

DIGITALES ARCHIV

ZBW – Leibniz-Informationszentrum Wirtschaft
ZBW – Leibniz Information Centre for Economics

Gintis, Herbert (Ed.)

Book

Epistemic game theory and modal logic

Reference: (2017). Epistemic game theory and modal logic. Basel, Switzerland : MDPI AG.

This Version is available at:

<http://hdl.handle.net/11159/781>

Kontakt/Contact

ZBW – Leibniz-Informationszentrum Wirtschaft/Leibniz Information Centre for Economics
Düsternbrooker Weg 120
24105 Kiel (Germany)
E-Mail: [rights\[at\]zbw.eu](mailto:rights[at]zbw.eu)
<https://www.zbw.eu/econis-archiv/>

Standard-Nutzungsbedingungen:

Dieses Dokument darf zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden. Sie dürfen dieses Dokument nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, aufführen, vertreiben oder anderweitig nutzen. Sofern für das Dokument eine Open-Content-Lizenz verwendet wurde, so gelten abweichend von diesen Nutzungsbedingungen die in der Lizenz gewährten Nutzungsrechte.

<https://zbw.eu/econis-archiv/termsfuse>

Terms of use:

This document may be saved and copied for your personal and scholarly purposes. You are not to copy it for public or commercial purposes, to exhibit the document in public, to perform, distribute or otherwise use the document in public. If the document is made available under a Creative Commons Licence you may exercise further usage rights as specified in the licence.

Article

Backward Induction *versus* Forward Induction Reasoning

Andres Perea

Department of Quantitative Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands; E-mail: a.perea@maastrichtuniversity.nl; Tel.: +31-43-3883922; Fax: +31-43-3884874

Received: 9 June 2010 / Accepted: 30 June 2010 / Published: 2 July 2010

Abstract: In this paper we want to shed some light on what we mean by backward induction and forward induction reasoning in dynamic games. To that purpose, we take the concepts of common belief in future rationality (Perea [1]) and extensive form rationalizability (Pearce [2], Battigalli [3], Battigalli and Siniscalchi [4]) as possible representatives for backward induction and forward induction reasoning. We compare both concepts on a conceptual, epistemic and an algorithm level, thereby highlighting some of the crucial differences between backward and forward induction reasoning in dynamic games.

Keywords: epistemic game theory; backward induction; forward induction; algorithms

1. Introduction

The ideas of *backward induction* and *forward induction* play a prominent role in the literature on dynamic games. Often, terms like backward and forward induction reasoning, and backward and forward induction concepts, are used to describe a particular pattern of reasoning in such games. But what exactly do we mean by backward induction and forward induction?

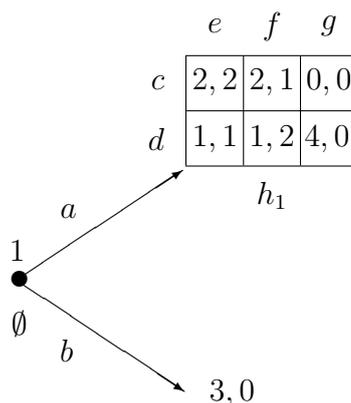
The literature offers no precise answer here. Only for the class of dynamic games with perfect information there is a clear definition of backward induction (based on Zermelo [5]), but otherwise there is no consensus on how to precisely formalize backward and forward induction. In fact, various authors have presented their own, personal interpretation of these two ideas. Despite this variety, there seems to be a common message in the authors' definitions of backward and forward induction, which can be described as follows:

Backward induction represents a pattern of reasoning in which a player, at every stage of the game, only reasons about the opponents' *future* behavior and beliefs, and not about choices that have been made in the past. So, he takes the opponents' past choices for granted, but does not draw any new conclusions from these.

In contrast, *forward induction* requires a player, at every stage, to think critically about the observed past choices by his opponents. He should always try to find a plausible reason for why his opponents have made precisely these choices in the past, and he should use this to possibly reconsider his belief about the opponents' future, present, and unobserved past choices.

In order to illustrate these two ideas, let us consider the game in Figure 1. So, at the beginning of the game, \emptyset , player 1 chooses between a and b . If he chooses b , the game ends and the players' utilities are 3 and 0. If he chooses a , the game moves to information set h_1 where players 1 and 2 simultaneously choose from $\{c, d\}$ and $\{e, f, g\}$ respectively.

Figure 1. Backwards induction and forward induction may lead to opposite choices.



If player 1 believes that player 2 chooses rationally, then he anticipates that player 2 will not choose g , and will therefore choose b at the beginning. But suppose now that the game actually reaches h_1 , and that player 2 must make a choice there. What should player 2 believe or do at h_1 ?

According to *backward induction*, player 2 should at h_1 only reason about player 1's behavior and beliefs at h_1 , and take his past choice a for granted. In that case, it is reasonable for player 2 to believe that player 1 still believes that player 2 will not choose g . Hence, player 2 will expect player 1 to choose c at h_1 , and player 2 would thus go for strategy e .

According to *forward induction*, player 2 should at h_1 try to make sense of player 1's past choice a . So, what reason could player 1 have to choose a , and not b , at the beginning? The only plausible reason could be that player 1 actually believed that player 2, with sufficiently high probability, would choose g . But if that were the case, then player 2 must conclude at h_1 that player 1 will choose d there, since that is his only chance to obtain more than 3. So, player 2 should then go for f , and not e .

So we see that backward and forward induction lead to opposite choices for player 2 in this game: backward induction leads to e , whereas forward induction leads to f . The crucial difference between the two ideas is that under backward induction, player 2 should at h_1 not draw any new conclusions from

player 1's past choice a . Under forward induction, player 2 should at h_1 use player 1's past choice a to form a new belief about player 1, namely that he believes that player 2 will (with sufficiently high probability) go for g .

Up to this stage we have only given some very broad, and rather unprecise, descriptions of backward and forward induction. Sometimes such descriptions are enough to analyze a game, like the one in Figure 1, but for other games it may be necessary to have a precise description of these two ideas, or to have them incorporated formally in some concept.

The literature actually offers a broad spectrum of formal concepts—most of these being *equilibrium* concepts—that incorporate the ideas of backward and forward induction. Historically, the equilibrium concept of *sequential equilibrium* (Kreps and Wilson [6]) is regarded as an important backward induction concept. Its main condition, *sequential rationality*, states that a player, at every stage, should believe that his opponents will make optimal choices at *present* and *future* information sets, given their beliefs there. No strategic conditions are imposed on beliefs about the opponents' *past* behavior. As such, sequential equilibrium only requires a player to reason about his opponents' future behavior, and may therefore be seen as a backward induction concept.

A problem with this concept, however, is that it incorporates an *equilibrium condition* which is hard to justify if the game is only played once, especially when there is no communication between the players before the game starts (see Bernheim [7] for a similar critique to Nash equilibrium). The equilibrium condition entails that a player believes that his opponents are correct about his own beliefs, and that he believes that two different players share the same belief about an opponent's future behavior. Aumann and Brandenburger [8], Asheim [9] and Perea [10] discuss similar conditions that lead to Nash equilibrium. Another drawback of the sequential equilibrium concept—and this is partially due to the equilibrium condition—is that the backward induction reasoning is somewhat *hidden* in the definition of sequential equilibrium, and not explicitly formulated as such.

In contrast, a backward induction concept that does not impose such equilibrium conditions, and which is very explicit about the backward induction reasoning being used, is *common belief in future rationality* (Perea [1]). It is a belief-based concept which states that a player should always believe that his opponents will choose rationally now and in the future, that a player always believes that every opponent always believes that his opponents will choose rationally now and in the future, and so on. No other conditions are imposed. The concept is closely related to sequential rationalizability (Dekel, Fudenberg and Levine [11,12] and Asheim and Perea [13]) and to backwards rationalizability (Penta [14]). (See Perea [1] for more details on this). Moreover, sequential equilibrium constitutes a refinement of common belief in future rationality, the main difference being the equilibrium condition that sequential equilibrium additionally imposes. In a sense, common belief in future rationality can be regarded as a backward induction concept that is similar to, but more basic and transparent than, sequential equilibrium. It can therefore serve as a basic representative of the idea of backward induction in general dynamic games, and we will use it as such in this paper.

Let us now turn to forward induction reasoning. In the 1980's and 1990's, forward induction has traditionally be modeled by equilibrium refinements. Some of these have been formulated as refinements of sequential equilibrium, by restricting the players' beliefs about the opponents' *past* behavior as well. Examples are justifiable sequential equilibrium (McLennan [15]), forward induction equilibrium

(Cho [16]) and stable sets of beliefs (Hillas [17]) for general dynamic games, and the intuitive criterion (Cho and Kreps [18]) and its various refinements for signaling games. By doing so, these authors have actually been incorporating a forward induction argument *inside* a backward induction concept, namely sequential equilibrium. So in a sense we are combining backward induction and forward induction in one and the same concept, and the result is a concept which does not purely represent the idea of backward induction nor forward induction.

As a consequence, these forward induction refinements of sequential equilibrium may fail to select intuitive forward induction strategies in certain games. Reconsider, for instance, the game in Figure 1. We have seen above that a natural forward induction argument uniquely selects strategy f for player 2. However, any forward induction refinement of sequential equilibrium necessarily selects only e for player 2. The reason is that e is the only sequential equilibrium strategy for player 2, so refining the sequential equilibrium concept will not be of any help here.

There are other forward induction equilibrium concepts that are *not* refinements of sequential equilibrium. Examples are stable sets of equilibria (Kohlberg and Mertens [19]), explicable equilibrium (Reny [20]) and Govindan and Wilson's [21] definition of forward induction—the latter two concepts being refinements of *weak* sequential equilibrium (Reny [20]) rather than sequential equilibrium.

As before, a problem with these equilibrium refinements is that it incorporates an equilibrium assumption which is problematic from an epistemic viewpoint. Moreover, the forward induction reasoning in these concepts is often not as transparent as it could be, partially due to this equilibrium assumption. In addition, the example in Figure 1 shows that in order to define a “pure” forward induction concept, we must step outside sequential equilibrium, and in fact step outside any backward induction concept, and simply build a new concept “from scratch”.

This is exactly what Pearce [2] did when he presented his *extensive form rationalizability* concept. The main idea is that a player, at each of his information sets, asks whether this information set could have been reached by rational¹ strategy choices by the opponents. If so, then he must believe that his opponents indeed *do* play rational strategies. In that case, he also asks whether this same information set could also have been reached by opponents who do not only choose rationally themselves, but who also believe that the other players choose rationally as well. If so, then he *must* believe that his opponents believe that the other players choose rationally as well. Iterating this argument finally leads to extensive form rationalizability.

This concept has many appealing properties. First, it is purely based on some very intuitive forward induction arguments, and not incorporated into some existing backward induction concept. In that sense, it is a very pure forward induction concept. Also, it has received a very appealing epistemic foundation in the literature (Battigalli and Siniscalchi [4]), and there is nowadays an easy elimination procedure supporting it (Shimoji and Watson [22]). So, the concept is attractive on an intuitive, an epistemic, and a practical level. That is why we will use this concept in this paper as a possible, appealing representative of the idea of forward induction.

The main objective of this paper is to compare the concept of common belief in future rationality—as a representative of backward induction reasoning—with the concept of extensive form

¹Here, by a *rational* strategy we mean a strategy that is optimal, at every information set, for *some* probabilistic belief about the opponents' strategy choices.

rationalizability—as a representative of forward induction reasoning. By doing so, we hope this paper will contribute towards better understanding the differences and similarities between backward induction and forward induction reasoning in dynamic games.

The outline of this paper is as follows. In Section 2 we formally present the class of dynamic games we consider, and develop an epistemic model for such games in order to formally represent the players' belief hierarchies. In Section 3 we define the concept of common belief in future rationality and present an elimination procedure, *backward dominance*, that supports it. Section 4 presents the concept of extensive form rationalizability, discusses Battigalli and Siniscalchi's [4] epistemic foundation, and presents Shimoji and Watson's [22] *iterated conditional dominance* procedure that supports it. In Section 5 we explicitly compare the two concepts with each other on a conceptual, epistemic and an algorithmic level.

2. Model

In this section we formally present the class of dynamic games we consider, and explain how to build an epistemic model for such dynamic games.

2.1. Dynamic Games

As we expect the reader to be familiar with the model of a dynamic game (or, extensive form game), we only list the relevant ingredients and introduce some pieces of notation. By I we denote the set of players, by X the set of non-terminal histories (or nodes) and by Z the set of terminal histories. By \emptyset we denote the beginning (or root) of the game. For every player i , we denote by H_i the collection of information sets for that player. Every information set $h \in H_i$ consists of a set of non-terminal histories. At every information set $h \in H_i$, we denote by $C_i(h)$ the set of choices (or actions) for player i at h . We assume that all sets mentioned above are *finite*, and hence we restrict to finite dynamic games in this paper. Finally, for every terminal history z and player i , we denote by $u_i(z)$ the utility for player i at z . As usual, we assume that there is *perfect recall*, meaning that a player never forgets what he previously did, and what he previously knew about the opponents' past choices.

We explicitly allow for *simultaneous moves* in the dynamic game. That is, we allow for non-terminal histories at which several players make a choice. Formally, this means that for some non-terminal histories x there may be different players i and j , and information sets $h_i \in H_i$ and $h_j \in H_j$, such that $x \in h_i$ and $x \in h_j$. In this case, we say that the information sets h_i and h_j are *simultaneous*. Explicitly allowing for simultaneous moves is important in this paper, especially for describing the concept of *common belief in future rationality*. We will come back to the issue of simultaneous moves in Section 3, when we formally introduce common belief in future rationality.

Say that an information set h follows some other information set h' if there are histories $x \in h$ and $y \in h'$ such that y is on the unique path from the root to x . Finally, we say that information set h *weakly follows* h' if either h follows h' , or h and h' are simultaneous.

2.2. Strategies

A strategy for player i is a complete choice plan, prescribing a choice at each of his information sets that can possibly be reached by this choice plan. Formally, for every $h, h' \in H_i$ such that h precedes h' , let $c_i(h, h')$ be the choice at h for player i that leads to h' . Note that $c_i(h, h')$ is unique by perfect recall. Consider a subset $\hat{H}_i \subseteq H_i$, not necessarily containing all information sets for player i , and a function s_i that assigns to every $h \in \hat{H}_i$ some choice $s_i(h) \in C_i(h)$. We say that s_i possibly reaches an information set h if at every $h' \in \hat{H}_i$ preceding h we have that $s_i(h') = c_i(h', h)$. By $H_i(s_i)$ we denote the collection of player i information sets that s_i possibly reaches. A strategy for player i is a function s_i , assigning to every $h \in \hat{H}_i \subseteq H_i$ some choice $s_i(h) \in C_i(h)$, such that $\hat{H}_i = H_i(s_i)$.

Notice that this definition slightly differs from the standard definition of a strategy in the literature. Usually, a strategy for player i is defined as a mapping that assigns to every information set $h \in H_i$ some available choice—also to those information sets h that cannot be reached by s_i . The definition of a strategy we use corresponds to what Rubinstein [23] calls a *plan of action*. One can also interpret it as the equivalence class of strategies (in the classical sense) that are outcome-equivalent. Hence, taking for every player the set of strategies as we use it corresponds to considering the pure strategy reduced normal form. However, for the concepts and results in this paper it does not make any difference which notion of strategy we use.

For a given information set h , denote by $S_i(h)$ the set of strategies for player i that possibly reach h . By $S_{-i}(h)$ we denote the strategy profiles for i 's opponents that possibly reach h , that is, $s_{-i} \in S_{-i}(h)$ if there is some $s_i \in S_i(h)$ such that (s_i, s_{-i}) reaches some history in h . By $S(h)$ we denote the set of strategy profiles $(s_i)_{i \in I}$ that reach some history in h . By perfect recall we have that $S(h) = S_i(h) \times S_{-i}(h)$ for every player i and every information set $h \in H_i$.

2.3. Epistemic Model

We now wish to model the players' beliefs in the game. At every information set $h \in H_i$, player i holds a belief about (a) the opponents' strategy choices, (b) the beliefs that the opponents have, at their information sets, about the other players' strategy choices, (c) the beliefs that the opponents have, at their information sets, about the beliefs their opponents have, at their information sets, about the other players' strategy choices, and so on. A possible way to represent such conditional belief hierarchies is as follows.

(Epistemic model) Consider a dynamic game Γ . An epistemic model for Γ is a tuple $M = (T_i, b_i)_{i \in I}$ where

- (a) T_i is a compact topological space, representing the set of types for player i ,
- (b) b_i is a function that assigns to every type $t_i \in T_i$, and every information set $h \in H_i$, a probability distribution $b_i(t_i, h) \in \Delta(S_{-i}(h) \times T_{-i})$.

Recall that $S_{-i}(h)$ represents the set of opponents' strategy combinations that possibly reach h . By $T_{-i} := \prod_{j \neq i} T_j$ we denote the set of opponents' type combinations. For a topological space X , we denote by $\Delta(X)$ the set of probability distributions on X with respect to the Borel σ -algebra. So, if there are more than two players in the game, we allow the players to hold *correlated* beliefs about the opponents' strategy choices (and types) at each of their information sets.

This model can be seen as an extension of the epistemic model in Ben-Porath [24], which was constructed specifically for games with perfect information. A similar model can also be found in Battigalli and Siniscalchi [25]. It is an *implicit* model, since we do not write down the belief hierarchies for the types explicitly, but these can rather be *derived* from the model. Namely, at every information set $h \in H_i$ a type t_i holds a conditional probabilistic belief $b_i(t_i, h)$ about the opponents' strategies and types. In particular, type t_i holds conditional beliefs about the opponents' strategies. As every opponent's type holds conditional beliefs about the other players' strategies, every type t_i holds at every $h \in H_i$ also a conditional belief about the opponents' conditional beliefs about the other players' strategy choices. And so on. So, in this way we can derive for every type the associated infinite conditional belief hierarchy. Since a type may hold different beliefs at different histories, a type may, during the game, revise his belief about the opponents' strategies, but also about the opponents' conditional beliefs.

To formally describe the concept of *common strong belief in rationality*, we need epistemic models that are *complete*, which means that every possible belief hierarchy must be present in the model.

(Complete epistemic model) An epistemic model $M = (T_i, b_i)_{i \in I}$ is complete if for every conditional belief vector $(b_i(h))_{h \in H_i}$ in $\prod_{h \in H_i} \Delta(S_{-i}(h) \times T_{-i})$ there is some type $t_i \in T_i$ with $b_i(t_i, h) = b_i(h)$ for every $h \in H_i$.

So, a complete epistemic model must necessarily be infinite. Battigalli and Siniscalchi [25] have shown that a complete epistemic model always exists for finite dynamic games, such as the ones we consider in this paper.

3. Common Belief in Future Rationality

We now present the concept of *common belief in future rationality* (Perea [1]), which is a typical backward induction concept. The idea is that a player always believes that (a) his opponents will choose rationally now and in the future, (b) his opponents always believe that their opponents will choose rationally now and in the future, and so on. After giving a precise epistemic formulation of this concept, we describe an algorithm, *backward dominance*, that supports it, and we illustrate this algorithm by means of an example.

3.1. Epistemic Formulation

We first define what it means for a strategy s_i to be optimal for a type t_i at a given information set h . Consider a type t_i , a strategy s_i and an information set $h \in H_i(s_i)$ that is possibly reached by s_i . By $u_i(s_i, t_i | h)$ we denote the expected utility from choosing s_i under the conditional belief that t_i holds at h about the opponents' strategy choices.

(Optimality at a given information set) Consider a type t_i , a strategy s_i and a history $h \in H_i(s_i)$. Strategy s_i is optimal for type t_i at h if $u_i(s_i, t_i | h) \geq u_i(s'_i, t_i | h)$ for all $s'_i \in S_i(h)$.

Remember that $S_i(h)$ is the set of player i strategies that possibly reach h . We can now define belief in the opponents' future rationality.

(Belief in the opponents' future rationality) Consider a type t_i , an information set $h \in H_i$, and an opponent $j \neq i$. Type t_i believes at h in j 's future rationality if $b_i(t_i, h)$ only assigns positive probability to j 's strategy-type pairs (s_j, t_j) where s_j is optimal for t_j at every $h' \in H_j(s_j)$ that weakly follows

h . Type t_i believes in the opponents' future rationality if at every $h \in H_i$, type t_i believes in every opponent's future rationality.

So, to be precise, a type that believes in the opponents' future rationality believes that every opponent chooses rationally now (if the opponent makes a choice at a simultaneous information set), and at every information set that follows. As such, the correct terminology would be "belief in the opponents' present and future rationality", but we stick to "belief in the opponents' future rationality" as to keep the name short.

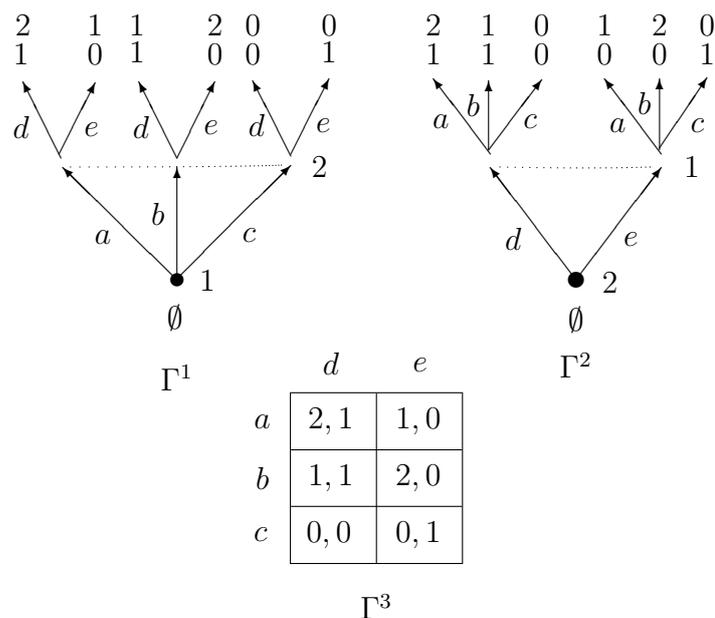
Next, we formalize the requirement that a player should not only believe in the opponents' future rationality, but should also always believe that every opponent believes in his opponents' future rationality, and so on.

(Common belief in future rationality) Type t_i expresses common belief in future rationality if (a) t_i believes in the opponents' future rationality, (b) t_i assigns, at every information set, only positive probability to opponents' types that believe in their opponents' future rationality, (c) t_i assigns, at every information set, only positive probability to opponents' types that, at every information set, only assign positive probability to opponents' types that believe in the opponents' future rationality, and so on.

Finally, we define those strategies that can rationally be chosen under common belief in future rationality. We say that a strategy s_i is rational for a type t_i if s_i is optimal for t_i at every $h \in H_i(s_i)$. In the literature, this is often called *sequential rationality*. We say that strategy s_i can rationally be chosen under common belief in future rationality if there is some epistemic model $M = (T_i, b_i)_{i \in I}$, and some type $t_i \in T_i$, such that t_i expresses common belief in future rationality, and s_i is rational for t_i .

For the concept of *common belief in future rationality*, it is crucial how we model the chronological order of moves in the game! Consider, for instance, the three games in Figure 2.

Figure 2. Chronological order of moves matters for "common belief in future rationality".



In game Γ^1 player 1 moves before player 2, in game Γ^2 player 2 moves before player 1, and in game Γ^3 both players choose simultaneously. In Γ^1 and Γ^2 , the second mover does not know which

choice has been made by the first mover. So, all three games represent a situation in which both players choose in complete ignorance of the opponent's choice. Since the utilities in the games are identical, one can argue that these three games are in some sense "equivalent". In fact, the three games above only differ by applying the transformation of *interchange of decision nodes*², as defined by Thompson [26]. However, for the concept of *common belief in future rationality* it crucially matters which of the three representations Γ^1 , Γ^2 or Γ^3 we choose.

In the game Γ^1 , common belief in future rationality does not restrict player 2's belief at all, as player 1 moves before him. So, player 2 can rationally choose d and e under common belief in future rationality here. On the other hand, player 1 may believe that player 2 chooses d or e under common belief in future rationality, and hence player 1 himself may rationally choose a or b under common belief in future rationality.

In the game Γ^2 , common belief in future rationality does not restrict player 1's beliefs as he moves after player 2. Hence, player 1 may rationally choose a or b under common belief in future rationality. Player 2 must therefore believe that player 1 will either choose a or b in the future, and hence player 2 can only rationally choose d under common belief in future rationality.

In the game Γ^3 , finally, player 1 can only rationally choose a , and player 2 can only rationally choose d under common belief in future rationality. Namely, if player 2 believes in player 1's (present and) future rationality, then player 2 believes that player 1 does not choose c , since player 1 moves at the same time as player 2. Therefore, player 2 can only rationally choose d under common belief in future rationality. If player 1 believes in player 2's (present and) future rationality, and believes that player 2 believes in player 1's (present and) future rationality, then player 1 believes that player 2 chooses d , and therefore player 1 can only rationally choose a under common belief in future rationality.

Hence, the precise order of moves is very important for the concept of common belief in future rationality! In particular, this concept is *not invariant* with respect to Thompson's [26] transformation of *interchange of decision nodes*. We will come back to this issue in Section 5.3.

3.2. Algorithm

Perea [1] presents an algorithm, *backward dominance*, that selects exactly those strategies than can rationally be chosen under common belief in future rationality. The algorithm proceeds by successively eliminating, at every information set, some strategies for the players. In the first round we eliminate, at every information set, those strategies for player i that are strictly dominated at a present or future information set for player i . In every further round k we eliminate, at every information set, those strategies for player i that are strictly dominated at a present or future information set h for player i , given the opponents' strategies that have survived until round k at that information set h . We continue until we cannot eliminate anything more.

In order to formally state the backward dominance procedure, we need the following definitions. Consider an information set $h \in H_i$ for player i , a subset $D_i \subseteq S_i(h)$ of strategies for player i that possibly reach h , and a subset $D_{-i} \subseteq S_{-i}(h)$ of strategy combinations for i 's opponents possibly reaching h . Then, (D_i, D_{-i}) is called a *decision problem* for player i at h , and we say that player i

²For a formal description of this transformation, the reader may consult Thompson [25], Elmes and Reny [26] or Perea [27].

is *active* at this decision problem. Note that several players may be active at the same decision problem, since several players may make a simultaneous move at the associated information set. Within a decision problem (D_i, D_{-i}) for player i , a strategy $s_i \in D_i$ is called *strictly dominated* if there is some randomized strategy $\mu_i \in \Delta(D_i)$ such that $u_i(\mu_i, s_{-i}) > u_i(s_i, s_{-i})$ for all $s_{-i} \in D_{-i}$. A decision problem at h is said to weakly follow an information set h' if h weakly follows h' . For a given information set $h \in H_i$, the *full* decision problem at h is the decision problem $(S_i(h), S_{-i}(h))$ where no strategies have been eliminated yet.

Backward dominance procedure

Initial step. For every information set h , let $\Gamma^0(h)$ be the full decision problem at h .

Inductive step. Let $k \geq 1$, and suppose that the decision problems $\Gamma^{k-1}(h)$ have been defined for every information set h . Then, at every information set h delete from the decision problem $\Gamma^{k-1}(h)$ those strategy combinations that involve a strategy s_i of some player i that is strictly dominated within some decision problem $\Gamma^{k-1}(h')$ for player i that weakly follows h . This yields the new decision problems $\Gamma^k(h)$. Continue this procedure until no further strategies can be eliminated in this way.

Say that a strategy s_i *survives the backward dominance procedure* if s_i is in $\Gamma^k(\emptyset)$ for every k . That is, s_i is never eliminated in the decision problem at the beginning of the game, \emptyset . Since we only have finitely many strategies in the game, and the decision problems can only become smaller at every step, this procedure must converge after finitely many steps. Perea [1] has shown that the algorithm always yields a nonempty set of strategies at every information set, and that the set of strategies surviving the algorithm is exactly the set of strategies that can rationally be chosen under common belief in future rationality. Combining these two insights then guarantees that common belief in future rationality is always possible—for every player we can always construct a type that expresses common belief in future rationality.

Note that the backward dominance procedure can be alternatively formulated as follows: If at a given decision problem $\Gamma^{k-1}(h)$ for player i strategy s_i is strictly dominated, then we eliminate s_i at $\Gamma^{k-1}(h)$ and at all decision problems $\Gamma^{k-1}(h')$ that come before h —that is, we eliminate s_i from h *backwards*. So, we can say that the backward dominance procedure, which characterizes the backward induction concept of common belief in future rationality, works by *backward elimination*. This, in turn, very clearly explains the word *backward* in backward induction concept.

3.3. Example

We will now illustrate the backward dominance procedure by means of an example. Consider again the game in Figure 1. At the beginning of the procedure we start with two decision problems, namely the full decision problem $\Gamma^0(\emptyset)$ at \emptyset where only player 1 is active, and the full decision problem $\Gamma^0(h_1)$ at h_1 where both players are active. These decision problems can be found in Table 1.

Table 1. The full decision problems in Figure 1.

Player 1 active				Players 1 and 2 active			
$\Gamma^0(\emptyset)$	<i>e</i>	<i>f</i>	<i>g</i>	$\Gamma^0(h_1)$	<i>e</i>	<i>f</i>	<i>g</i>
<i>(a, c)</i>	2, 2	2, 1	0, 0	<i>(a, c)</i>	2, 2	2, 1	0, 0
<i>(a, d)</i>	1, 1	1, 2	4, 0	<i>(a, d)</i>	1, 1	1, 2	4, 0
<i>b</i>	3, 0	3, 0	3, 0				

The backward dominance procedure does the following: In the first round, we eliminate from $\Gamma^0(\emptyset)$ strategy *(a, c)* as it is strictly dominated by *b* at player 1’s decision problem $\Gamma^0(\emptyset)$, and we eliminate from $\Gamma^0(\emptyset)$ and $\Gamma^0(h_1)$ strategy *g* as it is strictly dominated by *e* and *f* at player 2’s decision problem $\Gamma^0(h_1)$. In the second round, we eliminate from $\Gamma^1(\emptyset)$ strategy *(a, d)* as it strictly dominated by *b* at $\Gamma^1(\emptyset)$, and we eliminate strategy *(a, d)* also from $\Gamma^1(h_1)$ as it is strictly dominated by *(a, c)* at $\Gamma^1(h_1)$. In the third round, finally, we eliminate from $\Gamma^2(\emptyset)$ and $\Gamma^2(h_1)$ strategy *f*, as it is strictly dominated by *e* in $\Gamma^2(h_1)$. So, only strategies *b* and *e* remain at \emptyset . Hence, only strategies *b* and *e* can rationally be chosen under common belief in future rationality.

4. Extensive Form Rationalizability

We next turn to *extensive form rationalizability* (Pearce [2], Battigalli [3], Battigalli and Siniscalchi [4]), which is a typical forward induction concept. The idea is as follows: At every information set the corresponding player first asks whether this information set can be reached if his opponents would all choose rationally, that is, would choose optimally for some vectors of conditional beliefs. If so, then at that information set he must only assign positive probability to rational opponents’ strategies. In that case, he then asks: Can this information set also be reached by opponents’ strategies that are optimal if the opponents believe, whenever possible, that their opponents choose rationally? If so, then at that information set he must only assign positive probability to such opponents’ strategies. And so on. So, in a sense, at every information set the associated player looks for the highest degree of mutual belief in rationality that makes reaching this information set possible, and his beliefs at that information set should reflect this highest degree. We first provide a precise epistemic formulation of this concept, and then present an algorithm, *iterated conditional dominance*, that supports it. We finally illustrate the algorithm by means of an example.

4.1. Epistemic Formulation

The starting point in extensive form rationalizability is that a player, whenever possible, must believe that his opponents choose rationally. That is, if player *i* is at information set $h \in H_i$, he first asks whether *h* could have been reached by rational opponents’ strategies. If so, then at *h* he must assign positive probability *only* to rational opponents’ strategies. We say that player *i* *strongly believes in the opponents’ rationality* (Battigalli and Siniscalchi [4]).

In order to formalize this idea within an epistemic model, we must make sure that there are “enough” types in the model. To be more precise, if for a given information set $h \in H_i$ there is a rational strategy

s_j for opponent j that possibly reaches h , then there must be a type for player j inside the model for which s_j is optimal. Consider, for instance, the game in Figure 1, and suppose that our epistemic model would contain only one type for player 1, which believes that player 2 will choose e . Then, on the one hand, there is a rational strategy for player 1 that reaches h_1 , namely (a, d) . But the epistemic model does not contain a type for player 1 for which (a, d) is optimal. So, in this case, we must make sure that the epistemic model contains at least one type for player 1 for which (a, d) is optimal.

To guarantee this, it is enough to have a *complete* epistemic model. Namely, a complete model contains all possible belief hierarchies, so every potentially optimal strategy can be rationalized by at least one type in this model.

To formally define strong belief in the opponents' rationality, we need the following piece of notation: For every player i , information set $h \in H_i$, and subset of opponents' types $\tilde{T}_{-i} \subseteq T_{-i}$, let

$$(S_{-i}(h) \times \tilde{T}_{-i})^{rat} := \{(s_j, t_j)_{j \neq i} \in S_{-i}(h) \times \tilde{T}_{-i} : s_j \text{ rational for } t_j \text{ for all } j \neq i\}$$

Recall that strategy s_j is rational for type t_j if s_j is optimal for t_j at every information set in $H_j(s_j)$.

(Strong belief in the opponents' rationality) Consider a complete epistemic model $M = (T_i, b_i)_{i \in I}$. A type t_i strongly believes in the opponents' rationality if at every $h \in H_i$ with $(S_{-i}(h) \times T_{-i})^{rat} \neq \emptyset$, it holds that $b_i(t_i, h)((S_{-i}(h) \times T_{-i})^{rat}) = 1$.

That is, if for every opponent j there is strategy s_j leading to h and a type for which s_j is rational, then type t_i must at h only consider strategy-type pairs (s_j, t_j) where s_j is rational for type t_j . Let us define by T_i^1 the set of types $t_i \in T_i$ that strongly believe in the opponents' rationality.

Now, suppose that player i is at $h \in H_i$, and that $(S_{-i}(h) \times T_{-i})^{rat} \neq \emptyset$. So, h could have been reached by rational opponents' strategies. The next question that extensive form rationalizability asks is: Could h have been reached by opponents' strategies s_j that are optimal for opponents' types t_j that strongly believe in their opponents' rationality? If so, then player i at h should only consider such pairs (s_j, t_j) . In other words, if $(S_{-i}(h) \times T_{-i}^1)^{rat} \neq \emptyset$, then player i must at h only consider opponents' strategy-type combinations in $(S_{-i}(h) \times T_{-i}^1)^{rat}$. By iterating this argument, we arrive at the following recursive definition of *common strong belief in rationality* (Battigalli and Siniscalchi [4]).

(Common strong belief in rationality) Consider a complete epistemic model $M = (T_i, b_i)_{i \in I}$. Let $T_i^0 := T_i$ for every player i . For every $k \geq 1$ and every player i , let T_i^k contain those types $t_i \in T_i^{k-1}$ such that at every $h \in H_i$ with $(S_{-i}(h) \times T_{-i}^{k-1})^{rat} \neq \emptyset$, it holds that $b_i(t_i, h)((S_{-i}(h) \times T_{-i}^{k-1})^{rat}) = 1$. A type t_i expresses common strong belief in rationality if $t_i \in T_i^k$ for all k .

We say that strategy s_i can *rationally be chosen under common strong belief in rationality* if there is some complete epistemic model $M = (T_i, b_i)_{i \in I}$, and some type $t_i \in T_i$ expressing common strong belief in rationality, such that s_i is rational for t_i .

4.2. Algorithm

The concept of *extensive form rationalizability* has originally been proposed in Pearce [2] by means of an iterated reduction procedure. Later, Battigalli [3] has simplified this procedure and has shown that it delivers the same output as Pearce's procedure. Both procedures refine at every round the sets of strategies and conditional beliefs of the players. Battigalli and Siniscalchi [4] have shown that common strong belief in rationality selects exactly the extensive form rationalizable strategies for every player.

In this section we will consider yet another procedure leading to extensive form rationalizability, namely the *iterated conditional dominance* procedure developed by Shimoji and Watson [22]. The reason is that this procedure is closer to the backward dominance algorithm for common belief in future rationality, and therefore easier to compare.

The iterated conditional dominance procedure, like the backward dominance procedure, iteratedly removes strategies from decision problems. However, the criteria for removing a strategy in a particular decision problem are different. Remember that in the backward dominance procedure we remove a strategy for player i in the decision problem at h whenever it is strictly dominated in some decision problem for player i that *weakly follows* h . In the iterated conditional dominance procedure we remove a strategy for player i at the decision problem at h if there is some decision problem for player i , *not necessarily weakly following* h , at which it is strictly dominated. So, in the iterated conditional dominance procedure we would remove strategy s_i at h also if it is strictly dominated at some decision problem for player i that comes *before* h . Formally, the procedure can be formulated as follows.

Iterated conditional dominance procedure

Initial step. For every information set h , let $\Gamma^0(h)$ be the full decision problem at h .

Inductive step. Let $k \geq 1$, and suppose that the decision problems $\Gamma^{k-1}(h)$ have been defined for every information set h . Then, at every information set h delete from the decision problem $\Gamma^{k-1}(h)$ those strategy combinations that involve a strategy s_i for some player i that is strictly dominated within some decision problem $\Gamma^{k-1}(h')$ for player i , **not necessarily weakly following** h . This yields the new decision problems $\Gamma^k(h)$. Continue this procedure until no further strategies can be eliminated in this way.

A strategy s_i is said to survive this procedure if $s_i \in \Gamma^k(\emptyset)$ for all k . Shimoji and Watson [22] have shown that this procedure delivers exactly the set of extensive form rationalizable strategies. Hence, by Battigalli and Siniscalchi's [4] result, the iterated conditional dominance procedure selects exactly those strategies that can rationally be chosen under common strong belief in rationality.

Note that in the iterated conditional dominance procedure, it is possible that at a given decision problem $\Gamma^{k-1}(h)$ *all* strategies of a player i will be eliminated in step k —something that can never happen in the backward dominance procedure. Consider, namely, some information set $h \in H_i$, and some information set h' following h . Then, it is possible that within the decision problem $\Gamma^{k-1}(h)$, all strategies for player i in $\Gamma^{k-1}(h')$ are strictly dominated. In that case, we would eliminate in $\Gamma^{k-1}(h')$ all remaining strategies for player i ! Whenever this occurs, it is understood that at every further step nothing can be eliminated from the decision problem at h' anymore.

The iterated conditional dominance procedure thus has the following property: If at a given decision problem $\Gamma^{k-1}(h)$ for player i the strategy s_i is strictly dominated, then we eliminate s_i at $\Gamma^{k-1}(h)$, and at all decision problems $\Gamma^{k-1}(h')$ that come *before* and *after* it—that is, we eliminate s_i from h *backwards* and *forward*. So this algorithm, which characterizes the forward induction concept of extensive form rationalizability, proceeds by *backward and forward elimination*. From this perspective, the name “forward induction” is actually a bit misleading, as it would suggest the concept to work only in a forward fashion. This is not true: Extensive form rationalizability, when considered algorithmically, works both backwards and forward.

4.3. Example

To illustrate the iterated conditional dominance procedure, consider again the game from Figure 1, with its full decision problems $\Gamma^0(\emptyset)$ and $\Gamma^0(h_1)$ as depicted in Table 1. The iterated conditional dominance procedure works as follows here:

In the first round, we eliminate strategy (a, c) from $\Gamma^0(\emptyset)$ and $\Gamma^0(h_1)$ as it is strictly dominated by b at player 1's decision problem $\Gamma^0(\emptyset)$, and we eliminate from $\Gamma^0(\emptyset)$ and $\Gamma^0(h_1)$ strategy g as it is strictly dominated by e and f at player 2's decision problem $\Gamma^0(h_1)$. In the second round, we eliminate (a, d) from $\Gamma^1(\emptyset)$ and $\Gamma^1(h_1)$ as it is strictly dominated by b at $\Gamma^1(\emptyset)$, and we eliminate e from $\Gamma^1(\emptyset)$ and $\Gamma^1(h_1)$ as it is strictly dominated by f in $\Gamma^1(h_1)$. This only leaves strategies b and f at \emptyset , and hence only strategies b and f can rationally be chosen under extensive form rationalizability.

Recall that the backward dominance procedure uniquely selected strategies b and e , and hence only e can rationally be chosen by player 2 under common belief in future rationality. So, we see that both procedures (and hence their associated epistemic concepts) lead to unique but different strategy choices for player 2 in this example.

The crucial difference between both concepts lies in how player 2 at h_1 explains the surprise that player 1 has not chosen b . Under common belief in future rationality, player 2 believes at h_1 that player 1 has simply made a mistake, but he still believes that player 1 will choose rationally at h_1 , and he still believes that player 1 believes that he will not choose g at h_1 . So, player 2 believes at h_1 that player 1 will choose (a, c) , and therefore player 2 will choose e at h_1 . Under extensive form rationalizability, player 2 believes at h_1 that player 1's decision not to choose b was a rational decision, but this is only possible if player 2 believes at h_1 that player 1 believes that player 2 will irrationally choose g at h_1 (with sufficiently high probability). In that case, player 2 will believe at h_1 that player 1 will go for (a, d) , and therefore player 2 will choose f at h_1 .

5. Comparison Between the Concepts

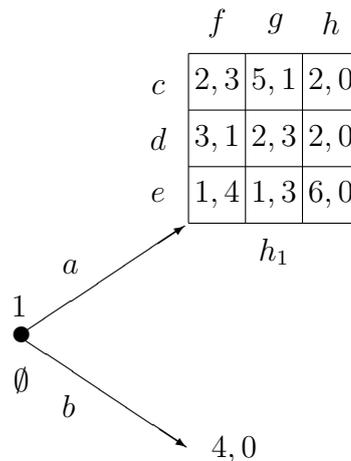
In this section we will compare the concepts of *common belief in future rationality* and *extensive form rationalizability (common strong belief in rationality)* on a conceptual, epistemic, algorithmic and behavioral level.

5.1. Conceptual Comparison: The Role of Rationality Orderings

An appealing way to look at extensive form rationalizability is by means of *rationality orderings* over strategies (Battigalli [29]). The idea is that for every player i we have an ordered partition (S_i^1, \dots, S_i^K) of his strategy set, where S_i^1 represents the set of "most rational" strategies, S_i^K the set of "least rational" strategies, and every strategy in S_i^k is deemed "more rational" than every strategy in S_i^{k+1} . At every information set $h \in H_i$, player i then looks for the most rational opponents' strategies that reach h , and assigns positive probability only to such opponents' strategies. Important is that these rationality orderings are *global*, that is, the players always use the *same* rationality orderings over the opponents' strategies to form their conditional beliefs.

To illustrate this, consider the game in Figure 3.

Figure 3. Rationality orderings.



Among player 1’s strategies, (a, c) , (a, e) and b are optimal for some belief, whereas (a, d) is not optimal for any belief. So, (a, d) may be considered the “least rational” strategy for player 1. So, for player 1 we have the “tentative” rationality ordering

$$RO_1^1 = (\{(a, c), (a, e), b\}, \{(a, d)\})$$

For player 2, strategies f and g are optimal at h_1 for some belief, whereas h is not optimal at h_1 for any belief. Hence, for player 2 we have the tentative rationality ordering

$$RO_2^1 = (\{f, g\}, \{h\})$$

Now, if player 1 believes that player 2 does not choose his least rational strategy h , then only (a, c) and b can be optimal, and not (a, e) . So, we obtain a refined rationality ordering

$$RO_1^2 = (\{(a, c), b\}, \{(a, e)\}, \{(a, d)\})$$

for player 1. Similarly, if player 2 believes at h_1 that player 1 will not choose his least rational strategy (a, d) , then only f can be optimal, and not g . So, for player 2 we obtain the refined rationality ordering

$$RO_2^2 = (\{f\}, \{g\}, \{h\})$$

But then, if player 1 believes that player 2 will choose his most rational strategy f , then player 1 will choose b . So, the final rationality orderings for the players are

$$RO_1 = (\{b\}, \{(a, c)\}, \{(a, e)\}, \{(a, d)\}) \text{ and } RO_2 = (\{f\}, \{g\}, \{h\})$$

Hence, if player 2 finds himself at h_1 , he believes that player 1 has chosen the most rational strategy that reaches h_1 , which is (a, c) . Player 2 must therefore choose f , and player 1, anticipating on this, should choose b . This is exactly what extensive form rationalizability does for this game.

Important is that both players agree on these specific rationality orderings RO_1 and RO_2 , and that player 2 uses the rationality ordering RO_1 throughout the game, in particular at h_1 to form his conditional belief there.

In contrast, the concept of common belief in future rationality *cannot* be described in terms of rationality orderings, or at least not by *global* rationality orderings that are used throughout the whole game. Consider again the game in Figure 3. Common belief in future rationality reasons as follows here: Player 1, at \emptyset and at h_1 , must believe that player 2 chooses rationally at h_1 , and hence must believe that player 2 will not choose h . Player 2, at h_1 , believes that player 1, at h_1 , believes in 2's rationality at h_1 . Hence, player 2 believes at h_1 that player 1 believes that 2 will not choose h . Also, player 2 believes at h_1 that player 1 chooses rationally at h_1 , and hence player 2 believes at h_1 that player 1 will not choose (a, e) . No further conclusions can be drawn at h_1 . Hence, player 2 can choose f or g at h_1 . But then, player 1 can choose (a, c) or b . So, common belief in future rationality selects strategies (a, c) and b for player 1, and strategies f and g for player 2.

Can this reasoning be represented by global rationality orderings on the players' strategies? The answer is "no". Suppose, namely, that such rationality orderings RO_1 and RO_2 would exist. As common belief in future rationality selects only the strategies (a, c) and b for player 1, strategies (a, c) and b should both be most rational under RO_1 . But then, if player 2 is at h_1 , he should conclude that player 1 has chosen (a, c) , as it is the most rational strategy under RO_1 that reaches h_1 . Consequently, player 2 should choose f , rendering (a, c) a suboptimal strategy for player 1. This, however, would contradict RO_1 , where (a, c) is considered to be a most rational strategy for player 1. Hence, there is no global rationality ordering on strategies that supports common belief in future rationality.

Rather, under common belief in future rationality, player 2 *changes* his rationality ordering over 1's strategies as the game proceeds. At the beginning, player 2 deems (a, c) and b more rational than (a, d) and (a, e) . However, at h_1 player 2 deems (a, c) and (a, d) "equally" rational, as under common belief in future rationality player 2 may at h_1 believe that player 1 chooses (a, c) or (a, d) .

5.2. Epistemic Comparison

By definition, a type t_i for player i is said to express *common belief in future rationality* if it always believes in the opponents' future rationality, always only assigns positive probability to opponents' types that always believe in their opponents' future rationality, always only assigns positive probability to opponents' types that always only assign positive probability to other players' types that always believe in their opponents' future rationality, and so on. As a consequence, type t_i always only assigns positive probability to opponents' types that express common belief in future rationality too. Hence, every type that expresses common belief in future rationality believes at every stage of the game that each of his opponents expresses common belief in future rationality as well. We may thus say that the concept of common belief in future rationality is "closed under belief".

Formally, "closed under belief" can be defined in the following way. Consider some epistemic model with sets of types T_i for every player i . Let $\hat{T}_i \subseteq T_i$ be a subset of types for every player i . Then, the combination $(\hat{T}_i)_{i \in I}$ of subsets of types is said to be *closed under belief* if for every player i , every type $t_i \in \hat{T}_i$, and every information set $h \in H_i$, the conditional belief $b_i(\hat{t}_i, h)$ only assigns positive probability to opponents' types t_j that are in \hat{T}_j . So if we take an epistemic model with sets of types $(T_i)_{i \in I}$, and define $T_i^{cbfr} \subseteq T_i$ to be the subset of types for player i that express common belief in future rationality, then the combination $(T_i^{cbfr})_{i \in I}$ of type subsets expressing common belief in future rationality is *closed under belief* in the sense above.

The same cannot be said about *common strong belief in rationality*—the epistemic foundation for *extensive form rationalizability*. Consider for instance the game from Figure 1. If player 2's type t_2 strongly believes in player 1's rationality, then t_2 must at h_1 believe that player 1 has rationally chosen (a, d) . More precisely, type t_2 must at h_1 only assign positive probability to strategy-type pairs $((a, d), t_1)$ for player 1 where (a, d) is optimal for type t_1 . This, however, can only be the case if t_1 assigns positive probability to player 2's irrational strategy g . But then, t_1 does certainly not strongly believe in 2's rationality. So we see that a type t_2 for player 2 who strongly believes in 1's rationality, must at h_1 necessarily assign positive probability to a type t_1 for player 1 who does *not* strongly believe in 2's rationality. In particular, a type t_2 for player 2 that expresses common strong belief in rationality, must at h_1 attach positive probability to a player 1 type t_1 that does *not* express common strong belief in rationality. Hence, the concept of common strong belief in rationality is certainly not closed under belief.

The latter is not surprising, as it follows from the very character of common strong belief in rationality. As we have seen in the previous subsection, this concept orders the players' strategies, and also types, from "most rational" to "least rational". Most rational are the types that express common strong belief in rationality, and least rational are the types that do not even strongly believe in the opponents' rationality, and there may be some subclasses in between. The idea of common strong belief in rationality is that at every information set, the corresponding player searches for the "most rational" opponents' types that could have been responsible for reaching this information set, and these opponents' types do not necessarily express common strong belief in rationality. In fact, typically these opponents' types are "less rational" than the "most rational types around", which are the ones expressing common strong belief in rationality. So, it is no surprise that the concept of common strong belief in rationality is not "closed under belief".

The fact that common belief in future rationality is closed under belief, and common strong belief in rationality is not, is also reflected in the *completeness* of the epistemic model needed for these two concepts. Note that for defining common strong belief in rationality we required a *complete* epistemic model (meaning that every possible belief hierarchy is present in the model), whereas for common belief in future rationality we did not. In fact, for the concept of common belief in future rationality a model with finitely many types is enough (see Perea [1]). So why do we have this difference?

The reason is that under common strong belief in rationality, player i must ask at every information set $h \in H_i$ whether h *could* have been reached by opponents' strategies that are optimal for *some* beliefs. To answer this question, he must consider *all* possible opponents' types—also those that do not express common strong belief in rationality—and see whether some of these types would support strategy choices that could lead to h . So, a complete epistemic model is needed here.

In contrast, under common belief in future rationality it is sufficient for player i to only consider opponents' types that express common belief in future rationality as well. In other words, there is no need for player i to step outside the sets of types expressing common belief in future rationality, and that is why we do not need a complete epistemic model here.

5.3. Algorithmic Comparison

In Sections 3 and 4 we have described two elimination procedures, backward dominance and iterated conditional dominance, that respectively lead to common belief in future rationality and extensive form rationalizability. A natural question is: Does the order and speed in which we eliminate strategies from the decision problems matter for the eventual result of these procedures? The answer is that it does not matter for the backward dominance procedure (see Perea [1]), whereas the order and speed of elimination is crucial for the iterated conditional dominance procedure.

Consider, namely, the game from Figure 1. Suppose that, in the first round of the iterated conditional dominance procedure, we would only eliminate strategy g , but not (a, c) , from $\Gamma^0(\emptyset)$ and $\Gamma^0(h_1)$. Then, in $\Gamma^1(h_1)$, strategy (a, d) is strictly dominated by (a, c) . Suppose that in round 2 we would only eliminate strategy (a, d) from $\Gamma^1(\emptyset)$ and $\Gamma^1(h_1)$. Suppose that in round 3 we would eliminate strategy f for player 2 at \emptyset and h_1 , as it has become strictly dominated at h_1 . Suppose that, finally, we would eliminate (a, c) at \emptyset and h_1 . So, for player 2 only strategy e would survive the procedure in this case. Recall, however, that if we eliminate “all that we can” at every round of the iterated conditional dominance procedure, then only strategy f would survive for player 2. Hence, the order and speed of elimination affects the outcome of the iterated conditional dominance procedure—it is absolutely crucial to eliminate at every round, and at every information set, all strategies we can.

Now, why is the order and speed of elimination relevant for the iterated conditional dominance procedure, but not for the backward dominance procedure? The reason has to do with rationality orderings as we have discussed them above. We have seen that extensive form rationalizability can be described by global rationality orderings on the players’ strategies, ranking them from “most rational” to “least rational”. At every information set, the corresponding player identifies the most rational opponents’ strategies that reach this information set, and assigns positive probability only to these strategies. For this construction to work, it is essential that all players agree on these specific rationality orderings.

The iterated conditional dominance procedure in fact *generates* these rationality orderings: All strategies that do not survive the first round are deemed “least rational”. All strategies that survive the first round, but not the second round, are deemed “second least rational” and so on. Finally, the strategies that survive all rounds are deemed “most rational”. So, this procedure does not only deliver the extensive form rationalizable strategies, it also delivers the rationality orderings on players’ strategies that *support* extensive form rationalizability. Since it is crucial that players agree on these rationality orderings, players must agree on the strategies that are eliminated at every round of the procedure: If at a certain round not all strategies that *could* be eliminated are in fact eliminated, then this would lead to a “coarser” rationality ordering in that round, which in turn could lead to completely different rationality orderings in the end.

This problem cannot occur for backward dominance: If at a certain information set a strategy that *could* have been eliminated is not in fact eliminated, then it will be eliminated at some later round anyhow. So, even if players would disagree on the order and speed of elimination, it would not affect their final strategy choices in the game.

We have seen in Section 3 that the concept of common belief in future rationality is sensitive to the transformation of *interchange of decision nodes*, as defined by Thompson [26]. This can be seen very clearly from its associated algorithm—the backward dominance procedure. In this algorithm, namely, whenever a strategy s_i is strictly dominated at a decision problem $\Gamma^{k-1}(h)$ for player i , we eliminate it from $\Gamma^{k-1}(h)$ and all decision problems $\Gamma^{k-1}(h')$ that come *before* h , but we do not eliminate it at decision problems $\Gamma^{k-1}(h')$ that come *after* h . By applying the transformation of *interchange of decision nodes*, we may interchange the chronological order of two information sets h and h' . So, before the transformation h comes before h' , whereas after the transformation h' comes before h . Hence it is possible that *before* the transformation we eliminate s_i at h because it is strictly dominated at h' , whereas *after* the transformation we can no longer do so because h now comes after h' . That is, the transformation of *interchange of decision nodes* may have important consequences for the output of the backward dominance procedure—and hence for the concept of common belief in future rationality.

It can be verified that the transformation of *interchange of decision nodes* has *no* consequences for the concept of extensive form rationalizability. This is most easily seen by studying the associated algorithm—the iterated conditional dominance procedure. In that procedure, whenever a strategy s_i is strictly dominated at a decision problem $\Gamma^{k-1}(h)$ for player i , we eliminate it at *all* decision problems $\Gamma^{k-1}(h')$ in the game. Hence, the precise chronological order of the information sets does not play a role, only the structure of the various decision problems $\Gamma^{k-1}(h)$ in the game. Since the transformation of *interchange of decision nodes* does not change this structure of the decision problems $\Gamma^{k-1}(h)$ in the game, it easily follows that the iterated conditional dominance procedure—and hence the concept of extensive form rationalizability—is invariant under the transformation of *interchange of decision nodes*.

5.4. Behavioral Comparison

In this section we ask whether there is any logical relationship between the strategy choices selected by common belief in future rationality, and those selected by extensive form rationalizability. The answer is “no”. This can already be concluded from the example in Figure 1. There, we have seen that common belief in future rationality uniquely selects strategy e for player 2, whereas extensive form rationalizability uniquely selects strategy f for this player. Hence, in this example both concepts yield completely opposite strategy selections for player 2.

There are other examples where common belief in future rationality is more restrictive than extensive form rationalizability, and yet other examples where it is exactly the other way around. Consider, for instance, the game from Figure 3. There, common belief in future rationality yields strategy choices (a, c) and b for player 1, and strategy choices f and g for player 2. Extensive form rationalizability, on the other hand, uniquely selects strategies b and f . So here extensive form rationalizability is more restrictive.

Now, replace in the example in Figure 1 the outcome $3, 0$ by $5, 0$. Then, common belief in future rationality would select strategy b for player 1, and strategy e for player 2, whereas extensive form rationalizability would select strategy b for player 1, and strategies e and f for player 2. So here common belief in future rationality is more restrictive.

Note, however, that in each of these examples the set of *outcomes* induced by extensive form rationalizability is always a subset of the set of outcomes induced by common belief in future rationality.

My conjecture is that this is true in general, but I could not find a formal proof yet. (In fact we know that it is true for all generic games with perfect information – see the paragraph below). So I leave this here as an interesting open problem.

An important special class of dynamic games, both for theory and applications, is the class of games with *perfect information*. These are games where at every stage only one player moves, and he always observes the choices made by others so far. Such a game is called *generic* if, for every player i and every information set $h \in H_i$, two different choices at h always lead to outcomes with different utilities for i .

In Perea [1] it has been shown that for the class of generic dynamic games with perfect information, the concept of common belief in future rationality leads to the unique backward induction strategies for the players. Battigalli [3] has proved that extensive form rationalizability, and hence common strong belief in rationality, leads to the backward induction *outcome*, but *not* necessarily to the backward induction *strategies*, in such games. As a consequence, for generic games with perfect information both concepts lead to the same outcome, namely the backward induction outcome, but not necessarily to the same strategies for the players.

References

1. Perea, A. Belief in the opponents' future rationality. Maastricht University: Maastricht, The Netherlands, 2010 (manuscript); Available online: <http://www.personeel.unimaas.nl/a.perea/> (accessed on 1 June 2010).
2. Pearce, D.G. Rationalizable strategic behavior and the problem of perfection. *Econometrica* **1984**, *52*, 1029–1050.
3. Battigalli, P. On rationalizability in extensive games. *J. Econ. Theor.* **1997**, *74*, 40–61.
4. Battigalli, P.; Siniscalchi, M. Strong belief and forward induction reasoning. *J. Econ. Theor.* **2002**, *106*, 356–391.
5. Zermelo, E. Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. In *Proceedings Fifth International Congress of Mathematicians*, Cambridge, UK, 22–28 August 1912; pp. 501–504.
6. Kreps, D.M.; Wilson, R. Sequential equilibria. *Econometrica* **1982**, *50*, 863–94.
7. Bernheim, B.D. Rationalizable strategic behavior. *Econometrica* **1984**, *52*, 1007–1028.
8. Aumann, R.; Brandenburger, A. Epistemic conditions for Nash equilibrium. *Econometrica* **1995**, *63*, 1161–1180.
9. Asheim, G.B. The Consistent preferences approach to deductive reasoning in games. In *Theory and Decision Library*; Springer: Dordrecht, The Netherlands, 2006.
10. Perea, A. A one-person doxastic characterization of Nash strategies. *Synthese* **2007**, *158*, 251–271 (*Knowledge, Rationality & Action*, 341–361).
11. Dekel, E.; Fudenberg, D.; Levine, D.K. Payoff information and self-confirming equilibrium. *J. Econ. Theor.* **1999**, *89*, 165–185.
12. Dekel, E.; Fudenberg, D.; Levine, D.K. Subjective uncertainty over behavior strategies: A correction. *J. Econ. Theor.* **2002**, *104*, 473–478.

13. Asheim, G.B.; Perea, A. Sequential and quasi-perfect rationalizability in extensive games. *Game. Econ. Behav.* **2005**, *53*, 15–42.
14. Penta, A. Robust dynamic mechanism design. University of Pennsylvania: Philadelphia, PA, USA, 2009 (manuscript).
15. McLennan, A. Justifiable beliefs in sequential equilibria. *Econometrica* **1985**, *53*, 889–904.
16. Cho, I.-K. A refinement of sequential equilibrium. *Econometrica* **1987**, *55*, 1367–1389.
17. Hillas, J. Sequential equilibria and stable sets of beliefs. *J. Econ. Theor.* **1994**, *64*, 78–102.
18. Cho, I.-K.; Kreps, D.M. Signaling games and stable equilibria. *Quart. J. Econ.* **1987**, *102*, 179–221.
19. Kohlberg, E.; Mertens, J.-F. On the strategic stability of equilibria. *Econometrica* **1986**, *54*, 1003–1038.
20. Reny, P.J. Backward induction, normal form perfection and explicable equilibrium. *Econometrica* **1992**, *60*, 627–649.
21. Govindan, S.; Wilson, R. On forward induction. *Econometrica* **2009**, *77*, 1–28.
22. Shimoji, M.; Watson, J. Conditional dominance, rationalizability, and game forms. *J. Econ. Theor.* **1998**, *83*, 161–195.
23. Rubinstein, A. Comments on the interpretation of game theory. *Econometrica* **1991**, *59*, 909–924.
24. Ben-Porath, E. Rationality, Nash equilibrium and backwards induction in perfect-information games. *Rev. Econ. Stud.* **1997**, *64*, 23–46.
25. Battigalli, P.; Siniscalchi, M. Hierarchies of conditional beliefs and interactive epistemology in dynamic games. *J. Econ. Theor.* **1999**, *88*, 188–230.
26. Thompson, F.B. Equivalence of games in extensive form. The Rand Corporation: Santa Monica, CA, USA, 1952; Discussion Paper RM 759.
27. Elmes, S.; Reny, P.J. On the strategic equivalence of extensive form games. *J. Econ. Theor.* **1994**, *62*, 1–23.
28. Perea, A. *Rationality in Extensive Form Games*; In *Theory and Decision Library, Series C*; Kluwer Academic Publishers: London, UK, 2001.
29. Battigalli, P. Strategic rationality orderings and the best rationalization principle. *Game. Econ. Behav.* **1996**, *13*, 178–200.

Article

The Role of Monotonicity in the Epistemic Analysis of Strategic Games

Krzysztof R. Apt^{1,*} and Jonathan A. Zvesper²

¹ Centrum for Mathematics and Computer Science (CWI), University of Amsterdam, Science Park 123, 1098 XG Amsterdam, The Netherlands

² Oxford University, Computing Laboratory, Wolfson Building, Parks Road, Oxford OX1 3QD, UK

* Author to whom correspondence should be addressed; E-Mail: apt@cw.nl; Tel: 31-20-592-4057; Fax: 31-20-592-4199.

Received: 23 July 2010; in revised form: 18 September 2010 / Accepted: 19 September 2010 /

Published: 8 October 2010

Abstract: It is well-known that in finite strategic games true common belief (or common knowledge) of rationality implies that the players will choose only strategies that survive the iterated elimination of strictly dominated strategies. We establish a general theorem that deals with monotonic rationality notions and arbitrary strategic games and allows to strengthen the above result to arbitrary games, other rationality notions, and transfinite iterations of the elimination process. We also clarify what conclusions one can draw for the customary dominance notions that are not monotonic. The main tool is Tarski's Fixpoint Theorem.

Keywords: true common beliefs; arbitrary games; monotonicity; Tarski's Fixpoint Theorem

1. Introduction

1.1. Contributions

In this paper we provide an epistemic analysis of arbitrary strategic games based on possibility correspondences. We prove a general result that is concerned with monotonic program properties¹ used by the players to select optimal strategies.

¹The concept of a monotonic property is introduced in Section 2.

More specifically, given a belief model for the initial strategic game, denote by $\mathbf{RAT}(\phi)$ the property that each player i uses a property ϕ_i to select his strategy ('each player i is ϕ_i -rational'). We establish in Section 3 the following general result:

Assume that each property ϕ_i is monotonic. The set of joint strategies that the players choose in the states in which $\mathbf{RAT}(\phi)$ is a true common belief is included in the set of joint strategies that remain after the iterated elimination of the strategies that for player i are not ϕ_i -optimal.

In general, transfinite iterations of the strategy elimination are possible. For some belief models the inclusion can be reversed.

This general result covers the usual notion of rationalizability in finite games and a 'global' version of the iterated elimination of strictly dominated strategies used in [1] and studied for arbitrary games in [2]. It does not hold for the 'global' version of the iterated elimination of weakly dominated strategies. For the customary, 'local' version of the iterated elimination of strictly dominated strategies we justify in Section 4 the statement

true common belief (or common knowledge) of rationality implies that the players will choose only strategies that survive the iterated elimination of strictly dominated strategies

for arbitrary games and transfinite iterations of the elimination process. Rationality refers here to the concept studied in [3]. We also show that the above general result yields a simple proof of the well-known version of the above result for finite games and strict dominance by a mixed strategy.

The customary, local, version of strict dominance is non-monotonic, so the use of monotonic properties has allowed us to provide epistemic foundations for a non-monotonic property. However, weak dominance, another non-monotonic property, remains beyond the reach of this approach. In fact, we show that in the above statement we cannot replace strict dominance by weak dominance. A mathematical reason is that its global version is also non-monotonic, in contrast to strict dominance, the global version of which is monotonic. To provide epistemic foundations of weak dominance the only currently known approaches are [4] based on lexicographic probability systems and [5] based on a version of the 'all I know' modality.

1.2. Connections

The relevance of monotonicity in the context of epistemic analysis of finite strategic games has already been pointed out in [6]. The distinction between local and global properties is from [7] and [8].

To show that for some belief models an equality holds between the set of joint strategies chosen in the states in which $\mathbf{RAT}(\phi)$ is true common belief and the set of joint strategies that remain after the iterated elimination of the strategies that for player i are not ϕ_i -rational requires use of transfinite ordinals. This complements the findings of [9] in which transfinite ordinals are used in a study of limited rationality, and [10], where a two-player game is constructed for which the ω_0 (the first infinite ordinal) and $\omega_0 + 1$ iterations of the rationalizability operator of [3] differ.

In turn, [11] show that arbitrary ordinals are necessary in the epistemic analysis of arbitrary strategic games based on partition spaces. Further, as shown in [2], the global version of the iterated elimination of strictly dominated strategies, when used for arbitrary games, also requires transfinite iterations of the underlying operator.

Finally, [12] invokes Tarski's Fixpoint Theorem, in the context of what the author calls "general systems", and uses this to prove that the set of rationalizable strategies in a finite non-cooperative game is the largest fixpoint of a certain operator. That operator coincides with the global version of the elimination of never-best-responses.

Some of the results presented here were initially reported in a different presentation, in [13].

2. Preliminaries

2.1. Strategic Games

Given n players ($n > 1$) by a **strategic game** (in short, a **game**) we mean a sequence $(S_1, \dots, S_n, p_1, \dots, p_n)$, where for all $i \in \{1, \dots, n\}$

- S_i is the non-empty set of **strategies** available to player i ,
- p_i is the **payoff function** for the player i , so $p_i : S_1 \times \dots \times S_n \rightarrow \mathcal{R}$, where \mathcal{R} is the set of real numbers.

We denote the strategies of player i by s_i , possibly with some superscripts. We call the elements of $S_1 \times \dots \times S_n$ **joint strategies**. Given a joint strategy s we denote the i th element of s by s_i , write sometimes s as (s_i, s_{-i}) , and use the following standard notation:

- $s_{-i} := (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$,
- $S_{-i} := S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_n$.

Given a finite non-empty set A we denote by ΔA the set of probability distributions over A and call any element of ΔS_i a **mixed strategy** of player i .

In the remainder of the paper we assume an initial strategic game

$$H := (H_1, \dots, H_n, p_1, \dots, p_n)$$

A **restriction** of H is a sequence (G_1, \dots, G_n) such that $G_i \subseteq H_i$ for all $i \in \{1, \dots, n\}$. Some of G_i s can be the empty set. We identify the restriction (H_1, \dots, H_n) with H . We shall focus on the complete lattice that consists of the set of all restrictions of the game H ordered by the componentwise set inclusion:

$$(G_1, \dots, G_n) \subseteq (G'_1, \dots, G'_n) \text{ iff } G_i \subseteq G'_i \text{ for all } i \in \{1, \dots, n\}$$

So in this lattice H is the largest element in this lattice.

2.2. Possibility Correspondences

In this and the next subsection we essentially follow the survey of [14]. Fix a non-empty set Ω of **states**. By an **event** we mean a subset of Ω .

A **possibility correspondence** is a mapping from Ω to the powerset $\mathcal{P}(\Omega)$ of Ω . We consider three properties of a possibility correspondence P :

- (i) for all ω , $P(\omega) \neq \emptyset$,
- (ii) for all ω and ω' , $\omega' \in P(\omega)$ implies $P(\omega') = P(\omega)$,
- (iii) for all ω , $\omega \in P(\omega)$.

If the possibility correspondence satisfies properties (i) and (ii), we call it a **belief correspondence** and if it satisfies properties (i)–(iii), we call it a **knowledge correspondence**.² Note that each knowledge correspondence P yields a partition $\{P(\omega) \mid \omega \in \Omega\}$ of Ω .

Assume now that each player i has at its disposal a possibility correspondence P_i . Fix an event E . We define

$$\Box E := \Box^1 E := \{\omega \in \Omega \mid \forall i \in \{1, \dots, n\} P_i(\omega) \subseteq E\}$$

by induction on $k \geq 1$

$$\Box^{k+1} E := \Box \Box^k E$$

and finally

$$\Box^* E := \bigcap_{k=1}^{\infty} \Box^k E$$

If all P_i s are belief correspondences, we usually write B instead of \Box and if all P_i s are knowledge correspondences, we usually write K instead of \Box . When $\omega \in B^* E$, we say that the event E is **common belief in the state** ω and when $\omega \in K^* E$, we say that the event E is **common knowledge in the state** ω .

An event F is called **evident** if $F \subseteq \Box F$. That is, F is evident if for all $\omega \in F$ we have $P_i(\omega) \subseteq F$ for all $i \in \{1, \dots, n\}$. In what follows we shall use the following alternative characterizations of common belief and common knowledge based on evident events:

$$\omega \in \Box^* E \text{ iff for some evident event } F \text{ we have } \omega \in F \subseteq \Box E \tag{1}$$

where $\Box = B$ or $\Box = K$ (see [16], respectively Proposition 4 on page 180 and Proposition on page 174), and

$$\omega \in K^* E \text{ iff for some evident event } F \text{ we have } \omega \in F \subseteq E \tag{2}$$

([17], page 1237).

2.3. Models for Games

We now relate these considerations to strategic games. Given a restriction $G := (G_1, \dots, G_n)$ of the initial game H , by a **model** for G we mean a set of states Ω together with a sequence of functions $\bar{s}_i : \Omega \rightarrow G_i$, where $i \in \{1, \dots, n\}$. We denote it by $(\Omega, \bar{s}_1, \dots, \bar{s}_n)$.

In what follows, given a function f and a subset E of its domain, we denote by $f(E)$ the range of f on E and by $f|E$ the restriction of f to E .

By the **standard model** \mathcal{M} for G we mean the model in which

²Note that the notion of a belief has two meanings in the literature on epistemic analysis of strategic games, so also in this paper. From the context it is always clear which notion is used. In the modal logic terminology a belief correspondence is a frame for the modal logic KD45 and a knowledge correspondence is a frame for the modal logic S5, see, e.g. [15].

- $\Omega := G_1 \times \dots \times G_n$
- $\bar{s}_i(\omega) := \omega_i$, where $\omega = (\omega_1, \dots, \omega_n)$

So the states of the standard model for G are exactly the joint strategies in G , and each \bar{s}_i is a projection function. Since the initial game H is given, we know the payoff functions p_1, \dots, p_n . So in the context of H the standard model is an alternative way of representing a restriction of H .

Given a (not necessarily standard) model $\mathcal{M} := (\Omega, \bar{s}_1, \dots, \bar{s}_n)$ for a restriction G and a sequence of events $\bar{E} = (E_1, \dots, E_n)$ in \mathcal{M} (i.e., of subsets of Ω) we define

$$G_{\bar{E}} := (\bar{s}_1(E_1), \dots, \bar{s}_n(E_n))$$

and call it the **restriction of G to \bar{E}** . When each E_i equals E we write G_E instead of $G_{\bar{E}}$.

Finally, we extend the notion of a model for a restriction G to a **belief model** for G by assuming that each player i has a belief correspondence P_i on Ω . If each P_i is a knowledge correspondence, we refer then to a **knowledge model**. We write each belief model as

$$(\Omega, \bar{s}_1, \dots, \bar{s}_n, P_1, \dots, P_n)$$

2.4. Operators

Consider a fixed complete lattice (D, \subseteq) with the largest element \top . In what follows we use ordinals and denote them by α, β, γ . Given a, possibly transfinite, sequence $(G_\alpha)_{\alpha < \gamma}$ of elements of D we denote their join and meet respectively by $\bigcup_{\alpha < \gamma} G_\alpha$ and $\bigcap_{\alpha < \gamma} G_\alpha$.

Let T be an operator on (D, \subseteq) , i.e., $T : D \rightarrow D$.

- We call T **monotonic** if for all $G, G', G \subseteq G'$ implies $T(G) \subseteq T(G')$, and **contracting** if for all G , $T(G) \subseteq G$.
- We say that an element G is a **fixpoint** of T if $G = T(G)$ and a **post-fixpoint** of T if $G \subseteq T(G)$.
- We define by transfinite induction a sequence of elements T^α of D , where α is an ordinal, as follows:
 - $T^0 := \top$,
 - $T^{\alpha+1} := T(T^\alpha)$,
 - for all limit ordinals β , $T^\beta := \bigcap_{\alpha < \beta} T^\alpha$.
- We call the least α such that $T^{\alpha+1} = T^\alpha$ the **closure ordinal** of T and denote it by α_T . We call then T^{α_T} the **outcome of** (iterating) T and write it alternatively as T^∞ .

So an outcome is a fixpoint reached by a transfinite iteration that starts with the largest element. In general, the outcome of an operator does not need to exist but we have the following classic result due to [18].³

³We use here its ‘dual’ version in which the iterations start at the largest and not at the least element of a complete lattice.

Tarski's Fixpoint Theorem Every monotonic operator T on (D, \subseteq) has an outcome, i.e., T^∞ is well-defined. Moreover,

$$T^\infty = \nu T = \cup\{G \mid G \subseteq T(G)\}$$

where νT is the largest fixpoint of T .

In contrast, a contracting operator does not need to have a largest fixpoint. But we have the following obvious observation.

Note 1. Every contracting operator T on (D, \subseteq) has an outcome, i.e., T^∞ is well-defined. \square

In Section 4 we shall need the following lemma, that modifies the corresponding lemma from [8] from finite to arbitrary complete lattices.

Lemma 1. Consider two operators T_1 and T_2 on (D, \subseteq) such that

- for all G , $T_1(G) \subseteq T_2(G)$,
- T_1 is monotonic,
- T_2 is contracting.

Then $T_1^\infty \subseteq T_2^\infty$.

Proof. We first prove by transfinite induction that for all α

$$T_1^\alpha \subseteq T_2^\alpha \tag{3}$$

By the definition of the iterations we only need to consider the induction step for a successor ordinal. So suppose the claim holds for some α . Then by the first two assumptions and the induction hypothesis we have the following string of inclusions and equalities:

$$T_1^{\alpha+1} = T_1(T_1^\alpha) \subseteq T_1(T_2^\alpha) \subseteq T_2(T_2^\alpha) = T_2^{\alpha+1}$$

This shows that for all α (3) holds. By Tarski's Fixpoint Theorem and Note 1 the outcomes of T_1 and T_2 exist, which implies the claim. \square

2.5. Iterated Elimination of Non-Rational Strategies

In this paper we are interested in analyzing situations in which each player pursues his own notion of rationality and this information is common knowledge or true common belief. As a special case we cover then the usually analyzed situation in which all players use the same notion of rationality.

Given player i in the initial strategic game $H := (H_1, \dots, H_n, p_1, \dots, p_n)$ we formalize his notion of rationality using an **optimality property** $\phi(s_i, G_i, G_{-i})$ that holds between a strategy $s_i \in H_i$, a set G_i of strategies of player i and a set G_{-i} of joint strategies of his opponents. Intuitively, $\phi(s_i, G_i, G_{-i})$ holds if s_i is an 'optimal' strategy for player i within the restriction $G := (G_i, G_{-i})$, assuming that he uses the property ϕ_i to select optimal strategies. In Section 4 we shall provide several natural examples of such properties.

We say that the property ϕ_i used by player i is **monotonic** if for all $G_{-i}, G'_{-i} \subseteq H_{-i}$ and $s_i \in H_i$

$$G_{-i} \subseteq G'_{-i} \text{ and } \phi(s_i, H_i, G_{-i}) \text{ imply } \phi(s_i, H_i, G'_{-i})$$

So monotonicity refers to the situation in which the set of strategies of player i is set to H_i and the set of joint strategies of player i 's opponents is increased.

Each sequence of properties $\phi := (\phi_1, \dots, \phi_n)$ determines an operator T_ϕ on the restrictions of H defined by

$$T_\phi(G) := G'$$

where $G := (G_1, \dots, G_n)$, $G' := (G'_1, \dots, G'_n)$, and for all $i \in \{1, \dots, n\}$

$$G'_i := \{s_i \in G_i \mid \phi_i(s_i, H_i, G_{-i})\}$$

Note that in defining the set of strategies G'_i we use in the second argument of ϕ_i the set H_i of player's i strategies in the *initial* game H and not in the *current* restriction G . This captures the idea that at every stage of the elimination process player i analyzes the status of each strategy in the context of his initial set of strategies.

Since T_ϕ is contracting, by Note 1 it has an outcome, *i.e.*, T_ϕ^∞ is well-defined. Moreover, if each ϕ_i is monotonic, then T_ϕ is monotonic and by Tarski's Fixpoint Theorem its largest fixpoint νT_ϕ exists and equals T_ϕ^∞ . Finally, G is a fixpoint of T_ϕ iff for all $i \in \{1, \dots, n\}$ and all $s_i \in G_i$, $\phi_i(s_i, H_i, G_{-i})$ holds.

Intuitively, $T_\phi(G)$ is the result of removing from G all strategies that are not ϕ_i -rational. So the outcome of T_ϕ is the result of the iterated elimination of strategies that for player i are not ϕ_i -rational.

3. Two Theorems

We now assume that each player i employs some property ϕ_i to select his strategies, and we analyze the situation in which this information is true common belief or common knowledge. To determine which strategies are then selected by the players we shall use the T_ϕ operator.

We begin by fixing a belief model $(\Omega, \bar{s}_1, \dots, \bar{s}_n, P_1, \dots, P_n)$ for the initial game H . Given an optimality property ϕ_i of player i we say that player i is **ϕ_i -rational in the state** ω if $\phi_i(\bar{s}_i(\omega), H_i, (G_{P_i(\omega)})_{-i})$ holds. Note that when player i believes (respectively, knows) that the state is in $P_i(\omega)$, the set $(G_{P_i(\omega)})_{-i}$ represents his belief (respectively, his knowledge) about other players' strategies. That is, $(H_i, (G_{P_i(\omega)})_{-i})$ is the restriction he believes (respectively, knows) to be relevant to his choice.

Hence $\phi_i(\bar{s}_i(\omega), H_i, (G_{P_i(\omega)})_{-i})$ captures the idea that if player i uses ϕ_i to select his strategy in the game he considers relevant, then in the state ω he indeed acts 'rationally'.

To reason about common knowledge and true common belief we introduce the event

$$\mathbf{RAT}(\phi) := \{\omega \in \Omega \mid \text{each player } i \text{ is } \phi_i\text{-rational in } \omega\}$$

and consider the following two events constructed out of it: $K^*\mathbf{RAT}(\phi)$ and $\mathbf{RAT}(\phi) \cap B^*\mathbf{RAT}(\phi)$. We then focus on the corresponding restrictions $G_{K^*\mathbf{RAT}(\phi)}$ and $G_{\mathbf{RAT}(\phi) \cap B^*\mathbf{RAT}(\phi)}$.

So strategy s_i is an element of the i th component of $G_{K^*\mathbf{RAT}(\phi)}$ if $s_i = \bar{s}_i(\omega)$ for some $\omega \in K^*\mathbf{RAT}(\phi)$. That is, s_i is a strategy that player i chooses in a state in which it is common knowledge that each player j is ϕ_j -rational, and similarly for $G_{\mathbf{RAT}(\phi) \cap B^*\mathbf{RAT}(\phi)}$.

The following result then relates for arbitrary strategic games the restrictions $G_{\mathbf{RAT}(\phi) \cap B^* \mathbf{RAT}(\phi)}$ and $G_{K^* \mathbf{RAT}(\phi)}$ to the outcome of the iteration of the operator T_ϕ .

Theorem 1.

(i) Suppose that each property ϕ_i is monotonic. Then for all belief models for H

$$G_{\mathbf{RAT}(\phi) \cap B^* \mathbf{RAT}(\phi)} \subseteq T_\phi^\infty$$

(ii) Suppose that each property ϕ_i is monotonic. Then for all knowledge models for H

$$G_{K^* \mathbf{RAT}(\phi)} \subseteq T_\phi^\infty$$

(iii) For some standard knowledge model for H

$$T_\phi^\infty \subseteq G_{K^* \mathbf{RAT}(\phi)}$$

So part (i) (respectively, (ii)) states that true common belief (respectively, common knowledge) of ϕ_i -rationality of each player i implies that the players will choose only strategies that survive the iterated elimination of non- ϕ -rational strategies.

Proof.

(i) Fix a belief model $(\Omega, \overline{s}_1, \dots, \overline{s}_n, P_1, \dots, P_n)$ for H . Take a strategy s_i that is an element of the i th component of $G_{\mathbf{RAT}(\phi) \cap B^* \mathbf{RAT}(\phi)}$. Thus we have $s_i = \overline{s}_i(\omega)$ for some state ω such that $\omega \in \mathbf{RAT}(\phi)$ and $\omega \in B^* \mathbf{RAT}(\phi)$. The latter implies by (1) that for some evident event F

$$\omega \in F \subseteq \{\omega' \in \Omega \mid \forall i \in \{1, \dots, n\} P_i(\omega') \subseteq \mathbf{RAT}(\phi)\} \tag{4}$$

Take now an arbitrary $\omega' \in F \cap \mathbf{RAT}(\phi)$ and $i \in \{1, \dots, n\}$. Since $\omega' \in \mathbf{RAT}(\phi)$, it holds that player i is ϕ_i -rational in ω' , i.e., $\phi_i(\overline{s}_i(\omega'), H_i, (G_{P_i(\omega')})_{-i})$ holds. But F is evident, so $P_i(\omega') \subseteq F$. Moreover by (4) $P_i(\omega') \subseteq \mathbf{RAT}(\phi)$, so $P_i(\omega') \subseteq F \cap \mathbf{RAT}(\phi)$. Hence $(G_{P_i(\omega')})_{-i} \subseteq (G_{F \cap \mathbf{RAT}(\phi)})_{-i}$ and by the monotonicity of ϕ_i we conclude that $\phi_i(\overline{s}_i(\omega'), H_i, (G_{F \cap \mathbf{RAT}(\phi)})_{-i})$ holds.

By the definition of T_ϕ this means that $G_{F \cap \mathbf{RAT}(\phi)} \subseteq T_\phi(G_{F \cap \mathbf{RAT}(\phi)})$, i.e. $G_{F \cap \mathbf{RAT}(\phi)}$ is a post-fixpoint of T_ϕ . But T_ϕ is monotonic since each property ϕ_i is. Hence by Tarski's Fixpoint Theorem $G_{F \cap \mathbf{RAT}(\phi)} \subseteq T_\phi^\infty$. But $s_i = \overline{s}_i(\omega)$ and $\omega \in F \cap \mathbf{RAT}(\phi)$, so we conclude by the above inclusion that s_i is an element of the i th component of T_ϕ^∞ . This proves the claim.

(ii) By the definition of common knowledge for all events E we have $K^*E \subseteq E$. Hence for all ϕ we have $K^* \mathbf{RAT}(\phi) \subseteq \mathbf{RAT}(\phi) \cap K^* \mathbf{RAT}(\phi)$ and consequently $G_{K^* \mathbf{RAT}(\phi)} \subseteq G_{\mathbf{RAT}(\phi) \cap K^* \mathbf{RAT}(\phi)}$.

So part (ii) follows from part (i).

(iii) Suppose $T_\phi^\infty = (G_1, \dots, G_n)$. Consider the event $F := G_1 \times \dots \times G_n$ in the standard model for H . Then $G_F = T_\phi^\infty$. Define each possibility correspondence P_i by

$$P_i(\omega) := \begin{cases} F & \text{if } \omega \in F \\ \Omega \setminus F & \text{otherwise} \end{cases}$$

Each P_i is a knowledge correspondence (also when $F = \emptyset$ or $F = \Omega$) and clearly F is an evident event.

Take now an arbitrary $i \in \{1, \dots, n\}$ and an arbitrary state $\omega \in F$. Since T_ϕ^∞ is a fixpoint of T_ϕ and $\bar{s}_i(\omega) \in G_i$ we have $\phi_i(\bar{s}_i(\omega), H_i, (T_\phi^\infty)_{-i})$, so by the definition of P_i we have $\phi_i(\bar{s}_i(\omega), H_i, (G_{P_i(\omega)})_{-i})$. This shows that each player i is ϕ_i -rational in each state $\omega \in F$, i.e., $F \subseteq \mathbf{RAT}(\phi)$.

Since F is evident, we conclude by (2) that in each state $\omega \in F$ it is common knowledge that each player i is ϕ_i -rational, i.e., $F \subseteq K^*\mathbf{RAT}(\phi)$. Consequently

$$T_\phi^\infty = G_F \subseteq G_{K^*\mathbf{RAT}(\phi)}$$

□

Items (i) and (ii) show that when each property ϕ_i is monotonic, for all belief models of H it holds that the joint strategies that the players choose in the states in which each player i is ϕ_i -rational and it is common belief that each player i is ϕ_i -rational (or in which it is common knowledge that each player i is ϕ_i -rational) are included in those that remain after the iterated elimination of the strategies that are not ϕ_i -rational.

Note that monotonicity of the ϕ_i properties was not needed to establish item (iii).

By instantiating the ϕ_i 's with specific properties we get instances of the above result that refer to specific definitions of rationality. This will allow us to relate the above result to the ones established in the literature. Before we do this we establish a result that identifies a large class of properties ϕ_i for which Theorem 1 does not apply.

Theorem 2. *Suppose that a joint strategy $s \notin T_\phi^\infty$ exists such that*

$$\phi_i(s_i, H_i, (\{s_j\}_{j \neq i}))$$

holds all $i \in \{1, \dots, n\}$. Then for some knowledge model for H the inclusion

$$G_{K^*\mathbf{RAT}(\phi)} \subseteq T_\phi^\infty$$

does not hold.

Proof. We extend the standard model for H by the knowledge correspondences P_1, \dots, P_n where for all $i \in \{1, \dots, n\}$, $P_i(\omega) = \{\omega\}$. Then for all ω and all $i \in \{1, \dots, n\}$

$$G_{P_i(\omega)} = (\{\bar{s}_1(\omega)\}, \dots, \{\bar{s}_n(\omega)\})$$

Let $\omega' := s$. Then for all $i \in \{1, \dots, n\}$, $G_{P_i(\omega')} = (\{s_1\}, \dots, \{s_n\})$, so by the assumption each player i is ϕ_i -rational in ω' , i.e., $\omega' \in \mathbf{RAT}(\phi)$. By the definition of P_i s the event $\{\omega'\}$ is evident and $\omega' \in K^*\mathbf{RAT}(\phi)$. So by (1) $\omega' \in K^*\mathbf{RAT}(\phi)$. Consequently $s = (\bar{s}_1(\omega'), \dots, \bar{s}_n(\omega')) \in G_{K^*\mathbf{RAT}(\phi)}$.

This yields the desired conclusion by the choice of s . □

4. Applications

We now analyze to what customary game-theoretic properties the above two results apply. By a *belief* of player i about the strategies his opponents play given the set G_{-i} of their joint strategies we mean one of the following possibilities:

- a joint strategy of the opponents of player i , i.e., $s_{-i} \in G_{-i}$, called a **point belief**,
- or, in the case the game is finite, a joint mixed strategy of the opponents of player i (i.e., $(m_1, \dots, m_{i-1}, m_{i+1}, \dots, m_n)$, where $m_j \in \Delta G_j$ for all $j \neq i$), called an **independent belief**,
- or, in the case the game is finite, an element of ΔG_{-i} , called a **correlated belief**.

In the second and third case the payoff function p_i can be lifted in the standard way to an **expected payoff** function $p_i : H_i \times \mathcal{B}_i(G_{-i}) \rightarrow \mathcal{R}$, where $\mathcal{B}_i(G_{-i})$ is the corresponding set of beliefs of player i held given G_{-i} .

We use below the following abbreviations, where $s_i, s'_i \in H_i$ and G_{-i} is a set of the strategies of the opponents of player i :

- (**strict dominance**) $s'_i \succ_{G_{-i}} s_i$ for $\forall s_{-i} \in G_{-i} p_i(s'_i, s_{-i}) > p_i(s_i, s_{-i})$
- (**weak dominance**) $s'_i \succ_{G_{-i}}^w s_i$ for $\forall s_{-i} \in G_{-i} p_i(s'_i, s_{-i}) \geq p_i(s_i, s_{-i}) \wedge \exists s_{-i} \in G_{-i} p_i(s'_i, s_{-i}) > p_i(s_i, s_{-i})$

In the case of finite games the relations $\succ_{G_{-i}}$ and $\succ_{G_{-i}}^w$ between a mixed strategy and a pure strategy are defined in the same way.

We now introduce natural examples of the optimality notion.

- $sd_i(s_i, G_i, G_{-i}) \equiv \neg \exists s'_i \in G_i s'_i \succ_{G_{-i}} s_i$
- (assuming H is finite) $msd_i(s_i, G_i, G_{-i}) \equiv \neg \exists m'_i \in \Delta G_i m'_i \succ_{G_{-i}} s_i$
- $wd_i(s_i, G_i, G_{-i}) \equiv \neg \exists s'_i \in G_i s'_i \succ_{G_{-i}}^w s_i$
- (assuming H is finite) $mwd_i(s_i, G_i, G_{-i}) \equiv \neg \exists m'_i \in \Delta G_i m'_i \succ_{G_{-i}}^w s_i$
- $br_i(s_i, G_i, G_{-i}) \equiv \exists \mu_i \in \mathcal{B}_i(G_{-i}) \forall s'_i \in G_i p_i(s_i, \mu_i) \geq p_i(s'_i, \mu_i)$

So sd_i and wd_i are the customary notions of strict and weak dominance and msd_i and mwd_i are their counterparts for the case of dominance by a mixed strategy. Note that the notion br_i of best response, comes in three ‘flavours’ depending on the choice of the set $\mathcal{B}_i(G_{-i})$ of beliefs.

Consider now the iterated elimination of strategies as defined in Subsection 2.5, so *with* the repeated reference by player i to the strategy set H_i . For the optimality notion sd_i such a version of iterated elimination was studied in [2], for mwd_i it was used in [4], while for br_i it corresponds to the rationalizability notion of [3].

In [10], [2] and [7] examples are provided showing that for the properties sd_i and br_i in general transfinite iterations (i.e., iterations beyond ω_0) of the corresponding operator are necessary to reach the outcome. So to establish for them part (iii) of Theorem 1 transfinite iterations of the T_ϕ operator are necessary.

The following lemma holds.

Lemma 2. *The properties sd_i , msd_i and br_i are monotonic.*

Proof. Straightforward. □

So Theorem 1 applies to the above three properties. In contrast, Theorem 1 does not apply to the remaining two properties wd_i and mwd_i , since, as indicated in [8], the corresponding operators T_{wd} and T_{mwd} are not monotonic, and hence the properties wd_i and mwd_i are not monotonic.

In fact, the desired inclusion does not hold and Theorem 2 applies to these two optimality properties. Indeed, consider the following game:

	<i>L</i>	<i>R</i>
<i>U</i>	1, 1	0, 1
<i>D</i>	1, 0	1, 1

Then the outcome of iterated elimination for both wd_i and mwd_i yields $G := (\{D\}, \{R\})$. Further, we have $wd_1(U, \{U, D\}, \{L\})$ and $wd_2(L, \{L, R\}, \{U\})$, and analogously for mwd_1 and mwd_2 .

So the joint strategy (U, L) satisfies the conditions of Theorem 2 for both wd_i and mwd_i . Note that this game also furnishes an example for non-monotonicity of wd_i since $wd_1(U, \{U, D\}, \{L, R\})$ does not hold.

This shows that the optimality notions wd_i and mwd_i cannot be justified in the used epistemic framework as ‘stand alone’ concepts of rationality.

5. Consequences of Common Knowledge of Rationality

In this section we show that common knowledge of rationality is sufficient to entail the customary iterated elimination of strictly dominated strategies. We also show that weak dominance is not amenable to such a treatment.

Given a sequence of properties $\phi := (\phi_1, \dots, \phi_n)$, we introduce an operator U_ϕ on the restrictions of H defined by

$$U_\phi(G) := G',$$

where $G := (G_1, \dots, G_n)$, $G' := (G'_1, \dots, G'_n)$, and for all $i \in \{1, \dots, n\}$

$$G'_i := \{s_i \in G_i \mid \phi_i(s_i, G_i, G_{-i})\}.$$

So when defining the set of strategies G'_i we use in the second argument of ϕ_i the set G_i of player’s i strategies in the *current* restriction G . That is, $U_\phi(G)$ determines the ‘locally’ ϕ -optimal strategies in G . In contrast, $T_\phi(G)$ determines the ‘globally’ ϕ -optimal strategies in G , in that each player i must consider all of his strategies s'_i that occur in his strategy set H_i in the *initial game* H .

So the ‘global’ form of optimality coincides with rationality, as introduced in Subsection 2.5, while the customary definition of iterated elimination of strictly (or weakly) dominated strategies refers to the iterations of the appropriate instantiation of the ‘local’ U_ϕ operator.

Note that the U_ϕ operator is non-monotonic for all non-trivial optimality notions ϕ_i such that $\phi_i(s_i, \{s_i\}, (\{s_j\}_{j \neq i}))$ for all joint strategies s , so in particular for br_i, sd_i, msd_i, wd_i and mwd_i . Indeed, given s let G_s denote the corresponding restriction in which each player i has a single strategy s_i .

Each restriction G_s is a fixpoint of U_ϕ . By non-triviality of ϕ_i s we have $U_\phi(H) \neq H$, so for each restriction G_s with s including an eliminated strategy the inclusion $U_\phi(G_s) \subseteq U_\phi(H)$ does not hold, even though $G_s \subseteq H$. In contrast, as we saw, by virtue of Lemma 2 the T_ϕ operator is monotonic for br_i, sd_i and msd_i .

First we establish the following consequence of Theorem 1. When each property ϕ_i equals br_i , we write here $\mathbf{RAT}(br)$ and similarly with U_{sd} .

Corollary 1.

(i) For all belief models

$$G_{\mathbf{RAT}(br) \cap B^* \mathbf{RAT}(br)} \subseteq U_{sd}^\infty$$

(ii) for all knowledge models

$$G_{K^* \mathbf{RAT}(br)} \subseteq U_{sd}^\infty$$

where in both situations we use in br_i the set of point beliefs.

Proof.

(i) By Lemma 2 and Theorem 1(i) $G_{\mathbf{RAT}(br) \cap B^* \mathbf{RAT}(br)} \subseteq T_{br}^\infty$ Each best response to a joint strategy of the opponents is not strictly dominated, so for all restrictions G

$$T_{br}(G) \subseteq T_{sd}(G)$$

Also, for all restrictions G , $T_{sd}(G) \subseteq U_{sd}(G)$. So by Lemma 1 $T_{br}^\infty \subseteq U_{sd}^\infty$, which concludes the proof.

(ii) By part (i) and the fact that $K^* \mathbf{RAT}(br) \subseteq \mathbf{RAT}(br)$. □

Part (ii) formalizes and justifies in the epistemic framework used here the often used statement:

common knowledge of rationality implies that the players will choose only strategies that survive the iterated elimination of strictly dominated strategies

for games with *arbitrary strategy sets* and *transfinite iterations* of the elimination process, and where best response means best response to a point belief.

In the case of finite games Theorem 1 implies the following result. For the case of independent beliefs it is implicitly stated in [19], explicitly formulated in [20] (see [14, page 181]) and proved using Harsanyi type spaces in [21].

Corollary 2. Assume the initial game H is finite.

(i) For all belief models for H

$$G_{\mathbf{RAT}(br) \cap B^* \mathbf{RAT}(br)} \subseteq U_{msd}^\infty,$$

(ii) for all knowledge models for H

$$G_{K^* \mathbf{RAT}(br)} \subseteq U_{msd}^\infty,$$

where in both situations we use in br_i either the set of point beliefs or the set of independent beliefs or the set of correlated beliefs.

Proof. The argument is analogous as in the previous proof but relies on a subsidiary result and runs as follows.

(i) Denote respectively by brp_i , bri_i and brc_i the best response property w.r.t. *point*, *independent* and *correlated* beliefs of the opponents. Below ϕ stands for either brp , bri or brc .

By Lemma 2 and Theorem 1 $G_{\mathbf{RAT}(\phi) \cap B^* \mathbf{RAT}(\phi)} \subseteq T_\phi^\infty$. Further, for all restrictions G we have both $T_\phi(G) \subseteq U_\phi(G)$ and $U_{br}(G) \subseteq U_{bri}(G) \subseteq U_{brc}(G)$. So by Lemma 1 $T_\phi^\infty \subseteq U_{brc}^\infty$. But by the result of [22], (page 60) (that is a modification of the original result of [23]), for all restrictions G we have $U_{brc}(G) = U_{msd}(G)$, so $U_{brc}^\infty = U_{msd}^\infty$, which yields the conclusion.

(ii) By (i) and the fact that $K^* \mathbf{RAT}(br) \subseteq \mathbf{RAT}(br)$. □

Finally, let us clarify the situation for the remaining two optimality notions, wd_i and mwd_i . For them the inclusions of Corollaries 1 and 2 do not hold. Indeed, it suffices to consider the following initial game H :

	L	R
U	1, 0	1, 0
D	1, 0	0, 0

Here every strategy is a best response but D is weakly dominated by U . So both U_{wd}^∞ and U_{mwd}^∞ are proper subsets of T_{br}^∞ . On the other hand by Theorem 1(iii) for some standard knowledge model for H we have $G_{K^* \mathbf{RAT}(br)} = T_{br}^\infty$. So for this knowledge model neither $G_{K^* \mathbf{RAT}(br)} \subseteq U_{wd}^\infty$ nor $G_{K^* \mathbf{RAT}(br)} \subseteq U_{mwd}^\infty$ holds.

Acknowledgements

We thank one of the referees for useful comments. We acknowledge helpful discussions with Adam Brandenburger, who suggested Corollaries 1 and 2, and with Giacomo Bonanno who, together with a referee of [7], suggested to incorporate common beliefs in the analysis. Joe Halpern pointed us to [16]. This paper was previously sent for consideration to another major game theory journal, but ultimately withdrawn because of different opinions with the referee. We would like to thank the referee and associate editor of that journal for their comments and help provided.

References

1. Milgrom, P.; Roberts, J. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica* **1990**, *58*, 1255–1278.
2. Chen, Y.C.; Long, N.V.; Luo, X. Iterated strict dominance in general games. *Games Econ. Behav.* **2007**, *61*, 299–315.
3. Bernheim, B.D. Rationalizable strategic behavior. *Econometrica* **1984**, *52*, 1007–1028.
4. Brandenburger, A.; Friedenberg, A.; Keisler, H. Admissibility in games. *Econometrica* **2008**, *76*, 307–352.
5. Halpern, J.; Pass, R. A logical characterization of iterated admissibility. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge (TARK XII)*; Stanford University,

- Standord, CA, USA, 6–8 July 2009; Available online: <http://arxiv.org/abs/0906.4326> (accessed on 30 September 2010).
6. van Benthem, J. Rational dynamics and epistemic logic in games. *Int. Game Theory Rev.* **2007**, *9*, 13–45.
 7. Apt, K.R. The many faces of rationalizability. *The B.E. J. Theoretical Econ.* **2007**, *7*, Article 18.; Available online: <http://arxiv.org/abs/cs.GT/0608011> (accessed on 30th September 2010).
 8. Apt, K.R. Relative Strength of Strategy Elimination Procedures. *Econ. Bull.* **2007**, *3*, 1–9; Available online: <http://www.economicsbulletin.com/> (accessed on 30th September 2010).
 9. Lipman, B.L. How to decide how to decide how to . . . : Modeling limited rationality. *Econometrica* **1991**, *59*, 1105–1125.
 10. Lipman, B.L. A note on the implications of common knowledge of rationality. *Games Econ. Behav.* **1994**, *6*, 114–129.
 11. Heifetz, A.; Samet, D. Knowledge spaces with arbitrarily high rank. *Games Econ. Behav.* **1998**, *22*, 260–273.
 12. Luo, X. General systems and ϕ -stable sets—A formal analysis of socioeconomic environments. *J. Math. Econ.* **2001**, *36*, 95–109.
 13. Apt, K.R. Epistemic analysis of strategic games with arbitrary strategy sets. In *Proceedings of the 11th Conference on Theoretical Aspects of Reasoning about Knowledge (TARK XI)*, Brussels, Belgium, June 25-27th, 2007; Available online: <http://portal.acm.org> (accessed on 30th September 2010).
 14. Battigalli, P.; Bonanno, G. Recent results on belief, knowledge and the epistemic foundations of game theory. *Res. Econ.* **1999**, *53*, 149–225.
 15. Blackburn, P.; de Rijke, M.; Venema, Y. *Modal Logic*; Cambridge University Press: Cambridge, UK, 2001.
 16. Monderer, D.; Samet, D. Approximating common knowledge with common beliefs. *Games Econ. Behav.* **1989**, *1*, 170–190.
 17. Aumann, R. Agreeing to disagree. *Ann. Stat.* **1976**, *4*, 1236–1239.
 18. Tarski, A. A lattice-theoretic fixpoint theorem and its applications. *Pacific J. Math.* **1955**, *5*, 285–309.
 19. Brandenburger, A.; Dekel, E. Rationalizability and correlated equilibria. *Econometrica* **1987**, *55*, 1391–1402.
 20. Stalnaker, R. On the evaluation of solution concepts. *Theor. Decis.* **1994**, *37*, 49–73.
 21. Brandenburger, A.; Friedenberg, A. Intrinsic correlation in games. *J. Econ. Theor.* **2008**, *141*, 28–67.
 22. Osborne, M.J.; Rubinstein, A. *A Course in Game Theory*; The MIT Press: Cambridge, MA, USA, 1994.
 23. Pearce, D.G. Rationalizable strategic behavior and the problem of perfection. *Econometrica* **1984**, *52*, 1029–1050.

© 2010 by the authors; licensee MDPI, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>).

Article

Consistent Beliefs in Extensive Form Games

Paulo Barelli ^{1,2}

¹ Department of Economics, University of Rochester, 214 Harkness Hall, Rochester, NY 14627, USA; E-Mail: paulo.barelli@rochester.edu; Tel.: 1-585-275-8075; Fax: 1-585-256-2309

² Insper Institute of Education and Research, Rua Quatá, 300 - Vila Olímpia 04546-042, São Paulo, Brazil

Received: 1 July 2010; in revised form: 26 September 2010 / Accepted: 15 October 2010 /

Published: 20 October 2010

Abstract: We introduce consistency of beliefs in the space of hierarchies of conditional beliefs (Battigalli and Siniscalchi) and use it to provide epistemic conditions for equilibria in finite multi-stage games with observed actions.

Keywords: hierarchies of conditional beliefs; epistemic conditions; common belief; correlated subgame perfect equilibrium

1. Introduction

Battigalli and Siniscalchi [1] constructed the space of hierarchies of conditional beliefs and used it to provide epistemic foundations for solution concepts in dynamic games. We consider the question of consistency of beliefs in the space of hierarchies of conditional beliefs. In the space of hierarchies of beliefs, Aumann [2], Aumann and Brandenburger [3] and Barelli [4], among others, have used consistency of beliefs to provide epistemic foundations for solution concepts in games in normal form. Here we provide an analogous analysis for multi-stage games with observable actions, in the corresponding space of hierarchies of conditional beliefs. In particular, we show that consistency of beliefs and extensive form rationality provide epistemic foundations for correlated subgame perfect equilibrium (correlated SPE), and these two conditions, plus a notion of constancy of conjectures, provide epistemic foundations for subgame perfect equilibrium (SPE).¹

¹For simplicity we deal only with finite multi-stage games with with observed actions, so sequential rationality is well captured by subgame perfection; the analysis can be generalized to include incomplete information and/or more complex information structures, where sequential equilibrium is the relevant equilibrium concept to capture sequential rationality.

The following simple example helps understand the ideas involved. Consider the standard Battle of Sexes game, with the payoff matrix below,

	<i>F</i>	<i>O</i>
<i>F</i>	2, 1	0, 0
<i>O</i>	0, 0	1, 2

The story is that the players decide simultaneously where to meet (either at a football game, *F*, or an opera house, *O*), and each player would rather go to the same place as the other, but has a preference for one venue over the other. Let $A_i = \{F, O\}$ for $i = 1, 2$ and $A = A_1 \times A_2$. A correlated equilibrium for such a simultaneous move game is a Nash equilibrium of the game augmented by some payoff irrelevant state space, which is understood by both players. Consider, for instance, that each player chooses *F* if the weather is good, and chooses *O* otherwise (that is, they go to the outdoor event if the weather is good, and to the indoor event if the weather is not good). It is clear that such a strategy is a Nash equilibrium of the game augmented by the state space {good weather, weather not good}: if the other player uses the strategy, it is in the given player's interest use it as well (if the weather is good (not good), a given player knows that the other will go to the football game (opera house), and will do well to go there too). Let p be the probability of the weather being good. Then the pair of strategies above gives rise to the distribution of joint actions $\eta \in \Delta(A)$ given by $\eta(F, F) = p$ and $\eta(O, O) = 1 - p$, and it is without loss to focus directly on such distributions in describing a correlated equilibrium. It suffices that, for each $a_i \in A_i$, the expected payoff of a_i given $\eta(a_i, \cdot) \in \Delta(A_j)$, $j \neq i$, is not smaller than the expected payoff of any other action a'_i , for $i = 1, 2$.

Now consider that the players play the game twice. That is, the players play the game once, observe its outcome, play it again, and get the sum of the payoffs obtained in each round. Let $H = \{\emptyset\} \cup A$ denote the set of histories. The empty history represents the first round, and each of the four joint strategies in A represents a possible second round. Recall that a SPE is a Nash equilibrium of the entire game that induces Nash equilibria at each subgame. Analogously, a correlated SPE is a correlated equilibrium of the entire game that induces correlated equilibria at each subgame. It can be described as follows. Let $\eta \in \Delta(A)$ be a correlated equilibrium of the original Battle of the Sexes game, like the η described above. A correlated SPE is a list of probability distributions $(\nu_h)_{h \in H}$ with $\nu_h \in \Delta(A)$ for each $h \in H$, where ν_h a correlated equilibrium for the continuation game at history $h \in A$ and ν_\emptyset a correlated equilibrium of the one shot game given by the first round outcome and the contingent second round outcome, given ν_h with $h \in A$. That is, each of the four continuation games is simply the original Battle of the Sexes game played after the first round. So a correlated equilibrium for a continuation game is a probability distribution $\eta \in \Delta(A)$. In the first round, on the other hand, each joint action gives rise to a (potentially) different continuation strategy. So it is not a simple stage game as the games in the second round. But it can be viewed as an one-shot game, with payoffs given by the sum of what is obtained in the first round and of the conditional payoffs in the second round, given the correlated equilibria of the four potential continuation games. Then, for instance, $\nu_h = \eta$ for all $h \in H$ is a correlated SPE, because $\nu_a (= \eta)$ is a correlated equilibrium of the continuation game after history $h = a$ for each $a \in A$, and given the four continuation correlated moves $(\nu_a)_{a \in A}$, $\nu_\emptyset (= \eta)$ is a correlated equilibrium of the game with payoffs $u_i(a) + u_i(\eta)$, where u_i is player i 's stage game payoff and $u_i(\eta)$ is the expected payoff given η (so, in

particular, the expected payoff given ν_\emptyset is simply $u_i(\eta) + u_i(\eta)$. More complex correlated SPE involving different correlated continuation strategies can be constructed analogously.

Likewise, let $\eta \in \Delta(A_1) \times \Delta(A_2)$ be a joint distribution associated with a Nash equilibrium of the stage game. For instance, $\eta(F, F) = \eta(O, O) = \frac{2}{9}$, $\eta(F, O) = \frac{4}{9}$ and $\eta(O, F) = \frac{1}{9}$, which is the joint distribution associated with the Nash equilibrium of the original Battle of the Sexes game in non degenerate mixed strategies. Then a list $(\nu_h)_{h \in H}$ with $\nu_h = \eta$ for all $h \in H$ is a SPE of the game, for the same reason as above. More complex SPE with different Nash equilibria of the continuation games can be constructed analogously.

Now let's perform an epistemic analysis on the game, that is, an analysis of knowledge and beliefs of the players. In order to do so, we append a type structure (T_1, T_2, g_1, g_2) with $g_{i,h} \in \Delta(S \times T_j)$ for each $h \in H$, where $S = S_1 \times S_2$ with $S_i = \{F, O\}^H$ for $i = 1, 2$. The beliefs of a type t_i , $(g_{i,h}(t_i))_{h \in H}$ form a conditional probability system (CPS), (the formal definitions are provided below). A "state" for a player is a strategy-type pair (s_i, t_i) , describing the player's strategy choice and beliefs. Epistemic statements can now be stated in terms of the states of the players. For instance, let $S_i(h)$ be player i 's set of strategies consistent with history $h \in H$. Let $\eta = \eta(\cdot | S_j(h))_{h \in H}$, where $\eta(\cdot | S_j(h)) \in \Delta(S_j(h))$ for each $h \in H$. We say that s_i is a best response to η , written $s_i \in r_i(\eta)$, if s_i maximizes the expected utility with respect to $\eta(\cdot | S_j(h))$ for every history h consistent with s_i . And we say that the strategy-type pair $(s_i, t_i) \in S_i \times T_i$ is rational if $s_i \in r_i((\text{marg}_{S_j} g_{i,h}(t_i))_{h \in H})$. Statements like "rationality is common knowledge among the players" can be described by a type structure where in each state $(s, t) \in S_1 \times S_2 \times T_1 \times T_2$ both players are rational. Note that a type of a player determines the conditional beliefs at every history, and rationality captures sequentially rational choices, after every history (given the conditional beliefs).

Assume that the beliefs of the players are consistent in the following sense. There is a CPS $(\mu_h)_{h \in H}$ with $\mu_h \in \Delta(S(h) \times T)$ for each $h \in H$, such that $g_{i,h}(t_i)(E \times T_j) = \mu_h(E \times T | t_i)$ for all $E \subset S$, $t_i \in T_i$ and $i = 1, 2$. The idea is analogous to action-consistency in Barelli [4], which is a generalization of the standard common prior assumption. Because strategies are in principle verifiable entities, we can conceive of an outside observer offering bets on S , conditional on each history, where the payouts of the bets are measured in utils. The two players will be in a *no-bets situation* if there does not exist a bet that yields a sure gain to an outsider. In Barelli [4] it is shown that this is equivalent to consistency of beliefs, as defined above.

Now, if consistency and rationality obtain at every $(s, t) \in S \times T$,² then we can identify a correlated SPE $(\nu_h)_{h \in H}$ from the CPS $(\mu_h)_{h \in H}$ by putting $\nu_h(a) = \mu_h(\{s : s_h = a\} \times T)$, for all $a \in A$. Indeed, if it is the case that beliefs are consistent and the CPS $(\mu_h)_{h \in H}$ satisfies $\mu_h(\{s : s_h = (F, F)\} \times T) = p$ and $\mu_h(\{s : s_h = (O, O)\} \times T) = 1 - p$ for every $h \in H$, then it is straightforward to verify that rationality is obtained at every state, and that we obtain the correlated SPE described above. Indeed, rationality implies that no player wants to deviate from the recommended action, as required in a correlated SPE, and $(\nu_h)_{h \in H}$ is exactly the correlated SPE above. Other correlated SPE are analogously obtained as we vary the consistent CPS $(\mu_h)_{h \in H}$. The key observation here is that, under consistency, rationality ensures that the system of inequalities defining a correlated SPE is met.

²More precisely, if throughout the support of the CPS $(\mu_h)_{h \in H}$ defined above we have rational strategy-type pairs.

If instead $\mu_h(\{s : s_h = (F, F)\} \times T) = \mu_h(\{s : s_h = (O, O)\} \times T) = \frac{2}{9}$, $\mu_h(\{s : s_h = (F, O)\} \times T) = \frac{4}{9}$ and $\mu_h(\{s : s_h = (O, F)\} \times T) = \frac{1}{9}$ for every $h \in H$, then we again have rationality at every state, and the SPE described above is obtained. As in Aumann and Brandenburger [3] and Barelli [4], the key observation is that constancy of conjectures in the support of the CPS $(\mu_h)_{h \in H}$ ensures that $(\mu_h(\{s : s_h = a\} \times T))_{a \in A}$ is the product of its marginals, just as above. So rationality, consistency and constancy of conjectures in the support of the CPS are sufficient conditions for a SPE. It is important to note that constancy of conjectures is implied by (but does not imply) conjectures being commonly known among the players.

2. Set Up

The set up is as in Battigalli and Siniscalchi [1]. Let X be a Polish space, and let \mathcal{A} be its Borel sigma algebra. Let $\mathcal{B} \in \mathcal{A}$ be a countable collection of clopen sets, with $\emptyset \notin \mathcal{B}$. The collection \mathcal{B} represents the relevant hypotheses. A CPS on $(X, \mathcal{A}, \mathcal{B})$ is a mapping $\mu(\cdot|\cdot) : \mathcal{A} \times \mathcal{B} \rightarrow [0, 1]$ satisfying: (i) $\mu(B|B) = 1$ for all $B \in \mathcal{B}$, (ii) $\mu(\cdot|B) \in \Delta(X)$, and (iii) for all $A \in \mathcal{A}, B, C \in \mathcal{B}$, if $A \subset B \subset C$ then $\mu(A|B)\mu(B|C) = \mu(A|C)$.³ The set of CPSs on $(X, \mathcal{A}, \mathcal{B})$ is a closed subset of $[\Delta(X)]^{\mathcal{B}}$, and it denoted by $\Delta^{\mathcal{B}}(X)$.

Consider a finite multi-stage game G with observable actions (Fudenberg and Tirole [5], Chap. 3). Let \mathcal{H} be the set of histories and let S_i be the set of strategies $s_i : \mathcal{H} \rightarrow A_i$, where A_i is the set of all possible actions for player $i \in I$, and $s_i(h) \in A_i(h)$ for each $h \in \mathcal{H}$, where $A_i(h)$ is the set of actions available at h . Let $u_i : S \rightarrow \mathcal{R}$ denote player i 's utility function, with $S = \times_{i \in I} S_i$. As usual, we use $A_{-i} = \times_{j \neq i} A_j$ and $A = \times_{i \in I} A_i$ (likewise for other sets, like T_i, T_{-i} and T below.)

A correlated equilibrium of a finite normal form game $(A_i, u_i)_{i \in I}$ is a probability distribution $\eta \in \Delta(A)$ satisfying

$$\sum_{a_{-i} \in A_{-i}} [u_i(a_i, a_{-i}) - u_i(a'_i, a_{-i})]\eta(a) \geq 0$$

for all $i \in I$ and all $a_i, a'_i \in A_i$. The interpretation is the one provided in the Introduction: the players use some external random device to peg their actions to, and assuming that the other players follow the recommended choices with the implied likelihoods, a given player has no incentive to deviate from his/her recommended choices. Because any such equilibrium generates a probability distribution over the joint actions, it is convenient to focus directly on such distributions (in the same way that a mixed strategy Nash equilibrium is defined directly on distributions, and not on the random variables that generate the distributions).

A correlated SPE of a finite multi-stage game with observed actions is given by $\nu = (\nu_h)_{h \in \mathcal{H}}$, with $\nu_h \in \Delta(A(h))$, which induces correlated equilibria at every subgame. That is, for each history h we have a continuation game $G(h)$ where the payoffs are defined for the histories that are consistent with h . Given h , we have a continuation correlated strategy $\nu|h$, given by the restriction of ν to histories consistent with h . Then a correlated SPE is ν such that $\nu|h$ is a correlated equilibrium of $G(h)$ for every $h \in \mathcal{H}$. Standard dynamic programming arguments show that this is equivalent to the description provided in the Introduction. An SPE is a correlated SPE ν with $\nu_h \in \times_{i \in I} \Delta(A_i(h))$ for each $h \in \mathcal{H}$.

³ $\Delta(X)$ denotes the space of probability measures on (X, \mathcal{A}) .

Let $S_i(h)$ be player i 's set of strategies consistent with history $h \in \mathcal{H}$, and let $\mathcal{H}(s_i)$ be the set of histories consistent with s_i . The relevant hypotheses for the players are thus the collection $\mathcal{B} = \{S(h) : h \in \mathcal{H}\}$. For a given player i , the hypotheses that are consistent with i 's strategies are $\mathcal{B}_i = \{S_i(h) : h \in \mathcal{H}\}$. As in Battigalli and Siniscalchi [1], we simplify notation by writing $\Delta^{\mathcal{B}_i}(\cdot)$ and $\Delta^{\mathcal{B}}(\cdot)$ as $\Delta^{\mathcal{H}}(\cdot)$.

In order to perform an epistemic analysis, we append a type structure to the game, describing the beliefs of the players. A type space is a tuple $\mathcal{T} = (T_i, g_i)_{i \in I}$ with $g_i : T_i \rightarrow \Delta^{\mathcal{H}}(S \times T_{-i})$ for each $i \in I$. Again, to simplify notation we write $(g_{i,h}(t_i))_{h \in \mathcal{H}} \in \Delta^{\mathcal{H}}(S \times T_{-i})$ instead of $(g_{i,S(h)}(t_i))_{S(h) \in \mathcal{B}} \in \Delta^{\mathcal{B}}(S \times T_{-i})$.

Let $\eta = \eta(\cdot | S_{-i}(h))_{h \in \mathcal{H}} \in \Delta^{\mathcal{H}}(S_{-i})$. We say that s_i is a best response to η , written $s_i \in r_i(\eta)$, if for all $h \in \mathcal{H}(s_i)$ and $s'_i \in S_i(h)$, we have

$$\sum_{s_{-i} \in S_{-i}(h)} [u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})] \eta(s_{-i} | S_{-i}(h)) \geq 0$$

We then say that a strategy-type pair (s_i, t_i) is rational if $s_i \in r_i((\text{marg}_{S_{-i}} g_{i,h}(t_i))_{h \in \mathcal{H}})$, and if $s_i \in S_i(h)$ then $g_{i,h}(t_i)(\{s_i\} \times S_{-i} \times T_{-i}) = 1$. We say that player i is rational at state $(s, t) \in S \times T$ if the strategy-type pair $(s_i, t_i) \in S_i \times T_i$ is rational.

A CPS $\mu \in \Delta^{\mathcal{H}}(S \times T)$ is called a consistent prior if

$$\mu_h(E \times T) = \int_{T_i} g_{i,h}(t_i)(E \times T_{-i}) \text{marg}_{T_i} \mu_h(dt_i)$$

for all $i \in I$, all $h \in \mathcal{H}$ and all $E \subset S$. It then follows that $g_{i,h}(t_i)(E \times T_{-i}) = \mu_h(E \times T | t_i)$ for all $i \in I$, all $h \in \mathcal{H}$, all $E \subset S$ and $\text{marg}_{T_i} \mu_h$ -a.e. t_i . Let $\text{supp } \mu = \bigcup_{h \in \mathcal{H}} \text{supp } \mu_h$ denote the support of the consistent prior. As advanced above, consistency is founded on players being in a *no-bets situation*, that is, a situation where an outside observer cannot make a sure gain on the group of players by offering bets on the strategy choices of the players. Proposition 5.3 in Barelli [4] establishes the equivalence between consistency and no-bets, and the reader is referred to that paper for further details.

For the sake of comparison with the literature, consider a finite normal form game $G = (A_i, u_i)_{i \in I}$ and a type space $\mathcal{T} = (T_i, \lambda_i)_{i \in I}$, with $\lambda_i(t_i) \in \Delta(A \times T_{-i})$ capturing hierarchies of beliefs. Player i is rational at state (a, t) if a_i is a best response to his conjecture $\text{marg}_{A_{-i}} \lambda_i(t_i)$ and $\lambda_i(t_i)(\{a_i\} \times A_{-i} \times T_{-i}) = 1$. A common prior is a probability measure $p \in \Delta(A \times T)$ such that $\lambda_i(t_i) = p(\cdot | t_i)$ for $\text{marg}_{T_i} p$ -a.e. t_i . An action-consistent prior is a probability measure $\pi \in \Delta(A \times T)$ such that $\text{marg}_A \lambda_i(t_i) = \text{marg}_A \pi(\cdot | t_i)$ for $\text{marg}_{T_i} \pi$ -a.e. t_i . Aumann [2] showed that, when there is a common prior, common knowledge of rationality implies that players play a correlated equilibrium. Aumann and Brandenburger [3] showed that common knowledge of rationality and of conjectures and the existence of a common prior are sufficient conditions for players to play a Nash equilibrium. These results were extended in Barelli [4] with the use of action-consistency in the place of common prior, rationality in the support of the action-consistent prior in the place of common knowledge of rationality and constancy of conjectures in the support of the action-consistent prior in the place of common knowledge of conjectures.

Note that the notion of consistency used here is much more demanding than using action-consistency in the normal form of the game. Consistency requires that players be at a no-bets situation after every

history $h \in \mathcal{H}$, whereas action-consistency allows for players to not be at a no-bets situation after histories that are not compatible with the strategy profiles in the support of the action-consistent prior. Players are required to be aware of a potential outsider at every counterfactual that they envisage while choosing their strategies.

3. Results

For a given consistent prior μ , let $\nu = (\nu_h)_{h \in \mathcal{H}}$ be given by $\nu_h(a) = \text{marg}_{S_i} \mu_h(s : s_h = a)$ for each $a \in A(h)$, so that $\nu_h \in \Delta(A(h))$. We have:

Proposition 1. *Let G be a finite multi-stage game with observed actions, and let \mathcal{T} be a type space associated with G . Assume that there exists a consistent prior $\mu \in \Delta^{\mathcal{H}}(S \times T)$ such that player i is rational at all $(s, t) \in \text{supp } \mu$, for every $i \in I$. Then ν defined above is a correlated SPE.*

Proof. By consistency, we have $\text{marg}_{S_{-i}} g_{i,h}(t_i) = \text{marg}_{S_{-i}} \mu_h(\cdot | t_i)$ for every $i \in I, h \in \mathcal{H}$ and $t_i \in \text{supp } \text{marg}_{T_i} \mu_h$. By rationality we then have for each $i \in I$ and every $h \in \mathcal{H}(s_i)$

$$\sum_{s_{-i} \in S_{-i}(h)} [u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})] \text{marg}_{S_{-i}} \mu_h(s_{-i} | t_i) \geq 0$$

for every (s_i, t_i) and every $s'_i \in S_i(h)$. Let $\eta = (\eta_h)_{h \in \mathcal{H}}$ with $\eta_h \in \Delta(S(h))$ be given by

$$\eta_h(s) = \int_{T_i(s_i)} \text{marg}_{S_{-i}} \mu_h(s_{-i} | t_i) \text{marg}_{T_i} \mu_h(dt_i)$$

where $T_i(s_i) = \{t'_i \in T_i : (s_i, t'_i) \text{ is rational}\}$, so that

$$\sum_{s_{-i} \in S_{-i}(h)} [u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})] \eta_h(s) \geq 0$$

for every $i \in I, s_i, s'_i$ in $S_i(h)$ and $h \in \mathcal{H}(s_i)$. Now notice that the restriction of ν to a history $h \in \mathcal{H}$, $\nu|h$, is the behavioral representation of η_h . By Kuhn's Theorem, the distribution over final outcomes induced by η_h is the same as that induced by $\nu|h$, so $\nu|h$ is a correlated equilibrium of the continuation game $G(h)$ for every history $h \in \mathcal{H}$, and we are done.

Let $\phi_{i,h}(t_i) = \text{marg}_{S_{-i}} g_{i,h}(t_i)$ denote the conjecture of type t_i at $h \in \mathcal{H}$. We have:

Proposition 2. *Let G be a finite multi-stage game with observed actions, and let \mathcal{T} be a type space associated with G . Assume that there exists a consistent prior $\mu \in \Delta^{\mathcal{H}}(S \times T)$ such that (i) player i is rational at all $(s, t) \in \text{supp } \mu$ for every $i \in I$ and (ii) $\phi_{i,h}(t_i) = \phi_{i,h}(t'_i)$ for every $i \in I$ for every $t_i, t'_i \in \text{supp } \text{marg}_{T_i} \mu_h$, for each $h \in \mathcal{H}$. Then ν defined above is a SPE.*

Proof. Fix $h \in H$ and let $\phi_{i,h}$ be player i 's constant conjecture in the support of $\text{marg}_{T_i} \mu_h$. By consistency and rationality, we have for each $s \in S(h)$

$$\text{marg}_S \mu_h(s) = \text{marg}_{T_i} \mu_h(T_i(s_i)) \phi_{i,h}(s_{-i})$$

where $T_i(s_i)$ is as in Proposition 1. Hence

$$\text{marg}_S \mu_h = \text{marg}_{S_i} \mu_h \otimes \text{marg}_{S_{-i}} \mu_h$$

Now induction in the number of players shows that $\text{marg}_S \mu_h = \otimes_{i \in I} \text{marg}_{S_i} \mu_h$, and a fortiori $\nu_h \in \times_{i \in I} \Delta(A_i(h))$, for all $h \in \mathcal{H}$. The result then follows from Proposition 1.

4. Conclusion

The results in section 3 tell us the following: under consistency, rationality yields correlated SPE and adding constancy of conjectures to these two conditions yield SPE. These results are analogous to the results in Aumann [2], Aumann and Brandenburger [3] and Barelli [4]. As in the latter, beliefs are required to be consistent only at events that are potentially observable by an outsider, who could in principle force beliefs to be consistent by offering bets on the observable events. Rationality and constancy of conjectures have to hold in the support of the consistent prior. Because rationality and/or constancy of conjectures are implied by (but do not imply) rationality and/or conjectures being commonly known among the players, we have that rationality need not be common knowledge for players to play a correlated SPE, and neither do conjectures have to be common knowledge for players to play a SPE.

References

1. Battigalli, P.; Siniscalchi, M. Hierarchies of Conditional Beliefs and Interactive Epistemology in Dynamic Games. *J. Econ. Theory* **1999**, *88*, 188–230.
2. Aumann, R. Correlated Equilibrium as an Expression of Bayesian Rationality. *Econometrica* **1987**, *55*, 1–18.
3. Aumann, R.; Brandenburger, A. Epistemic Conditions for Nash Equilibrium. *Econometrica* **1995**, *63*, 1161–1180.
4. Barelli, P. Consistency of Beliefs and Epistemic Conditions for Nash and Correlated Equilibria. *Game. Econ. Behav.* **2009**, *67*, 363–375.
5. Fudenberg, D.; Tirole, J. *Game Theory*; The MIT Press: Cambridge, MA, USA, 1992.

© 2010 by the authors; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license <http://creativecommons.org/licenses/by/3.0/>.

Article

A Modal Logic of Epistemic Games

Emiliano Lorini * and François Schwarzentruber *

IRIT-CNRS, Université de Toulouse, 118 route de Narbonne, 31062 Toulouse, France

* Author to whom correspondence should be addressed; E-Mails: lorini@irit.fr, schwarze@irit.fr.

Received: 11 June 2010; in revised form: 7 September 2010 / Accepted: 22 October 2010 /

Published: 2 November 2010

Abstract: We propose some variants of a multi-modal of joint action, preference and knowledge that support reasoning about epistemic games in strategic form. The first part of the paper deals with games with complete information. We first provide syntactic proofs of some well-known theorems in the area of interactive epistemology that specify some sufficient epistemic conditions of equilibrium notions such as Nash equilibrium and Iterated Deletion of Strictly Dominated Strategies (IDSDS). Then, we present a variant of the logic extended with dynamic operators of Dynamic Epistemic Logic (DEL). We show that it allows to express the notion IDSDS in a more compact way. The second part of the paper deals with games with weaker forms of complete information. We first discuss several assumptions on different aspects of perfect information about the game structure (e.g., the assumption that a player has perfect knowledge about the players' strategy sets or about the preference orderings over strategy profiles), and show that every assumption is expressed by a corresponding logical axiom of our logic. Then we provide a proof of Harsanyi's claim that all uncertainty about the structure of a game can be reduced to uncertainty about payoffs. Sound and complete axiomatizations of the logics are given, as well as some complexity results for the satisfiability problem.

Keywords: game theory; modal logic; dynamic epistemic logic

1. Introduction

The aim of this article is to propose a modal logic framework that allows to reason about epistemic games in strategic form. In this kind of games players decide what to do according to some general

principles of rationality while being uncertain about several aspects of the interaction such as other agents' choices, other agents' preferences, *etc.*

While epistemic games have been extensively studied in economics (in the so-called interactive epistemology area, see e.g., [1–6]) and while there have been few analyses of epistemic games in modal logic (see, e.g., [3,7–9]), no modal logic approach to epistemic games has been proposed up to now that addresses all the following issues at the same time:

- to provide a *logic*, and a corresponding *formal semantics*, which is sufficiently general:
 - to express solution concepts like Nash Equilibrium or Iterated Deletion of Strictly Dominated Strategies (IDSDS) and to derive syntactically the epistemic and rationality conditions on which such solution concepts are based,
 - to study epistemic games both with *complete* information and with *incomplete* information;
- to prove its *soundness* and *completeness*;
- to study its computational properties like *decidability* and *complexity*.

In this article, we try to fill this gap by proposing some variants of a multi-modal logic of joint action, preference and knowledge interpreted on a Kripke-style semantics, that allow to represent both epistemic games with complete information and epistemic games with different forms of incomplete information about the game structure. We give sound and complete axiomatizations of all these logics as well as some complexity results for the satisfiability problem.

The article is organized in two parts: the first part is focused on strategic games with complete information, while the second one extends the analysis to strategic games with incomplete information.

In Section 2 we present a modal logic, called \mathcal{MLEG} (*Modal Logic of Epistemic Games*), that supports reasoning about epistemic games with complete information in which an agent can only have uncertainty about other agents' current choices. A complete axiomatization and complexity results for this logic are given.

Section 3 is devoted to the analysis in \mathcal{MLEG} of the epistemic conditions of Nash equilibrium and IDSDS. We use the logic \mathcal{MLEG} in order to provide syntactic proofs of some well-known theorems in the area of interactive epistemology such as the theorem that specifies some sufficient epistemic conditions of Nash equilibrium in terms of players' rationality and knowledge about other players' choices, and the theorem characterizing IDSDS in terms of common knowledge of rationality.

In Section 4 we make \mathcal{MLEG} dynamic by extending it with constructions of Dynamic Epistemic Logic (DEL) [10–12], and we show that this dynamic version of \mathcal{MLEG} allows to express the notion IDSDS in a more compact way than in the static \mathcal{MLEG} . A complete axiomatization for this dynamic extension of the logic \mathcal{MLEG} is given.

In Section 5 we show how our logical framework can be easily adapted in order to study strategic interaction with incomplete information about the game structure. In Section 6 we discuss several assumptions on different aspects of complete information about the game structure (e.g., the assumption that a player has perfect knowledge about the players' strategy sets or about the players' preference ordering over strategy profiles). We show that every assumption is expressed by a corresponding

logical axiom. Consequently, a class of epistemic games characterized by a specific aspect of complete information about the game structure corresponds to a specific variant of the logic \mathcal{MLEG} . We present some complexity results for these variants of the logic \mathcal{MLEG} , that are interesting because they show how complexity of our logic varies from games with complete information to games with incomplete information.

We also provide a formal proof of Harsanyi's claim that all uncertainty about the structure of a game can be reduced to uncertainty about payoffs. The novel aspect of our contribution is that we prove Harsanyi's claim in a purely qualitative setting with no probabilities, while existing proofs are given in a quantitative setting with probabilities (see, e.g., [13]).

Proofs of theorems are given in an Annex at the end of the article.

Before concluding this introduction, we would like to emphasize an aspect of our work that could be interesting for a game-theorist.

As the logics presented in this paper are sound and complete, they allow to study strategic interaction both at the semantic level and at the syntactic level. In this sense, they provide a formal framework which unifies two approaches traditionally opposed by authors working in the area of formal interactive epistemology: the semantic approach and the syntactic approach.¹ However, it is worth noting that syntactic derivations of various results concerning the epistemic foundations of game theory are not interesting in itself. Instead, this kind of analysis is useful to identify specific features that are important for the foundations of game theory, for example whether certain assumptions on the players' knowledge are indeed necessary to prove results concerning the epistemic conditions of equilibrium notions such as Nash equilibrium and IDSDS. Typical assumptions on players' knowledge are for example the assumption that knowledge is positively and negatively introspective (i.e. if I know that φ is true then I know that I know that φ is true, and if I do not know that φ is true then I know that I do not know that φ is true), the factivity principle that knowing that φ implies that φ is true, or the assumption that a player has perfect knowledge about some aspects of the game such as the players' strategy sets and the players' preference ordering over strategy profiles.

2. A Logic of Joint Actions, Knowledge and Preferences

We present in this section the multi-modal logic \mathcal{MLEG} (*Modal Logic of Epistemic Games*) integrating the concepts of joint action, knowledge and preference. This logic supports reasoning about epistemic games in strategic form in which an agent might be uncertain about the current choices of the other agents.

2.1. Syntax

The syntactic primitives of \mathcal{MLEG} are the finite set of agents $Agnt$, the set of atomic formulas Atm , a nonempty finite set of atomic action names $Act = \{a_1, a_2, \dots, a_{|Act|}\}$. Non-empty sets of agents are called *coalitions* or *groups*, noted C_1, C_2, \dots . We note $2^{Agnt^*} = 2^{Agnt} \setminus \{\emptyset\}$ the set of coalitions.

To every agent $i \in Agnt$ we associate the set Act_i of all possible ordered pairs agent/action $i:a$, that is, $Act_i = \{i:a \mid a \in Act\}$. Besides, for every coalition C we note Δ_C the set of all joint actions of this

¹See [2] for an analysis of the relation between the two approaches.

coalition, that is, $\Delta_C = \prod_{i \in C} Act_i$. Elements in Δ_C are C -tuples noted $\alpha_C, \beta_C, \gamma_C, \delta_C, \dots$. If $C = Agt$, we write Δ instead of Δ_{Agt} . Elements in Δ are also called strategy profiles. Given $\delta \in \Delta$, we note δ_i the element in δ corresponding to agent i . Moreover, for notational convenience, we write $\delta_{-i} = \delta_{Agt \setminus \{i\}}$.

The language $\mathcal{L}_{\mathcal{M}\mathcal{L}\mathcal{E}\mathcal{G}}$ of the logic $\mathcal{M}\mathcal{L}\mathcal{E}\mathcal{G}$ is given by the following rule:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \vee \varphi) \mid [\delta_C]\varphi \mid \Box\varphi \mid K_i\varphi \mid [\text{good}]_i\varphi$$

where p ranges over Atm , i ranges over Agt , and δ_C ranges over $\bigcup_{C \in 2^{Agt^*}} \Delta_C$. The classical Boolean connectives $\wedge, \rightarrow, \leftrightarrow, \top$ (tautology) and \perp (contradiction) are defined from \vee and \neg in the usual manner. We also follow the standard rules for omission of parentheses.

The formula $[\delta_C]\varphi$ reads “if coalition C chooses the joint action δ_C then φ holds”. Therefore, $[\delta_C]\perp$ reads “coalition C does not choose the joint action δ_C ”.

\Box is a necessity operator which enables to quantify over possible joint actions of all agents, that is, over the strategy profiles of the current game (the terms “joint actions of all agents” and “strategy profiles” are supposed here to be synonymous). $\Box\varphi$ reads “ φ holds for every alternative strategy profile of the current game”, or simply “ φ is necessarily true”.

Operators K_i are standard epistemic modal operators. Construction $K_i\varphi$ is read as usual “agent i knows that φ is true”, whereas the construction $[\text{good}]_i\varphi$ is read “ φ is true in all worlds which are for agent i at least as good as the current one concerning the strategy profile that is chosen”. We define $\langle \text{good} \rangle_i\varphi$ as an abbreviation of $\neg[\text{good}]_i\neg\varphi$. Operators $[\text{good}]_i$ are used in $\mathcal{M}\mathcal{L}\mathcal{E}\mathcal{G}$ to define agents’ preference orderings over the strategy profiles of the current game. Similar operators are studied by [14] (see Section 3.1 for a discussion).

We use $EK_C\varphi$ as an abbreviation of $\bigwedge_{i \in C} K_i\varphi$, i.e. every agent in C knows φ (if $C = \emptyset$ then $EK_C\varphi$ is equivalent to \top). Then we define by induction $EK_C^k\varphi$ for every natural number $k \in \mathbf{N}$:

$$EK_C^0\varphi \stackrel{\text{def}}{=} \varphi$$

and for all $k \geq 1$,

$$EK_C^k\varphi \stackrel{\text{def}}{=} EK_C(EK_C^{k-1}\varphi)$$

We define for all natural numbers $n \in \mathbf{N}$, $MK_C^n\varphi$ as an abbreviation of $\bigwedge_{1 \leq k \leq n} EK_C^k\varphi$. $MK_C^n\varphi$ expresses C ’s mutual knowledge that φ up to n iterations, i.e., everyone in C knows φ , everyone in C knows that everyone in C knows φ , and so on until level n .

Finally, $\langle \delta_C \rangle\varphi$ abbreviates $\neg[\delta_C]\neg\varphi$, $\diamond\varphi$ abbreviates $\neg\Box\neg\varphi$ and $\widehat{K}_i\varphi$ abbreviates $\neg K_i\neg\varphi$. $\diamond\varphi$ means “ φ is possibly true”. Therefore $\langle \delta_C \rangle\varphi$ reads “coalition C chooses the joint action δ_C and φ holds”, and $\langle \delta_C \rangle\top$ simply reads “coalition C chooses the joint action δ_C ”.

The operator \diamond and the operators $\langle \delta_C \rangle$ can be combined in order to express what a coalition of agents can do. In particular, $\diamond\langle \delta_C \rangle\top$ has to be read “coalition C can choose the joint action δ_C ”. For the individual case, $\diamond\langle i:a \rangle\top$ has to be read “agent i can choose action a ” or also “action a is in the strategy set (action repertoire) of agent i ”. Furthermore, $\diamond\langle \delta \rangle\top$ is read “coalition Agt can choose the joint action (strategy profile) δ ” or also “ δ is a strategy profile of the current game”.

2.2. Semantics

In this subsection, we introduce a Kripke-style possible world semantics of our logic $\mathcal{ML}\mathcal{EG}$.

Definition 1 ($\mathcal{ML}\mathcal{EG}$ -frames). $\mathcal{ML}\mathcal{EG}$ -frames are tuples $F = \langle W, \sim, R, E, \preceq \rangle$ where:

- W is a nonempty set of possible worlds or states;
- \sim is an equivalence relation on W ;
- R is a collection of total functions $R_C : W \rightarrow \Delta_C$ one for every coalition $C \in 2^{Agt^*}$, mapping every world in W to a joint action of the coalition such that:
 - C1** $R_C(w) = \delta_C$ if and only if for every $i \in C$, $R_i(w) = \delta_i$,²
 - C2** if for every $i \in Agt$ there is v_i such that $w \sim v_i$ and $R_i(v_i) = \delta_i$ then there is a v such that $w \sim v$ and $R_{Agt}(v) = \delta$;
- $E : Agt \rightarrow 2^{W \times W}$ maps every agent i to an equivalence relation E_i on W such that:
 - C3** if wE_iv , then $R_i(w) = i:a$ if and only if $R_i(v) = i:a$,
 - C4** if wE_iv then $w \sim v$;
- $\preceq : Agt \rightarrow 2^{W \times W}$ maps every agent i to a reflexive, transitive relation \preceq_i on W such that:
 - C5** if $w \preceq_i v$ then $w \sim v$,
 - C6** if $w \sim v$ and $w \sim u$ then $v \preceq_i u$ or $u \preceq_i v$.

$R_C(w) = \delta_C$ means that coalition C performs the joint action δ_C at world w .

If $w \sim v$ then w and v correspond to alternative strategy profiles of the same game. For short, we say that v is alternative to w . Given a world w , we use the notation $\sim(w) = \{v \mid w \sim v\}$ to denote the equivalence class made up of those worlds corresponding to alternative strategy profiles of the game of which w is one of the strategy profile. Consider e.g., $Agt = \{1, 2\}$ and $Act = \{c, d\}$. In the frame in Figure 1 we have $w_1 \sim w_2$. This means that the joint action performed at w_1 (viz. $\langle 1:c, 1:c \rangle$) and the one performed at w_2 (viz. $\langle 1:c, 1:d \rangle$) are alternative strategy profiles of the same game defined by the equivalence class $\sim(w_1) = \{w_1, w_2, w_3, w_4\}$.

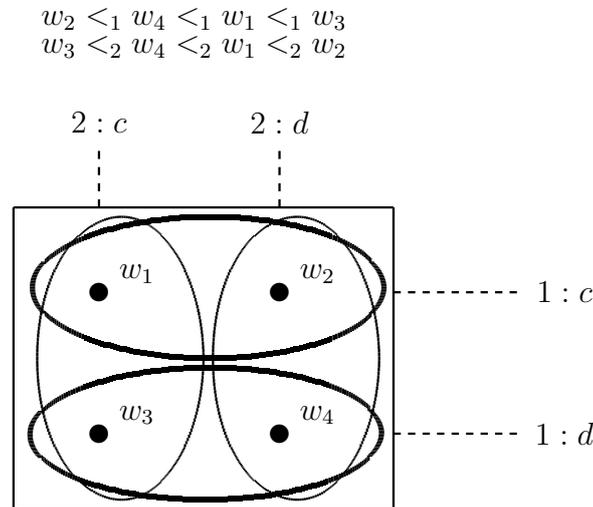
For every $C \subseteq Agt$, if there exists $v \in \sim(w)$ such that C performs δ_C at v then we say that δ_C is possible at w (or δ_C can be performed at w).

wE_iv means that, for agent i , world v is (epistemically) possible at w , whilst $w \preceq_i v$ means that for agent i , world v is at least as good as world w . We write $w =_i v$ iff $w \preceq_i v$ and $v \preceq_i w$, and $w <_i v$ iff $w \preceq_i v$ and not $v \preceq_i w$.

Let us discuss the semantic constraints in Definition 1.

²Note that for notational convenience we write R_i instead of $R_{\{i\}}$.

Figure 1. The equivalence class $\{w_1, w_2, w_3, w_4\}$ represents the Prisoner’s Dilemma game [15] between two players 1 and 2 (action c stands for ‘cooperate’ and action d stands for ‘defect’). Thick ellipses are epistemic relations for 1, thin ellipses are epistemic relations for 2 (both 1 and 2 are uncertain about the other’s action).



According to Constraint **C1**, at world w coalition C chooses the joint action δ_C if and only if, every agent i in C chooses the action δ_i at w . In other words, a certain joint action is performed by a coalition if and only if every agent in the coalition does his part of the joint action. According to the Constraint **C2**, if all individual actions in a joint action δ are possible at world w , then their simultaneous occurrence is also possible at world w .

Constraint **C3** just says that an agent knows what he has decided to do. This is a standard assumption in interactive epistemology and epistemic analysis of games (see [3] for instance).

We suppose complete information about the specification of the game, including the players’ strategy sets (or action repertoires) and the players’ preference ordering over strategy profiles. This assumption is formally expressed by the Constraint **C4**: if world v is epistemically possible for agent i at w , then w and v correspond to alternative strategy profiles of the same game. Complete information about the structure of the game is a standard assumption in game theory. In Section 5, this assumption will be relaxed in order to deal with realistic situations in which an agent might be uncertain about his own utility and other agents’ utilities associated to a certain strategy profile, as well as about his own action repertoire and other agents’ action repertoires.

Finally, we have two constraints over the relations \preceq_i . We suppose that a world v is for agent i at least as good as w only if v is a world which is possible at w , i.e., only if v and w correspond to alternative strategy profiles of the same game (Constraint **C5**). Furthermore, we suppose that every agent has a complete preference ordering over the strategy profiles of the current game (Constraint **C6**).

REMARK. Note that in the case of complete information (Constraint **C4**) the relation \sim is superfluous because all other relations are included into \sim (Constraints **C5** and **C6**). So in this case we can suppose \sim to be the universal relation $W \times W$ and \square to be the well-known universal modality.³ We decided to

³ The universal modality has been used since the dawning of modal logic (it is just a plain old S5 operator [16]). More recently, it has been used in several modal logic analysis of preferences and games (see, e.g., [14,17,18]).

introduce the relation \sim in this part of the paper in order to be able to generalize the definition of model in the case of incomplete information (see Section 5). In fact, in the case of games with incomplete information, a player can imagine alternative games and there is no one-to-one correspondence between models and games (*i.e.*, every model does not necessarily correspond to a unique strategic game). Therefore, the relation \sim can no longer be supposed to be the universal relation.

Definition 2 ($\mathcal{ML}\mathcal{EG}$ -models). $\mathcal{ML}\mathcal{EG}$ -models are couples $F = \langle F, \pi \rangle$ where:

- F is a $\mathcal{ML}\mathcal{EG}$ -frame;
- $\pi : Atm \rightarrow 2^W$ is a valuation function.

The truth conditions for Boolean operators and for operators $[\delta_C]$, \Box , K_i and $[\text{good}]_i$ are:

- $M, w \models p$ iff $w \in \pi(p)$;
- $M, w \models \neg\varphi$ iff not $M, w \models \varphi$;
- $M, w \models \varphi \vee \psi$ iff $M, w \models \varphi$ or $M, w \models \psi$;
- $M, w \models [\delta_C]\varphi$ iff if $R_C(w) = \delta_C$ then $M, w \models \varphi$;
- $M, w \models \Box\varphi$ iff $M, v \models \varphi$ for all v such that $w \sim v$;
- $M, w \models K_i\varphi$ iff $M, v \models \varphi$ for all v such that wE_iv ;
- $M, w \models [\text{good}]_i\varphi$ iff $M, v \models \varphi$ for all v such that $w \preceq_i v$.

A formula φ is *true in an $\mathcal{ML}\mathcal{EG}$ -model M* iff $M, w \models \varphi$ for every world w in M . A formula φ is *$\mathcal{ML}\mathcal{EG}$ -valid* (noted $\models \varphi$) iff φ is true in all $\mathcal{ML}\mathcal{EG}$ -models. A formula φ is *$\mathcal{ML}\mathcal{EG}$ -satisfiable* iff $\neg\varphi$ is not $\mathcal{ML}\mathcal{EG}$ -valid.

2.3. Axiomatization and Complexity Results

We call $\mathcal{ML}\mathcal{EG}$ the logic that is axiomatized by the principles given in Figure 2.

Note that the principles of modal logic S5 for the operator \Box are: the four axiom schemas (K) $(\Box\varphi \wedge \Box(\varphi \rightarrow \psi)) \rightarrow \Box\psi$, (T) $\Box\varphi \rightarrow \varphi$, (4) $\Box\varphi \rightarrow \Box\Box\varphi$, (B) $\varphi \rightarrow \Box\Diamond\varphi$, and the rule of inference (Necessitation) $\frac{\varphi}{\Box\varphi}$. The principles of modal logic S5 for the operators K_i are: the four axiom schemas (K) $(K_i\varphi \wedge K_i(\varphi \rightarrow \psi)) \rightarrow K_i\psi$, (T) $K_i\varphi \rightarrow \varphi$, (4) $K_i\varphi \rightarrow K_iK_i\varphi$, (B) $\varphi \rightarrow K_i\widehat{K}_i\varphi$, and the rule of inference (Necessitation) $\frac{\varphi}{K_i\varphi}$. The principles of modal logic S4 for the operators $[\text{good}]_i$ are: the three axiom schemas (K) $([\text{good}]_i\varphi \wedge [\text{good}]_i(\varphi \rightarrow \psi)) \rightarrow [\text{good}]_i\psi$, (T) $[\text{good}]_i\varphi \rightarrow \varphi$, (4) $[\text{good}]_i\varphi \rightarrow [\text{good}]_i[\text{good}]_i\varphi$, and the rule of inference (Necessitation) $\frac{\varphi}{[\text{good}]_i\varphi}$.

Note also that Axiom **Indep** is the $\mathcal{ML}\mathcal{EG}$ counterpart of the so-called *axiom of independence of agents* of STIT logic (the logic of *Seeing to it that*) [19]. This axiom enables to express the basic game theoretic assumption that the set of strategy profiles of a game in strategic form is the cartesian product of the sets of individual actions for the agents in *Agt*.

Figure 2. Axiomatization of $\mathcal{ML}\mathcal{EG}$

(CPL)	All principles of classical propositional logic
(S5 $_{\square}$)	All principles of modal logic S5 for \square
(S5 $_{K_i}$)	All principles of modal logic S5 for every K_i
(S4 $_{[\text{good}]_i}$)	All principles of modal logic S4 for every $[\text{good}]_i$
(Def $_{[\delta_C]}$)	$[\delta_C] \varphi \leftrightarrow (\langle \delta_C \rangle \top \rightarrow \varphi)$
(JointAct)	$\langle \delta_C \rangle \top \leftrightarrow \bigwedge_{i \in C} \langle \delta_i \rangle \top$
(Active)	$\bigvee_{\delta_C \in \Delta_C} \langle \delta_C \rangle \top$
(Single)	$\langle \delta_C \rangle \top \rightarrow [\delta'_C] \perp$ if $\delta_C \neq \delta'_C$
(Indep)	$\left(\bigwedge_{i \in \text{Agt}} \diamond \langle \delta_i \rangle \top \right) \rightarrow \diamond \langle \delta \rangle \top$
(Aware)	$\langle i:a \rangle \top \rightarrow K_i \langle i:a \rangle \top$
(Incl $_{[\text{good}]_i, \square}$)	$\square \varphi \rightarrow [\text{good}]_i \varphi$
(PrefConnect)	$(\diamond \varphi \wedge \diamond \psi) \rightarrow (\diamond(\varphi \wedge \langle \text{good} \rangle_i \psi) \vee \diamond(\psi \wedge \langle \text{good} \rangle_i \varphi))$
(CompleteInfo)	$\square \varphi \rightarrow K_i \varphi$
(ModusPonens)	$\frac{\varphi, \varphi \rightarrow \psi}{\psi}$

We write $\vdash_{\mathcal{ML}\mathcal{EG}} \varphi$ if φ is a theorem of $\mathcal{ML}\mathcal{EG}$, that is, if φ can be deduced by applying the axioms and the rules of inference of the logic $\mathcal{ML}\mathcal{EG}$.

As the following Theorem 1 highlights, we can prove that the logic $\mathcal{ML}\mathcal{EG}$ is sound and complete with respect to the class of $\mathcal{ML}\mathcal{EG}$ -models.

Theorem 1. *$\mathcal{ML}\mathcal{EG}$ is determined by the class of $\mathcal{ML}\mathcal{EG}$ -models.*

Moreover we can prove a result about complexity of the satisfiability problem of the logic $\mathcal{ML}\mathcal{EG}$, that is, the complexity of the problem of deciding whether a given $\mathcal{ML}\mathcal{EG}$ formula φ is $\mathcal{ML}\mathcal{EG}$ -satisfiable or not. Here, we give a lower-bound and upper-bound for the complexity of the satisfiability problem (for more information about complexity theory, the reader may refer to [20]):

Theorem 2. *If the number of agents is greater or equal to 2, the satisfiability problem of $\mathcal{ML}\mathcal{EG}$ is EXPTIME-hard and in NEXPTIME.*

We conjecture that the satisfiability problem of $\mathcal{ML}\mathcal{EG}$ is EXPTIME-complete. Indeed, we think that this can be proved by the argument used for proving that the satisfiability problem of S5 $_2$ with common knowledge [21] and the satisfiability problem of propositional dynamic logic (PDL) [22] are EXPTIME-complete.

2.4. A Variant of with Joint Determinism

We present here a variant of $\mathcal{ML}\mathcal{E}\mathcal{G}$ where models have an additional constraint of determinism for the joint actions of all agents: different worlds in an equivalence class $\sim(w)$ correspond to the occurrences of *different* strategy profiles:

CD if $w \sim v$ and $R_{Agt}(w) = \delta$ and $R_{Agt}(v) = \delta$, then $w = v$;

A $\mathcal{ML}\mathcal{E}\mathcal{G}$ -model satisfying the constraint **CD** is called a $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -model.

The axiom corresponding to the Constraint **CD** is:

(JointDet)
$$(\langle \delta \rangle \top \wedge \varphi) \rightarrow \Box(\langle \delta \rangle \top \rightarrow \varphi)$$

We call $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ the logic that is axiomatized by the principles given in Figure 2 plus axiom **JointDet**.

Theorem 3. $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ is determined by the class of $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models.

Although the Constraint **CD** excludes pure uncertainty about uncertainty (*i.e.*, cases where the same profile is played at two states on the basis of different information), it is interesting because it allows to establish a connection between our logical framework and Coalition Logic (CL) [23,24], where it is assumed that if every agent in Agt opts for an action the next state of the world is uniquely determined. As shown in [25], the logic $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ extended by the operator *next* of linear temporal logic (LTL) embeds Coalition Logic (CL) [23,24]. In particular, if we extend $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ by the temporal operator X (where $X\varphi$ means “ φ will be true in the next state”) CL cooperation modalities of the form $[C]$ can be reconstructed in our logic $\mathcal{ML}\mathcal{E}\mathcal{G}$ as follows.

$$tr([C]\varphi) = \bigvee_{\delta \in \Delta} (\langle \delta_C \rangle \top \wedge \Box(\langle \delta_C \rangle \top \rightarrow X\varphi))$$

That is, the CL expression “coalition C can enforce an outcome state satisfying φ ” (noted $[C]\varphi$) is translated in our logic as “there exists a joint action δ_C of the agents in C such that the agents in C can perform δ_C , and necessarily if the agents in C perform δ_C then φ will be true in the next state, no matter what the agents outside C do”.

REMARK. It is worth noting that, while $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ embeds Coalition Logic, the basic logic $\mathcal{ML}\mathcal{E}\mathcal{G}$ embeds Chellas’ STIT logic with agents and groups [26], under the hypothesis that the number of agents’ choices is bounded (see [27] for more details). In fact, differently from Coalition Logic, in STIT joint actions of all agents are not necessarily deterministic. STIT logic has formulas of the form $[C \text{ cstit}: \varphi]$ that are read “group C sees to it that φ ”. The translation of STIT modalities of the form $[C \text{ cstit}:]$ into $\mathcal{ML}\mathcal{E}\mathcal{G}$ would be the following:

$$tr([C \text{ cstit}: \varphi]) = \bigvee_{\delta \in \Delta} (\langle \delta_C \rangle \top \wedge \Box(\langle \delta_C \rangle \top \rightarrow \varphi))$$

That is, the STIT expression “group C sees to it that φ ” is translated into \mathcal{DLA} as “there exists a joint action δ_C of the agents in C such that the agents in C perform δ_C , and necessarily if the agents in C perform δ_C then φ will be true, no matter what the agents outside C do”.

The constraint of joint determinism **CD** is also useful for complexity reasons. Indeed, if we add **CD** to our logic, the complexity of the satisfiability problem drops to NP.

Theorem 4. *The satisfiability problem of $\mathcal{ML}\mathcal{EG}^{det}$ is NP-complete.*

3. A Logical Account of Epistemic Games

This section is devoted to the analysis in the modal logic $\mathcal{ML}\mathcal{EG}$ of the epistemic aspects of strategic games. We first consider the basic game-theoretic concepts of best response and Nash equilibrium, and their relationships with the notion of epistemic rationality assumed in classical game theory. Finally, we provide an analysis of Iterated Deletion of Strictly Dominated Strategies (IDSDS).

3.1. Best Response and Nash Equilibrium

The modal operators $[good]_i$ and \Box enable to capture in $\mathcal{ML}\mathcal{EG}$ a notion of comparative goodness over formulas of the kind “ φ is for agent i at least as good as ψ ”, noted $\psi \leq_i \varphi$:

$$\psi \leq_i \varphi \stackrel{\text{def}}{=} \Box (\psi \rightarrow \langle good \rangle_i \varphi)$$

According to the previous definition, φ is for agent i at least as good as ψ if and only if, for every world v corresponding to a strategy profile of the current game in which ψ is true, there is a world u corresponding to a strategy profile of the current game in which φ is true and which is for agent i at least as good as world v . We can prove that $\psi \leq_i \varphi$ is a total preorder. Indeed, the formulas $\psi \leq_i \psi$ (reflexivity), $(\varphi_1 \leq_i \varphi_2) \wedge (\varphi_2 \leq_i \varphi_3) \rightarrow (\varphi_1 \leq_i \varphi_3)$ (transitivity) and $(\varphi_1 \leq_i \varphi_2) \vee (\varphi_2 \leq_i \varphi_1)$ (connectedness, also called completeness) are valid in $\mathcal{ML}\mathcal{EG}$. We define the corresponding strict ordering over formulas:

$$\psi <_i \varphi \stackrel{\text{def}}{=} (\psi \leq_i \varphi) \wedge \neg(\varphi \leq_i \psi)$$

Formula $\psi <_i \varphi$ has to read “ φ is for agent i strictly better than ψ ”. Finally, we define a notion of comparative goodness over strategy profiles and the corresponding strict ordering over strategy profiles:

$$\delta \leq_i \delta' \stackrel{\text{def}}{=} \langle \delta \rangle \top \leq_i \langle \delta' \rangle \top \text{ and } \delta <_i \delta' \stackrel{\text{def}}{=} (\delta \leq_i \delta') \wedge \neg(\delta' \leq_i \delta)$$

Formula $\delta \leq_i \delta'$ has to be read “strategy profile δ' is for agent i at least as good as strategy profile δ ” and formula $\delta <_i \delta'$ has to be read “strategy profile δ' is for agent i strictly better than strategy profile δ ”.

Some basic concepts of game theory can be expressed in $\mathcal{ML}\mathcal{EG}$ in terms of comparative goodness. We first consider *best response*. Agent i 's action a is said to be a best response to the other agents' joint action δ_{-i} , noted $\text{BR}(i:a, \delta_{-i})$, if and only if i cannot improve his utility by deciding to do something different from a while the others choose the joint action δ_{-i} , that is:

$$\text{BR}(i:a, \delta_{-i}) \stackrel{\text{def}}{=} \bigwedge_{b \in \text{Act}} (\langle i:b, \delta_{-i} \rangle \leq_i \langle i:a, \delta_{-i} \rangle)$$

REMARK. Note that the definition of best response $\text{BR}(i:a, \delta_{-i})$ given above only works for a complete preference relation \preceq_i . To see why, suppose that $\text{Agt} = \{i, j\}$ and $\text{Act} = \{a, b\}$, and consider the model $\mathcal{M} = \langle W, \sim, R, E, \preceq, \pi \rangle$ such that $W = \{w_1, w_2\}$, $\sim = W \times W$, $R_{\text{Agt}}(w_1) = \langle i:a, j:a \rangle$, $R_{\text{Agt}}(w_2) = \langle i:b, j:a \rangle$ and $E_i = E_j = \preceq_i = \preceq_j = \{(w_1, w_1), (w_2, w_2)\}$. Here the relation \preceq_i does not satisfy the constraint **C6** (i.e., \preceq_i is not complete). Intuitively, $\text{BR}(i:a, j:a)$ should be true at any world of \mathcal{M} because if i plays a while j plays a , he does not improve his utility by playing b . Nevertheless, as

$\langle i:b,j:a \rangle \leq_i \langle i:a,j:a \rangle$ is false at any world of \mathcal{M} , we have that $\text{BR}(i:a,j:a)$ is also false at every world of \mathcal{M} .

Given a certain strategic game, the strategy profile (or joint action) δ is said to be a *Nash equilibrium* if and only if for every agent $i \in \text{Agt}$, i 's action δ_i is a best response to the other agents' joint action δ_{-i} :

$$\text{Nash}(\delta) \stackrel{\text{def}}{=} \bigwedge_{i \in \text{Agt}} \text{BR}(\delta_i, \delta_{-i})$$

From Axiom **CompleteInfo**, S5 for \square , Axioms K and T for K_i the five $\mathcal{ML}\mathcal{E}\mathcal{G}$ -theorems in Lemma 1 are provable. They express complete information about the players' preferences ordering over strategy profiles, complete information about the existence of a Nash equilibrium, and complete information about the players' action repertoires. Surprisingly formulas of Lemma 1 are kinds of introspection properties, but they are provable without axioms of positive and negative introspections for knowledge: (4) $K_i\varphi \rightarrow K_iK_i\varphi$ or (5) $\neg K_i\varphi \rightarrow K_i\neg K_i\varphi$.

Proposition 1. For all $i \in \text{Agt}$ and $n \in \mathbb{N}$:

- $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} (\psi \leq_i \varphi) \leftrightarrow \text{MK}_{\text{Agt}}^n(\psi \leq_i \varphi)$
- $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} (\psi <_i \varphi) \leftrightarrow \text{MK}_{\text{Agt}}^n(\psi <_i \varphi)$
- $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Nash}(\delta) \leftrightarrow \text{MK}_{\text{Agt}}^n \text{Nash}(\delta)$
- $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \diamond \langle \delta_i \rangle \top \leftrightarrow \text{MK}_{\text{Agt}}^n \diamond \langle \delta_i \rangle \top$
- $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \square [\delta_i] \perp \leftrightarrow \text{MK}_{\text{Agt}}^n \square [\delta_i] \perp$

It has to be noted that weak preference operators $[\text{good}]_i$, used here to define some basic concepts of game theory, have been studied before by [14,17] and by [18], where complete axiomatizations for different kinds of preference logics and for a combination of preference logic with epistemic logic are given. In [14,17] Liu proposed a complete modal logic of knowledge and preference extended by dynamic operators of knowledge update and preference upgrade in the style of dynamic epistemic logic (DEL). In [18] van Benthem *et al.* studied different variants of preference logic that allow to express different readings of *ceteris paribus* preferences [28]. They first present a basic modal logic of weak and strict preference which allows to express the “all other things being normal” reading of *ceteris paribus* preferences. Then they present a more general modal logic in which modal operators of weak and strict preference are relativized to sets of formulas representing conditions to be kept *equal*. They show that this logic allows to express the “all other things being equal” reading of *ceteris paribus* preferences and to characterize the notion of Nash equilibrium as a preference for a given strategy profile for a game, given that *others* keep the same strategy.⁴ One of the main contribution of our work is to propose a modal logic which integrates the notion of weak preference studied by Liu and van Benthem *et al.* with notions of action and knowledge, and which provides a suitable framework for a logical analysis of epistemic strategic games both with perfect information and with weaker forms of perfect information. The latter are the subject of the second part of the paper (Sections 5 and 6).

⁴The *normality* reading of *ceteris paribus* preferences expresses preferences which hold under certain normal conditions, whereas the *equality* reading expresses preferences which hold when certain facts are kept constant.

3.2. Epistemic Rationality

The following $\mathcal{ML}\mathcal{EG}$ formula characterizes a notion of rationality which is commonly supposed in the epistemic analysis of games (see, e.g., [4,7]):

$$\bigwedge_{a,b \in Act} \left(\langle i:a \rangle \top \rightarrow \bigvee_{\delta \in \Delta} \left(\widehat{K}_i \langle \delta_{-i} \rangle \top \wedge (\langle \delta_{-i}, i:b \rangle \leq_i \langle \delta_{-i}, i:a \rangle) \right) \right)$$

This means that an agent i is rational if and only if, if he chooses a particular action a then for every alternative action b , there exists a joint action δ_{-i} of the other agents that he considers possible such that, playing a while the others play δ_{-i} is for i at least as good as playing b while the others play δ_{-i} . This means that epistemic rationality simply consists in not choosing a strategy that is strictly dominated within the agent's set of epistemic alternatives.

As formula $\delta \leq_i \delta'$ and formula $K_i(\delta \leq_i \delta')$ are equivalent in $\mathcal{ML}\mathcal{EG}$, the previous definition of rationality can be rewritten in the following equivalent form:

$$\text{Rat}_i \stackrel{\text{def}}{=} \bigwedge_{a,b \in Act} \left(\langle i:a \rangle \top \rightarrow \bigvee_{\delta \in \Delta} \left(\widehat{K}_i \langle \delta_{-i} \rangle \top \wedge K_i (\langle \delta_{-i}, i:b \rangle \leq_i \langle \delta_{-i}, i:a \rangle) \right) \right)$$

Theorem 5. For all $i \in \text{Agt}$:

$$(5a) \quad \vdash_{\mathcal{ML}\mathcal{EG}} \text{Rat}_i \leftrightarrow K_i \text{Rat}_i$$

$$(5b) \quad \vdash_{\mathcal{ML}\mathcal{EG}} \neg \text{Rat}_i \leftrightarrow K_i \neg \text{Rat}_i$$

Theorem 5 highlights that the concepts of rationality and irrationality are introspective. That is, an agent i is (resp. is not) epistemically rational if and only if he knows this. The syntactic proof of Theorem 5a given in the annex shows that it can be proved either by means of Axioms K, T, 4 and 5 for knowledge or by means of Axioms K, T and 5 for knowledge and a principle of introspection over preferences of the form “ $(\delta \leq_i \delta') \rightarrow K_i(\delta \leq_i \delta')$ ”. Theorem 5b is provable from Theorem 5a by means of Axioms T and 5 for knowledge.

The following theorem specifies some sufficient epistemic conditions of Nash equilibrium: if all agents are rational and every agent knows the choices of the other agents, then the selected strategy profile is a Nash equilibrium. This theorem has been stated for the first time by Aumann & Brandeburger [1,5].

Theorem 6. For all $\delta \in \Delta$:

$$\vdash_{\mathcal{ML}\mathcal{EG}} \left(\left(\bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \wedge \bigwedge_{i \in \text{Agt}} K_i \langle \delta_{-i} \rangle \top \right) \rightarrow \text{Nash}(\delta)$$

The syntactic proof of Theorem 6 in the annex at the end of the paper shows that it can be proved just by means of Axioms K and T for the epistemic modal operators. Axiom **CompleteInfo**, positive and negative introspection for knowledge ($K_i\varphi \rightarrow K_i K_i\varphi$ and $\neg K_i\varphi \rightarrow K_i \neg K_i\varphi$) are not needed for the proof.

3.3. Iterated Deletion of Strictly Dominated Strategies

A strategy a for agent i is a *strictly dominated strategy*, noted $\text{SD}^{\leq 0}(i:a)$, if and only if, if a can be performed then there is another strategy b such that, no matter what joint action δ_{-i} the other agents

choose, playing b is for i strictly better than playing a :

$$SD^{\leq 0}(i:a) \stackrel{\text{def}}{=} \diamond \langle i:a \rangle \top \rightarrow \bigvee_{b \in Act} \left(\diamond \langle i:b \rangle \top \wedge \bigwedge_{\delta \in \Delta} (\diamond \langle \delta_{-i} \rangle \top \rightarrow (\langle \delta_{-i}, i:a \rangle <_i \langle \delta_{-i}, i:b \rangle)) \right)$$

An example of strictly dominated strategy is cooperation in the Prisoner Dilemma (PD) game: whether one's opponent chooses to cooperate or defect, defection yields a higher payoff than cooperation. Therefore, a rational player will never play a dominated strategy. So when trying to predict the behavior of rational players, we can rule out all strictly dominated strategies. The so-called Iterated Deletion of Strictly Dominated Strategies (IDSDS) (or iterated strict dominance) [15] is a procedure that starts with the original game and, at each step, for every player i removes from the game all i 's strictly dominated strategies, thereby generating a subgame of the original game, and that repeats this process again and again. IDSDS can be inductively characterized in our logic $\mathcal{ML}\mathcal{EG}$ by defining a concept of strict dominance in the subgame of depth at most n , noted $SD^{\leq n}(i:a)$. For every $n \geq 1$:

$$SD^{\leq n}(i:a) \stackrel{\text{def}}{=} \neg SD^{\leq n-1}(i:a) \rightarrow \bigvee_{b \in Act} \left(\neg SD^{\leq n-1}(i:b) \wedge \bigwedge_{\delta \in \Delta} (\neg SD^{\leq n-1}(\delta_{-i}) \rightarrow (\langle \delta_{-i}, i:a \rangle <_i \langle \delta_{-i}, i:b \rangle)) \right)$$

where $SD^{\leq k}(\delta_C)$ is defined as follows

$$SD^{\leq k}(\delta_C) \stackrel{\text{def}}{=} \bigvee_{i \in C} SD^{\leq k}(\delta_i)$$

for every $k \geq 0$ and for every δ_C . According to this definition, a is a strictly dominated strategy for agent i in a subgame of depth at most n , noted $SD^{\leq n}(i:a)$, if and only if, if a is not strictly dominated for i in all subgames of depth $k < n$ then there is another strategy b such that b is not strictly dominated for i in all subgames of depth $k < n$ and, no matter what joint action δ_{-i} the other agents choose, if the elements in δ_{-i} are not dominated in all subgames of depth $k < n$ then playing b is for i strictly better than playing a . In other terms $SD^{\leq n}(i:a)$ means that strategy $i:a$ does not survive after n rounds of IDSDS. We can prove by recurrence on n that the length of the formula $SD^{\leq n}(\delta)$ is

$$O(|Act||Agt|^{2n+1})$$

where $O(\dots)$ is the ‘‘Big Oh Notation’’ [20], $|Act|$ is the number of actions and $|Agt|$ is the number of agents and n is the number of rounds of IDSDS. That is, the length of the formula $SD^{\leq n}(\delta)$ is exponential in n . In Section 4, we are going to extend the language of $\mathcal{ML}\mathcal{EG}$ in order to capture the concept of IDSDS with a more compact formula.

As the following $\mathcal{ML}\mathcal{EG}$ -theorems highlight, the truth of $SD^{\leq n}(i:a)$ depends on the game but does not depend on the world where the formula is evaluated.

Proposition 2. For all $a \in Act$, for all $n \geq 0$, we have:

- $\vdash_{\mathcal{ML}\mathcal{EG}} SD^{\leq n}(i:a) \leftrightarrow \Box SD^{\leq n}(i:a)$;
- $\vdash_{\mathcal{ML}\mathcal{EG}} \neg SD^{\leq n}(i:a) \leftrightarrow \Box \neg SD^{\leq n}(i:a)$.

The following Theorem 7 is the qualitative version of a probabilistic- based result of Stalnaker [29] who has been the first to use probabilistic Kripke structures in order to characterize the IDSDS procedure in terms of common knowledge of rationality (see [3,30] for some recent discussion of Stalnaker’s results). A similar result was also proved, with differing degrees of formality, by Bernheim [31], Pearce [32], Brandenburger & Dekel [33], and Tan & Werlang [34]. Note that Stalnaker’s proof is purely semantic. According to the Theorem 7, if there is mutual knowledge of rationality among the players to n levels and the agents play the strategy profile δ then, for every agent i , δ_i survives IDSDS until the subgame of depth $n+1$.

Theorem 7. For all positive integer n , for all $\delta \in \Delta$, we have:

$$\vdash_{\mathcal{ML}\mathcal{EG}} \left(\left(\text{MK}_{\text{Agt}}^n \bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \wedge \langle \delta \rangle \top \right) \rightarrow \neg \text{SD}^{\leq n}(\delta)$$

(note that $\neg \text{SD}^{\leq n}(\delta)$ is just the abbreviation of $\bigwedge_{i \in \text{Agt}} \neg \text{SD}^{\leq k}(\delta_i)$).

The syntactic proof of Theorem 7 given in the annex shows that, although Axiom **CompleteInfo**, positive and negative introspection for knowledge ($\text{K}_i\varphi \rightarrow \text{K}_i\text{K}_i\varphi$ and $\neg\text{K}_i\varphi \rightarrow \text{K}_i\neg\text{K}_i\varphi$) are not needed for the proof, we need to assume that a player has complete information about the players’ strategy sets as well as about the players’ preference ordering over strategy profiles.

Table 1 summarizes the sufficient conditions for the syntactic proof of Theorem 7 together with the sufficient conditions for the syntactic proofs of Theorems 5 and 6. It highlights an interesting aspect of our syntactic analysis of games based on modal logic: the fact that we can easily verify whether certain assumptions about knowledge and information over the game structure are indeed necessary to prove results concerning the epistemic foundations of game theory.

It has to be noted that Theorem 7 provides only one direction of the characterization result for the IDSDS procedure as formulated in the game-theoretic literature, according to which IDSDS is fully characterized by the epistemic condition of common knowledge of rationality between the players (see, e.g., [3,29,30]).

The other direction states approximately that, for every strategic game, if δ is the strategy profile that is chosen and that survives to the infinite procedure IDSDS, then *there is (a state in) an epistemic model* for that game in which the profile δ is played and the players have common knowledge of rationality.

This statement is formally expressed by the following theorem (a similar result is proved in [3]).

Theorem 8. Consider an arbitrary $\mathcal{ML}\mathcal{EG}$ -model $M = \langle W, \sim, R, E, \preceq, \pi \rangle$, a world w in M and $\delta \in \Delta$ such that for all positive integers n we have $M, w \models \langle \delta \rangle \top \wedge \neg \text{SD}^{\leq n}(\delta)$. Then, there is a model $M' = \langle W, \sim, R, E', \preceq, \pi \rangle$ such that for all positive integers n we have $M', w \models \text{MK}_{\text{Agt}}^n \bigwedge_{i \in \text{Agt}} \text{Rat}_i$.

The idea is that, for every strategic game, if δ is the strategy profile of this game which is chosen and that survives after n rounds of the procedure IDSDS, for all positive integers n , then it is possible to find an “epistemic configuration” for the players which satisfies common knowledge of rationality between the players. In other words, if δ is a strategy profile of a given strategic game that survives after all rounds of the procedure IDSDS then it is always possible to justify the choice of δ by the fact that the players have common knowledge of rationality.

Table 1. Some sufficient conditions for Theorems 5, 6, 7.

Assumptions about knowledge operators	Assumptions about information over game structure	Result
KT45	none	Theorem 5
KT5	introspection over preferences: $(\delta \leq_i \delta') \rightarrow K_i(\delta \leq_i \delta')$	Theorem 5
KT	none	Theorem 6
KT	complete information about players' preference ordering over strategy profiles: $(\delta \leq_j \delta') \rightarrow K_i(\delta \leq_j \delta')$ $(\delta <_j \delta') \rightarrow K_i(\delta <_j \delta')$ complete information about players' strategy sets: $\Diamond\langle j:a \rangle \top \rightarrow K_i \Diamond\langle j:a \rangle \top$ $\Box[j:a] \perp \rightarrow K_i \Box[j:a] \perp$	Theorem 7

Note that Theorem 8 is different from stating that the logic $\mathcal{ML}\mathcal{E}\mathcal{G}$ cannot prove the negation of conjunctions of this type: a given profile δ is played, survives after n rounds of IDSDS and rationality of players is common knowledge up to degree n . By the completeness of $\mathcal{ML}\mathcal{E}\mathcal{G}$ this just means that the conjunction $(MK_{Agt}^n \bigwedge_{i \in Agt} Rat_i) \wedge \langle \delta \rangle \top \wedge \neg SD^{\leq n}(\delta)$ is consistent in $\mathcal{ML}\mathcal{E}\mathcal{G}$. Indeed, the latter trivially holds, as we can always exhibit the trivial model $M = \langle W, \sim, R, E, \preceq, \pi \rangle$ such that $W = \{w\}$, $R_{Agt}(w) = \delta$ and $E'_i(w) = \{w\}$ for every $i \in Agt$, and in which $M, w \models (MK_{Agt}^n \bigwedge_{i \in Agt} Rat_i) \wedge \langle \delta \rangle \top \wedge \neg SD^{\leq n}(\delta)$ holds for every n .

3.4. Discussion: Related Works on Modal Logic Analysis of Epistemic Games under Complete Information

Although several modal logics of games in strategic forms have been proposed in recent times (see, e.g., [25,35]), few modal logics of epistemic games under complete information exist. Among them we should mention [3,8,9,36]. Let us compare our modal logic $\mathcal{ML}\mathcal{E}\mathcal{G}$ with some of these alternative approaches.

De Bruin [8] has developed a logical framework which enables to reason about the epistemic aspects of strategic games and of extensive games. His system deals with several game-theoretic concepts like the concepts of knowledge, rationality, Nash equilibrium, iterated strict dominance, backward induction. Nevertheless, de Bruin's approach differs from ours in several respects. First of all, our logical approach to epistemic games is *minimalistic* since it relies on few primitive concepts: knowledge, action, historical necessity and preference. All other notions such Nash equilibrium, rationality, iterated strict dominance are defined by means of these four primitive concepts. On the contrary, in de Bruin's logic all those notions are atomic propositions managed by a *ad hoc* axiomatization (see, e.g., ([8], pp. 61,65) where

special propositions for rationality and iterated strict dominance are introduced). Secondly, we provide a semantics and a complete axiomatization for our logic of epistemic games. De Bruin's approach is purely syntactic: no model-theoretic analysis of games is proposed nor completeness result for the proposed logic is given. Finally, de Bruin does not provide any complexity results for his logic while we provide complexity results for the satisfiability problem of our logic.

In [36] van der Hoek & Pauly investigate how modal logics can be used to describe and reason about games. They show how epistemic logic can be combined with constructions expressing agents' preferences over strategy profiles in order to study the epistemic aspects of strategic games and to define a concept of rationality similar to the one discussed in Section 3.2. Although van der Hoek & Pauly discuss the combination of action, preference and epistemics for the analysis of epistemic games they do not provide a unified modal logic framework combining operators for knowledge, for preference and for action with a complete axiomatization and with a study of its computational properties like decidability and complexity. The latter is one of the main contribution of our work.

Roy [9] has recently proposed a modal logic integrating preference, knowledge and intention. In his approach every world in a model is associated to a nominal which directly refers to a strategy profile in a strategic game. This approach is however limited in expressing formally the structure of a strategic game. In particular, in Roy's logic there is no principle like the $\mathcal{ML}\mathcal{EG}$ Axiom **Indep** explaining how possible actions δ_i of individual agents are combined to form a strategy profile δ of the current game. Another limitation of Roy's approach is that it does not allow to express the concept of (weak) rationality that we have been able to define in Section 3.2 (see [9], pp. 101). As discussed in the previous sections this is a crucial concept in interactive epistemology since it is used for giving epistemic justifications of several solution concepts like Nash equilibrium and IDSDS (see Theorems 6 and 7).

Bonanno [3] combines modal operators for belief and common belief with constructions expressing agents' preferences over individual actions and strategy profiles, and applies them to the semantic characterization of solution concepts like Iterated Deletion of Strictly Dominated Strategies (IDSDS) and Iterated Deletion of Inferior Profiles (IDIP). As in [9], in Bonanno's logic every world in a model corresponds to a strategy profile of the current game. Although this logic allows to express the concept of weak rationality, it is not sufficiently general to enable to express in the object language solution concepts like Nash equilibrium and IDSDS (note that the latter is defined by Bonanno only in the metalanguage).

It is to be noted that, differently from $\mathcal{ML}\mathcal{EG}$, most modal logics of epistemic games in strategic form (including Roy's logic and Bonanno's logic) postulate a one-to-one correspondence between models and games (i.e. every model of the logic corresponds to a unique strategic game, and worlds in the model are all strategy profiles of this game). Such an assumption is quite restrictive since it prevents from analyzing in the logic games with incomplete information about the game structure in which an agent can imagine alternative games. We will show in Section 5 that this is something we can do in our logical framework by removing Axiom **CompleteInfo** from $\mathcal{ML}\mathcal{EG}$.

4. Game Transformation

We provide in this section an alternative and more compact characterization of the procedure IDSDS in our logic $\mathcal{ML}\mathcal{EG}$. To this aim, we introduce special events whose effect is to transform the current game by removing certain strategies from it. In particular, these special events can be used to delete

a strictly dominated strategy from the current game. These special events are similar to the notion of announcement in Dynamic Epistemic Logic (DEL) [10–12].

$\mathcal{L}_{\mathcal{GT}}$ is the set of *game transformation formulas* and is defined by the following rule:

$$\chi ::= \Box\psi \rightarrow [i:a] \perp \mid \chi \wedge \chi$$

where $\psi \in \mathcal{L}_{\mathcal{MLEG}}$, $i \in \text{Agt}$ and $a \in \text{Act}$. Thus, game transformation formulas are of the form ‘if property ψ necessarily holds in the current game, then action a should not be performed by agent i ’.

\mathcal{GT} is the set of *game transformation events* and is defined as $\mathcal{GT} = \{\chi! \mid \chi \in \mathcal{L}_{\mathcal{GT}}\}$.

We extend the \mathcal{MLEG} language with dynamic operators of the form $[\chi!]$ with $\chi! \in \mathcal{GT}$. The formula $[\chi!]\varphi$ has to be read ‘ φ holds, after the occurrence of the game transformation event $\chi!$ ’. We call $\mathcal{MLEG}^{\mathcal{GT}}$ the extended logic. The truth condition for $[\chi!]\varphi$ is:

$$M, w \models [\chi!]\varphi \text{ iff if } M, w \models \chi \text{ then } M^\chi, w \models \varphi$$

with $M^\chi = \langle W^\chi, \sim^\chi, R^\chi, E^\chi, \preceq^\chi, \pi^\chi \rangle$ and:

$$\begin{aligned} W^\chi &= \{u \mid u \in W \text{ and } M, u \models \chi\} \\ \sim^\chi &= \sim \cap (W^\chi \times W^\chi) \\ \text{for every } C \in 2^{\text{Agt}^*}, R_C^\chi &= R_C|_{W^\chi} \\ \text{for every } i \in \text{Agt}, E_i^\chi &= E_i \cap (W^\chi \times W^\chi) \\ \text{for every } i \in \text{Agt}, \preceq_i^\chi &= \preceq_i \cap (W^\chi \times W^\chi) \\ \text{for every } p \in \text{Atm}, \pi^\chi(p) &= \pi(p) \cap W^\chi \end{aligned}$$

Thus, an event $\chi!$ removes from the model M all worlds in which χ is false. Every epistemic relations E_i , every preference orderings \preceq_i , every function R_C , and the valuation π are restricted to the worlds in which χ is true.

In the resulting structure M^χ , the relations \sim^χ , $R_{\delta_C}^\chi$, E_i^χ , \preceq_i^χ verify the constraints of Definition 1. This result is summed up in the following theorem.

Theorem 9. *Let $\chi \in \mathcal{L}_{\mathcal{GT}}$. If M is a \mathcal{MLEG} model then M^χ is a \mathcal{MLEG} model.*

REMARK. The syntactic restriction on game transformation formulas is given in order to ensure that the updated model M^χ is still a \mathcal{MLEG} model. In fact, Theorem 9 does not hold if we allow χ to be any formula in $\mathcal{L}_{\mathcal{MLEG}}$. For instance suppose M is a \mathcal{MLEG} model such that $W = \{w, v, u, z\}$, $\sim(w) = \{w, v, u, z\}$, $R_{\{i,j\}}(w) = \langle i:a, j:a \rangle$, $R_{\{i,j\}}(v) = \langle i:a, j:b \rangle$, $R_{\{i,j\}}(u) = \langle i:b, j:a \rangle$ and $R_{\{i,j\}}(z) = \langle i:b, j:b \rangle$. If $\chi = \langle i:a \rangle \top \vee \langle j:b \rangle \top$ then the updated model M^χ is no longer a \mathcal{MLEG} -model because it does not satisfy the constraint **C2**.

We have reduction axioms for the dynamic operators $[\chi!]$.

Theorem 10. *The following schemata are valid in the logic $\mathcal{ML}\mathcal{EG}^{\mathcal{GT}}$.*

- R1.** $[\chi!]p \leftrightarrow (\chi \rightarrow p)$
- R2.** $[\chi!]\neg\varphi \leftrightarrow (\chi \rightarrow \neg[\chi!]\varphi)$
- R3.** $[\chi!](\varphi_1 \wedge \varphi_2) \leftrightarrow ([\chi!]\varphi_1 \wedge [\chi!]\varphi_2)$
- R4.** $[\chi!]\Box\varphi \leftrightarrow (\chi \rightarrow \Box[\chi!]\varphi)$
- R5.** $[\chi!]\mathbf{K}_i\varphi \leftrightarrow (\chi \rightarrow \mathbf{K}_i[\chi!]\varphi)$
- R6.** $[\chi!][\text{good}]_i\varphi \leftrightarrow (\chi \rightarrow [\text{good}]_i[\chi!]\varphi)$
- R7.** $[\chi!][\delta_C]\varphi \leftrightarrow (\langle\delta_C\rangle\top \rightarrow [\chi!]\varphi)$

The principles **R1.-R7.** are called reduction axioms because, read from left to right, they reduce the complexity of those operators in a formula. In particular the principles **R1.-R7.** explains how to transform any formula φ of the language with dynamic operators in a formula without dynamic operators. More generally, we have an axiomatization result.

Theorem 11. *The logic $\mathcal{ML}\mathcal{EG}^{\mathcal{GT}}$ is completely axiomatized by the axioms and inference rules of $\mathcal{ML}\mathcal{EG}$ together with the schemata of Theorem 10 together with the following rule of replacement of proved equivalence:*

$$\frac{\psi_1 \leftrightarrow \psi_2}{\varphi \leftrightarrow \varphi[\psi_1 := \psi_2]}$$

where $\varphi[\psi_1 := \psi_2]$ is the formula φ in which we have replaced all occurrences of ψ_1 by ψ_2 .

Now, consider the following formula:

$$\chi_{\text{SD}} \stackrel{\text{def}}{=} \bigwedge_{i \in \text{Agt}, a \in \text{Act}} (\Box \text{SD}^{\leq 0}(i:a) \rightarrow [i:a]\perp)$$

where $\text{SD}^{\leq 0}(i:a)$ has been defined in Subsection 3.3. The effect of the game transformation event $\chi_{\text{SD}}!$ is to delete from every game $\sim(w)$ in the model M all worlds in which a strictly dominated strategy is played by some agent.

As the following Theorem 12 highlights, the procedure IDSDS that we have characterized in Section 3.3 in the static $\mathcal{ML}\mathcal{EG}$ can be characterized in a more compact way in $\mathcal{ML}\mathcal{EG}^{\mathcal{GT}}$. Suppose δ is the selected strategy profile. Then, δ survives IDSDS until the subgame of depth $n+1$ if and only if, the event $\chi_{\text{SD}}!$ can occur $n+1$ times in sequence.

Theorem 12. *For all $\delta \in \Delta$, for all $n \geq 0$,*

$$\vdash_{\mathcal{ML}\mathcal{EG}^{\mathcal{GT}}} \langle\delta\rangle\top \rightarrow (\neg \text{SD}^{\leq n}(\delta) \leftrightarrow \langle\chi_{\text{SD}}!\rangle^{n+1}\top).$$

The above theorem says that if δ is performed, then the formula $\neg \text{SD}^{\leq n}(\delta)$, defined in Subsection 3.3, whose length is exponential in n and the more compact formula $\langle\chi_{\text{SD}}!\rangle^{n+1}\top$ are equivalent. Indeed the length of the formula $\langle\chi_{\text{SD}}!\rangle^{n+1}\top$ is $O(n(|\text{Agt}||\text{Act}|)^2)$ where n is the number of IDSDS rounds, $|\text{Agt}|$ is the number of agents and $|\text{Act}|$ is the maximal number of actions.

We conjecture that there is no formula $\varphi \in \mathcal{ML}\mathcal{EG}$ more compact than $\neg \text{SD}^{\leq n}(\delta)$ such that $\vdash_{\mathcal{ML}\mathcal{EG}} \langle\delta\rangle\top \rightarrow (\neg \text{SD}^{\leq n}(\delta) \leftrightarrow \varphi)$. If our conjecture is true, Theorem 12 would imply that the representation

of IDSDS in $\mathcal{MLEG}^{\mathcal{GT}}$ is necessarily more succinct than the representation of IDSDS in \mathcal{MLEG} (i.e., there is no representation of IDSDS in \mathcal{MLEG} that is equally or more succinct than the representation of IDSDS in $\mathcal{MLEG}^{\mathcal{GT}}$). The latter is indeed a variant of the result given in [37] showing that S5 with public announcements is more succinct than S5.

Finally, here is a compact reformulation of Theorem 7 in $\mathcal{MLEG}^{\mathcal{GT}}$:

Theorem 13. For all $n \geq 0$, $\vdash_{\mathcal{MLEG}^{\mathcal{GT}}} \left(\text{MK}_{\text{Agt}}^n \bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \rightarrow \langle \chi_{\text{SD}}! \rangle^{n+1} \top$.

It has to be noted that the approach to game dynamics based on Dynamic Epistemic Logic (DEL) proposed here is inspired by [7] in which strategic equilibrium is defined by fixed-points of operations of repeated announcement of suitable epistemic statements and rationality assertions. However, the analysis of epistemic games proposed in [7] is mainly semantical and the author does not provide a full-fledged modal language for epistemic games which allows to express in the object language solution concepts like Nash Equilibrium or IDSDS, and the concept of rationality. Moreover, van Benthem's analysis does not include any completeness result for the proposed framework and there is no proposal of reduction axioms for a combination of DEL with a static logic of epistemic games. On the contrary, these two aspects are central in our analysis.

5. Incomplete Information

We here consider a more general class of games that includes strategic games with incomplete information about the game structure including the players' strategy sets (or action repertoires) and the players' preference ordering over strategy profiles. This kind of games have been explored in the past by Harsanyi [38]. A more recent analysis is given by [39].

We are interested here in verifying whether the results obtained in Sections 3.2 and 3.3 can be generalized to this kind of games, that is:

1. Are rationality of every player and every agent's knowledge about other agents' choices still sufficient to ensure that the selected strategy profile is a Nash equilibrium in a strategic game with incomplete information about the game structure?
2. Is mutual knowledge of rationality among the players still sufficient to ensure that the selected strategy profile survives iterated deletion of dominated strategies in a strategic game with incomplete information about the game structure?

To answer these questions, we have to remove Axiom **CompleteInfo** of the form $\Box\varphi \rightarrow K_i\varphi$ from \mathcal{MLEG} and the corresponding semantic constraint **C4** from the definition of \mathcal{MLEG} frames expressing the hypothesis of complete information about the game structure. We call \mathcal{MLEG}^* the resulting logic and \mathcal{MLEG}^* -models the resulting class of models. Then we have to check whether Theorems 6 and 7 given in Sections 3.2 and 3.3 are still derivable in \mathcal{MLEG}^* .

We have a positive answer to the previous first question. Indeed, the formula

$$\left(\left(\bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \wedge \bigwedge_{i \in \text{Agt}} K_i \langle \delta_{-i} \rangle \top \right) \rightarrow \text{Nash}(\delta)$$

is derivable in \mathcal{MLEG}^* . But we have a negative answer to the second question. Indeed, the following formula is invalid in \mathcal{MLEG}^* for every $\delta \in \Delta$ and for every $n \in \mathbf{N}$ such that $n > 0$:

agent 2 and are for 2 strictly better than the situations in which agent 1 successfully burgles the apartment (worlds w_2 and w_4).

At world w_1 agent 1 enters by the main door while agent 2 follows the security procedure. This is a world in the model M in which agent 1 has some uncertainty. Indeed, in this world agent 1 can imagine the alternative game defined by the equivalence class $\sim(w_5) = \{w_5, w_6, w_7, w_8\}$ in which he enters by the main door while agent 2 does nothing (world w_6). We suppose that in such a game, even if agent 2 follows the security procedure, agent 1 will succeed in burgling his apartment (*i.e.*, the equivalence class $\sim(w_5) = \{w_5, w_6, w_7, w_8\}$ represents the situation in which the security procedure does not work). This is the reason why the four strategy profiles $\langle 1:main, 2:skip \rangle$, $\langle 1:back, 2:skip \rangle$, $\langle 1:main, 2:proc \rangle$ and $\langle 1:back, 2:proc \rangle$ are equally preferable for the two agents.

Concerning the automated reasoning aspects of the logic $\mathcal{ML}\mathcal{EG}^*$, we conjecture that the complexity of its satisfiability problem is PSPACE. Indeed, we think that it is possible to build a tableau method using only a polynomial amount of memory for the satisfiability problem of $\mathcal{ML}\mathcal{EG}^*$. In other terms, we conjecture that if we move from $\mathcal{ML}\mathcal{EG}$ to $\mathcal{ML}\mathcal{EG}^*$, the complexity decreases from EXPTIME-hard to PSPACE. We here provide an unsurprising lower-bound for the complexity of $\mathcal{ML}\mathcal{EG}^*$.

Theorem 14. *The satisfiability problem of a given formula φ in a $\mathcal{ML}\mathcal{EG}^*$ -model is PSPACE-hard.*

The situation is different when we add the Axiom **JointDet** of joint determinism discussed in Section 2.4. Let us call $\mathcal{ML}\mathcal{EG}^{det*}$ the logic resulting from adding Axiom **JointDet** of joint determinism to the logic $\mathcal{ML}\mathcal{EG}^*$ and $\mathcal{ML}\mathcal{EG}^{det*}$ -models the class of models resulting from adding the corresponding Constraint **CD** to $\mathcal{ML}\mathcal{EG}^*$ -models. While the complexity of the satisfiability problem for $\mathcal{ML}\mathcal{EG}^{det}$ was NP-complete, it increases to PSPACE-complete for $\mathcal{ML}\mathcal{EG}^{det*}$. More precisely:

Theorem 15. • *If $\text{card}(Agt) = 1$ and $\text{card}(Act) = 1$ then the satisfiability problem of a given formula φ in a $\mathcal{ML}\mathcal{EG}^{det*}$ -model is NP-complete.*

- *If $\text{card}(Agt) \geq 2$ or $\text{card}(Act) \geq 2$ then the satisfiability problem of a given formula φ in a $\mathcal{ML}\mathcal{EG}^{det*}$ -model is PSPACE-complete.*

6. Weaker Forms of Complete Information

In the previous section, we have removed Axiom **CompleteInfo** of the form $\Box\varphi \rightarrow K_i\varphi$ from the logic $\mathcal{ML}\mathcal{EG}$ to obtain a new logic $\mathcal{ML}\mathcal{EG}^*$ in which agents may have incomplete information about all aspects of the game they play, including the players' strategy sets (or action repertoires) and the players' preference ordering over strategy profiles.

Nevertheless, in some cases we would like to suppose that agents have complete information about some specific aspects of the game they play. For example, we would like to suppose that:

1. an agent has complete information about his strategy sets even though he may have incomplete information about other agents' strategy sets or,
2. that an agent has complete information about the strategy set of every agent even though he may have incomplete information about agents' preference ordering over strategy profiles.

The former assumption applies to the scenario in which a robber enters a bank, approaches the bank teller and demands money waving a gun. In this situation the bank teller has complete information about his strategy set: he knows that he can either sound the alarm or do nothing. But the bank teller does not know the robber's strategy set, as he is not sure whether the robber's gun is loaded or not (*i.e.*, the bank teller does not know whether the robber is able to kill him by shooting). The latter assumption applies to a card game like Poker. In Poker a player has complete information about every player's strategy set, as he knows that a given point in the game a player has the option to check (if no bet is in front of him), bet, or fold. However, a Poker player has incomplete information about other players' preference ordering over strategy profiles, as he cannot see other players' cards.

In other terms, we would like to build variants of $\mathcal{ML}\mathcal{EG}^*$ in which some formulas of Lemma 1 are derivable.

6.1. Complete Information about Strategy Sets

We show here how to relax the Axiom **CompleteInfo** in order to express the assumption of complete information about strategy sets without necessarily assuming complete information over the payoffs.

If we replace Axiom **CompleteInfo** by the following axiom schemas:

$$(\mathbf{CompleteInfoStrategy}_i) \quad \Diamond\langle i:a \rangle \top \rightarrow K_i \Diamond\langle i:a \rangle \top$$

for all $i \in \mathit{Agt}$ and $a \in \mathit{Act}$, then every agent i has complete information about his strategy set. That is, if an agent i can perform an action a then agent i knows that he can perform action a . Axiom **CompleteInfoStrategy_i** corresponds to the following semantic constraint on models. For every $i \in \mathit{Agt}$ and $a \in \mathit{Act}$:

$$\mathbf{C7} \quad \text{if } wE_i u \text{ and there is } v \text{ such that } w \sim v \text{ and } i:a = R_i(v) \text{ then, there is } z \text{ such that } u \sim z \text{ and } i:a = R_i(z).$$

If we replace Axiom **CompleteInfo** by the following axiom schemas:

$$(\mathbf{CompleteInfoStrategy}_{i,j}) \quad \Diamond\langle j:a \rangle \top \rightarrow K_i \Diamond\langle j:a \rangle \top$$

for all $i, j \in \mathit{Agt}$ and $a \in \mathit{Act}$, then an agent i has complete information about the strategy sets of every agent. That is, if an agent j can perform an action a then every agent i knows that agent j can perform action a . Axiom **CompleteInfoStrategy_{i,j}** corresponds to the following semantic constraint on models. For every $i, j \in \mathit{Agt}$ and $a \in \mathit{Act}$:

$$\mathbf{C8} \quad \text{if } wE_i u \text{ and is } v \text{ such that } w \sim v \text{ and } j:a = R_j(v) \text{ then, there is } z \text{ such that } u \sim z \text{ and } j:a = R_j(z).$$

Obviously Axiom **CompleteInfoStrategy_{i,j}** is more general than Axiom **CompleteInfoStrategy_i**, that is, **CompleteInfoStrategy_{i,j}** implies **CompleteInfoStrategy_i**. It is also worth noting that the previous Axiom **CompleteInfoStrategy_{i,j}** together with Axiom **Indep** and Axiom **JointAct** imply $\Diamond\langle \delta \rangle \top \rightarrow K_i \Diamond\langle \delta \rangle \top$. The latter means that if δ is a strategy profile of the current game then every agent knows this.

REMARK. Note also that **CompleteInfoStrategy_i** and **CompleteInfoStrategy_{i,j}** are respectively equivalent to $\Box[i:a] \perp \rightarrow K_i \Box[i:a] \perp$ and $\Box[j:a] \perp \rightarrow K_i \Box[j:a] \perp$ thanks to Axiom 5 for the epistemic operators K_i .

In the sequel we call \mathcal{MLEG}^{**} the logic that results from adding the previous Axiom **CompleteInfoStrategy** $_{i,j}$ to the logic \mathcal{MLEG}^* discussed in Section 5, and we call \mathcal{MLEG}^{**} -models the corresponding models that results from adding the semantic constraint **C8** to the \mathcal{MLEG}^* -models. \mathcal{MLEG}^{**} is the logic of epistemic strategic games in which the only uncertainty is about agents' preference ordering over strategy profiles.

6.2. An Analysis of the Harsanyi Transformation

We conclude our discussion about games with incomplete information by shedding light on Harsanyi's claim that all uncertainty about the structure of a game can be reduced to uncertainty about payoffs [38].

Harsanyi proposed a way of transforming a game with uncertainty over both the payoffs and the strategy choices of the players into a game with no strategy-set uncertainty, without affecting the epistemic implications. In particular, Harsanyi proposed a way of reducing all kinds of imperfect information about the structure of a game to imperfect information about the strategy choices without affecting the rationality or irrationality of a player.

The basic idea of Harsanyi's transformation is that having a strategy with a highly undesirable payoff is for a player equivalent to not having the strategy at all. Suppose we start with a game with uncertainty over both the payoffs and the strategy choices of the players. This means that, some player i has a strategy a in his strategy set and another player j does not know this or some player i does not have a strategy a in his strategy set and another player j does not know this. To eliminate player's j strategy-set uncertainty is sufficient to add the strategy a to the set of strategies which in player j 's opinion are included in player i 's strategy space and to assign the lowest possible payoff to the new strategy profiles in which player i chooses strategy a . If player j is rational then the game transformation does not affect his choice, as his decision is not affected by the highly undesirable options that have been added by the game transformation.

We here provide a formal proof of Harsanyi's claim in a purely qualitative setting with no probabilities. See [13] for a formal proof of Harsanyi's claim in a quantitative setting using *interactive belief systems* à la Aumann & Brandeburger [1] with probabilities.

Let us start with a model $M = \langle W, \sim, R, E, \preceq, \pi \rangle$ of the logic \mathcal{MLEG}^* in which players may have incomplete information about all aspects of the game. We want to show that we can build a corresponding model $M' = \langle W', \sim', R', E', \preceq', \pi' \rangle$ of the logic \mathcal{MLEG}^{**} in which players can only have incomplete information about payoffs and which satisfies the same formulas Rat_i as M .

Let $[W] = \{\sim(w) \mid w \in W\}$ be the partition of W induced by the equivalence relation \sim . We note S_1, S_2, \dots the elements of $[W]$. Let $\Pi_M = \{\delta \mid \text{for every } \delta_i \text{ there is } u \in W \text{ such that } M, u \models \langle \delta_i \rangle \top\}$ be the set of strategy profiles with respect to the model M .

The model M' can be defined as follows.

- for every $S_i \in [W]$, $S'_i = S_i \cup \{w_i^\delta \mid \delta \in \Pi_M \text{ and there is no } v \in S_i \text{ such that } M, v \models \langle \delta \rangle \top\}$
- $W' = \bigcup_{S_i \in [W]} S'_i$
- for every $w, v \in W'$, $w \sim' v$ if and only if there is $S_i \in [W]$ such that $w, v \in S'_i$

- for every $C \in 2^{Agt^*}$ and $w \in W$, $R'_C(w) = R_C(w)$
- for every $C \in 2^{Agt^*}$ and $w_i^\delta \in W' \setminus W$, $R'_C(w) = \delta_C$
- for every $i \in Agt$ and $w \in W$, $E'_i(w) = E_i(w)$
- for every $i \in Agt$ and $w_i^\delta \in W' \setminus W$, $E'_i(w_i^\delta) = \{w_i^\delta\}$
- for every $i \in Agt$ and $w \in W$, $\preceq'_i(w) = \preceq_i(w)$
- for every $i \in Agt$ and $w_i^\delta \in W' \setminus W$, $\preceq'_i(w_i^\delta) = \sim(w_i^\delta)$
- for every $p \in Atm$, $\pi'(p) = \pi(p) \cup \{w_i^\delta \mid w_i^\delta \in W'\}$

Note that the crucial condition is the definition of the preference relation \preceq'_i for all $w_i^\delta \in W' \setminus W$. Following Harsanyi's intuition, we assign the lowest possible payoff to the worlds (and to the corresponding strategy profiles) that have been added through the game transformation in order to remove strategy-set uncertainty and that did not exist in the initial model M .

It is straightforward to check that the new model M' is indeed a $\mathcal{ML}\mathcal{E}\mathcal{G}^{**}$ -model without strategy-set uncertainty.

The following Theorem 16 is a formal characterization of Harsanyi's claim. It says that an agent i is rational at a given world w of model M if and only if, agent i is rational at world w of model M' which is obtained by removing strategy-set uncertainty from model M .

Theorem 16. *For every $w \in W$ and for every $i \in Agt$, $M, w \models \text{Rat}_i$ if and only if $M', w \models \text{Rat}_i$.*

The following corollary of Theorem 16 highlights that Harsanyi transformation does not affect common knowledge about the rationality or irrationality of a player.

Corollary 1. *For every $w \in W$, for every $i \in Agt$ and for every $C \in 2^{Agt^*}$, we have $M, w \models \text{MK}_C^n \text{Rat}_i$ if and only if $M', w \models \text{MK}_C^n \text{Rat}_i$.*

7. Conclusions

We have presented a multi-modal logic that enables to reason about epistemic games in strategic form. This logic, called $\mathcal{ML}\mathcal{E}\mathcal{G}$ (*Modal Logic of Epistemic Games*), integrates the concepts of joint action, preference and knowledge. We have shown that $\mathcal{ML}\mathcal{E}\mathcal{G}$ provides a highly flexible formal framework for the analysis of the epistemic aspects of strategic interaction. Indeed, $\mathcal{ML}\mathcal{E}\mathcal{G}$ can be easily adapted in order to integrate different assumptions on players' knowledge about the structure of a game.

Directions for future research are manifold. In this article (Section 3.2) we only considered the notion of individualistic rationality assumed in classical game theory: an agent decides to perform a certain action only if the agent believes that this action is a best response to what he expects the others will do. Our plan is to extend the present modal logic analysis of epistemic games to other forms of rationality such as fairness and reciprocity [40]. According to these notions of rationality, rational agents are not necessarily self-interested but they also consider the benefits of their choices for the group. Moreover, their decisions can be affected by their beliefs about other agents' willingness to act for the well-being of the group. In [41] we did some first steps into this direction.

Another aspect we intend to investigate in the future is a generalization of our approach to mixed strategies. Indeed, at the current stage the multi-modal logic $\mathcal{ML}\mathcal{E}\mathcal{G}$ only enables to reason about pure strategies. To this aim, we will have to extend $\mathcal{ML}\mathcal{E}\mathcal{G}$ by modal operators of probabilistic beliefs as the ones studied by [42,43]. We also postpone to future work an analysis of the epistemic conditions of Bayesian equilibrium in the resulting logical framework.

Acknowledgements

We would like to thank the anonymous reviewer for his very helpful comments.

This research is supported by the french project “Social ties in economics: experiments and theory” financed by the Agence Nationale de la Recherche (ANR), “Jeunes Chercheuses et Jeunes Chercheurs” program.

References

1. Aumann, R.J.; Brandenburger, A. Epistemic Conditions for Nash Equilibrium. *Econometrica* **1995**, *63*, 1161–1180.
2. Aumann, R. Interactive Epistemology I: Knowledge. *Int. J. Game Theory* **1999**, *28*, 263–300.
3. Bonanno, G. A syntactic approach to rationality in games with ordinal payoffs. *Proceedings of the 8th Conference on Logic and the Foundations of Game and Decision Theory (LOFT 2008)*, Amsterdam, The Netherlands, 3–5 July, 2008; Amsterdam University Press: Amsterdam, The Netherlands; Texts in Logic and Games Series, pp. 59–86.
4. Battigalli, P.; Bonanno, G. Recent Results on Belief, Knowledge and the Epistemic Foundations of Game Theory. *Res. Econ.* **1999**, *53*, 149–225.
5. Brandenburger, A. Knowledge and Equilibrium in Games. *J. Econ. Perspect.* **1992**, *6*, 83–101.
6. Gintis, H. *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*; Princeton University Press: Princeton, NJ, USA, 2009.
7. Van Benthem, J. Rational Dynamics and Epistemic Logic in Games. *Int. Game Theory Rev.* **2007**, *9*, 13–45.
8. De Bruin, B. Explaining Games: On the Logic of Game Theoretic Explanations. Ph.D thesis; University of Amsterdam: Amsterdam, The Netherlands, 2004.
9. Roy, O. Thinking before Acting: Intentions, Logic, Rational Choice. Ph.D thesis; University of Amsterdam: Amsterdam, The Netherlands, 2008.
10. Van Ditmarsch, H.P.; van der Hoek, W.; Kooi, B. *Dynamic Epistemic Logic*; Springer: Dordrecht, The Netherlands, 2007.
11. Baltag, A.; Moss, L.S. Logics for Epistemic Programs. *Synthese* **2004**, *139*, 165–224.
12. Gerbrandy, J.; Groeneveld, W. Reasoning about Information Change. *J. Logic Lang. Info.* **1997**, *6*, 147–196.
13. Hu, H.; Stuart, H.W. An Epistemic Analysis of the Harsanyi Transformation. *Int. J. Game Theory* **2001**, *30*, 517–525.
14. van Benthem, J.; Liu, F. Dynamic Logic of Preference Upgrade. *J. Appl. Non-Classical Logics* **2007**, *17*, 157–182.

15. Osborne, M.J.; Rubinstein, A. *A Course in Game Theory*; MIT Press: Cambridge, MA, USA, 1994.
16. Lewis, C.I.; Langford, C.H. *Symbolic Logic*; Dover Publications: New York, NY, USA, 1932.
17. Liu, F. Changing for the Better: Preference Dynamics and Agent Diversity. Ph.D thesis; University of Amsterdam: Amsterdam, The Netherlands, 2008.
18. Van Benthem, J.; Girard, P.; Roy, O. Everything Else Being Equal: A Modal Logic for *Ceteris Paribus* Preferences. *J. Phil. Logic* **2009**, *38*, 83–125.
19. Belnap, N.; Perloff, M.; Xu, M. *Facing the Future: Agents and Choices in Our Indeterminist World*; Oxford University Press: New York, NY, USA, 2001.
20. Papadimitriou, C.H. *Computational Complexity*; Addison Wesley: Reading, MA, USA, 1994.
21. Halpern, J.Y.; Moses, Y. A Guide to Completeness and Complexity for Modal Logics of Knowledge and Belief. *Artif. Intell.* **1992**, *54*, 319–379.
22. Pratt, V. A Near-Optimal Method for Reasoning about Action. *J. Comput. Syst. Sci.* **1980**, *20*, 231–254.
23. Pauly, M. Logic for Social Software. Ph.D thesis; University of Amsterdam: Amsterdam, The Netherlands, 2001.
24. Pauly, M. A Modal Logic for Coalitional Power in Games. *J. Logic Comput.* **2002**, *12*, 149–166.
25. Lorini, E. A Dynamic Logic of Agency II: Deterministic DLA, Coalition Logic, and Game Theory. *J. Logic Lang. Info.* **2010**, *19*, 327–351.
26. Horty, J.F. *Agency and Deontic Logic*; Oxford University Press: Oxford, UK, 2001.
27. Herzig, A.; Lorini, E. A Dynamic Logic of Agency I: STIT, Abilities and Powers. *J. Logic Lang. Info.* **2010**, *19*, 89–121.
28. Von Wright, G.H. *The Logic of Preference*; Edinburgh University Press: Edinburgh, UK, 1963.
29. Stalnaker, R. On the Evaluation of Solution Concepts. *Theor. Decis.* **1994**, *37*, 49–73.
30. Board, O. Knowledge, Beliefs, and Game-Theoretic Solution Concepts. *Oxford Rev. Econ. Policy* **2002**, *18*, 418–432.
31. Bernheim, D. Axiomatic Characterizations of Rational Choice in Strategic Environments. *Scand. J. Econ.* **1986**, *88*, 473–488.
32. Pearce, D. Rationalizable Strategic Behavior and the Problem of Perfection. *Econometrica* **1984**, *52*, 1029–1050.
33. Brandenburger, A.; Dekel, E. Rationalizability and Correlated Equilibria. *Econometrica* **1987**, *55*, 1391–1402.
34. Tan, T.; Werlang, S. The Bayesian Foundation of Solution Concepts of Games. *J. Econ. Theory* **1988**, *45*, 370–391.
35. Van der Hoek, W.; Jamroga, W.; Wooldridge, M. A Logic for Strategic Reasoning. *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2005)*, Utrecht, The Netherlands, 25–29 July, 2005; ACM Press: New York, NY, USA, 2005; pp. 157–164.
36. Van der Hoek, W.; Pauly, M. Modal Logic for Games and Information. In *Handbook of Modal Logic (3)*; Blackburn, P., Van Benthem, J., Wolter, F., Eds.; Elsevier: Amsterdam, The Netherlands, 2006.

37. Lutz, C. Complexity and Succinctness of Public Announcement Logic. *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)*, Hakodate, Japan, 8–12 May, 2006; ACM Press: New York, USA, 2006; pp. 137–143.
38. Harsanyi, J.C. Games with Incomplete Information Played by ‘Bayesian’ Players. *Manag. Sci.* **1967**, *14*, 159–182.
39. Halpern, J.Y.; Rego, L. Generalized Solution Concepts in Games with Possibly Unaware Players. *Proceedings of the Eleventh Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2007)*, Brussels, Belgium, 25–27 June, 2007; pp. 253–262.
40. Fehr, E.; Schmidt, K.M. Theories of Fairness and Reciprocity: Evidence and Economic Applications. In *Advances in Economics and Econometrics*; Cambridge University Press: Cambridge, UK, 2003.
41. Lorini, E. A Logical Account of Social Rationality in Strategic Games. *Proceedings of the 9th Conference on Logic and the Foundations of Game and Decision Theory (LOFT 2010)*, Toulouse, France, 5–7 July, 2010.
42. Halpern, J. *Reasoning about Uncertainty*; MIT Press: Cambridge, MA, USA, 2003.
43. Fagin, R.; Halpern, J. Reasoning about Knowledge and Probability. *J. Assn. Comput. Mach.* **1994**, *41*, 340–367.
44. Blackburn, P.; de Rijke, M.; Venema, Y. *Modal Logic*; Cambridge University Press: Cambridge, UK, 2001.
45. Sahlqvist, H. Completeness and Correspondence in the First and Second Order Semantics for Modal Logics. *Proceedings of the Third Scandinavian Logic Symposium* **1975**, *82*, 110–143.
46. Chen, C.; Lin, I. The Complexity of Propositional Modal Theories and the Complexity of Consistency of Propositional Modal Theories. *Logical Found. Comput. Sci.* **1994**, *813*, 69–80.
47. Chellas, B.F. *Modal Logic: An Introduction*; Cambridge University Press: Cambridge, UK, 1980.
48. Savitch, W. Relationships between Nondeterministic and Deterministic Tape Complexities*. *J. Comput. Syst. Sci.* **1970**, *4*, 177–192.

A. ANNEX: Proofs of Some Theorems

A.1. Proof of Theorems 1 and 3

\mathcal{MLEG} is determined by the class of \mathcal{MLEG} -models. \mathcal{MLEG}^{det} is determined by the class of \mathcal{MLEG}^{det} -models.

Proof. We only provide a sketch of the proof of Theorem 3. The proof of Theorem 1 is a straightforward adaptation of the proof of Theorem 3. It is sufficient to remove the constraint (S6) from the following definition 3.

It is straightforward to show that all axioms in Figure 2 are valid and that the rules of inference preserve validity in the class of \mathcal{MLEG}^{det} -models. The other part of the proof is shown using two major steps.

Step 1. We provide an alternative semantics for $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ in terms of standard Kripke models whose semantic conditions correspond one-to-one to the axioms in Table 2. The definition of Kripke $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models is the following one.

Definition 3 (Kripke $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -model). *Kripke $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models are tuples $M = \langle W, \sim, R, E, \preceq, \pi \rangle$ where:*

- W is a nonempty set of possible worlds or states;
- \sim is an equivalence relation on W ;
- $R : \bigcup_{C \in 2^{Agt^*}} \Delta_C \longrightarrow 2^{W \times W}$ maps every joint action δ_C to a transition relation $R_{\delta_C} \subseteq W \times W$ between possible worlds such that:
 - S1 $R_{\delta_C}(w) \neq \emptyset$ if and only if, for every $i \in C$ $R_{\delta_i}(w) \neq \emptyset$,
 - S2 if $R_{\delta_C}(w) \neq \emptyset$ then $R_{\delta_C}(w) = \{w\}$,
 - S3 $\bigcup_{\delta_C \in \Delta_C} R_{\delta_C}(w) \neq \emptyset$,
 - S4 if $\delta_C \neq \delta'_C$ then $R_{\delta_C}(w) = \emptyset$ or $R_{\delta'_C}(w) = \emptyset$,
 - S5 if for every $i \in Agt$ there is v_i such that $w \sim v_i$ and $R_{\delta_i}(v_i) \neq \emptyset$ then there is a v such that $w \sim v$ and $R_\delta(v) \neq \emptyset$;
 - S6 if $w \sim v$ and $R_\delta(w) \neq \emptyset$ and $R_\delta(v) \neq \emptyset$, then $w = v$;
- $E : Agt \longrightarrow W \times W$ maps every agent i to an equivalence relation E_i on W such that:
 - S6 if $(w, v) \in E_i$, then $i:a = R_i(w)$ if and only if $i:a = R_i(v)$,
 - S7 if wE_iv then $w \sim v$;
- $\preceq : Agt \longrightarrow W \times W$ maps every agent i to a reflexive, transitive relation \preceq_i on W such that:
 - S8 if $w \preceq_i v$ then $w \sim v$,
 - S9 if $w \sim v$ and $w \sim v'$ then $v \preceq_i v'$ or $v' \preceq_i v$;
- $\pi : Atm \longrightarrow 2^W$ is a valuation function.

Truth conditions of $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ formulas in Kripke $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models are again standard for atomic formulas and the Boolean operators. The truth conditions for Boolean operators and for operators \square , K_i and $[\text{good}]_i$ are the ones of Section 2.2. The truth condition for operators $[\delta_C]$ are:

- $M, w \models [\delta_C] \varphi$ iff $M, v \models \varphi$ for all $v \in R_{\delta_C}(w)$.

It is a routine task to prove that the axiomatic system of the logic $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ given in Table 2 is sound and complete with respect to this class of Kripke $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models via the Sahlqvist theorem, cf. [44, Th. 2.42]. Indeed all axioms in Table 2 are in the so-called Sahlqvist class [45]. Thus, they are all expressible as first-order conditions on Kripke models and are complete with respect to the defined model classes.

Step 2. The second step shows that the semantics in terms of $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models of Definition 2 and the semantics in terms of Kripke $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models of Definition 3 are equivalent. As the logic

$\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ is sound and complete for the class of Kripke $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models and is sound for the class of $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models, we have that for every $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ formula φ , if φ is valid in the class of Kripke $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models then φ is valid in the class of $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models. Consequently, for every $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ formula φ , if φ is satisfiable in the class of $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models then φ is satisfiable in the class of Kripke $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models. Therefore, in this second step we just need to show that for every $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ formula φ , if φ is satisfiable in the class of Kripke $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models then φ is satisfiable in the class of $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models.

Suppose φ is satisfiable in the class of Kripke $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -models. This means that there is a Kripke $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -model $M = \langle W, \sim, R, E, \preceq, \pi \rangle$ and world w such that $M, w \models \varphi$. We can now build a $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -model $M' = \langle W', R', E', \preceq', \pi' \rangle$ which satisfies φ . The model M' is defined as follows:

- $W' = W$;
- for every $C \in 2^{Agt^*}$ and $v \in W'$, $R'_C(v) = \delta_C$ if and only if $R_{\delta_C}(v) \neq \emptyset$;
- for every $i \in Agt$, $E'_i = E_i$;
- for every $i \in Agt$, $\preceq'_i = \preceq_i$;
- $\pi' = \pi$.

By induction on the structure of φ , it is just a trivial exercise to show that we have $M', w \models \varphi$.

A.2. Proof of Theorem 2

Proof. Let us start to prove that the satisfiability problem of $\mathcal{ML}\mathcal{E}\mathcal{G}$ is EXPTIME-hard when $card(Agt) \geq 2$. Let us consider two distinct agents $i, j \in Agt$. Let us consider a modal formula φ made of operators K_i, K_j and \Box . It is easy to check that the following two statements are equivalent:

- φ is satisfiable in the logic where K_i and K_j are S5-operators and \Box is the universal modality;
- φ is satisfiable in $\mathcal{ML}\mathcal{E}\mathcal{G}$.

So we have a reduction from the satisfiability problem of $\mathcal{ML}\mathcal{E}\mathcal{G}$ to the satisfiability problem of S5₂ plus universal modality \Box .

But the satisfiability problem of S5₂ plus the universal modality \Box is EXPTIME-hard as it is the case for the satisfiability problem of K plus the universal modality [46]. Indeed, we can reduce the satisfiability problem of S5₂ plus the universal modality \Box to the satisfiability problem of K plus the universal modality by translating a formula of K plus the universal modality into S5₂ plus universal modality. Let x be an extra proposition. The translation works as follows:

- $tr_i(\blacksquare\psi) = x \wedge K_i[\neg x \rightarrow tr_j(\psi)]$ where \blacksquare is the K-operator;
- $tr_j(\blacksquare\psi) = \neg x \wedge K_j[x \rightarrow tr_i(\psi)]$ where \blacksquare is the K-operator;
- for all $a \in \{i, j\}$, $tr_a = \Box[tr_i(\psi) \vee tr_j(\psi)]$ where \Box is the universal operator.

And φ is satisfiable in K plus universal modality iff $tr_i(\varphi)$ is satisfiable in $S5_2$ plus the universal modality. So $\mathcal{ML}\mathcal{EG}$ is EXPTIME-hard.

Now let us prove that the satisfiability problem of $\mathcal{ML}\mathcal{EG}$ is NEXPTIME. We are going to prove that we can make a filtration of any $\mathcal{ML}\mathcal{EG}$ model, preserving both the semantic constraints of Definition 1 and the truth of formulas (see [47] for a general introduction to the filtration method in modal logic). Let us consider a $\mathcal{ML}\mathcal{EG}$ -model $M = \langle W, \sim, R, E, \preceq, \pi \rangle$ where we suppose \sim to be the universal modality, without loss of generality. As usual, we consider a formula φ , the set $\Gamma = SF(\varphi)$ of all subformulas of φ and the equivalence relation \equiv over W defined by $w \equiv u$ iff for all $\psi \in \Gamma$, $M, w \models \psi$ iff $M, u \models \psi$. We note $|w|$ the equivalence class of \equiv containing w . Let us define $M' = \langle W', \sim, R', E', \preceq', \pi' \rangle$ by:

- $W' = \{|w| \mid w \in W\}$;
- $\sim = W' \times W'$;
- $R_C(|w|) = R_C(w)$;
- $|w|E_i|u|$ iff for all formulas $K_i\psi \in \Gamma$, $M, w \models K_i\psi$ iff $M, u \models K_i\psi$ and $R_i(w) = R_i(u)$;
- $|w| \preceq'_i |u|$ iff for all formulas $[\text{good}]_i\psi \in \Gamma$, $M, w \models [\text{good}]_i\psi$ implies $M, u \models [\text{good}]_i\psi$;
- $\pi'(p) = \{|w| \mid w \in \pi(p) \text{ and } p \text{ appears in } \varphi\}$.

We leave the reader checking that M' is well-defined, M' satisfies the constraints of Definition 1 and that if $M, w \models \varphi$ then $M', |w| \models \varphi$.

This filtration implies that if a formula φ is satisfiable, then it is satisfiable in a model of size $O(2^{|\varphi|})$ where $|\varphi|$ is the length of the formula φ . A possible algorithm for solving the satisfiability of φ may be as follows:

- Guess non-deterministically a $\mathcal{ML}\mathcal{EG}$ -model $M = \langle F, \pi \rangle$ whose size is bounded by $O(2^{|\varphi|})$ where π only gives truthness of propositions occurring in φ ;
- Guess non-deterministically a world w of M ;
- Check if $M, w \models \varphi$.

This algorithm non-deterministically runs in exponential time. So the satisfiability problem of $\mathcal{ML}\mathcal{EG}$ is in NEXPTIME.

A.3. Proof of Theorem 4

Proof. The satisfiability problem of $\mathcal{ML}\mathcal{EG}^{det}$ is clearly NP-hard because it is a conservative extension of the classical propositional logic whose satisfiability problem is NP-complete (Cook's Theorem [20]).

Now let us prove it is in NP. Clearly if a formula φ is $\mathcal{ML}\mathcal{EG}^{det}$ -satisfiable, there exists a $\mathcal{ML}\mathcal{EG}^{det}$ -model $F = \langle F, \pi \rangle$ whose size is bounded by $card(Act)^{card(Agt)}$. Here is a non-deterministic algorithm to check if a given formula φ is satisfiable:

- Guess non-deterministically a $\mathcal{ML}\mathcal{EG}^{det}$ -model $M = \langle F, \pi \rangle$ whose size is bounded by $card(Act)^{card(Agt)}$ where π only gives truthness of propositions occurring in φ ;

- Guess non-deterministically a world w of M ;
- Check if $M, w \models \varphi$.

This algorithm non-deterministically runs in polynomial time. So the satisfiability problem of $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ is in NP.

A.4. Proof of Theorem 5a

For all $i \in \text{Agt}$, we have: $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Rat}_i \leftrightarrow \text{K}_i \text{Rat}_i$

Lemma 1. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} [i:a]\perp \rightarrow \text{K}_i[i:a]\perp$

Proof. 1. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \bigvee_{b \in \text{Act}} \langle i:b \rangle \top$;

from **Active**;

2. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} [i:a]\perp \rightarrow [i:a]\perp \wedge \bigvee_{b \in \text{Act}} \langle i:b \rangle \top$

by 1. and Boolean principles;

3. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} [i:a]\perp \wedge \bigvee_{b \in \text{Act}} \langle i:b \rangle \top \rightarrow \bigvee_{b \neq a} \langle i:b \rangle \top$

by Boolean principles;

4. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \langle i:b \rangle \top \rightarrow \text{K}_i \langle i:b \rangle \top$ if $b \neq a$

by **Aware**;

5. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \bigvee_{b \neq a} \langle i:b \rangle \top \rightarrow \bigvee_{b \neq a} \text{K}_i \langle i:b \rangle \top$

by 4. and Boolean principles;

6. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \langle i:b \rangle \top \rightarrow [i:a]\perp$ if $b \neq a$; by **Single**;

7. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{K}_i(\langle i:b \rangle \top \rightarrow [i:a]\perp)$ if $b \neq a$

by Necessitation of K_i from 6;

8. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{K}_i \langle i:b \rangle \top \rightarrow \text{K}_i [i:a]\perp$ if $b \neq a$

by Axiom K for K_i plus **Modus Ponens** from 7;

9. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \bigvee_{b \neq a} \text{K}_i \langle i:b \rangle \top \rightarrow \text{K}_i [i:a]\perp$

by Boolean principles from 8.

10. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} [i:a]\perp \rightarrow \text{K}_i [i:a]\perp$

by 2, 3, 5 and 9.

Now let us prove Theorem 5a. We give here a version of the proof that uses Axioms K, T, 4 and 5 for epistemic modal operators.

Proof. 1. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Rat}_i \leftrightarrow \bigwedge_{a,b \in \text{Act}} (\langle i:a \rangle \top \rightarrow \bigvee_{\beta \in \Delta} (\widehat{\text{K}}_i \langle \beta_{-i} \rangle \top \wedge \text{K}_i(\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle)))$

by Definition of Rat_i ;

$$2. \widehat{K}_i \langle \beta_{-i} \rangle \top \leftrightarrow K_i \widehat{K}_i \langle \beta_{-i} \rangle \top$$

by Axiom 5 for K_i ;

$$3. K_i(\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle) \leftrightarrow K_i K_i(\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle)$$

by Axiom 4 for K_i plus Boolean principles;

$$4. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Rat}_i \rightarrow \bigwedge_{a,b \in \text{Act}} (\langle i:a \rangle \top \rightarrow \bigvee_{\beta \in \Delta} (K_i \widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i K_i(\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle)))$$

by 1, 2, 3 and Boolean principles;

$$5. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} K_i \varphi \wedge K_i \psi \dots \leftrightarrow K_i(\varphi \wedge \psi \dots)$$

by modal logic K principles;

$$6. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} [i:a] \perp \rightarrow K_i [i:a] \perp$$

by Lemma 1;

$$7. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Rat}_i \rightarrow \bigwedge_{a,b \in \text{Act}} (K_i [i:a] \perp \vee \bigvee_{\beta \in \Delta} K_i(\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i(\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle)))$$

by 4, 5, 6 and Boolean principles;

$$8. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} K_i \varphi \vee K_i \psi \dots \rightarrow K_i(\varphi \vee \psi \dots)$$

by modal logic K principles;

$$9. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Rat}_i \rightarrow K_i \underbrace{\bigwedge_{a,b \in \text{Act}} ([i:a] \perp \vee \bigvee_{\beta \in \Delta} (\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i(\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle)))}_{\text{Rat}_i}$$

by 7 and 8;

$$10. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} K_i \text{Rat}_i \rightarrow \text{Rat}_i$$

by Axiom T of K_i ;

$$11. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Rat}_i \leftrightarrow K_i \text{Rat}_i$$

by 9 and 10.

We give another version of the proof of Theorem 5a that uses Axioms K , T and 5 for epistemic modal operators and introspection over preferences “ $(\delta \leq_i \delta') \rightarrow K_i(\delta \leq_i \delta')$ ”.

$$\text{Proof. } 1. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Rat}_i \leftrightarrow \left(\bigwedge_{a,b \in \text{Act}} (\langle i:a \rangle \top \rightarrow \bigvee_{\beta \in \Delta} (\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i(\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle))) \right)$$

by Definition of Rat_i ;

$$2. \widehat{K}_i \langle \beta_{-i} \rangle \top \leftrightarrow K_i \widehat{K}_i \langle \beta_{-i} \rangle \top$$

by Axiom 5 for K_i ;

$$3. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Rat}_i \leftrightarrow \bigwedge_{a,b \in \text{Act}} (\langle i:a \rangle \top \rightarrow \bigvee_{\beta \in \Delta} (K_i \widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i(\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle)))$$

by 1 and 2;

4. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} K_i\varphi \wedge K_i\psi \dots \leftrightarrow K_i(\varphi \wedge \psi \dots)$
by modal logic **K** principles;
5. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} [i:a] \perp \rightarrow K_i [i:a] \perp$
by Lemma 1;
6. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Rat}_i \leftrightarrow \bigwedge_{a,b \in \text{Act}} (K_i [i:a] \perp \vee \bigvee_{\beta \in \Delta} K_i (\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle)))$
by 2, 4, 5 and Boolean principles;
7. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} K_i\varphi \vee K_i\psi \dots \rightarrow K_i(\varphi \vee \psi \dots)$
by modal logic **K** principles;
8. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Rat}_i \rightarrow K_i \bigwedge_{a,b \in \text{Act}} ([i:a] \perp \vee \bigvee_{\beta \in \Delta} (\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle)))$
9. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle) \rightarrow K_i(\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle)$
by Lemma 1 (or introspection over preferences);
10. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Rat}_i \rightarrow K_i \underbrace{\bigwedge_{a,b \in \text{Act}} ([i:a] \perp \vee \bigvee_{\beta \in \Delta} (\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i(\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle)))}_{\text{Rat}_i}$
by 8 and 9;
11. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} K_i \text{Rat}_i \rightarrow \text{Rat}_i$
by Axiom T of K_i ;
12. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Rat}_i \leftrightarrow K_i \text{Rat}_i$
by 10 and 11.

A.5. Proof of Theorem 6

For all $\delta \in \Delta$, we have:

$$\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \left(\left(\bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \wedge \bigwedge_{i \in \text{Agt}} K_i \langle \delta_{-i} \rangle \top \right) \rightarrow \text{Nash}(\delta)$$

Proof. 1. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \langle \delta_{-i} \rangle \top \rightarrow [\beta_{-i}] \perp$ if $\beta_{-i} \neq \delta_{-i}$

by **Single**;

$$2. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} K_i (\langle \delta_{-i} \rangle \top \rightarrow [\beta_{-i}] \perp) \text{ if } \beta_{-i} \neq \delta_{-i}$$

by necessitation of K_i ;

$$3. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} (K_i \langle \delta_{-i} \rangle \top \rightarrow K_i [\beta_{-i}] \perp) \text{ if } \beta_{-i} \neq \delta_{-i}$$

by 2, axiom **K** of K_i plus **Modus Ponens**;

$$4. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} K_i \langle \delta_{-i} \rangle \top \wedge \widehat{K}_i \langle \beta_{-i} \rangle \top \leftrightarrow \perp \text{ if } \beta_{-i} \neq \delta_{-i}$$

by 3 and Boolean principles;

5. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} K_j \langle \delta_{-j} \rangle \top \rightarrow \langle \delta_{-j} \rangle \top$
by axiom T of **SS** $_{K_i}$.
6. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \langle \delta_{-j} \rangle \top \rightarrow \langle \delta_i \rangle \top$ if $j \neq i$
by **JointAct**;
7. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} K_j \langle \delta_{-j} \rangle \top \rightarrow \langle \delta_i \rangle \top$ for $j \neq i$
by 5 and 6;
8. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \left(\left(\bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \wedge \bigwedge_{i \in \text{Agt}} K_i \langle \delta_{-i} \rangle \top \right)$
 $\rightarrow \bigwedge_{b \in \text{Act}} \left(\langle \delta_i \rangle \top \rightarrow \bigvee_{\beta \in \Delta} \left(\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i (\langle i:b, \beta_{-i} \rangle \top \leq_i (\langle \delta_i, \beta_{-i} \rangle \top)) \right) \right) \wedge \langle \delta_i \rangle \top \wedge K_i \langle \delta_{-i} \rangle \top$
by Boolean principles and 7;
9. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \bigwedge_{b \in \text{Act}} \left(\langle \delta_i \rangle \top \rightarrow \bigvee_{\beta \in \Delta} \left(\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i (\langle i:b, \beta_{-i} \rangle \top \leq_i (\langle \delta_i, \beta_{-i} \rangle \top)) \right) \right) \wedge \langle \delta_i \rangle \top \wedge$
 $K_i \langle \delta_{-i} \rangle \top$
 $\rightarrow \bigwedge_{b \in \text{Act}} \bigvee_{\beta \in \Delta} \left(\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i (\langle i:b, \beta_{-i} \rangle \top \leq_i (\langle \delta_i, \beta_{-i} \rangle \top)) \right) \wedge K_i \langle \delta_{-i} \rangle \top$
by Boolean principle “ $A \wedge (A \rightarrow B) \rightarrow B$ ”;
10. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \bigwedge_{b \in \text{Act}} \bigvee_{\beta \in \Delta} \left(\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i (\langle i:b, \beta_{-i} \rangle \top \leq_i (\langle \delta_i, \beta_{-i} \rangle \top)) \right) \wedge K_i \langle \delta_{-i} \rangle \top$
 $\rightarrow \bigwedge_{b \in \text{Act}} \bigvee_{\beta \in \Delta} \left(\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i \langle \delta_{-i} \rangle \top \wedge K_i (\langle i:b, \beta_{-i} \rangle \top \leq_i (\langle \delta_i, \beta_{-i} \rangle \top)) \right)$
by distributivity of \wedge over $\bigvee_{\beta \in \Delta}$;
11. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \bigwedge_{b \in \text{Act}} \bigvee_{\beta \in \Delta} \left(\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i \langle \delta_{-i} \rangle \top \wedge K_i (\langle i:b, \beta_{-i} \rangle \top \leq_i (\langle \delta_i, \beta_{-i} \rangle \top)) \right)$
 $\rightarrow \bigwedge_{b \in \text{Act}} K_i (\langle i:b, \delta_{-i} \rangle \top \leq_i (\langle \delta_i, \delta_{-i} \rangle \top))$
by 4 plus Boolean principles;
12. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} K_i (\langle i:b, \delta_{-i} \rangle \top \leq_i (\langle \delta_i, \delta_{-i} \rangle \top)) \rightarrow (\langle i:b, \delta_{-i} \rangle \top \leq_i (\langle \delta_i, \delta_{-i} \rangle \top))$
by Axiom T of K_i ;
13. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \bigwedge_{b \in \text{Act}} K_i (\langle i:b, \delta_{-i} \rangle \top \leq_i (\langle \delta_i, \delta_{-i} \rangle \top)) \rightarrow \underbrace{\bigwedge_{b \in \text{Act}} (\langle i:b, \delta_{-i} \rangle \top \leq_i (\langle \delta_i, \delta_{-i} \rangle \top))}_{\text{BR}(\delta_i, \delta_{-i})}$
by 12 and Boolean principles;
14. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \left(\left(\bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \wedge \bigwedge_{i \in \text{Agt}} K_i \langle \delta_{-i} \rangle \top \right) \rightarrow \text{BR}(\delta_i, \delta_{-i})$
by 8, 9, 10, 11 and 13;
15. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \left(\left(\bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \wedge \bigwedge_{i \in \text{Agt}} K_i \langle \delta_{-i} \rangle \top \right) \rightarrow \text{Nash}(\delta)$
by 14 and Boolean principles.

A.6. Proof of Theorem 7

For all $\delta \in \Delta$, $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \left(\left(\text{MK}_{\text{Agt}}^n \bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \wedge \langle \delta_C \rangle \top \right) \rightarrow \neg \text{SD}^{\leq n}(\delta_C)$

Proof.

Lemma 2. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{SD}^{\leq n}(i:a) \rightarrow \text{K}_j \text{SD}^{\leq n}(i:a)$

Proof. The proof of the lemma consists in proving by induction on n that $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{SD}^{\leq n}(i:a) \rightarrow \text{K}_j \text{SD}^{\leq n}(i:a)$ and $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \neg \text{SD}^{\leq n}(i:a) \rightarrow \text{K}_j \neg \text{SD}^{\leq n}(i:a)$. We leave the proof of these two $\mathcal{ML}\mathcal{E}\mathcal{G}$ -theorems based on Lemma 1 to the reader.

Basic case $n = 0$

Here we prove $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \left(\left(\bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \wedge \langle \delta_C \rangle \top \right) \rightarrow \neg \text{SD}^{\leq 0}(\delta_C)$.

1. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \langle \delta_C \rangle \top \rightarrow \bigwedge_{i \in C} \langle \delta_i \rangle \top$
by Axiom **JointAct**
2. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} K_i(\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i : \delta_i \rangle) \rightarrow (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i : \delta_i \rangle)$
by Axiom T of K_i ;
3. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \left(\left(\bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \wedge \langle \delta_C \rangle \top \right)$
 $\rightarrow \langle \delta_i \rangle \top \wedge \bigwedge_{b \in \text{Act}} \left(\bigvee_{\beta \in \Delta} \left(\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i : \delta_i \rangle) \right) \right)$ if $c \in C$
By 1, 2, axiom T for K_i and Boolean principles;
4. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \widehat{K}_i \langle \beta_{-i} \rangle \top \rightarrow \diamond \langle \beta_{-i} \rangle \top$
by **CompleteInfo** or (Lemma 1 considered as axioms plus axiom T for \square);
5. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \langle \delta_i \rangle \top \rightarrow \diamond \langle \delta_i \rangle \top$
by Axiom T for \square ;
6. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \langle \delta_i \rangle \top \wedge \bigwedge_{b \in \text{Act}} \left(\bigvee_{\beta \in \Delta} \left(\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i : \delta_i \rangle) \right) \right)$
 $\rightarrow \diamond \langle \delta_i \rangle \top \wedge \bigwedge_{b \in \text{Act}} \left(\bigvee_{\beta \in \Delta} (\diamond \langle \beta_{-i} \rangle \top \wedge \langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i : \delta_i \rangle) \right)$
by 4, 5 and Boolean principles;
7. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \diamond \langle \delta_i \rangle \top \wedge \bigwedge_{b \in \text{Act}} \left(\bigvee_{\beta \in \Delta} (\diamond \langle \beta_{-i} \rangle \top \wedge \langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i : \delta_i \rangle) \right)$
 $\rightarrow \underbrace{\diamond \langle \delta_i \rangle \top \wedge \bigwedge_{b \in \text{Act}} \diamond \langle i:b \rangle \top}_{\neg \text{SD}^{\leq 0}(\delta_i)} \rightarrow \left(\bigvee_{\beta \in \Delta} (\diamond \langle \beta_{-i} \rangle \top \wedge \langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i : \delta_i \rangle) \right)$
by Boolean principles;
8. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \left(\left(\bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \wedge \langle \delta_C \rangle \top \right) \rightarrow \neg \text{SD}^{\leq 0}(\delta_i)$ if $c \in C$
by 3, 6, 7 and Boolean principles;
9. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \left(\left(\bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \wedge \langle \delta_C \rangle \top \right) \rightarrow \neg \text{SD}^{\leq 0}(\delta_C)$
by 8 and Boolean principles.

Inductive case

let $n \in \mathbf{N}$ and let us prove that if the theorem 7 is true for all $k \leq n$ then it is true for $n + 1$.

1. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \left(\left(\text{MK}_{Agt}^{n+1} \bigwedge_{i \in Agt} \text{Rat}_i \right) \wedge \langle \delta_C \rangle \top \right) \rightarrow \left(\left(\text{MK}_{Agt}^n \bigwedge_{i \in Agt} \text{Rat}_i \right) \wedge \langle \delta_C \rangle \top \right)$
by Axiom T for K_i plus Boolean principles;
2. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \left(\left(\text{MK}_{Agt}^n \bigwedge_{i \in Agt} \text{Rat}_i \right) \wedge \langle \delta_C \rangle \top \right) \rightarrow \neg \text{SD}^{\leq n}(\delta_C)$
by induction;
3. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \neg \text{SD}^{\leq n}(\delta) \rightarrow \neg \text{SD}^{\leq n}(\delta_i)$
by Definition of $\neg \text{SD}^{\leq n}(\delta)$ and Boolean principles;
4. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \left(\left(\text{MK}_{Agt}^{n+1} \bigwedge_{i \in Agt} \text{Rat}_i \right) \wedge \langle \delta_C \rangle \top \right) \rightarrow \neg \text{SD}^{\leq n}(\delta_i)$
by 1, 2, 3;
5. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \left(\left(\text{MK}_{Agt}^{n+1} \bigwedge_{i \in Agt} \text{Rat}_i \right) \wedge \langle \delta_C \rangle \top \right) \rightarrow \text{Rat}_i \wedge K_i \text{MK}_{Agt}^n \text{Rat}_i \wedge \langle \delta_C \rangle \top$;
by Boolean principles
6. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \text{Rat}_i \wedge K_i \text{MK}_{Agt}^n \text{Rat}_i \wedge \langle \delta_C \rangle \top$
 $\rightarrow \bigwedge_{b \in Act} \bigvee_{\beta \in \Delta} (\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, \delta_i \rangle)) \wedge K_i \text{MK}_{Agt}^n \text{Rat}_i$ if $i \in C$
by **JointAct**, definition of Rat_i and Boolean principles;
7. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \bigwedge_{b \in Act} \bigvee_{\beta \in \Delta} (\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, \delta_i \rangle)) \wedge K_i \text{MK}_{Agt}^n \text{Rat}_i$
 $\rightarrow \bigwedge_{b \in Act} \bigvee_{\beta \in \Delta} (\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i \text{MK}_{Agt}^n \text{Rat}_i \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, \delta_i \rangle))$
by distributivity of \wedge over $\bigvee_{\beta \in \Delta}$;
8. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i \text{MK}_{Agt}^n \text{Rat}_i \rightarrow \widehat{K}_i (\langle \beta_{-i} \rangle \top \wedge \text{MK}_{Agt}^n \text{Rat}_i)$
by modal logic K principle “ $\widehat{K}_i A \wedge K_i B \rightarrow \widehat{K}_i (A \wedge B)$ ”;
9. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} (\langle \beta_{-i} \rangle \top \wedge \text{MK}_{Agt}^n \text{Rat}_i) \rightarrow \neg \text{SD}^{\leq n}(\beta_{-i})$
by induction;
10. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} K_i ((\langle \beta_{-i} \rangle \top \wedge \text{MK}_{Agt}^n \text{Rat}_i) \rightarrow \neg \text{SD}^{\leq n}(\beta_{-i}))$
by necessitation rule on 9;
11. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \widehat{K}_i (\langle \beta_{-i} \rangle \top \wedge \text{MK}_{Agt}^n \text{Rat}_i) \rightarrow \widehat{K}_i \neg \text{SD}^{\leq n}(\beta_{-i})$
by modal logic K principles applied on 10;
12. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \widehat{K}_i \neg \text{SD}^{\leq n}(\beta_{-i}) \rightarrow \neg \text{SD}^{\leq n}(\beta_{-i})$ by Lemma A.6;
13. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i \text{MK}_{Agt}^n \text{Rat}_i \rightarrow \neg \text{SD}^{\leq n}(\beta_{-i})$
from 8, 11 and 12;
14. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \bigwedge_{b \in Act} \bigvee_{\beta \in \Delta} (\widehat{K}_i \langle \beta_{-i} \rangle \top \wedge K_i \text{MK}_{Agt}^n \text{Rat}_i \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, \delta_i \rangle))$
 $\rightarrow \bigwedge_{b \in Act} \bigvee_{\beta \in \Delta} (\neg \text{SD}^{\leq n}(\beta_{-i}) \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, \delta_i \rangle))$
by 13 and Boolean principles;

15. $\vdash_{\mathcal{ML}\mathcal{EG}} \bigwedge_{b \in Act} \bigvee_{\beta \in \Delta} (\neg \text{SD}^{\leq n}(\beta_{-i}) \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:\delta_i \rangle))$
 $\rightarrow \bigwedge_{b \in Act} (\neg \text{SD}^{\leq n}(i:b) \rightarrow \bigvee_{\beta \in \Delta} (\neg \text{SD}^{\leq n}(\beta_{-i}) \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, \delta_i \rangle)))$
 by Boolean principles;
16. $\vdash_{\mathcal{ML}\mathcal{EG}} \left(\left(\text{MK}_{Agt}^{n+1} \bigwedge_{i \in Agt} \text{Rat}_i \right) \wedge \langle \delta_C \rangle \top \right)$
 $\rightarrow \bigwedge_{b \in Act} (\neg \text{SD}^{\leq n}(i:b) \rightarrow \bigvee_{\beta \in \Delta} (\neg \text{SD}^{\leq n}(\beta_{-i}) \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, \delta_i \rangle)))$ if $c \in C$
 by 5, 6, 7, 14, 15;
17. $\vdash_{\mathcal{ML}\mathcal{EG}} \left(\left(\text{MK}_{Agt}^{n+1} \bigwedge_{i \in Agt} \text{Rat}_i \right) \wedge \langle \delta_C \rangle \top \right) \rightarrow \neg \text{SD}^{\leq n+1}(\delta_i)$ if $c \in C$
 by 4 and 16;
18. $\vdash_{\mathcal{ML}\mathcal{EG}} \left(\left(\text{MK}_{Agt}^{n+1} \bigwedge_{i \in Agt} \text{Rat}_i \right) \wedge \langle \delta_C \rangle \top \right) \rightarrow \neg \text{SD}^{\leq n+1}(\delta_C)$
 by 17.

A.7. Proof of Theorem 8

Consider an arbitrary $\mathcal{ML}\mathcal{EG}$ -model $M = \langle W, \sim, R, E, \preceq, \pi \rangle$, a world w in M and $\delta \in \Delta$ such that for all positive integers n we have $M, w \models \langle \delta \rangle \top \wedge \neg \text{SD}^{\leq n}(\delta)$. Then, there is a model $M' = \langle W, \sim, R, E', \preceq, \pi \rangle$ such that for all positive integers n we have $M', w \models \text{MK}_{Agt}^n \bigwedge_{i \in Agt} \text{Rat}_i$.

Proof. The proof is based on the following Lemma 3.

Lemma 3. For all $\delta \in \Delta$, we have

$$\vdash_{\mathcal{ML}\mathcal{EG}} \neg \text{SD}^{\leq n}(i:a) \leftrightarrow$$

$$(\neg \text{SD}^{\leq n-1}(i:a) \wedge \bigwedge_{b \in Act} \bigvee_{\delta \in \Delta} (\neg \text{SD}^{\leq n-1}(\delta_{-i}) \wedge (\langle \delta_{-i}, i:b \rangle \leq_i \langle \delta_{-i}, i:a \rangle)))$$

In other words, player i 's strategy a survives after n rounds of IDSDS if and only if, a survives after $n - 1$ rounds of IDSDS and in the subgame of depth n , for every alternative strategy b of i , there is a joint action δ_{-i} of the other agents that survives after $n - 1$ rounds of IDSDS such that playing a while the others play δ_{-i} is for i at least as good as playing b while the others play δ_{-i} .

Lemma 3 ensures that the definition of $\text{SD}^{\leq n}(i:a)$ can be rewritten in the following shorter equivalent form:

$$\text{SD}^{\leq n}(i:a) \stackrel{\text{def}}{=} \neg \text{SD}^{\leq n-1}(i:a) \rightarrow$$

$$\bigvee_{b \in Act} \bigwedge_{\delta \in \Delta} (\neg \text{SD}^{\leq n-1}(\delta_{-i}) \rightarrow (\langle \delta_{-i}, i:a \rangle <_i \langle \delta_{-i}, i:b \rangle))$$

Let us consider w such that for all positive integers n , $M, w \models \langle \delta \rangle \top \wedge \neg \text{SD}^{\leq n}(\delta)$. We can now show how to build the accessibility relations E'_i of the model M' in such a way that for all positive integers n $M', w \models \text{MK}_{Agt}^n \bigwedge_{i \in Agt} \text{Rat}_i$. The construction goes as follows.

For all positive integers n , let A_n be the subset of all joint actions $\beta \in \Delta$ such that $M, w \models \neg \text{SD}^{\leq n}(\beta)$. As $\vdash_{\mathcal{ML}\mathcal{EG}} \neg \text{SD}^{\leq n+1}(\beta) \rightarrow \neg \text{SD}^{\leq n}(\beta)$, we have $A_{n+1} \subseteq A_n$. Let us define $A_\infty = \bigcap_{n \in \mathbb{N}} A_n$. As Δ is finite, there exists a positive integer n_0 such that $A_\infty = A_{n_0}$ and for all positive integers $n > n_0$,

$A_n = A_{n_0}$. Let Ω be the set of all worlds u such that $u \sim w$ and such that there exists $\beta \in A_\infty$ such that $M, u \models \langle \beta \rangle \top$. Note that $w \in \Omega$.

For all $i \in \text{Agt}$, we define E'_i as follows:

- for all $s, t \in W$, $sE'_i t$ iff either $s, t \in \Omega$ or $s = t$.

Now, let us prove that for all $i \in \text{Agt}$, for all $s \in \Omega$, we have $M', s \models \text{Rat}_i$. Let $a \in \text{Act}$ be such that $s \models \langle i:a \rangle \top$. As $s \in \Omega$, we have $M', s \models \neg \text{SD}^{\leq n_0+1}(i:a)$. By Lemma 3, it implies that for all $b \in \text{Act}$, there exists $\beta \in \Delta$ such that $M', s \models \neg \text{SD}^{\leq n_0}(\beta_{-i})$ and $\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle$. But by definition of E'_i , we have equivalence between $M', s \models \widehat{K}_i \langle \beta_{-i} \rangle \top$ and $M', s \models \neg \text{SD}^{\leq n_0}(\beta_{-i})$. So for all $s \in \Omega$, we have $M', s \models \text{Rat}_i$. As for all $i \in \text{Agt}$ we have $E'_i(w) = \Omega$, we obtain $M', w \models \text{MK}_{\text{Agt}}^n \bigwedge_{i \in \text{Agt}} \text{Rat}_i$ for all positive integers n .

A.8. Proof of Theorem 9

If M is a $\mathcal{ML}\mathcal{E}\mathcal{G}$ model then M^\times is a $\mathcal{ML}\mathcal{E}\mathcal{G}$ model.

Proof. It is just a routine to verify that \sim^\times and every E_i^\times are equivalence relations, every \preceq_i^\times is reflexive and transitive, and the model M^\times satisfies the semantic constraints **C1**, **C4**, **C5** and **C6**.

Let us prove that M^\times satisfies constraints **C2** and **C3**.

We first prove that M^\times satisfies constraint **C2**. We introduce the following useful notation. Suppose $\chi_1, \chi_2 \in \mathcal{L}_{\mathcal{GT}}$. Then, $\chi_2 \rightsquigarrow \chi_3$ iff there is $\chi_3 \in \mathcal{L}_{\mathcal{GT}}$ such that $\chi_1 = \chi_2 \wedge \chi_3$.

Now, suppose for every $i \in \text{Agt}$ there is v_i such that $v_i \sim^\times w$ and $R_i^\times(v_i) = \delta_i$. It follows that for every $i \in \text{Agt}$ there is v_i such that $v_i \sim w$ and $R_i(v_i) = \delta_i$. The latter implies that there is v such that $v \sim w$ and $R_\delta(v) \neq \emptyset$ (by the semantic constraint **C2**). Now, suppose for all v' if $v' \sim^\times w$ then $R_{\text{Agt}}^\times(v') = \delta$. It follows that: there is $i \in \text{Agt}$ and $\psi \in \mathcal{L}_{\mathcal{ML}\mathcal{E}\mathcal{G}}$ such that $\Box\psi \rightarrow [\delta_i] \perp \rightsquigarrow \chi$ and $M, v \models \Box\psi$. The latter implies that there is $i \in \text{Agt}$ and $\psi \in \mathcal{L}_{\mathcal{ML}\mathcal{E}\mathcal{G}}$ such that $\Box\psi \rightarrow [\delta_i] \perp \rightsquigarrow \chi$ and for all $v' \sim w$, $M, v' \models \Box\psi$. We conclude that there is no $v_i \sim^\times w$ such that $R_i^\times(v_i) = \delta_i$ which leads to a contradiction.

We now consider constraint **C3**. Suppose $wE_i^\times v$ and $R_i^\times(w) = i:a$. It follows that $wE_i v$ and $R_i(w) = i:a$ which implies $R_i(v) = i:a$, because M satisfies constraint **C3**. The latter implies $R_i^\times(v) = i:a$. Now, suppose $wE_i^\times v$ and $R_i^\times(v) = i:a$. It follows that $wE_i v$ and $R_i(v) = i:a$ which implies $R_i(w) = i:a$, because M satisfies constraint **C3**. The latter implies $R_i^\times(w) = i:a$.

A.9. Proof of Theorem 10

Proof. The proofs of **R1-R6** go as in Dynamic Epistemic Logic (DEL) (see [10]). We here prove **R7**.

$M, w \models [\chi!] [\delta_C] \varphi$,

IFF if $M, w \models \chi$ then $M^\times, w \models [\delta_C] \varphi$,

IFF if $M, w \models \chi$ then $M^\times, w \models \langle \delta_C \rangle \top \rightarrow \varphi$ (by Axiom **Def**_[\delta_C]),

IFF if $M, w \models \chi$ then $M^\times, w \models [\delta_C] \perp$ or $M^\times, w \models \varphi$,

IFF if $M^\times, w \models \langle \delta_C \rangle \top$ then, if $M, w \models \chi$ then $M^\times, w \models \varphi$,

IFF if $M^\times, w \models \langle \delta_C \rangle \top$ then, $M, w \models [\chi!] \varphi$,

IFF if $M, w \models \langle \delta_C \rangle \top$ then, $M, w \models [\chi!] \varphi$,

IFF if $M, w \models \langle \delta_C \rangle \top \rightarrow [\chi!] \varphi$.

A.10. Proof of Theorem 11

The logic $\mathcal{ML}\mathcal{EG}^{GT}$ is completely axiomatized by the axioms and inference rules of $\mathcal{ML}\mathcal{EG}$ together with the schemata of Theorem 10.

Proof. By means of the principles **R1-R7** in Theorem 10, it is straightforward to prove that for every $\mathcal{ML}\mathcal{EG}^{GT}$ formula there is an equivalent $\mathcal{ML}\mathcal{EG}$ formula. In fact, each reduction axiom **R2-R7**, when applied from the left to the right by means of the rule of replacement of proved equivalence, yields a simpler formula, where “simpler” roughly speaking means that the dynamic operator is pushed inwards. Once the dynamic operator attains an atom it is eliminated by the equivalence **R1**. Hence, the completeness of $\mathcal{ML}\mathcal{EG}^{GT}$ is a straightforward consequence of Theorem 1.

A.11. Proof of Theorem 12

For all $\delta \in \Delta$, for all $n \geq 0$,

$$\vdash_{\mathcal{ML}\mathcal{EG}^{GT}} \langle \delta \rangle \top \rightarrow (\neg \text{SD}^{\leq n}(\delta) \leftrightarrow \langle \chi_{\text{SD}}! \rangle^{n+1} \top).$$

Lemma 4. $\vdash_{\mathcal{ML}\mathcal{EG}} \bigvee_{a \in \text{Act}} \neg \text{SD}^{\leq 0}(i:a)$.

Proof. 1. $\vdash_{\mathcal{ML}\mathcal{EG}} \bigvee_{a_1 \in \text{Act}} \langle i:a_1 \rangle \top$

by **Active**;

$$2. \vdash_{\mathcal{ML}\mathcal{EG}} \bigvee_{a_1 \in \text{Act}} \diamond \langle i:a_1 \rangle \top$$

by 1 and **T** for \square and Boolean principles;

$$3. \vdash_{\mathcal{ML}\mathcal{EG}} \bigvee_{\beta \in \Delta} \langle \beta_{-i} \rangle \top$$

by **Active**;

$$4. \vdash_{\mathcal{ML}\mathcal{EG}} \bigvee_{\beta \in \Delta} \diamond \langle \beta_{-i} \rangle \top$$

by 3;

$$5. \vdash_{\mathcal{ML}\mathcal{EG}} \bigwedge_{a \in \text{Act}} \text{SD}^{\leq 0}(i:a) \rightarrow \bigwedge_{a \in \text{Act}} \text{SD}^{\leq 0}(i:a) \wedge \bigvee_{a \in \text{Act}} \diamond \langle i:a \rangle \top \wedge \bigvee_{\beta \in \Delta} \diamond \langle \beta_{-i} \rangle \top$$

by 2 and 4;

$$6. \vdash_{\mathcal{ML}\mathcal{EG}} \diamond \langle \beta_{-i} \rangle \top \wedge \text{SD}^{\leq 0}(i:a) \wedge \diamond \langle i:a \rangle \top \rightarrow \bigvee_{b \in \text{Act}} (\langle \beta_{-i}, i:a \rangle <_i \langle \beta_{-i}, i:b \rangle \wedge \diamond \langle i:b \rangle \top)$$

by Definition $\text{SD}^{\leq 0}(i:a)$ and Boolean principles;

$$7. \vdash_{\mathcal{ML}\mathcal{EG}} \diamond \langle \beta_{-i} \rangle \top \wedge \bigwedge_{a \in \text{Act}} \text{SD}^{\leq 0}(i:a) \wedge \diamond \langle i:a \rangle \top$$

$$\rightarrow \bigvee_{b \in \text{Act}} (\langle \beta_{-i}, i:a \rangle <_i \langle \beta_{-i}, i:b \rangle \wedge \diamond \langle \beta_{-i} \rangle \top \wedge \bigwedge_{a \in \text{Act}} \text{SD}^{\leq 0}(i:a) \wedge \diamond \langle i:b \rangle \top)$$

by 6 and Boolean principles to propagate $\bigwedge_{a \in \text{Act}} \text{SD}^{\leq 0}(i:a)$;

$$8. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \diamond\langle\beta_{-i}\rangle\top \wedge \bigwedge_{a \in Act} \text{SD}^{\leq 0}(i:a) \wedge \diamond\langle i:a\rangle\top$$

$$\rightarrow \bigvee_{b_1 \in Act} \bigvee_{b_2 \in Act} \cdots \bigvee_{b_n \in Act} \langle\beta_{-i}, i:a\rangle <_i \langle\beta_{-i}, i:b_1\rangle \langle\beta_{-i}, i:b_1\rangle <_i \langle\beta_{-i}, i:b_2\rangle \cdots$$

$$\langle\beta_{-i}, i:b_{n-1}\rangle <_i \langle\beta_{-i}, i:b_n\rangle \text{ for all } n \geq 1;$$

by 7 and Boolean principles (induction on n);

$$9. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \bigvee_{b_1 \in Act} \bigvee_{b_2 \in Act} \cdots \bigvee_{b_n \in Act} \langle\beta_{-i}, i:a\rangle <_i \langle\beta_{-i}, i:b_1\rangle \langle\beta_{-i}, i:b_1\rangle <_i \langle\beta_{-i}, i:b_2\rangle \cdots$$

$$\langle\beta_{-i}, i:b_{n-1}\rangle <_i \langle\beta_{-i}, i:b_n\rangle$$

$$\rightarrow \bigvee_{b \in Act} \langle\beta_{-i}, i:b\rangle <_i \langle\beta_{-i}, i:b\rangle \text{ if } n > \text{card}(Agt)$$

by Boolean principles and because $<_i$ is transitive (as $n > \text{card}(Agt)$, all sequence b_1, \dots, b_n are such that there exists $i \neq j$ such that $b_i = b_j$);

$$10. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \langle\beta_{-i}, i:b\rangle <_i \langle\beta_{-i}, i:b\rangle \rightarrow \perp$$

by Definition of $<_i$ and Boolean principles;

$$11. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \diamond\langle\beta_{-i}\rangle\top \wedge \bigwedge_{a \in Act} \text{SD}^{\leq 0}(i:a) \wedge \diamond\langle i:a\rangle\top \rightarrow \perp$$

by 8, 9, 10;

$$12. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \bigwedge_{a \in Act} \text{SD}^{\leq 0}(i:a) \wedge \bigvee_{a \in Act} \diamond\langle i:a\rangle\top \wedge \bigvee_{\beta \in \Delta} \diamond\langle\beta_{-i}\rangle\top \rightarrow \perp$$

by 11 and Boolean principles;

$$13. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \bigwedge_{a \in Act} \text{SD}^{\leq 0}(i:a) \rightarrow \perp$$

by 5 and 12;

$$14. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}} \bigvee_{a \in Act} \neg \text{SD}^{\leq 0}(i:a)$$

by 13 and Boolean principle.

Lemma 5. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \neg \text{SD}^{\leq n}(i:a) \wedge \chi_{\text{SD}} \leftrightarrow \langle\chi_{\text{SD}}!\rangle \neg \text{SD}^{\leq n-1}(i:a)$.

Proof. We prove it by induction. Let us consider the case $n = 0$ where $\neg \text{SD}^{\leq -1}(i:a) = \diamond\langle i:a\rangle\top$ by convention.

1. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \neg\text{SD}^{\leq 0}(i:a) \rightarrow \Diamond\langle i:a \rangle\top$
by Definition of $\neg\text{SD}^{\leq 0}(i:a)$ and Boolean principles;
2. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \bigwedge_{j \in \text{Agt}} \bigvee_{b \in \text{Act}} \neg\text{SD}^{\leq 0}(j:b)$
by Lemma 4 and Boolean principles;
3. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \bigwedge_{j \in \text{Agt}} \bigvee_{b \in \text{Act}} \neg\text{SD}^{\leq 0}(j:b) \leftrightarrow \bigwedge_{j \in \text{Agt}} \bigvee_{b \in \text{Act}} \neg\text{SD}^{\leq 0}(j:b) \wedge \Diamond\langle j:b \rangle\top$
by 1 and Boolean principles;
4. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \bigwedge_{j \in \text{Agt}} \bigvee_{b \in \text{Act}} \neg\text{SD}^{\leq 0}(j:b) \wedge \Diamond\langle j:b \rangle\top$
by 2 and 3;
5. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \neg\text{SD}^{\leq 0}(i:a) \wedge \chi_{\text{SD}} \leftrightarrow \neg\text{SD}^{\leq 0}(i:a) \wedge \chi_{\text{SD}} \wedge \Diamond\langle i:a \rangle\top$
by 1 and Boolean principles;
6. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \neg\text{SD}^{\leq 0}(i:a) \wedge \chi_{\text{SD}} \wedge \Diamond\langle i:a \rangle\top \leftrightarrow \neg\text{SD}^{\leq 0}(i:a) \wedge \chi_{\text{SD}} \wedge \Diamond\langle i:a \rangle\top \wedge \bigwedge_{j \in \text{Agt}} \bigvee_{b \in \text{Act}} \neg\text{SD}^{\leq 0}(j:b) \wedge \Diamond\langle j:b \rangle\top$
by 4 and Boolean principles;
7. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \neg\text{SD}^{\leq 0}(i:a) \wedge \chi_{\text{SD}} \wedge \Diamond\langle i:a \rangle\top \wedge \bigwedge_{j \in \text{Agt}} \bigvee_{b \in \text{Act}} \neg\text{SD}^{\leq 0}(j:b) \wedge \Diamond\langle j:b \rangle\top$
 $\rightarrow \chi_{\text{SD}} \wedge \bigvee_{\beta \in \Delta} (\neg\text{SD}^{\leq 0}(\beta_{-i}, i:a) \wedge \Diamond\langle \beta_{-i}, i:a \rangle\top)$
by Boolean principles and **Indep**;
8. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \chi_{\text{SD}} \wedge \bigvee_{\beta \in \Delta} (\neg\text{SD}^{\leq 0}(\beta_{-i}, i:a) \wedge \Diamond\langle \beta_{-i}, i:a \rangle\top)$
 $\rightarrow \chi_{\text{SD}} \wedge \bigvee_{\beta \in \Delta} \Diamond(\neg\text{SD}^{\leq 0}(\beta_{-i}, i:a) \wedge \langle \beta_{-i}, i:a \rangle\top)$
by Proposition 2 and **K**(\Box) principles;
9. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \chi_{\text{SD}} \wedge \bigvee_{\beta \in \Delta} \Diamond(\neg\text{SD}^{\leq 0}(\beta_{-i}, i:a) \wedge \langle \beta_{-i}, i:a \rangle\top)$
 $\rightarrow \chi_{\text{SD}} \wedge \bigvee_{\beta \in \Delta} \Diamond(\chi_{\text{SD}} \wedge \langle i:a \rangle\top)$
by definitions of χ_{SD} , $\neg\text{SD}^{\leq 0}(\beta_{-i}, i:a)$, **JointAct** and Boolean principles;
10. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle \chi_{\text{SD}}! \rangle \langle i:a \rangle\top \leftrightarrow \langle i:a \rangle\top \wedge \chi_{\text{SD}}$ by **R2.** and **R7.**
11. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \chi_{\text{SD}} \wedge \bigvee_{\beta \in \Delta} \Diamond(\chi_{\text{SD}} \wedge \langle i:a \rangle\top) \rightarrow \chi_{\text{SD}} \wedge \Diamond(\chi_{\text{SD}} \wedge \langle i:a \rangle\top)$;
by Boolean principles;
12. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \chi_{\text{SD}} \wedge \Diamond(\chi_{\text{SD}} \wedge \langle i:a \rangle\top) \leftrightarrow \chi_{\text{SD}} \wedge \Diamond(\langle \chi_{\text{SD}}! \rangle \langle i:a \rangle\top)$
by Boolean principles and 10;
13. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \chi_{\text{SD}} \wedge \Diamond(\langle \chi_{\text{SD}}! \rangle \langle i:a \rangle\top) \leftrightarrow \langle \chi_{\text{SD}}! \rangle \Diamond\langle i:a \rangle\top$.
by **R2.** and **R4.**;
14. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \neg\text{SD}^{\leq 0}(i:a) \wedge \chi_{\text{SD}} \rightarrow \langle \chi_{\text{SD}}! \rangle \neg\text{SD}^{\leq -1}(i:a)$
by 5, 6, 7, 8, 9, 11, 12, 13.

$$15. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle \chi_{SD}! \rangle \neg SD^{\leq -1}(i:a) \rightarrow \chi_{SD} \wedge \Diamond(\chi_{SD} \wedge \langle i:a \rangle \top)$$

by 13, 12;

$$16. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle i:a \rangle \top \wedge \chi_{SD} \rightarrow \neg SD^{\leq 0}(i:a) \wedge \chi_{SD};$$

by Definition of χ_{SD} and Boolean principles;

$$17. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \chi_{SD} \wedge \Diamond(\chi_{SD} \wedge \langle i:a \rangle \top) \rightarrow \chi_{SD} \wedge \Diamond \neg SD^{\leq 0}(i:a)$$

by 16, modal logic $K(\Box)$ principles and Boolean principles;

$$18. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \Diamond \neg SD^{\leq 0}(i:a) \wedge \chi_{SD} \rightarrow \neg SD^{\leq 0}(i:a) \wedge \chi_{SD}$$

by $\Diamond \neg SD^{\leq 0}(i:a) \leftrightarrow \neg SD^{\leq 0}(i:a)$;

$$19. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle \chi_{SD}! \rangle \neg SD^{\leq -1}(i:a) \rightarrow \neg SD^{\leq 0}(i:a) \wedge \chi_{SD}$$

by 15, 17, 18;

$$20. \vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \neg SD^{\leq 0}(i:a) \wedge \chi_{SD} \leftrightarrow \langle \chi_{SD}! \rangle \neg SD^{\leq -1}(i:a)$$

by 14 and 19.

Now let us consider the inductive case.

1. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle \chi_{SD!} \rangle \neg SD^{\leq n}(i:a) \leftrightarrow \langle \chi_{SD!} \rangle$
 $\left(\neg SD^{\leq n-1}(i:a) \wedge \bigwedge_{b \in Act} (\neg SD^{\leq n-1}(i:b) \rightarrow \bigvee_{\beta \in \Delta} (\neg SD^{\leq n}(\beta_{-i}) \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle))) \right)$
 by Definition of $\neg SD^{\leq n}(i:a)$;
2. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle \chi_{SD!} \rangle$
 $\left(\neg SD^{\leq n-1}(i:a) \wedge \bigwedge_{b \in Act} (\neg SD^{\leq n-1}(i:b) \rightarrow \bigvee_{\beta \in \Delta} (\neg SD^{\leq n-1}(\beta_{-i}) \wedge (\langle \beta_{-i}, i:b \rangle \leq_i \langle \beta_{-i}, i:a \rangle))) \right)$
 $\leftrightarrow (\langle \chi_{SD!} \rangle \neg SD^{\leq n-1}(i:a) \wedge \bigwedge_{b \in Act} (\langle \chi_{SD!} \rangle \neg SD^{\leq n-1}(i:b) \rightarrow$
 $\bigvee_{\beta \in \Delta} (\langle \chi_{SD!} \rangle \neg SD^{\leq n-1}(\beta_{-i}) \wedge \langle \chi_{SD!} \rangle (\langle \beta_{-i}, i:b \rangle \wedge \leq_i \langle \beta_{-i}, i:a \rangle)))$
 by Boolean principles and rules **R2.** and **R3.** (we can distribute $\langle \chi_{SD!} \rangle$ over Boolean connectives);
3. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle \chi_{SD!} \rangle (\langle \beta_{-i}, i:b \rangle \wedge \leq_i \langle \beta_{-i}, i:a \rangle) \leftrightarrow (\langle \beta_{-i}, i:b \rangle \wedge \leq_i \langle \beta_{-i}, i:a \rangle) \wedge \chi_{SD}$
 by Boolean principles and Axiom **R2., R3., R4.** and **R6.;**
4. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} (\langle \chi_{SD!} \rangle \neg SD^{\leq n-1}(i:a) \wedge \bigwedge_{b \in Act} (\langle \chi_{SD!} \rangle \neg SD^{\leq n-1}(i:b) \rightarrow$
 $\bigvee_{\beta \in \Delta} (\langle \chi_{SD!} \rangle \neg SD^{\leq n-1}(\beta_{-i}) \wedge \langle \chi_{SD!} \rangle (\langle \beta_{-i}, i:b \rangle \wedge \leq_i \langle \beta_{-i}, i:a \rangle)))$
 $\leftrightarrow \chi_{SD} \wedge \neg SD^{\leq n}(i:a) \wedge \bigwedge_{b \in Act} (\chi_{SD} \wedge \neg SD^{\leq n}(i:b) \rightarrow \bigvee_{\beta \in \Delta} (\chi_{SD} \wedge \neg SD^{\leq n}(\beta_{-i}) \wedge$
 $\chi_{SD} \wedge (\langle \beta_{-i}, i:b \rangle \wedge \leq_i \langle \beta_{-i}, i:a \rangle)))$
 by induction and 3;
5. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} [\chi_{SD} \wedge \neg SD^{\leq n}(i:a) \wedge \bigwedge_{b \in Act} (\chi_{SD} \wedge \neg SD^{\leq n}(i:b) \rightarrow \bigvee_{\beta \in \Delta} (\chi_{SD} \wedge$
 $\neg SD^{\leq n}(\beta_{-i}) \wedge \chi_{SD} \wedge (\langle \beta_{-i}, i:b \rangle \wedge \leq_i \langle \beta_{-i}, i:a \rangle)))] \leftrightarrow [\chi_{SD} \wedge$
 $\neg SD^{\leq n}(i:a) \wedge \underbrace{\bigwedge_{b \in Act} (\neg SD^{\leq n}(i:b) \rightarrow \bigvee_{\beta \in \Delta} (\neg SD^{\leq n}(\beta_{-i}) \wedge (\langle \beta_{-i}, i:b \rangle \wedge \leq_i \langle \beta_{-i}, i:a \rangle)))}_{\neg SD^{\leq n+1}(i:a)}]$
 by Boolean principles (we remove the multiple “ $\chi_{SD} \wedge$ ”);
6. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle \chi_{SD!} \rangle \neg SD^{\leq n}(i:a) \leftrightarrow \langle \chi_{SD!} \rangle \leftrightarrow \chi_{SD} \wedge \neg SD^{\leq n+1}(i:a)$
 by 1, 2, 4, 5.

Now let us finish the proof:

1. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \neg\text{SD}^{\leq n}(\delta) \rightarrow \neg\text{SD}^{\leq n-1}(\delta)$
by definition of $\neg\text{SD}^{\leq n}(\delta)$ and Boolean principles;
2. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \neg\text{SD}^{\leq n}(\delta) \rightarrow \neg\text{SD}^{\leq 0}(\delta)$
by 1 and Boolean principles (induction on n);
3. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \neg\text{SD}^{\leq 0}(\delta) \wedge \langle\delta\rangle\top \rightarrow \chi_{\text{SD}}$
by Boolean principles (see definition of χ_{SD});
4. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle\delta\rangle\top \wedge \chi_{\text{SD}} \leftrightarrow \langle\chi_{\text{SD}}!\rangle\langle\delta\rangle\top$
by rule **R7.**;
5. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle\delta\rangle\top \wedge \neg\text{SD}^{\leq n}(\delta) \wedge \chi_{\text{SD}} \rightarrow \langle\chi_{\text{SD}}!\rangle(\langle\delta\rangle\top \wedge \neg\text{SD}^{\leq n-1}(\delta) \wedge \chi_{\text{SD}})$
by 3, 4, Lemma 5, **R2.** and **R3.**;
6. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle\delta\rangle\top \wedge \neg\text{SD}^{\leq n}(\delta) \wedge \chi_{\text{SD}} \rightarrow \langle\chi_{\text{SD}}!\rangle^{n+1}(\langle\delta\rangle\top)$
by induction with 5;
7. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle\delta\rangle\top \wedge \langle\chi_{\text{SD}}!\rangle^{n+1}\top \rightarrow \langle\chi_{\text{SD}}!\rangle(\langle\delta\rangle\top \wedge \langle\chi_{\text{SD}}!\rangle^n\top)$
by rule **R2.** and **R3.**;
8. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle\delta\rangle\top \wedge \langle\chi_{\text{SD}}!\rangle^{n+1}\top \rightarrow \langle\chi_{\text{SD}}!\rangle^{n+1}\langle\delta\rangle\top$
by 7 and induction;
9. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle\chi_{\text{SD}}!\rangle^{n+1}\langle\delta\rangle\top \rightarrow \langle\chi_{\text{SD}}!\rangle^{n+1}\diamond\langle\delta\rangle\top$
by **R2.**, **R4.** and **T** for \square and Boolean principles;
10. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle\chi_{\text{SD}}!\rangle^{n+1}\diamond\langle\delta\rangle\top \rightarrow \neg\text{SD}^{\leq n}(\delta)$
by Lemma 5 and induction;
11. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle\delta\rangle\top \rightarrow (\neg\text{SD}^{\leq n}(\delta) \rightarrow \langle\chi_{\text{SD}}!\rangle^{n+1}\top)$
by 2, 3, 6;
12. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle\delta\rangle\top \rightarrow (\langle\chi_{\text{SD}}!\rangle^{n+1}\top \rightarrow \neg\text{SD}^{\leq n}(\delta))$
by 8, 9, 10;
13. $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \langle\delta\rangle\top \rightarrow (\langle\chi_{\text{SD}}!\rangle^{n+1}\top \leftrightarrow \neg\text{SD}^{\leq n}(\delta))$
by 11 and 12.

A.12. Proof of Theorem 13.

For all $n \geq 0$, $\vdash_{\mathcal{ML}\mathcal{E}\mathcal{G}^{GT}} \left(\text{MK}_{\text{Agt}}^n \bigwedge_{i \in \text{Agt}} \text{Rat}_i \right) \rightarrow \langle\chi_{\text{SD}}!\rangle^{n+1}\top$.

Proof. By Theorem 7, Theorem 12 and Boolean principles.

A.13. Proof of Theorem 14

The satisfiability problem of a given formula φ in a $\mathcal{ML}\mathcal{E}\mathcal{G}^*$ -model is PSPACE-hard.

Proof. Let us prove that the satisfiability problem of $\mathcal{ML}\mathcal{E}\mathcal{G}^*$ is PSPACE-hard. Let i be an agent. Let us consider a formula φ written only with atomic propositions and with modal operators \Box and K_i . We have equivalence between:

1. φ is satisfiable in a $\mathcal{ML}\mathcal{E}\mathcal{G}^*$ -model;
2. φ is satisfiable in a model of the logic $S5_2(\Box, K_i)$ (i.e. the fusion of the logic S5 for \Box and S5 for K_i).

Hence, we have reduced the satisfiability problem of a given formula φ in a $\mathcal{ML}\mathcal{E}\mathcal{G}^*$ -model to the satisfiability problem of a given formula φ of $S5_2(\Box, K_i)$ which is PSPACE-hard. So the satisfiability problem of a given formula φ in a $\mathcal{ML}\mathcal{E}\mathcal{G}^*$ -model is PSPACE-hard.

A.14. Proof of Theorem 15

- If $\text{card}(Agt) = 1$ and $\text{card}(Act) = 1$ then the satisfiability problem of a given formula φ in a $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$ -model is NP-complete.
- If $\text{card}(Agt) \geq 2$ or $\text{card}(Act) \geq 2$ the satisfiability problem of a given formula φ in a $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$ -model is PSPACE-complete.

Proof. We give here some hint for the proof. When there is only one agent and $\text{card}(Act) = 1$ then the games are trivial and reduced to singletons. In these settings, a $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$ -frame $F = \langle W, \sim, R, E, \preceq \rangle$ is such that \sim and \preceq_i for each agent i are equal to the relation $\{(w, w) \mid w \in W\}$. So the modal operators $[\text{good}]_i$ and \Box are superfluous. The operator $[\delta_C]$ can be treated as a proposition. Hence the logic is similar to the logic S5 which is NP. This is the main argument why when there is only one agent and $\text{card}(Act) = 1$ the logic $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$ is NP. NP-hardness is granted because $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$ is a conservative extension of Classical Propositional Logic.

Now let us prove that the satisfiability problem of a given formula φ in a $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$ -model is PSPACE-hard in other cases. First let us consider the case where $\text{card}(Agt) \geq 2$. Let us consider two distinct agents $i, j \in Agt$. Let φ be a formula written only with atomic propositions and with epistemic modal operators K_i and K_j . We have equivalence between:

1. φ is satisfiable in a $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$ -model;
2. φ is satisfiable in the logic $S5_2(K_i, K_j)$ (i.e., the fusion of the logic S5 for K_i and S5 for K_j).

The direction 1. \rightarrow 2. is straightforward and is already true with the assumption of the Axiom **CompleteInfo**. The direction 2. \rightarrow 1. comes from the fact that the Constraint **C4** (corresponding to the Axiom **CompleteInfo**) has disappeared. So we can easily transform a model of the epistemic modal logic into a $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$ -model. Note that in the case of the logic $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$, the direction 2. \rightarrow 1. is not true anymore. Indeed, it is not possible to transform a model of $S5_2(K_a, K_b)$ with more than

$card(Act)^{card(Agt)}$ worlds into a $\mathcal{ML}\mathcal{E}\mathcal{G}^{det}$ -model. Hence, we have reduced the satisfiability problem of a given formula φ in a $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$ -model into the satisfiability problem of a given formula ψ of $\text{S5}_2(K_i, K_j)$ which is PSPACE-hard. So the satisfiability problem of a given formula φ in a $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$ -model is PSPACE-hard.

Now let us consider the case where $Agt = \{i\}$ and $card(Act) \geq 2$. Let a and b be two distinct actions. We prove that we can reduce the satisfiability problem of a given formula φ in a $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$ -model to the satisfiability problem of K . Here is a possible translation:

- $tr_0(\blacksquare\psi) = i:a \wedge \Diamond K_i tr_1(\psi)$ where \blacksquare is the K -operator;
- $tr_1(\blacksquare\psi) = i:b \wedge \Diamond K_i tr_0(\psi)$ where \blacksquare is the K -operator;
- $tr_0(p) = i:a \wedge p$ for all propositions p ;
- $tr_1(p) = i:b \wedge p$ for all propositions p .

And φ is satisfiable in K iff $tr_0(\varphi)$ is satisfiable in $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$. Hence, the logic $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$ is also PSPACE-hard in this case.

Now we are going to prove that the satisfiability problem of $\mathcal{ML}\mathcal{E}\mathcal{G}^{det*}$ is PSPACE. We do not give all the details but we give the idea for a tableau method [21] for the logic $\mathcal{ML}\mathcal{E}\mathcal{G}^*$. The tableau method is a non-deterministic procedure. The creation of a model proceeds as follows:

- We start the procedure by guessing a “grid”, that is to say an equivalence class for the relation \sim of maximal size $card(Act)^{card(Agt)}$ and also its preference relation as in the algorithm of Theorem 2. We also choose non-deterministically a world w in this class.
- We adapt the classical tableau method rules for the epistemic modal logic [21], that is to say:
 - Suppose that a world w contains a formula of the form $K_i\psi$. Then we propagate the formula ψ in all nodes v such that wE_iv .
 - Suppose that a world w contains a formula of the form $\widehat{K}_i\psi$. Then we create an equivalence class for \sim , we choose a point v such that $R_i(v) = R_i(w)$ in this equivalence class and we propagate ψ in v .
- Suppose that a node w contains a formula $\Box\psi$. Then we propagate the formula ψ in all nodes v such that $v \sim w$;
- Suppose that a node w contains a formula $\Diamond\psi$. Then we choose non-deterministically a world v such that $v \sim w$ and we propagate ψ in v .
- Suppose that a node w contains a formula $[\text{good}]_i\varphi\psi$. Then we propagate the formula ψ in all nodes v such that $v \preceq_i w$;
- Suppose that a node w contains a formula $\langle \text{good} \rangle_i\psi$. Then we choose non-deterministically a world v such that $v \preceq_i w$ and we propagate ψ in v .

During the construction, we explore the structure in depth first so that we only need to have one branch in memory at each step. Thus, the algorithm is a non-deterministic procedure that uses only a polynomial amount of memory. So the satisfiability problem of $\mathcal{ML}\mathcal{EG}^{det*}$ is in NPSpace. According to the Savitch's theorem [48], it is in PSPACE.

A.15. Proof of Theorem 16

For every $w \in W$ and for every $i \in \text{Agt}$, $M, w \models \text{Rat}_i$ if and only if $M', w \models \text{Rat}_i$.

Proof. (\Rightarrow) We first prove the left-to-right direction. Suppose that $M, w \models \langle i:a \rangle \top$ and $M, w \models \text{Rat}_i$. The latter means that for every $b \in \text{Act}$ there is $\delta \in \Delta$ such that $M, w \models \widehat{K}_i \langle \delta_{-i} \rangle \top \wedge K_i (\langle \delta_{-i}, i:b \rangle \leq_i \langle \delta_{-i}, i:a \rangle)$.

The latter means that for every $b \in \text{Act}$ there is $\delta \in \Delta$ such that:

- A. there is $v \in E_i(w)$ such that $R_{-i}(v) = \delta_{-i}$ and
- B. for all $u, v \in W$ if $v \in E_i(w)$ and $u \sim v$ and $R_{\text{Agt}}(u) = \langle \delta_{-i}, i:b \rangle$ then there is $z \in W$ such that $u \preceq_i z$ and $R_{\text{Agt}}(z) = \langle \delta_{-i}, i:a \rangle$.

Consider an arbitrary $b \in \text{Agt}$. It follows that there is an element δ of Δ which satisfies the previous conditions A and B. By definition of $E'_i(w)$ and $R'_{-i}(v)$ we have that: if there is $v \in E_i(w)$ such that $R_{-i}(v) = \delta_{-i}$ then there is $v \in E'_i(w)$ such that $R'_{-i}(v) = \delta_{-i}$. Therefore, from item A, we conclude:

- C. there is $v \in E'_i(w)$ such that $R'_{-i}(v) = \delta_{-i}$.

Take two arbitrary worlds $u, v \in W'$ and suppose that $v \in E'_i(w)$ and $u \sim' v$ and $R'_{\text{Agt}}(u) = \langle \delta_{-i}, i:b \rangle$. By definition of M' , we have $v \in E_i(w)$ and $\langle \delta_{-i}, i:a \rangle \in \Pi_M$. Consider now the following three cases.

CASE 1. Suppose $u \in W$. Then, by definition of M' , we have that $u \sim v$ and $R_{\text{Agt}}(u) = \langle \delta_{-i}, i:b \rangle$. Therefore, from item B, it follows that there is $z \in W$ such that $u \preceq_i z$ and $R_{\text{Agt}}(z) = \langle \delta_{-i}, i:a \rangle$. From the latter, by definition of M' , we conclude that there is $z \in W'$ such that $u \preceq'_i z$ and $R'_{\text{Agt}}(z) = \langle \delta_{-i}, i:a \rangle$.

CASE 2. Suppose that $u \notin W$ and that there is $z \in W$ such that $z \sim v$ and $R_{\text{Agt}}(z) = \langle \delta_{-i}, i:a \rangle$. From the former, by definition of M' , it follows that $\preceq'_i(u) = \sim'(u)$. From the latter, by definition of M' , it follows that there is $z \in W'$ such that $z \sim' v$ and $R'_{\text{Agt}}(z) = \langle \delta_{-i}, i:a \rangle$. Therefore, we have that there is $z \in W'$ such that $u \preceq'_i z$ and $R'_{\text{Agt}}(z) = \langle \delta_{-i}, i:a \rangle$.

CASE 3. Suppose that $u \notin W$ and that there is no $z \in W$ such that $z \sim v$ and $R_{\text{Agt}}(z) = \langle \delta_{-i}, i:a \rangle$. From the former, by definition of M' , it follows that $\preceq'_i(u) = \sim'(u)$. From the latter, by definition of M' (and the fact that $\langle \delta_{-i}, i:a \rangle \in \Pi_M$), it follows that there is $z \in W'$ such that $z \sim' v$ and $R'_{\text{Agt}}(z) = \langle \delta_{-i}, i:a \rangle$ and $\preceq'_i(z) = \sim'(z)$. Therefore, we have that there is $z \in W'$ such that $u \preceq'_i z$ and $R'_{\text{Agt}}(z) = \langle \delta_{-i}, i:a \rangle$.

From the previous three cases, it follows that:

- D. for all $u, v \in W'$ if $v \in E'_i(w)$ and $u \sim' v$ and $R'_{\text{Agt}}(u) = \langle \delta_{-i}, i:b \rangle$ then there is $z \in W'$ such that $u \preceq'_i z$ and $R'_{\text{Agt}}(z) = \langle \delta_{-i}, i:a \rangle$.

From the items C and D we conclude that $M', w \models \text{Rat}_i$.

(\Leftarrow) Let us prove the right-to-left direction. Suppose that $M', w \models \langle i:a \rangle \top$ and $M', w \models \text{Rat}_i$. The latter means that for every $b \in \text{Act}$ there is $\delta \in \Delta$ such that $M', w \models \widehat{K}_i \langle \delta_{-i} \rangle \top \wedge K_i (\langle \delta_{-i}, i:b \rangle \leq_i \langle \delta_{-i}, i:a \rangle)$.

The latter means that for every $b \in \text{Act}$ there is $\delta \in \Delta$ such that:

E. there is $v \in E'_i(w)$ such that $R'_{-i}(v) = \delta_{-i}$ and

F. for all $u, v \in W'$ if $v \in E'_i(w)$ and $u \sim' v$ and $R'_{\text{Agt}}(u) = \langle \delta_{-i}, i:b \rangle$ then there is $z \in W'$ such that $u \preceq'_i z$ and $R'_{\text{Agt}}(z) = \langle \delta_{-i}, i:a \rangle$.

Consider an arbitrary $b \in \text{Agt}$. It follows that there is an element δ of Δ which satisfies the previous conditions E and F. By definition of $E'_i(w)$ and $R'_{-i}(v)$ we have that: if there is $v \in E'_i(w)$ such that $R'_{-i}(v) = \delta_{-i}$ then there is $v \in E_i(w)$ such that $R_{-i}(v) = \delta_{-i}$. Therefore, from item E, we conclude:

G. there is $v \in E_i(w)$ such that $R_{-i}(v) = \delta_{-i}$.

Take two arbitrary worlds $u, v \in W$ and suppose that $v \in E_i(w)$ and $u \sim v$ and $R_{\text{Agt}}(u) = \langle \delta_{-i}, i:b \rangle$. By Definition of $E'_i(w)$, \sim' and $R'_{\text{Agt}}(u)$ we have $v \in E'_i(w)$ and $u \sim' v$ and $R'_{\text{Agt}}(u) = \langle \delta_{-i}, i:b \rangle$. Thus, by the previous item F, there is $z \in W'$ such that $u \preceq'_i z$ and $R'_{\text{Agt}}(z) = \langle \delta_{-i}, i:a \rangle$. But as $u \in W$ we have $\preceq'_i(u) = \preceq_i(u)$. Thus $z \in W$ and $R'_{\text{Agt}}(z) = R_{\text{Agt}}(z)$.

It follows that:

H. for all $u, v \in W$ if $v \in E_i(w)$ and $u \sim v$ and $R_{\text{Agt}}(u) = \langle \delta_{-i}, i:b \rangle$ then there is $z \in W$ such that $u \preceq_i z$ and $R_{\text{Agt}}(z) = \langle \delta_{-i}, i:a \rangle$.

From the items G and H we conclude that $M, w \models \text{Rat}_i$.

A.16. Proof of Corollary 1

For every $w \in W$, for every $i \in \text{Agt}$ and for every $C \in 2^{\text{Agt}^*}$, we have $M, w \models \text{MK}_C^n \text{Rat}_i$ if and only if $M', w \models \text{MK}_C^n \text{Rat}_i$.

Proof. Define a world v to be C -reachable from world w in n steps (with $n \geq 1$), and note this $w E_{C,n} v$, if and only if there exist worlds w_0, \dots, w_n such that $w_0 = w$ and $w_n = v$ and for all $0 \leq k \leq n-1$, there exists $i \in C$ such that $w_k E_i w_{k+1}$. Define $E_{C,n}(w) = \{v \mid w E_{C,n} v\}$. We have $M, w \models \text{MK}_C^n \varphi$ if and only if $M, v \models \varphi$ for all $v \in E_{C,n}(w)$.

By definition of M' , we have $E_{C,n}(w) = E'_{C,n}(w)$ for all $w \in W$. Therefore, $M, w \models \text{MK}_C^n \text{Rat}_i$ if and only if $M, v \models \text{Rat}_i$ for all $v \in E'_{C,n}(w)$.

Moreover, according to Theorem 16, for every $v \in W$ we have $M, v \models \text{Rat}_i$ if and only if $M', v \models \text{Rat}_i$. Therefore, we have that $M, v \models \text{Rat}_i$ for all $v \in E'_{C,n}(w)$ if and only if $M', v \models \text{Rat}_i$ for all $v \in E'_{C,n}(w)$.

It follows that, $M, w \models \text{MK}_C^n \text{Rat}_i$ if and only if $M', w \models \text{MK}_C^n \text{Rat}_i$.

Article

Toward a Theory of Play: A Logical Perspective on Games and Interaction

Johan van Benthem^{1,2,*}, Eric Pacuit³ and Olivier Roy^{4,†}

¹ ILLC, University of Amsterdam, 1090 GE Amsterdam, The Netherlands

² Department of Philosophy, Stanford University, Stanford, CA 94305, USA

³ TiLPS, Tilburg University, Warandelaan 2, 5037 AB Tilburg, The Netherlands

⁴ Faculty of Philosophy, University of Groningen, 9700 AB Groningen, The Netherlands

† Current address: Ludwig-Maximilians-Universität München, Munich Center for Mathematical Philosophy, Ludwigstrasse 31, 80539 Munich, Germany.

* Author to whom correspondence should be addressed; E-Mail: J.vanBenthem@uva.nl.

Received: 25 November 2010; in revised form: 9 February 2011 / Accepted: 11 February 2011 /
Published: 16 February 2011

Abstract: Logic and game theory have had a few decades of contacts by now, with the classical results of epistemic game theory as major high-lights. In this paper, we emphasize a recent new perspective toward “logical dynamics”, designing logical systems that focus on the actions that change information, preference, and other driving forces of agency. We show how this dynamic turn works out for games, drawing on some recent advances in the literature. Our key examples are the long-term dynamics of information exchange, as well as the much-discussed issue of extensive game rationality. Our paper also proposes a new broader interpretation of what is happening here. The combination of logic and game theory provides a fine-grained perspective on information and interaction dynamics, and we are witnessing the birth of something new which is not just logic, nor just game theory, but rather a *Theory of Play*.

Keywords: dynamic epistemic logic; games; interaction

For many contemporary logicians, games and social interaction are important objects of investigation. *Actions*, *strategies* and *preferences* are central concepts in computer science and philosophical logic,

and their combination raises interesting questions of *definability*, *axiomatization* and *computational complexity* [1–4]. Epistemic game theory, c.f. [5], has added one more element to this mix, again familiar to logicians: the role of factual and higher-order information. This much is well-understood, and there are excellent sources, that we need not reproduce here, though we will recall a few basics in what follows.

In this paper we will take one step further, assuming that the reader knows the basics of logic and game theory. We are going to take a look at all these components from a dynamic logical perspective, emphasizing actions that make information flow, change beliefs, or modify preferences—in ways to be explained below. For us, understanding social situations as dynamic logical processes where the participants interactively revise their beliefs, change their preferences, and adapt their strategies is a step towards a more finely-structured theory of rational agency. In a simple phrase that sums it up, this joint off-spring “in the making” of logic and game theory might be called a *Theory of Play* instead of a theory of games.

The paper starts by laying down the main components of such a theory, a logical take on the dynamics of actions, preferences, and information (Sections 1 and 2). We then show that this perspective has already shed new light on the long-term dynamics of information exchange, Section 3, as well as on the question of extensive game rationality, Section 4. We conclude with general remarks on the relation between logic and game theory, pleading for cross-fertilization instead of competition. This paper is introductory and programmatic throughout. Our treatment is heavily based on evidence from a number of recent publications demonstrating a variety of new developments.

1. An Encounter Between Logic and Games

A first immediate observation is that games as they stand are natural models for many existing logical languages: epistemic, doxastic and preference logics, as well as conditional logics and temporal logics of action. We do not aim at encyclopedic description of these systems—[2] is a relatively up-to-date overview. This section just gives some examples setting the scene for our later more detailed dynamic-logic analyses.

1.1. Strategic Games

Even simple strategic games call for logical analysis, with new questions arising at once. To a logician, a game matrix is a semantic model of a rather special kind that invites the introduction of well-known languages. Recall the main components in the definition of a *strategic game* for a set of n players N : (1) a nonempty set A_i of actions for each $i \in N$, and (2) a utility function or preference ordering on the set of outcomes. For simplicity, one often identifies the outcomes with the set $S = \prod_{i \in N} A_i$ of *strategy profiles*. As usual, given a strategy profile $\sigma \in S$ with $\sigma = (a_1, \dots, a_n)$, σ_i denotes the i th projection (*i.e.*, $\sigma_i = a_i$) and σ_{-i} denotes the choices of all agents except agent i : $\sigma_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$.

Games as models. Now, from a logical perspective, it is natural to treat the set S of strategy profiles as a universe of “possible worlds”.¹ These worlds then carry three natural relations, that are entangled in various ways. For each $\sigma, \sigma' \in S$, define for each player $i \in N$:

- $\sigma \geq_i \sigma'$ iff player i prefers the outcome σ at least as much as outcome σ' ,
- $\sigma \sim_i \sigma'$ iff $\sigma_i = \sigma'_i$: this epistemic relation represents player i ’s “view of the game” at the *ex interim* stage where i ’s choice is fixed but the choices of the other players’ are unknown,
- $\sigma \approx_i \sigma'$ iff $\sigma_{-i} = \sigma'_{-i}$: this relation of “action freedom” gives the alternative choices for player i when the other players’ choices are fixed.²

This can all be packaged in a relational structure

$$\mathcal{M} = \langle S, \{\sim_i\}_{i \in N}, \{\approx_i\}_{i \in N}, \{\geq_i\}_{i \in N} \rangle$$

with S the set of strategy profiles and the relations just defined.

Matching modal game languages. The next question is what is the “right” logical language to reason about these structures? The goal here is not simply to formalize standard game-theoretic reasoning. That could be done in a number of ways, often in the first-order language of these relational models. Rather, the logician will aim for a *well-behaved* language, with a good balance between the level of formalization and other desirable properties, such as perspicuous axiomatization, low computational complexity of model checking and satisfiability, and the existence of an elegant meta-theory for the system. In particular, the above game models suggest the use of *modal languages*, whose interesting balance of expressive power and computational complexity has been well-researched over the last decades.³

Our first key component—players’ desires or *preferences*—has been the subject of logical analysis since at least the work of [10].⁴ Here is a modern take on preference logic [12,14]. A modal *betterness model* for a set N of players is a tuple $\mathcal{M} = \langle W, \{\geq_i\}_{i \in N}, V \rangle$ where W is a nonempty set of states, for each $i \in N$, $\geq_i \subseteq W \times W$ is a *preference ordering*, and V is a valuation function $V : \text{At} \rightarrow \wp(W)$ (At is a set of atomic propositions describing the ground facts about the situation being modeled). Precisely which properties \geq_i should have has been the subject of debate in philosophy: in this paper, we assume that the relation is reflexive and transitive. For each \geq_i , the corresponding *strict* preference ordering is written $>_i$.

A modal language to describe betterness models uses modalities $\langle \geq_i \rangle \varphi$ saying that “there is a world at least as good as the current world satisfying φ ”, and likewise for strict preference:

- $\mathcal{M}, w \models \langle \geq_i \rangle \varphi$ iff there is a v with $v \geq_i w$ and $\mathcal{M}, v \models \varphi$
- $\mathcal{M}, w \models \langle >_i \rangle \varphi$ iff there is a v with $v \geq_i w$, $w \not\geq_i v$, and $\mathcal{M}, v \models \varphi$

¹We could also have more abstract worlds, carrying strategy profiles without being identical to them. This additional generality is common in epistemic game theory, see e.g. [6], but it is not needed in what follows.

²We have borrowed the appealing term “freedom” from [7].

³We cannot go into details of the modern modal paradigm here, but refer to the textbooks [8,9].

⁴See [11–13] for a contemporary discussion and references.

Standard techniques in modal model theory apply to definability and axiomatization in this modal preference language: we refer to ([9], Chapter 3) and [13] for details. Both [12] and [13] show how this language can also define “lifted” generic preferences between propositions, *i.e.*, properties of worlds.

Next, the full modal game language for the above models must also include modalities for the relations that we called the “view of the game” and the “action freedom”. But this is straightforward, as these are even closer to standard notions studied in epistemic and action logics.

Again, we start with a set At of atomic propositions that represent basic facts about the strategy profiles.⁵ Now, we add obvious modalities for the other two relations to get a full modal logic of strategic games:

- $\sigma \models [\sim_i]\varphi$ iff for all σ' , if $\sigma \sim_i \sigma'$ then $\sigma' \models \varphi$.
- $\sigma \models [\approx_i]\varphi$ iff for all σ' , if $\sigma \approx_i \sigma'$ then $\sigma' \models \varphi$.
- $\sigma \models \langle \geq_i \rangle \varphi$ iff there exists σ' such that $\sigma' \geq_i \sigma$ and $\sigma' \models \varphi$.
- $\sigma \models \langle >_i \rangle \varphi$ iff there is a σ' with $\sigma' \geq_i \sigma$, $\sigma \not\geq_i \sigma'$, and $\sigma' \models \varphi$

Some issues in modal game logic for strategic games. A language allows us to say things about structures. But what about a calculus of reasoning: what is the logic of our modal logic of strategic games? For convenience, we restrict attention to 2-player games. First, given the nature of our three relations, the separate logics are standard: modal **S4** for preference, and modal **S5** for epistemic outlook and action freedom. What is of greater interest, and logical delicacy, is the *interaction* of the three modalities. For instance, the following combination of two modalities makes φ true in each world of a game model:

$$[\sim_i][\approx_i]\varphi$$

Thus, the language also has a so-called “universal modality”. Moreover, this modality can be defined in two ways, since we also have that:

the equivalence $[\sim_i][\approx_i]\varphi \leftrightarrow [\approx_i][\sim_i]\varphi$ is valid⁶ in game models.

This validity depends on the geometrical “grid property” of game matrices that if one can go $x \sim_i y \approx_i z$, then there exists a point u with $x \approx_i u \sim_i z$.

This may look like a pleasant structural feature of matrices, but its logical effects are delicate. It is well-known that the general logic of such a bi-modal language on grid models is not decidable, and not even axiomatizable: indeed, it is “ Π_1^1 -complete”.⁷ In particular, satisfiability in grid-like models can encode computations of Turing machines on their successive rows, or alternatively, they can encode geometrical “tiling problems” whose complexity is known to be high. From a logical point of view, simple-looking strategic matrix games can be quite complex computational structures.

⁵For example, a proposition p_i^a might say “agent i plays action a ”.

⁶A formula is valid in a class of model whenever it is true at all states in all models of that class. C.f. [9, chap.1] for details.

⁷Cf. [15–18] for formal details behind the assertions in this paragraph.

However, there are two ways in which these complexity results can be circumvented. One is that we have mainly looked at finite games, where additional validities hold⁸—and then, the complexity may be lower. Determining the precise modal logic of finite game matrices appears to be an open problem.

Here is another interesting point. It is known that the complexity of such logics may go down drastically when we allow more models, in particular, models where some strategy profiles have been ruled out. One motivation for this move has to do with *dependence* and *independence* of actions.⁹ Full matrix models make players' actions independent, as reflected in the earlier grid property. By contrast, general game models omitting some profiles can represent dependencies between players' actions: changing a move for one may only be possible by changing a move for another. The general logic of game models allowing dependencies does not validate the above commutation law. Indeed, it is much simpler: being just multi-agent modal S5. Thus, complexity of logics matches interesting decisions on how we view players: as independent, or correlated.

Against this background of available actions, information, and freedom, the preference structure of strategic games adds further interesting features. One benchmark for modal game logics has been the definition of the strategy profiles that are in Nash Equilibrium. And this requires defining the usual notion of *best response* for a player. One can actually prove¹⁰ that best response is not definable in the language that we have so far. One extension that would do the job is taking an *intersection modality*:

$$\mathcal{M}, \sigma \models \langle \approx_i \cap >_i \rangle \varphi \text{ iff for each } \sigma' \text{ if } \sigma(\approx_i \cap >_i) \sigma' \text{ then } \mathcal{M}, \sigma' \models \varphi$$

Then the best response for player i is defined as $\neg \langle \approx_i \cap >_i \rangle \top$.

Questions of complexity and complete axiomatization then multiply. But we can also deal with preference structure in other ways. Introduce proposition letters "*Best(i)*" for players i saying that the profiles where they hold are best responses for i in the game model. Then one finds interesting properties of such models reflected in the logic. One example is that each finite game model has a cycle of points where (for simplicity assume there are only two players i and j):

$$\sigma \sim_i \sigma^1 \sim_j \sigma^2 \sim_i \dots \sim_i \sigma^n \sim_j \sigma$$

where $\sigma \models \text{Best}(i)$, $\sigma^1 \models \text{Best}(j)$, $\sigma^2 \models \text{Best}(i)$, \dots , $\sigma^n \models \text{Best}(j)$. Such loops represent subgames where all players are "strongly rational" in the sense of considering it possible that their current move is a best response to what their opponent is doing. Thus, the logic encodes basic game theory.¹¹

Our main point with this warm-up discussion for our logical Theory of Play is that the simple matrix pictures that one sees in a beginner's text on game theory *are* already models for quite sophisticated logics of action, knowledge and preference. Thus, games of even the simplest sort have hidden depths

⁸Cf. [4] for some concrete examples of modal "Gregorczyk axioms".

⁹Cf. [19] for this, and what follows. Game theorists have also studied correlations extensively, c.f. [20,21]. The precise relation between our logical and their probabilistic approaches to correlations is still to be investigated.

¹⁰We omit the simple modal "bisimulation"-based argument here.

¹¹Cf. [22] for technical details, including connections to epistemic "fixed-point logics" over game models, as well as applications to game solution procedures.

for logicians: there is much more to them than we might think, including immediate open problems for logical research.¹²

1.2. Extensive Games

Just like strategic games, interactive agency in the more finely-structured *extensive games* offers a natural meeting point with logic. We will demonstrate this with a case study of *Backwards Induction*, a famous benchmark at the interface, treated in a slightly novel way. Our treatment in this section will be rather classical, that is static and not information-driven. However, in Section 4 we return to the topic, giving it a dynamic, epistemic twist.

Dynamic logic of actions and strategies. The first thing to note is that the sequential structure of players' actions in an extensive game lends itself to logical analysis. A good system to use for this purpose is *propositional dynamic logic (PDL)*, originally designed to analyze programs and computation (see [27] for the original motivation and subsequent theory). Let Act be a set of primitive actions. An *action model* is a tuple $\mathcal{M} = \langle W, \{R_a \mid a \in \text{Act}\}, V \rangle$ where W is an abstract set of states, or stages in an extensive game, and for each $a \in \text{Act}$, $R_a \subseteq W \times W$ is a binary *transition relation* describing possible transition from states w to w' by executing the action a . On top of this atomic repertoire, the tree structure of extensive games supports complex action expressions, constructed by the standard regular operations of “indeterministic choice” (\cup), “sequential composition” ($;$) and “unbounded finitary iteration” ($*$: Kleene star):

$$\alpha := a \mid \alpha \cup \beta \mid \alpha; \beta \mid \alpha^*$$

This syntax recursively defines complex relations in action models:

- $R_{\alpha \cup \beta} := R_\alpha \cup R_\beta$
- $R_{\alpha; \beta} := R_\alpha \circ R_\beta$
- $R_{\alpha^*} := \bigcup_{n \geq 0} R_\alpha^n$. $R_\alpha^0 = \text{Id}$ (the identity relation) and $R_\alpha^{n+1} = R_\alpha^n \circ R_\alpha$.

The key dynamic modality $[\alpha]\varphi$ now says that “after the move described by the program expression α is taken, φ is true”:

$$\mathcal{M}, w \models [\alpha]\varphi \text{ iff for each } v, \text{ if } wR_\alpha v \text{ then } \mathcal{M}, v \models \varphi$$

PDL has been used for describing solution concepts on extensive games by many authors [2,4,28]. An extended discussion of logics that can explicitly define strategies in extensive games is found in [29].

Adding preferences: the case of Backwards Induction. As before, a complete logical picture must bring in players' preferences on top of *PDL*, along the lines of our earlier modal preference logic. To show how this works, we consider a key pilot example: the Backwards Induction (*BI*) algorithm. This procedure marks each node of an extensive game tree with *values* for the players (assuming that distinct end nodes have different utility values):¹³

¹²For further illustrations of logics on strategic games, cf. [23–26].

¹³In what follows, we shall mainly work with *finite games*, though current dynamic and temporal logics can also deal with infinite games.

BI Algorithm: At end nodes, players already have their values marked. At further nodes, once all daughters are marked, the player to move gets her maximal value that occurs on a daughter, while the other, non-active player gets his value on that maximal node.

The resulting strategy for a player selects the successor node with the highest value. The resulting set of moves for all players (still a function on nodes given our assumption on end nodes) is the “*bi* strategy”.

Relational strategies and set preference. But to a logician, a strategy is best viewed as a subrelation of the total *move* relation. It is an advice to restrict one’s next choice in some way, similar to the more general situation where our *plans* constrain our choices. Mathematically, this links up with the usual way of thinking about programs and procedures in computational logic, in terms of the elegant algebra of relations and its logic *PDL* as defined earlier.

When the above algorithm is modified to a relational setting—we can now drop assumptions about unicity at end-points—we find an interesting new feature: special assumptions about players. For instance, it makes sense to take a minimum value for the passive player at a node over all highest-value moves for the active player. But this is a worst-case assumption: my counter-player does not care about my interests after her own are satisfied. But we might also assume that she does, choosing a maximal value for me among her maximum nodes. This highlights an important feature: *solution methods* are not neutral, they encode significant assumptions about players.

One interesting way of understanding the variety that arises here has to do with the earlier modal preference logic. We might say in general that the driving idea of *Rationality* behind relational *BI* is the following:

I do not play a move when I have another whose outcomes I prefer.

But preferences between moves that can lead to different sets of outcomes call for a notion of “lifting” the given preference on end-points of the game to sets of end-points. As we said before, this is a key topic in preference logic, and here are many options: the game-theoretic rationality behind *BI* has a choice point. One popular version in the logical literature is this:

$$\forall y \in Y \exists x \in X \ x <_i y$$

This says that we choose a move with the highest maximal value that can be achieved. A more demanding notion of preference for a set *Y* over *X* in the logical literature [10] is the $\forall\forall$ clause that

$$\forall y \in Y \forall x \in X \ x <_i y$$

Here is what relational *BI* looks like when we follow the latter stipulation, which makes *Rationality* less demanding, and hence the method more cautious:

First mark all moves as *active*. Call a move *a dominated* if it has a sibling move all of whose reachable endpoints via active nodes are preferred by the current player to all reachable endpoints via *a* itself. The second version of the *BI* algorithm works in stages:

At each stage, mark dominated moves in the $\forall\forall$ sense of preference as *passive*, leaving all others active.

Here “reachable endpoints” by a move are all those that can be reached via a sequence of moves that are still active at this stage.

We will analyze just this particular algorithm in our logics to follow, but our methods apply much more widely.

Defining Backwards Induction in logic. Many logical definitions for the *BI* strategy have been published [cf. again the survey in 2, Section 3]. Here is a modal version combining the logics of action and preferences presented earlier—significantly, involving operator commutations between these:

Theorem 1.1 ([30]). *For each extensive game form, the strategy profile σ is a backward induction solution iff σ is played at the root of a tree satisfying the following modal axiom for all propositions p and players i :*

$$(turn_i \wedge \langle \sigma^* \rangle (end \wedge p)) \rightarrow [move_i] \langle \sigma^* \rangle (end \wedge \langle \geq_i \rangle p)$$

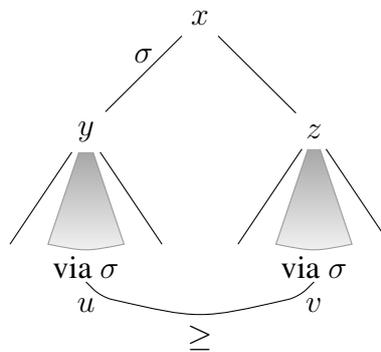
Here $move_i = \bigcup_{a \text{ is an } i\text{-move}} a$, $turn_i$ is a propositional variable saying that it is i ’s turn to move, and end is a propositional variable true at only end nodes. Instead of a proof, we merely develop the logical notions involved a bit further.

The meaning of the crucial axiom follows by a modal frame correspondence ([9], Chapter 3).¹⁴ Our notion of Rationality reappears:

Fact 1.2. *A game frame makes $(turn_i \wedge [\sigma^*](end \rightarrow p)) \rightarrow [move_i] \langle \sigma^* \rangle (end \wedge \langle pref_i \rangle p)$ true for all i at all nodes iff the frame has this property for all i :*

RAT: *No alternative move for the current player i guarantees outcomes via further play using σ that are all strictly better for i than all outcomes resulting from starting at the current move and then playing σ all the way down the tree.*

A typical picture to keep in mind here, and also later on in this paper, is this:



¹⁴“Game frames” here are extensive games extended with one more binary relation σ .

More formally, *RAT* is this *confluence property* for action and preference:

$$\text{CF} \quad \bigwedge_i \forall x \forall y ((\text{turn}_i(x) \wedge x \sigma y) \rightarrow \\ (x \text{ move } y \wedge \forall z (x \text{ move } z \rightarrow \exists u \exists v (\text{end}(u) \wedge \text{end}(v) \wedge y \sigma^* v \wedge z \sigma^* u \wedge u \leq_i v)))$$

Now, a simple inductive proof on the depth of finite game trees shows for our cautious algorithm that:

Theorem 1.3. *BI is the largest subrelation S of the move relation in a game with (a) S has a successor at each intermediate node, (b) S satisfies CF.*

This result is not very deep, but it opens a door to a whole area of research.

The general view: fixed-point logics for game trees. We are now in the realm of a well-known logic of computation, viz. *first-order fixed-point logic* $LFP(FO)$ [31]. The above analysis really tells us:

Theorem 1.4. *The BI relation is definable as a greatest-fixed-point formula in the logic $LFP(FO)$.*

Here is the explicit definition in $LFP(FO)$:

$$BI(x, y) = \nu S. xy \cdot x \text{ move } y \wedge \bigwedge_i (\text{Turn}_i(x) \rightarrow \forall z (x \text{ move } z \rightarrow \\ \exists u \exists v (\text{end}(u) \wedge \text{end}(v) \wedge S.yv \wedge S.zu \wedge u \leq_i v)))$$

The crucial feature making this work is a typical logical point: the occurrences of the relation S in the property CF are *syntactically positive*, and this guarantees upward monotonic behaviour. We will not go into technical details of this connection here, except for noting the following.

Fixed-point formulas in computational logics like this express at the same time static definitions of the *bi* relation, and procedures computing it.¹⁵ Thus, fixed-point logics are an attractive language for extensive games, since they analyze both the statics and dynamics of game solution.

This first analysis of the logic behind extensive games already reveals the fruitfulness of putting together logical and game-theoretical perspectives. But it still leaves untouched the dynamics of deliberation and information flow that determine players' expectations and actual play as a game unfolds, an aspect of game playing that both game theorists and logicians have extensively studied in the last decades. In what follow we make these features explicit, deploying the full potential of the fine-grained Theory of Play that we propose.

¹⁵One can use the standard defining sequence for a greatest fixed-point, starting from the total *move* relation, and see that its successive decreasing approximation stages S^k are exactly the 'active move stages' of the above algorithm. This and related connections have been analyzed in greater mathematical detail in [32].

2. Information Dynamics

The background to the logical systems that follow is a move that has been called a “Dynamic Turn” in logic, making informational acts of inference, but also observations, or questions, into explicit first-class citizens in logical theory that have their own valid laws that can be brought out in the same mathematical style that has served standard logic so well for so long. The program has been developed in great detail in [19,33] drawing together a wide range of relevant literature, but we will only use some basic components here: single events of information change and, later on in this paper, longer-term interactive processes of information change. Towards the end of the paper, we will also briefly refer to other dynamic components of rational agency, with dynamic logics for acts of strategy change, or even preference change.

Players’ informational attitudes can be broadly divided into two categories: *hard* and *soft* information [34,35].¹⁶ *Hard information*, and its companion attitude, is information that is *veridical* and *not revisable*. This notion is intended to capture what agents are fully and correctly certain of in a given game situation. So, if an agent has hard information that some fact φ is true, then φ really is true. In absence of better terminology and following common usage in the literature, we use the term *knowledge* to describe this very strong type of informational attitude. By contrast, *soft information* is, roughly speaking, anything that is not “hard”: it is not necessarily veridical, and it is revisable in the presence of new information. As such, it comes much closer to *beliefs* or more generally attitudes that can be described as “regarding something as true” [36]. This section introduces some key logical systems for describing players’ hard and soft information in a game situation, and how this information can change over time.

2.1. Hard Information and Public Announcements

Recall that N is the set of players, and A a set of atomic sentences p describing ground facts, such as “player i choose action a ” or “the red card is on the table”. A non-empty set W of *worlds* or *states* then represent possible configurations of plays for a fixed game. Typically, players have hard information about the structure of the game—e.g., which moves are available, and what are their own preferences and choices, at least in the *ex interim* stage of analysis.

Static epistemic logic. Rather than directly representing agents’ information in terms of syntactic statements, in this paper, we use standard epistemic models for “semantic information” encoded by epistemic “*indistinguishability relations*”. Setting aside some conceptual subtleties for the purpose of exposition, we will assume that indistinguishability is an equivalence relation. Each agent has some “hard information” about the situation being modeled, and agents cannot distinguish between any two states that agree on this information. This is essentially what we called the player’s “view of the game” in Section 1. Technically, we then get well-known structures:

¹⁶ Note that the distinction “hard” *versus* “soft” information has to do with the way agents take an incoming new signal that some proposition is true. Despite a similarity in terms, this is orthogonal to the standard game-theoretic contrast between “perfect” and “imperfect” information, which is rather about how much players know about their position during a game. Players can receive both hard and soft information in both perfect and imperfect information games.

Definition 2.1. [Epistemic Model] An *epistemic model* $\mathcal{M} = \langle W, \{\sim_i\}_{i \in N}, V \rangle$ has a non-empty set of worlds W ; for each $i \in N$, $\sim_i \subseteq W \times W$ is reflexive, transitive and symmetric; and $V : \text{At} \rightarrow \wp(W)$ is a valuation map. \times

A simple modal language describes properties of these structures. Formally, \mathcal{L}_{EL} is the set of sentences generated by the grammar:

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_i\varphi$$

where $p \in \text{At}$ and $i \in N$. The propositional connectives $\rightarrow, \leftrightarrow, \vee$ are defined as usual, and the dual L_i of K_i is $\neg K_i \neg \varphi$. The intended interpretation of $K_i\varphi$ is “according to agent i ’s current (hard) information, φ is true” (in popular jargon, “ i knows that φ is true”). Here is the standard truth definition:

Definition 2.2. Let $\mathcal{M} = \langle W, \{\sim_i\}_{i \in N}, V \rangle$ be an epistemic model. For each $w \in W$, φ is true at state w , denoted $\mathcal{M}, w \models \varphi$, is defined by induction:

- $\mathcal{M}, w \models p$ iff $w \in V(p)$
- $\mathcal{M}, w \models \neg\varphi$ iff $\mathcal{M}, w \not\models \varphi$
- $\mathcal{M}, w \models \varphi \wedge \psi$ iff $\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models K_i\varphi$ iff for all $v \in W$, if $w \sim_i v$ then $\mathcal{M}, v \models \varphi$

We call φ *satisfiable* if there is a model $\mathcal{M} = \langle W, \{\sim_i\}_{i \in N}, V \rangle$ and $w \in W$ with $\mathcal{M}, w \models \varphi$, and say φ is valid in \mathcal{M} , denoted $\mathcal{M} \models \varphi$, if $\mathcal{M}, w \models \varphi$ for all $w \in W$. \times

Given the definition of the dual of K_i , it is easy to see that:

$$\mathcal{M}, w \models L_i\varphi \text{ iff there is a } v \in W \text{ such that } \mathcal{M}, v \models \varphi$$

This says that “ φ is consistent with agent i ’s current hard information”.

Information update. Now comes a simple concrete instance of the above-mentioned “Dynamic Turn”. Typically, hard information can *change*, and this crucial phenomenon can be added to our logic explicitly.

The most basic type of information change is a *public announcement* [37,38]. This is an event where some proposition φ (in the language of \mathcal{L}_{EL}) is made publicly available, in full view, and with total reliability. Clearly, the effect of such an event should be to *remove* all states that do not satisfy φ : new hard information shrinks a current range of uncertainty.

Definition 2.3. [Public Announcement.] Let $\mathcal{M} = \langle W, \{\sim_i\}_{i \in N}, V \rangle$ be an epistemic model and φ an epistemic formula. The model updated by the **public announcement of** φ is the structure $\mathcal{M}^\varphi = \langle W^\varphi, \{\sim_i^\varphi\}_{i \in N}, V^\varphi \rangle$ where $W^\varphi = \{w \in W \mid \mathcal{M}, w \models \varphi\}$, for each $i \in N$, $\sim_i^\varphi = \sim_i \cap W^\varphi \times W^\varphi$, and for all atomic proposition p , $V^\varphi(p) = V(p) \cap W^\varphi$. \times

Clearly, if \mathcal{M} is an epistemic model then so is \mathcal{M}^φ . The two models describe two different moments in time, with \mathcal{M} the current information state of the agents and \mathcal{M}^φ the information state after the information that φ is true has been incorporated in \mathcal{M} . This temporal dimension can be represented explicitly in our logical language:

Let \mathcal{L}_{PAL} extend \mathcal{L}_{EL} with expressions of the form $[\varphi]\psi$ with $\varphi \in \mathcal{L}_{EL}$. The intended interpretation of $[\varphi]\psi$ is “ ψ is true after the public announcement of φ ” and truth is defined as $\mathcal{M}, w \models [\varphi]\psi$ iff if $\mathcal{M}, w \models \varphi$ then $\mathcal{M}^\varphi, w \models \psi$.

Now, in the earlier definition of public announcement, we can also allow formulas from the extended language \mathcal{L}_{PAL} : the recursion will be in harmony. As an illustration, a formula like $\neg K_i\psi \wedge [\varphi]K_i\psi$ says that “agent i currently does not know ψ but after the announcement of φ , agent i knows ψ ”. So, the language of \mathcal{L}_{PAL} describes what is true both before and after the announcement while explicitly mentioning the informational event that achieved this.

While this is a broad extension of traditional conceptions of logic, standard methods still apply. A fundamental insight is that there is a strong logical relationship between what is true before and after an announcement, in the form of so-called *reduction axioms*:

Theorem 2.4. *On top of the static epistemic base logic, the following reduction axioms completely axiomatize the dynamic logic of public announcement:*

$[\varphi]p$	\leftrightarrow	$\varphi \rightarrow p, \text{ where } p \in \text{At}$
$[\varphi]\neg\psi$	\leftrightarrow	$\varphi \rightarrow \neg[\varphi]\psi$
$[\varphi](\psi \wedge \chi)$	\leftrightarrow	$[\varphi]\psi \wedge [\varphi]\chi$
$[\varphi][\psi]\chi$	\leftrightarrow	$[\varphi \wedge [\varphi]\psi]\chi$
$[\varphi]K_i\varphi$	\leftrightarrow	$\varphi \rightarrow K_i(\varphi \rightarrow [\varphi]\psi)$

Going from left to right, these axioms reduce syntactic complexity in a stepwise manner. This recursive style of analysis has set a model for the logical analysis of informational events generally. Thus, information dynamics and logic form a natural match.

2.2. Group Knowledge

Both game theorists and logicians have extensively studied a next phenomenon after the individual notions considered so far: *group knowledge and belief*.¹⁷ We assume that the reader is familiar with the relevant notions, recalling just the merest basics. For a start, the statement “everyone in the (finite) group $G \subseteq N$ knows φ ” can be defined as follows:

$$E_G\varphi := \bigwedge_{i \in G} K_i\varphi$$

Following [41]¹⁸, the intended interpretation of “it is common knowledge in G that φ ” ($C_G\varphi$) is the infinite conjunction:

$$\varphi \wedge E_G\varphi \wedge E_GE_G\varphi \wedge E_GE_GE_G\varphi \wedge \dots$$

In general, we need to add a new operator $C_G\varphi$ to the earlier epistemic language for this. It takes care of all iterations of knowledge modalities by inspecting all worlds reachable through finite sequences of

¹⁷[39] and [40] provide an extensive discussion.

¹⁸Cf. [42] for an alternative reconstruction.

epistemic accessibility links for arbitrary agents. Let $\mathcal{M} = \langle W, \{\sim_i\}_{i \in N}, V \rangle$ be an epistemic model, with $w \in W$. Truth of formulas of the form $C\varphi$ is defined by:

$$\mathcal{M}, w \models C_G\varphi \text{ iff for all } v \in W, \text{ if } wR_G^*v \text{ then } \mathcal{M}, v \models \varphi$$

where $R_G^* := (\bigcup_{i \in G} \sim_i)^*$ is the reflexive transitive closure of $\bigcup_{i \in G} \sim_i$. As for valid laws of reasoning, the complete epistemic logic of common knowledge expresses principles of “reflective equilibrium”, or mathematically, fixed-points: ¹⁹

- Fixed-Point Axiom: $C_G\varphi \rightarrow E_G C_G\varphi$
- Induction Axiom: $\varphi \wedge C_G(\varphi \rightarrow E_G\varphi) \rightarrow C_G\varphi$

Studying group knowledge is just a half-way station to a more general move in current logics of agency. Common knowledge is a notion of group information that is definable in terms of what the individuals know about each others. But taking collective agents—a committee, a scientific research community—seriously as logical actors in their own right brings us beyond this reductionist perspective.

Finally, what about dynamic logics for group modalities? Baltag, Moss and Solecki [44] proved that the extension of \mathcal{L}_{EL} with common knowledge and public announcement operators is strictly more expressive than with common knowledge alone. Nonetheless, a technical reduction axiom-style recursive analysis is still possible, as carried out in [45].

2.3. Soft Information and Soft Announcements

But rational agents are not just devices that keep track of hard information, and produce indubitable knowledge all the time. What seems much more characteristic of intelligent behaviour, as has been pointed out by philosophers and psychologists alike, is our creative learning ability of having *beliefs*, perhaps based on soft information, that overshoot the realm of correctness. And the dynamics of that is found in our skills in *revising* those beliefs when they turn out to be wrong. Thus, the dynamics of “correction” is just as important to rational agency as that of “correctness”.

Models of belief via plausibility. While there is an extensive literature on the theory of belief revision, starting with [46], truly logical models of the dynamics of beliefs, hard and soft information have only been developed recently. For a start, we need a static base, extending epistemic models with softer, revisable informational attitudes. One appealing approach is to endow epistemic ranges with a *plausibility ordering* for each agent: a pre-order (reflexive and transitive) $w \preceq_i v$ that says “player i considers world v at least as plausible as w .” As a convenient notation, for $X \subseteq W$, we set $Min_{\preceq_i}(X) = \{v \in W \mid v \preceq_i w \text{ for all } w \in X\}$, the set of minimal elements of X according to \preceq_i . The plausibility ordering \preceq_i represents which possible worlds an agent considers more likely, encoding soft information. Such models representing have been used by logicians [35,47,48], game theorists [49], and computer scientists [50,51]:

¹⁹Cf. [43] for an easy way of seeing why the next principles do the job.

Definition 2.5 (Epistemic-Doxastic Models). An **epistemic-doxastic model** is a tuple:

$$\mathcal{M} = \langle W, \{\sim_i\}_{i \in N}, \{\preceq_i\}_{i \in N}, V \rangle$$

where $\langle W, \{\sim_i\}_{i \in N}, V \rangle$ is an epistemic model and, for each $i \in N$, \preceq_i is a well-founded²⁰ reflexive and transitive relation on W satisfying, for all $w, v \in W$:

- plausibility implies possibility: if $w \preceq_i v$ then $w \sim_i v$.
- locally-connected: if $w \sim_i v$ then either $w \preceq_i v$ or $v \preceq_i w$.²¹ ×

These richer models can define many basic soft informational attitudes:

- **Belief:** $\mathcal{M}, w \models B_i \varphi$ iff for all $v \in \text{Min}_{\preceq_i}([w]_i)$, $\mathcal{M}, v \models \varphi$.
This is the usual notion of belief which satisfies standard properties,
- **Safe Belief:** $\mathcal{M}, w \models \Box_i \varphi$ iff for all v , if $v \preceq_i w$ then $\mathcal{M}, v \models \varphi$.
Thus, φ is safely believed if φ is true in *all* states the agent considers more plausible. This stronger notion of belief has also been called *certainty* by some authors ([52], Section 13.7).²²

Soft attitudes in terms of information dynamics. As noted above, a crucial feature of soft informational attitudes is that they are *defeasible* in light of new evidence. In fact, we can characterize these attitudes in terms of the type of evidence which can prompt the agent to adjust them. To make this precise, consider the natural notion of a *conditional belief* in an epistemic-doxastic model \mathcal{M} . We say i *believes φ given ψ* , denoted $B_i^\psi \varphi$, if

$$\mathcal{M}, w \models B_i^\psi \varphi \text{ iff for all } v \in \text{Min}_{\preceq_i}([\psi]_{\mathcal{M}} \cap [w]_i), \mathcal{M}, v \models \varphi$$

where $[\varphi]_{\mathcal{M}} = \{w \mid \mathcal{M}, w \models \varphi\}$ is the usual truth set of φ . So, ‘ B_i^ψ ’ encodes what agent i will believe upon receiving (possibly misleading) evidence that ψ is *true*.²³ Unlike beliefs, conditional beliefs may be inconsistent (*i.e.*, $B^\psi \perp$ may be true at some state). In such a case, agent i cannot (on pain of inconsistency) revise by ψ , but this will only happen if the agent has hard information that ψ is false. Indeed, $K \neg \varphi$ is logically equivalent to $B_i^\varphi \perp$ over the class of epistemic-doxastic models. This suggests the following dynamic characterization of hard information as unrevisable belief:

$$\mathcal{M}, w \models K_i \varphi \text{ iff } \mathcal{M}, w \models B_i^\psi \varphi \text{ for all } \psi$$

Safe belief can be similarly characterized by restricting the admissible evidence:

- $\mathcal{M}, w \models \Box_i \varphi$ iff $\mathcal{M}, w \models B_i^\psi \varphi$ for all ψ with $\mathcal{M}, w \models \psi$.
i.e., i safely believes φ iff i continues to believe φ given any true formula.

Baltag and Smets [55] give an elegant logical characterization of all these notions by adding the safe belief modality \Box_i to the epistemic language \mathcal{L}_{EL} .

²⁰Well-foundedness is only needed to ensure that for any set X , $\text{Min}_{\preceq_i}(X)$ is nonempty. This is important only when W is infinite—and there are ways around this in current logics. Moreover, the condition of connectedness can also be lifted, but we use it here for convenience.

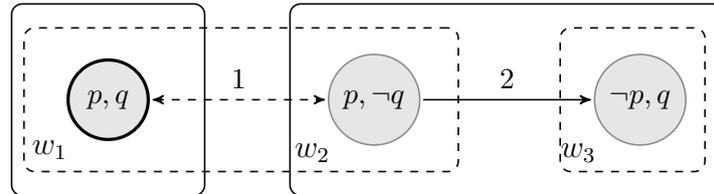
²¹We can even prove the following equivalence: $w \sim_i v$ iff $w \preceq_i v$ or $v \preceq_i w$.

²²Another notion is *Strong Belief*: $\mathcal{M}, w \models B_i^s \varphi$ iff there is a v with $w \sim_i v$ and $\mathcal{M}, v \models \varphi$ and $\{x \mid \mathcal{M}, x \models \varphi\} \cap [w]_i \preceq_i \{x \mid \mathcal{M}, x \models \neg \varphi\} \cap [w]_i$, where $[w]_i$ is the equivalence class of w under \sim_i . This has been studied by [53,54].

²³We can define belief $B_i \varphi$ as $B_i^\top \varphi$: belief in φ given a tautology.

Belief change under hard information. Let us now turn to the systematic logical issue of how beliefs change under new hard information, *i.e.*, the logical laws governing $[\varphi]B_i\psi$. One might think this is taken care of by conditional belief $B_i^\varphi\psi$, and indeed they are when ψ is a *ground formula* not containing any modal operators. But in general, they are different.

Example 2.6. [Dynamic Belief Change *versus* Conditional Belief] Consider state w_1 in the following epistemic-doxastic model:



In this model, the solid lines represent agent 2’s hard and soft information (the box is 2’s hard information \sim_2 and the arrow represent 2’s soft information \preceq_2) while the dashed lines represent 1’s hard and soft information. Reflexive arrows are not drawn to keep down the clutter in the picture. Note that at state w_1 , agent 2 *knows* p and q (e.g., $w_1 \models K_2(p \wedge q)$), and agent 1 believes p but not q ($w_1 \models B_1p \wedge \neg B_1q$). Now, although agent 1 does not *know* that agent 2 knows p , agent 1 does believe that agent 2 believes q ($w_1 \models B_1B_2q$). Furthermore, agent 1 maintains this belief *conditional on* p : $w_1 \models B_1^pB_2q$. However, public announcing the true fact p , removes state w_3 and so we have $w_1 \models [p]\neg B_1B_2q$. Thus a belief in ψ conditional on φ is *not* the same as a belief in ψ *after* the public announcement of φ . The reader is invited to check that $B_i^p(p \wedge \neg K_ip)$ is satisfiable but $[\!|p]B_i(p \wedge \neg K_ip)$ is not satisfiable.²⁴

The example is also interesting as the announcement of a *true* fact misleads agent 1 by forcing her to drop her belief that agent 2 believes q ([33], pg. 182). Despite these intricacies, the logical situation is clear: The dynamic logic of changes in absolute and conditional beliefs under public announcement is completely axiomatizable by means of the static base logic of belief over plausibility models plus the following complete reduction axiom:

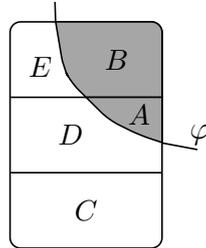
$$[\varphi]B_i^\psi\chi \leftrightarrow (\varphi \rightarrow B_i^{\varphi \wedge [\varphi]\psi}[\varphi]\chi)$$

Belief change under soft information. Public announcement assumes that agents treat the source of the incoming information as *infallible*. But in many scenarios, agents *trust* the source of the information up to a point. This calls for *softer* announcements, that can also be brought under our framework. We only make some introductory remarks: see ([33], Chapter 7) and [55] for more extensive discussion.

How to incorporate less-than-conclusive evidence that φ is true into an epistemic-doxastic model \mathcal{M} ? Eliminating worlds is too radical for that. It makes all updates irreversible. What we need for a soft announcement of a formula φ is thus not to eliminate worlds altogether, but rather *modify the plausibility ordering* that represents an agent’s current hard and soft information state. The goal is to *rearrange* all

²⁴The key point is stated in ([56], pg. 2): “ $B_i^\psi\varphi$ says that if agent i would learn φ , she would come to believe that ψ was true before the learning, while $[\!|\varphi]B_i\psi$ says that after learning φ , i would come to believe that ψ is the case (after the learning).” This observation will be of importance in our analysis of Agreement Theorems later on.

states in such a way that φ is believed, and perhaps other desiderata are met. There are many “policies” for doing this [57], but here, we only mention two, that have been widely discussed in the literature on belief revision. The following picture illustrates the setting:



Suppose the agent considers all states in C as least as plausible as all states in $A \cup D$, which she, in turns, considers at least as plausible as all states in $B \cup E$. If the agent gets evidence in favor of φ from a source that she barely trusts. How is she to update her plausibility ordering?

Perhaps the most ubiquitous policy is *conservative upgrade*, which lets the agent only tentatively accept the incoming information φ by making the best φ the new minimal set and keeping the old plausibility ordering the same on all other worlds. In the above picture a conservative upgrade with φ results in the new ordering $A \prec_i C \prec_i D \prec_i B \cup E$. The general logical idea here is this: “plausibility upgrade is model reordering”.²⁵ This view can be axiomatized in a dynamic logic in the same style as we did with earlier scenarios ([33], Chapter 7 for details).

In what follows, we will focus on a more radical policy for belief upgrade, between the soft conservative upgrade and hard public announcements. The idea behind such *radical upgrade* is to move *all* φ worlds ahead of all other worlds, while keeping the order inside these two zones the same. In the picture above, a radical upgrade by φ would result in $A \prec_i B \prec_i C \prec_i D \prec_i E$.

The precise definition of radical upgrades goes as follow. Let $[[\varphi]]_i^w = \{x \mid \mathcal{M}, x \models \varphi\} \cap [w]_i$ (where $[w]_i$ is the equivalence class of w under \sim_i) denote this set of φ worlds:

Definition 2.7 (Radical Upgrade.). *Given an epistemic-doxastic model $\mathcal{M} = \langle W, \{\sim_i\}_{i \in N}, \{\preceq_i\}_{i \in N}, V \rangle$ and a formula φ , the radical upgrade of \mathcal{M} with φ is the model $\mathcal{M}^{\uparrow\varphi} = \langle W^{\uparrow\varphi}, \{\sim_i^{\uparrow\varphi}\}_{i \in N}, \{\preceq_i^{\uparrow\varphi}\}_{i \in N}, V^{\uparrow\varphi} \rangle$ with $W^{\uparrow\varphi} = W$, for each i , $\sim_i^{\uparrow\varphi} = \sim_i$, $V^{\uparrow\varphi} = V$ and finally, for all $i \in N$ and $w \in W^{\uparrow\varphi}$:*

- for all $x \in [[\varphi]]_i^w$ and $y \in [[\neg\varphi]]_i^w$, set $x \prec_i^{\uparrow\varphi} y$,
- for all $x, y \in [[\varphi]]_i^w$, set $x \preceq_i^{\uparrow\varphi} y$ iff $x \preceq_i y$, and
- for all $x, y \in [[\neg\varphi]]_i^w$, set $x \preceq_i^{\uparrow\varphi} y$ iff $x \preceq_i y$. ✕

A logical analysis of this type of information change uses modalities $[\uparrow_i\varphi]\psi$ meaning “after i ’s radical upgrade of φ , ψ is true”, interpreted as follows:

$$\mathcal{M}, w \models [\uparrow_i\varphi]\psi \text{ iff } \mathcal{M}^{\uparrow_i\varphi}, w \models \psi. \text{ } ^{26}$$

Here is how belief revision under soft information can be treated:

²⁵The most general dynamic point is this: “Information update is model transformation”.
²⁶Conservative upgrade is the special case of radical upgrade with the modal formula $best_i(\varphi, w) := Min_{\preceq_i}([w]_i \cap \{x \mid \mathcal{M}, x \models \varphi\})$.

Theorem 2.8. *The dynamic logic of radical upgrade is completely axiomatized by the complete static epistemic-doxastic base logic plus, essentially, the following recursion axiom for conditional beliefs:*

$$[\uparrow\varphi]B^\psi\chi \leftrightarrow (L(\varphi \wedge [\uparrow\varphi]\psi) \wedge B^{\varphi \wedge [\uparrow\varphi]\psi}[\uparrow\varphi]\chi) \vee (\neg L(\varphi \wedge [\uparrow\varphi]\psi) \wedge B^{[\uparrow\varphi]\psi}[\uparrow\varphi]\chi)$$

This result is from [58], and its proof shows how revision policies as plausibility transformations really give agents not just new beliefs, but also new conditional beliefs – a point sometimes overlooked in the literature.

2.4. The General Logical Dynamics Program

Our logical treatment of update with hard and soft information reflects a general methodology, central to the Theory of Play that we advocate here. Information dynamics is about steps of model transformation, either in their the universe of worlds, or their relational structure, or both.

Other dynamic actions and events. These methods work much more generally than we are able to show here, including model update with information that may be partly private, but also for various other relevant actions, such as *inference* manipulating finer syntactic information, or *questions* modifying a current agenda of issues for investigation. These methods even extend beyond the agents’ informational attitudes, such as the dynamics of preferences expressing their “evaluation” of the world.²⁷

From local to global dynamics. One further important issue is this. Most information flow only makes sense in a longer-term temporal setting, where agents can pursue goals and engage in strategic interaction. This is the realm of *epistemic-doxastic temporal logics* that describe a “Grand Stage” of histories unfolding over time. By now, there are several studies linking up between the dynamic logics of local informational step that we have emphasized, and abstract long-term temporal logics. We refer to [33,59] for these new developments, that are leading to complete logics of information dynamics with “protocols” and what may be called *procedural information* that agents have about the process they are in. Obviously, this perspective is very congenial to extensive games, and in the rest of this paper, it will return in many places, though always concretely.²⁸

3. Long-term Information Dynamics

We now discuss a first round of applications of the main components of the Theory of Play outlined in the previous sections. We leave aside games for the moment, and concentrate on the dynamic of information in interaction. These applications have in common that they use single update steps, but then iterate them, according to what might be called “protocols” for conversation, learning, or other relevant processes. It is the resulting limit behavior that will mainly occupy us in this section.

We first consider agreement theorems, well known to game theorists, showing how repeated conditioning and public announcements lead to consensus in the limit. This opens the door a general

²⁷See [12] on dynamic logics for agents’ *preference changes* between worlds, triggered by commands or other actions with evaluative or moral force.

²⁸In terms of the cited literature, we will now engage in concrete “logic of protocols”.

analysis of fixed-points of repeated attitude changes, raising new questions for logic as well as for interactive epistemology. Next we discuss underlying logical issues, including extensions to scenarios of belief merge and formation of group preferences in the limit. Finally we return to a concrete illustration: viz. learning scenarios, a fairly recent chapter in logical dynamics, at the intersection of logic, epistemology, and game theory.

3.1. Agreement Dynamics

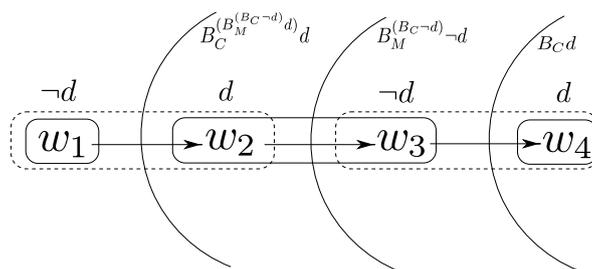
Agreement Theorems, introduced in [60], show that common knowledge of disagreement about posterior beliefs is impossible given a common prior. Various generalizations have been given to other informational attitudes, such as probabilistic common belief [61] and qualitative non-negatively introspective “knowledge” [62]. These results naturally suggest dynamic scenarios, and indeed [63] have shown that agreement can be dynamically reached by repeated Bayesian conditioning, given common prior beliefs.

The logical tools introduced above provide a unifying framework for these various generalizations, and allow to extend them to other informational attitudes. For the sake of conciseness, we will not cover static agreement results in this paper. The interested reader can consult [64,65].

For a start, we will focus on a comparison between agreements reached via conditioning and via public announcements, reporting the work of [65]. In the next section, we show how generalized scenarios of this sort can also deal with softer forms of information change, allowing for diversity in update policies within groups.

Repeated Conditioning Lead to Agreements. The following example, inspired by a recent Hollywood production, illustrates how agreements are reached by repeated belief conditioning:

Figure 1. Cobb and Mal on the window ledge.



Example 3.1. Cobb and Mal are standing on a window ledge, arguing whether they are dreaming or not. Cobb needs to convince Mal, otherwise dreadful consequences will ensue. For the sake of the example, let us assume that Cobb knows they are not dreaming, but Mal mistakenly believes that they are: state w_1 in Figure 1. The solid and dashed rectangles represent, respectively, Cobb’s and Mal’s hard information. The arrow is their common plausibility ordering.

With some thinking, Mal can come to agree with Cobb. The general procedure for achieving this goes as follows: A *sequence of simultaneous belief conditioning acts* starts with the agents’ simple belief about φ , i.e. for all i , the first element $\mathbb{B}_{1,i}$ in the sequence is $B_i\varphi$ if $\mathcal{M}, w \models B_i\varphi$, and $\neg B_i\varphi$ otherwise.

Agent i 's beliefs about φ at a successor stage are defined by taking her beliefs about φ , conditional upon learning the others' belief about φ at that stage. Formally, for two agents i, j then: $\mathbb{B}_{n+1,i} = B_i^{\mathbb{B}_{n,j}\varphi}\varphi$ if $\mathcal{M}, w \models B_i^{\mathbb{B}_{n,j}\varphi}\varphi$, and $\neg B_i^{\mathbb{B}_{n,j}\varphi}\varphi$ otherwise.²⁹

Following the zones marked with an arc in Figure 1, the reader can check that, at w_1 , Mal needs three rounds of conditioning to switch her belief about their waking, and thus reach an agreement with Cobb. Her belief stays the same upon learning that Cobb believes that they are not dreaming. Let us call this fact φ . The turning point occurs when she learns that Cobb would not change his mind even if he would learn φ . Conditional on this, she now believes that they are indeed not dreaming. Note that Cobb's beliefs stay unchanged throughout, since he knows the true state at the outset.

Iterated conditioning thus leads to agreement, given common priors. Indeed, conditioning induces a decreasing map from subsets to subsets, which guarantees the existence of a fixed points, where all agent's conditional beliefs stabilize. Once the agents have reached this fixed-point, they have eliminated all higher-order uncertainties concerning the posteriors beliefs about φ of the others. Their posteriors beliefs are now common knowledge:

Theorem 3.2 ([65]). *At the fixed-point n of a sequence of simultaneous conditioning acts on φ , for all $w \in W$ and $i \in I$, we have that:*

$$\mathcal{M}, w \models C_I\left(\bigwedge_{i \in I} \mathbb{B}_{n,i}\varphi\right)$$

The reader accustomed to static agreement theorems will see that we are now only a small step away from concluding that sequences of simultaneous conditionings lead to agreements, as it is indeed the case in our example. Since common prior and common belief of posteriors suffice for agreement, we get:

Corollary 3.3. *Take any sequence of conditioning acts for a formula φ , as defined above, in a finite model with common prior. At the fixed point of this sequence, either all agents believe φ or they all don't believe φ .*

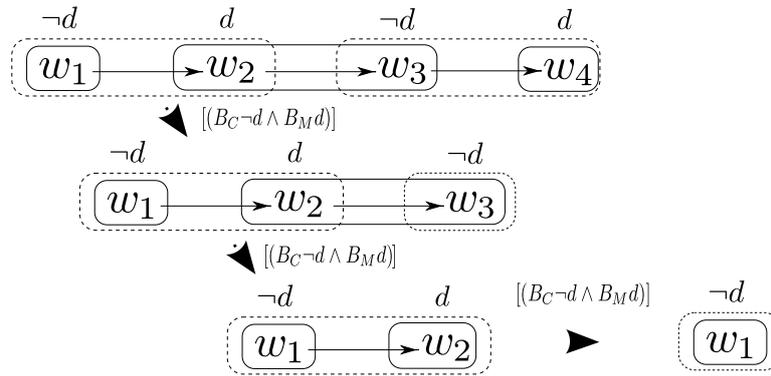
This recasts, in our logical framework, the result of [63], showing how “dialogs” lead to agreements. Still, belief conditioning has a somewhat private character.³⁰ In the example above, Cobb remains painfully uncertain of Mal's thinking process until he sees her changing her mind, that is until she makes the last step of conditioning. Luckily for Cobb, they can do better, as we will now proceed to show.

Repeated Public Announcements Lead to Agreements. Figure 2 shows another scenario, where Cobb and Mal publicly and repeatedly announce their beliefs at w_1 . They keep announcing the same thing, but each time, this induces important changes in both agents' higher-order information. Mal is led stepwise to realize that they are not dreaming, and crucially, Cobb also knows that Mal receives and processes this information. As the reader can check, at each step in the process, Mal's beliefs are common knowledge.

²⁹This definition is meant to fix intuition only. Full details on how to deal with *infinite scenarios*, here and later, are in the cited paper.

³⁰See the remarks on page 66 contrasting public announcement and belief conditioning.

Figure 2. Cobb and Mal’s discussion on the window ledge.



One again, Figure 2 exemplifies a general fact. We first define a *dialogue about φ* as a sequence of public announcements. Let \mathcal{M}, w be a finite pointed

epistemic-doxastic model.³¹ Now let $\mathbb{B}_{1,i}^w$, *i*’s *original belief state at w*, be $B_i\varphi$ if this formula holds at w , and $\neg B_i\varphi$, otherwise. Agent *i*’s $n + 1$ belief state, written $\mathbb{B}_{n+1,i}^w$, is defined as $[\bigwedge_{j \in I} \mathbb{B}_{n,j}^w \varphi] B_i\varphi$ if $\mathcal{M}, w \models [\bigwedge_{j \in I} \mathbb{B}_{n,j}^w \varphi] B_i\varphi$, and as $[\bigwedge_{j \in I} \mathbb{B}_{n,j}^w \varphi] \neg B_i\varphi$, otherwise. Intuitively, a dialogue about φ is a process in which all agents in a group publicly and repeatedly announce their posterior beliefs about φ , while updating with the information received in each round.

In dialogues, just like with belief conditioning, iterated public announcements induce decreasing maps between epistemic-doxastic models, and thus are bound to reach a fixed point, where no further discussion is needed. At this point, the protagonists are guaranteed to have reached consensus:

Theorem 3.4 ([65]). *At the fixed-point \mathcal{M}_n, w of a public dialogue about φ among agents in a group I :*

$$\mathcal{M}_n, w \models C_I \left(\bigwedge_{i \in I} \mathbb{B}_{n,i} \right)$$

Corollary 3.5 ([65]). *For any public dialogue about φ , if there is a common prior that is a well-founded plausibility order, then at the fixed-point \mathcal{M}_n, w , either all agents believe φ or all do not believe φ .*

As noted in the literature [63,64], the preceding dynamics of agreement is one of higher-order information. In the examples above, Mal’s information about the ground facts of dreaming or not dreaming, does not change until the very last round of conditioning or public announcement. The information she gets by learning about Cobb’s beliefs affects her higher-order beliefs, *i.e.*, what she believes about Cobb’s information. This importance of higher-order information flow is a general phenomenon, well-known to epistemic game theorists, which the present logical perspective treats in a unifying way.

Agreements and Dynamics: Further Issues. Here are a few points about the preceding scenarios that invite generalization. Classical agreement results require the agents to be “like-minded” [66]. Our analysis of agreement in dynamic-epistemic logic reveals that this like-mindedness extends beyond the common prior assumption: it also requires the agents to process the information they receive in the same

³¹Our analysis also applies to infinite models: see the cited papers.

way.³² One can easily find counter-examples to the agreement theorems when the update rule is not the same for all agents. Indeed, the issue of “agent diversity” is largely unexplored in our logics (but see [12] for an exception).

A final point is this. While agreement scenarios seem special, to us, they demonstrate a general topic, viz. how different parties in a conversation, say a “Skeptic” and an ordinary person, can modify their positions interactively. In the epistemological literature, this dynamic conversational feature has been neglected—and the above, though solving things in a general way, at least suggests that there might be interesting structure here of epistemological interest.

3.2. Logical Issues about Hard and Soft Limit Behavior

One virtue of our logical perspective is that we can study the above limit phenomena in much greater generality.

Hard information. For a start, for purely logical reasons, iterated public announcement of any formula φ in a model \mathcal{M} must stop at a limit model $\text{lim}(\mathcal{M}, \varphi)$ where φ has either become true throughout (it has become common knowledge), or its negation is true throughout.³³ This raises an intriguing open model-theoretic problem of telling, purely from syntactic form, when a given formula is uniformly “self-fulfilling” (the case where common knowledge is reached), or when “self-refuting” (the case where common knowledge is reached of the negation). Game-theoretic assertions of rationality tend to be self-fulfilling, as we shall see in Section 4 below. But there is no stigma attached to the self-refuting case: e.g., the ignorance assertion in the famous Muddy Children puzzle is self-refuting in the limit. Thus, behind our single scenarios, there is a whole area of limit phenomena that have not yet been studied systematically in epistemic logic.³⁴

In addition to definability, there is complexity and proof. Van Benthem [4] shows how announcement limit submodels can be defined in various known epistemic fixed-point logics, depending on the syntactic shape of φ . Sometimes the resulting formalisms are decidable, e.g., when the driving assertion φ has “existential positive form”, as in the mentioned Muddy Children puzzle, or simple rationality assertions in games.

But these scenarios are still quite special, in that the same assertion gets repeated. There is large variety of further long-term scenarios in the dynamic logic literature, starting from the “Tell All” protocols in [69–71] where agents tell each other all they know at each stage, turning the initial *distributed knowledge* of the group into explicit *common knowledge*.

Soft information. In addition to the limit dynamics of knowledge under hard information, there is the limit behavior of belief, making for more realistic dialog scenarios. This allows for more interesting phenomena in the earlier update sequences. An example is iterated hard information dovetailing agents’ opinions, flipping sides in the disagreement until the very last steps of the dialogue (cf. [33] and [72],

³²Thanks to Alexandru Baltag for pointing out this feature to us.

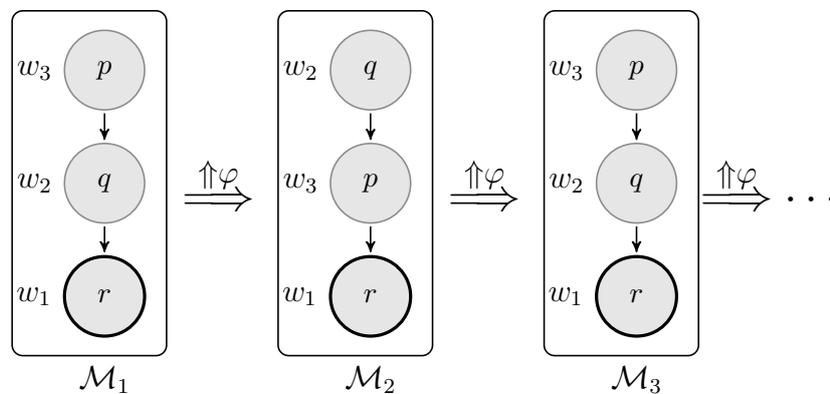
³³We omit some details with pushing the process through infinite ordinals. The final stage is discussed further in terms of “redundant assertions” in [67].

³⁴Even in the single-step case, characterizing “self-fulfilling” public announcements has turned out quite involved [68].

p.110-111). Such disagreement flips can occur until late in the exchange, but as we saw above, they are bound to stop at some point.

All these phenomena get even more interesting mathematically with dialogs involving soft announcements $[\uparrow\varphi]$, when limit behavior can be much more complex, as we will see in the next section. Some relevant observations can be found in [71], and in Section 4 below. First, there need not be convergence at all, the process can oscillate:

Example 3.6. Suppose that φ is the formula $(r \vee (B^{-r}q \wedge p)) \vee (B^{-r}p \wedge q)$ and consider the one agent epistemic-doxastic models pictured below. Since $\llbracket\varphi\rrbracket^{\mathcal{M}_1} = \{w_3, w_1\}$, we have $\mathcal{M}_1^{\uparrow\varphi} = \mathcal{M}_2$. Furthermore, $\llbracket\varphi\rrbracket^{\mathcal{M}_2} = \{w_2, w_1\}$, so $\mathcal{M}_2^{\uparrow\varphi} = \mathcal{M}_3$. Since, \mathcal{M}_3 is the same model as \mathcal{M}_1 , we have a cycle:



In line with this, players’ conditional beliefs may keep changing along the stages of an infinite dialog.³⁵ But still, there is often convergence at the level of agents’ absolute factual beliefs about that the world is like. Indeed, here is a result from [71]:

Theorem 3.7. *Every iterated sequence of truthful radical upgrades stabilizes all simple non-conditional beliefs in the limit.*

Belief and Preference Merge. Finally, we point at some further aspects of the topics raised here. Integrating agents’ orderings through some prescribed process has many similarities with other areas of research. One is *belief merge* where groups of agents try to arrive at a shared group plausibility order, either as a way of replacing individual orders, or as a way of creating a further group agent that is a most reasonable amalgam of the separate components. And this scenario is again much like those of *social choice theory*, where individual agents have to aggregate preference orders into some optimal public ordering. This naturally involves dynamic analysis of the processes of *deliberation* that lead to the eventual act of voting.³⁶ Thus, the technical issues raised in this section have much wider impact. We may be seeing the contours of a systematic logical study of conversation, deliberation and related social processes.

³⁵Infinite iteration of plausibility reordering is in general a non-monotonic process closer to philosophical theories of truth revision in the philosophical literature [73,74]. The technical theory developed on the latter topic in the 1980s may be relevant to our concerns here [75].

³⁶ Van Benthem [33], Chapter 12, elaborates this connection in more technical detail.

3.3. Learning

We conclude this section with one concrete setting where many of the earlier themes come together, viz. *formal learning theory*: see [76–78]. The paradigm we have in mind is identification in the limit of correct hypotheses about the world (cf. [79] on language learning), though formal learning theory in epistemology has also studied concrete learning algorithms for inquiry of various sorts.

The learning setting shows striking analogies with the dynamic-epistemic logics that we have presented in this paper. What follows is a brief summary of recent work in [80,81], to show how our logics link up with learning theory. For broader philosophical backgrounds in epistemology, we refer to [82]. The basic scenario of formal learning theory is one of an agent trying to formulate correct and informative hypotheses about the world, on the basis of an input stream of evidence (in general, an infinite history) whose totality describes what the world is like. At each finite stage of such a sequence, an agent outputs a current hypothesis about the world, which can be modified as new evidence comes in. Success of such a *learning function* in recognition can be of two kinds: either a correct hypothesis is identified uniformly on all histories by some finite stage (the strong notion of “finite identifiability”), or more weakly, each history reaches a point where a correct hypothesis is stated, but when that is may vary according to the history (“identifiability in the limit”). There is a rich mathematical theory of learning functions and what classes of hypotheses can, and cannot, be described by them.

Now, it is not hard to recognize many features here of the logical dynamics that we have discussed. The learning function outputs beliefs, that get revised as new hard information comes in (we think of the observation of the evidence stream as a totally reliable process). Indeed, it is possible to make very precise connections here. We can take the possible hypotheses as our possible worlds, each of which allows those evidence streams (histories of investigation) that satisfy that hypothesis. Then observing successive pieces of evidence is a form of public announcement allowing us to prune the space of worlds. The beliefs involved can be modeled as we did before, by a plausibility ordering on the set of worlds for the agent, which may be modified by successive observations.

On the basis of this simple analogy, [83] prove results like the following, making connections very tight:

Theorem 3.8. *Public announcement-style eliminative update is a universal method: for any learning function, there exists a plausibility order that encodes the successive learning states as current beliefs. The same is true, taking observations as events of soft information, for radical upgrade of plausibility orders.*

Theorem 3.9. *When evidence streams may contain a finite amount of errors, public announcement-style update is no longer a universal learning mechanisms, but radical upgrade still is.*

With these bridges in place, one can also introduce logical languages in the learning-theoretic universe. [80] show how many notions in learning theory then become expressible in dynamic-epistemic or epistemic-temporal languages, say convergence in the limit as necessary future truth of knowledge of a correct hypothesis about the world.³⁷ Thus, we seem to be witnessing the beginning of merges between dynamic logic, belief revision theory and learning theory.

³⁷The logical perspective can actually define many further refinements of learning desiderata, such as reaching future stages when the agent’s knowledge becomes introspective, or when her belief becomes correct, or known.

Such combinations of dynamic epistemic logic and learning theory also invite comparison with game theory. Learning, for instance, to coordinate on a Nash equilibrium in repeated games, has been extensively studied, with many positive and negative results—see, for example, [84].³⁸

This concludes our exploration of long-term information dynamics in our logical setting. We have definitely not exhausted all possible connections, but we hope to have shown how a general Theory of Play fits in naturally with many different areas, providing a common language between them.

4. Solution Dynamics on Extensive Games

We now return to game theory proper, and bring our dynamic logic perspective to bear on an earlier benchmark example: Backwards Induction. This topic has been well-discussed already by eminent authors, but we hope to add a number of new twists suggesting broader ramifications in the study of agency.

In the light of logical dynamics, the main interest of a solution concept is not its “outcome”, its set of strategy profiles, but rather its “process”, the way in which these outcomes are reached. Rationality seems largely a feature of procedures we follow, and our dynamic logics are well-suited to focus on that.

4.1. First Scenario: Iterated Announcement of Rationality

Here is a procedural line on Backwards Induction as a rational process. We can take *BI* to be a process of prior off-line *deliberation* about a game by players whose minds proceed in harmony, though they need not communicate in reality. The treatment that follows was proposed by [22] (which mainly deals with strategic games), and studied in much greater detail by [85].

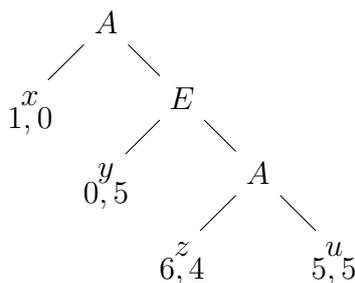
As we saw in Section 3, public announcements saying that some proposition φ is true transform an epistemic model \mathcal{M} into its submodel $\mathcal{M}_{|\varphi}$ whose domain consists of just those worlds in \mathcal{M} that satisfy φ . Now the driving assertion for the Backwards Induction procedure is the following assertion. It states essentially the notion of Rationality discussed in our static analysis of Section 1. As before, at a turn for player i , a move a is *dominated* by a sibling b (a move available at the same node) if every history through a ends worse, in terms of i 's preference, than every history through b :

“at the current node, no player ever chose a strictly dominated move coming here” (**rat**)

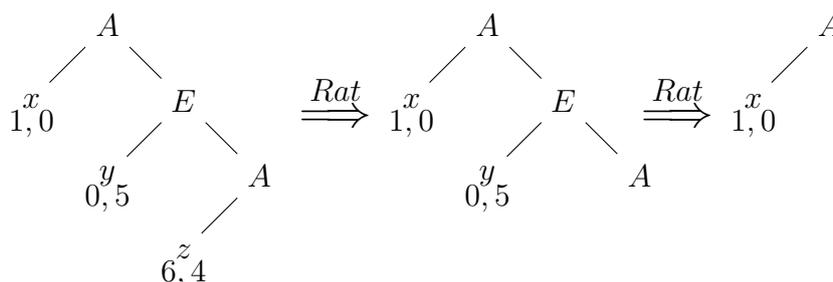
This makes an informative assertion about nodes in a game tree, that can be true or false. Thus, announcing this formula **rat** as a fact about the players will in general make the current game tree smaller. But then we get a dynamics of iteration as in our scenarios of Section 3. In the new smaller game tree, new nodes may become dominated, and hence announcing **rat** again (saying that it still holds after this round of deliberation) makes sense, and so on. As we have seen, this process must reach a limit:

Example [Solving games through iterated assertions of Rationality.] Consider a game with three turns, four branches, and pay-offs for A, E in that order:

³⁸Many of these results live in a probabilistic setting, but dynamic logic and probability is another natural connection, that we have to forego in this paper.



Stage 0 rules out u , the only point where **rat** fails, Stage 1 rules out z and the node above it (the new points where **rat** fails), and Stage 2 rules out y and the node above it. In the remaining game, Rationality reigns supreme:



We see how the *BI* solution emerges from the given game step by step. The general result follows from a simple correspondence between subrelations of the total *move* relation and sets of nodes ([85] has a precise proof with full details):

Theorem 4.1. *In any game tree \mathcal{M} , the model $(rat, \mathcal{M})^\#$ is the actual subtree computed by the *BI* procedure.*

The logical background here is just as we have seen earlier in our epistemic announcement dynamics. The actual *BI* play is the limit sub-model, where **rat** holds throughout. In terms of our earlier distinction, this means that Rationality is a “self-fulfilling” proposition: its announcement eventually makes it true everywhere, and hence *common knowledge* of rationality emerges in the process. Thus, the algorithmic definition of the *BI* procedure in Section 1 and our iterated announcement scenario amount to the same thing. One might say then that our deliberation scenario is just a way of “conversationalizing” a mathematical fixed-point computation. Still, it is of independent interest. Viewing a game tree as an logical model, we see how repeated announcement of Rationality eventually makes this property true throughout the remaining model: it has made itself into common knowledge.

4.2. Second Scenario: Belief and Soft Plausibility Upgrade

Many foundational studies in game theory view Rationality as choosing a best action *given what one believes* about the current and future behaviour of the players. An appealing alternative take on the *BI* procedure does not eliminate any nodes of the initial game, but rather endows it with “progressive expectations” on how the game will proceed. This is the plausibility dynamics that we studied in Section 3, now performing a *soft announcement* of **rat**, where the appropriate action is the “radical upgrade” studied earlier. The essential information produced by the algorithm is then in the binary

plausibility relations that it creates inductively for players among end nodes in the game, standing for complete histories or “worlds”:

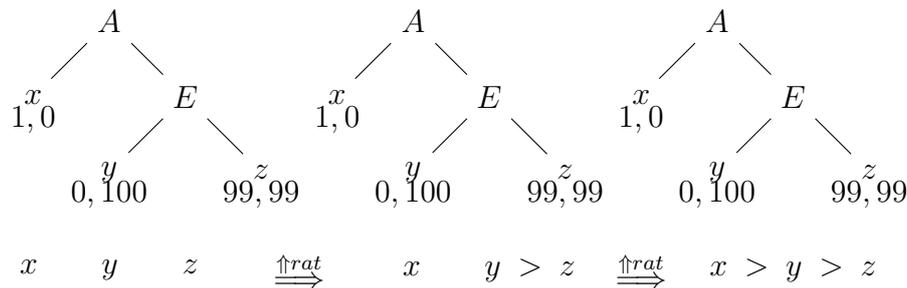
Example [The *BI* outcome in a soft light.] A soft scenario does not remove nodes but modifies the plausibility relation. To implement this, we start with all endpoints of the game tree incomparable.³⁹ Next, at each stage, we compare sibling nodes, using this notion:

A move x for player i dominates its sibling y in beliefs if the *most plausible* end nodes reachable after x along any path in the whole game tree are all better for the active player than all the *most plausible* end nodes reachable in the game after y .

Rationality* (rat^*) says no player plays a move that is dominated in beliefs. Now we perform essentially a radical upgrade $\uparrow rat^*$.⁴⁰

If game node x dominates node y in beliefs, make all end nodes reachable from x more plausible than those reachable from y , keeping the old order inside these zones.

This changes the plausibility order, and hence the pattern of dominance-in-belief, so that iteration makes sense. Here are the stages in our earlier example, where letters x, y, z stand for the end nodes of the game:



In the first game tree, going right is not yet dominated in beliefs for A by going left. rat^* only has bite at E 's turn, and an upgrade takes place that makes $(0, 100)$ more plausible than $(99, 99)$. After this upgrade, however, going right has now become dominated in beliefs, and a new upgrade takes place, making A 's going left most plausible. Here is the general result [33,85]:

Theorem 4.2. *On finite trees, the Backwards Induction strategy is encoded in the plausibility order for end nodes created by iterated radical upgrade with rationality-in-belief.*

Again this is “self-fulfilling”: at the end of the procedure, the players have acquired common belief in rationality. An illuminating way of proving this uses an idea from [86]:

Strategies as plausibility relations. Each sub-relation R of the total *move* relation induces a total plausibility order $ord(R)$ on endpoints of a game:

$x ord(R) y$ iff, looking up at the first node z where the histories of x, y diverged, if x was reached via an R move from z , then so is y .

³⁹Other versions of our scenario would rather make them equi-plausible.

⁴⁰We refer to [85] for technical details.

More generally, relational strategies correspond one-to-one with “move-compatible” total orders of endpoints. In particular, conversely, each such order \leq induces a strategy $rel(\leq)$. Now we can relate the computation in our upgrade scenario for belief and plausibility to the earlier relational algorithm for BI in Section 1:

Fact 4.3. *For any game tree \mathcal{M} and any k , $rel((\uparrow rat^*)^k, \mathcal{M}) = BI^k$.*

Thus, the algorithmic view of Backwards Induction and its procedural doxastic analysis in terms of forming beliefs amount to the same thing. Still, as with our iterated announcement scenario, the dynamic logical view has interesting features of its own. One is that it yields fine-structure to the plausibility relations among worlds that are usually taken as primitive in doxastic logic. Thus games provide an underpinning for the possible worlds semantics of belief that seems of interest per se.

4.3. Logical Dynamic Foundations of Game Theory

We have seen how several dynamic approaches to Backwards Induction amount to the same thing. To us, this means that the notion is logically stable. Of course, extensionally equivalent definitions can still have interesting intensional differences. For instance, the above analysis of strategy creation and plausibility change seems the most realistic description of the “entanglement” of belief and rational action in the behaviour of agents. But as we will discuss soon, a technical view in terms of fixed-point logics may be the best mathematical approach linking up with other areas.

No matter how we construe them, one key feature of our dynamic announcement and upgrade scenarios is this. Unlike the usual epistemic foundation results, common knowledge or belief of rationality is not assumed, but *produced* by the logic. This reflects our general view that rationality is primarily a property of procedures of deliberation or other logical activities, and only secondarily a property of outcomes of such procedures.

4.4. Logics of Game Solution: General Issues

Our analysis does not just restate existing game-theoretic results, it also raises new issues in the logic of rational agency. Technically, all that has been said in Sections 2 and 3 can be formulated in terms of existing *fixed-point logics* of computation, such as the modal “ μ -calculus” and the first-order fixed-point logic $LFP(FO)$. This link with a well-developed area of computational logic is attractive, since many results are known there, and we may use them to investigate game solution procedures that are quite different from Backwards Induction.⁴¹ But the analysis of game solutions also brings some new logical issues to this area.

Game solution and fragments of fixed-point logics. Game solution procedures need not use the full power of fixed-point languages for recursive procedures. It makes sense to use small decidable fragments where appropriate. Still, it is not quite clear right now what the best fragments are. In particular, our earlier analysis intertwines two different relations on trees: the *move* relation of action and computation, and the *preference* relations for players on endpoints. And the question is what happens to known properties of computational logics when we add such preference relations:

⁴¹See the dissertation [32] for details, linking up with computational logic.

The complexity of rationality. In combined logics of action and knowledge, it is well-known that apparently harmless assumptions such as Perfect Recall for agents make the validities undecidable, or non-axiomatizable, sometimes even Π_1^1 -complete [15]. The reason is that these assumptions generate commuting diagrams for actions *move* and epistemic uncertainty \sim satisfying a “confluence property”

$$\forall x \forall y ((x \text{ move } y \wedge y \sim z) \rightarrow \exists u (x \sim u \wedge u \text{ move } z))$$

These patterns serve as the basic grid cells in encodings of complex “tiling problems” in the logic.⁴² Thus, the logical theory of games for players with perfect memory is more complex than that of forgetful agents [15,18]. But now consider the non-epistemic property of rationality studied above, that mixes action and preference. Our key property *CF* in Section 1 had a confluence flavour, too, with a diagram involving action and preference:

$$\begin{aligned} \forall x \forall y ((\text{Turn}_i(x) \wedge x \sigma y) \rightarrow \forall z (z \text{ move } z \rightarrow \forall u ((\text{end}(u) \wedge y \sigma^* u) \\ \rightarrow \exists v (\text{end}(v) \wedge z \sigma^* v \wedge v \leq_i u)))) \end{aligned}$$

So, what is the complexity of fixed-point logics for players with this kind of regular behaviour? Can it be that Rationality, a property meant to make behaviour simple and predictable, actually makes its theory complex?

Zooming in and zooming out: modal logics of best action. The main trend in our analysis has been toward making dynamics explicit in richer logics than the usual epistemic-doxastic-preferential ones, in line with the program in [33]. But in logical analysis, there are always two opposite directions intertwined: getting at important reasoning patterns by making things more explicit, or rather, by making things less explicit!

In particular, in practical reasoning, we are often only interested in what are our *best actions* without all details of their justification. As a mathematical abstraction, it would then be good to extract a simple surface logic for reasoning with best actions, while hiding most of the machinery:

Can we axiomatize the modal logic of finite game trees with a *move* relation and its transitive closure, turns and preference relations for players, and a new relation *best* computed by Backwards Induction?

Further logical issues in our framework concern extensions to *infinite games*, games with *imperfect information*, and scenarios with *diverse agents*. See [12,72,87] for some first explorations.

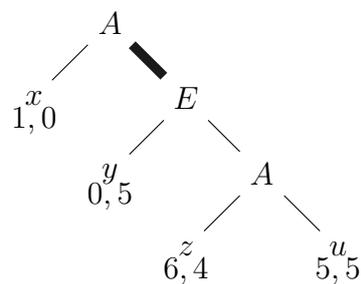
4.5. From Games to Their Players

We end by high-lighting a perhaps debatable assumption of our analysis so far. It has been claimed that the very Backwards Induction reasoning that ran so smoothly in our presentation, is incoherent when we try to “replay” it in the opposite order, when a game is actually played.⁴³

⁴²Recall our earlier remarks in Section 1 on the complexity of strategic games.

⁴³There is a large literature focused on this “paradox” of backwards induction which we do not discuss here. See, for example, [88].

Example [The ‘Paradox of Backwards Induction’.] Recall the style of reasoning toward a Backward Induction solution, as in our earlier simple scenario:



Backwards Induction tells us that A will go left at the start, on the basis of logical reasoning that is available to both players. But then, if A plays *right* (as marked by the thick black line) what should E conclude? Does not this mean that A is not following the *BI* reasoning, and hence that all bets are off as to what he will do later on in the game? It seems that the very basis for the computations in our earlier sections collapses.⁴⁴

Responses to this difficulty vary. Many game-theorists seem under-impressed. The characterization result of [89] assumes that players know that rationality prevails throughout.⁴⁵ One can defend this behaviour by assuming that the other player only makes isolated mistakes. Baltag, Smets and Zvesper [86] essentially take the same tack, deriving the *BI* strategy from an assumption of “stable true belief” in rationality, a gentler form of stubbornness stated in terms of dynamic-epistemic logic.

Players’ revision policies. We are more inclined toward the line of [91,92]. A richer analysis should add an account of the *types of agent* that play the game. In particular, we need to represent the *belief revision policies* of the players, that determine what they will do when making a surprising observation contradicting their beliefs in the course of a game. There are many different options for such policies in the above example, such as “It was just an error, and A will go back to being rational”, “ A is telling me that he wants me to go right, and I will be rewarded for that”, “ A is an automaton with a general rightward tendency”, and so on.⁴⁶ Our analysis so far has omitted this type of information about players of the game, since our algorithms made implicit uniform assumptions about their prior deliberation, as well as what they are going to do as the game proceeds.

This matching up of two directions of thought: *backwards* in “off-line dynamics” of deliberation, and *forwards* in “on-line dynamics” of playing the actual game, is a major issue in its own right, beyond specific scenarios. Belief revision policies and other features of players must come in as explicit components of the theory, in order to deal with the dynamics of how players update knowledge and revise beliefs as a game proceeds.

But all this is exactly what the logical dynamics of Section 2 is about. Our earlier discussion has shown how acts of information change and belief revision can enter logic in a systematic manner.

⁴⁴The drama is clearer in longer games, when A has many comebacks toward the right.

⁴⁵Samet [90] calls this “rationality no matter what”, a stubborn unshakable belief that players will act rationally later on, even if they have never done so up until now.

⁴⁶ One reaction to these surprise events might even be a switch to an entirely new style of reasoning about the game. That would require a more finely-grained *syntax-based* views of revision: cf. the discussion in [93].

Thus, once more, the richer setting that we need for a truly general theory of game solution is a perfect illustration for the general Theory of Play that we have advocated.

5. Conclusion

Logic and game theory form a natural match, since the structures of game theory are very close to being models of the sort that logicians typically study. Our first illustrations reviewed existing work on static logics of game structure, drawing attention to the fixed-point logic character of game solution methods. This suggests a broader potential for joining forces between game theory and computational logic, going beyond specific scenarios toward more general theory. To make this more concrete, we then presented the recent program of “logical dynamics” for information-driven agency, and showed how it throws new light on basic issues studied in game theory, such as agreement scenarios and game solution concepts.

What we expect from this contact is not the solution of problems afflicting game theory through logic, or vice versa, remedying the aches and pains of logic through game theory. Of course, game theorists may be led to new thoughts by seeing how a logician treats (or mistreats) their topics, and also, as we have shown, logicians may see interesting new open problems through the lense of game theory.

But fruitful human relations are usually not therapeutic: they lead to new facts, in the form of shared offspring. In particular, one broad trend behind much of what we have discussed here is this. Through the fine-structure offered by logic, we can see the dynamics of games as played in much more detail, making them part of a general analysis of agency that also occurs in many other areas, from “multi-agent systems” in computer science to social epistemology and the philosophy of action. It is our expectation that the offspring of this contact might be something new, neither fully logic nor game theory: a *Theory of Play*, rather than just a theory of games.

References and Notes

1. *Epistemic Logic and the Theory of Games and Decisions*; Bacharach, M., Gerard-Varet, L., Mongin, P., Shin, H., Eds.; Kluwer Academic Publishers: Dordrecht, The Netherlands, 1997.
2. Van der Hoek, W.; Pauly, M. Modal logic for games and information. In *Handbook of Modal Logic*; Blackburn, P., van Benthem, J., Wolter, F., Eds.; Elsevier: Amsterdam, The Netherlands, 2006; pp. 1077–1148.
3. Bonanno, G. Modal logic and game theory: Two alternative approaches. *Risk Decis. Policy* **2002**, *7*, 309–324.
4. Van Benthem, J. Games in dynamic epistemic logic. *Bull. Econ. Res.* **2001**, *53*, 219–248.
5. Brandenburger, A. The power of paradox: Some recent developments in interactive epistemology. *Int. J. Game Theory* **2007**, *35*, 465–492.
6. Aumann, R. Interactive epistemology I: Knowledge. *Int. J. Game Theory* **1999**, *28*, 263–300.
7. Seligman, J. Hybrid logic for analyzing games. Presented at the Workshop Door to Logic, Beijing, China, May 2010.
8. Van Benthem, J. *Modal Logic for Open Minds*; CSLI Publications: Stanford, CA, USA, 2010.

9. Blackburn, P.; de Rijke, M.; Venema, Y. *Modal Logic*; Cambridge University Press: Cambridge, UK, 2002.
10. Von Wright, G.H. *The Logic of Preference*; University of Edinburgh Press: Edinburgh, UK, 1963.
11. Hansson, S.O.; Grüne-Yanoff, T. Preferences. In *The Stanford Encyclopedia of Philosophy*, Spring 2009 ed.; Zalta, E.N., Ed.; Stanford University: Stanford, CA, USA, 2009.
12. Liu, F. Changing for the Better: Preference Dynamics and Agent Diversity. Ph.D Thesis, Institute for Logic, Language and Computation (DS-2008-02), University of Amsterdam: Amsterdam, The Netherlands, 2008.
13. Van Benthem, J.; Girard, P.; Roy, O. Everything Else Being Equal: A Modal Logic for *Ceteris Paribus* Preferences. *J. Phil. Logic* **2009**, *38*, 83–125.
14. Girard, P. Modal Logic for Belief and Preference Change. Ph.D Thesis, Stanford University: Stanford, CA, USA, 2008.
15. Halpern, J.; Vardi, M. The complexity of reasoning about knowledge and time. *J. Comput. Syst. Sci.* **1989**, *38*, 195–237.
16. Marx, M. Complexity of modal logic. In *Handbook of Modal Logic*; Blackburn, P., van Benthem, J., Wolter, F., Eds.; Elsevier: Amsterdam, The Netherlands, 2006; pp. 139–189.
17. Gabbay, D.; Kurucz, A.; Wolter, F.; Zakharyashev, M. *Many-Dimensional Modal Logics: Theory and Applications*; Elsevier: Amsterdam, The Netherlands, 2003.
18. Van Benthem, J.; Pacuit, E. The tree of knowledge in action: Towards a common perspective. In *Proceedings of Advances in Modal Logic Volume 6*, Queensland, Australia, 25–28 September 2006; Governatori, G., Hodkinson, I., Venema, Y., Eds.; King's College Press: London, UK, 2006; pp. 87–106.
19. Van Benthem, J. *Exploring Logical Dynamics*; CSLI Press: Stanford, CA, USA, 1996.
20. Aumann, R.J. Correlated equilibrium as an expression of bayesian rationality. *Econometrica* **1987**, *55*, 1–18.
21. Brandenburger, A.; Friedenberg, A. Intrinsic correlations in games. *J. Econ. Theor.* **2008**, *141*, 28–67.
22. Van Benthem, J. Rational dynamics and epistemic logic in games. *Int. J. Game Theory Rev.* **2007**, *9*, 13–45.
23. De Bruin, B. *Explaining Games: The Epistemic Programme in Game Theory*; Springer: Berlin, Germany, 2010.
24. Bonanno, G. A Syntactic Approach to Rationality in Games with Ordinal Payoffs. In *Logic and the Foundations of Game and Decision Theory*; Bonanno, G., van der Hoek, W., Wooldridge, M., Eds.; Amsterdam University Press: Amsterdam, The Netherlands, 2008; pp. 59–86.
25. Van Benthem, J. Open problems in logic and games. In *We Will Show Them! Essays in Honour of Dov Gabbay*; Artemov, S., Barringer, H., d'Avila Garcez, A., Lamb, L., Woods, J., Eds.; College Publications: London, UK, 2005.
26. Lorini, E.; Schwarzentruher, F.; Herzig, A. Epistemic games in modal logic: Joint actions, knowledge and preferences all together. Proceedings of LORI-II: Workshop on Logic, Rationality and Interaction, Chongqing, China, 2009; He, X., Horty, J., Pacuit, E., Eds.; Springer-Verlag: Berlin, Germany, 2009; pp. 212–226.

27. Harel, D.; Kozen, D.; Tiuryn, J. *Dynamic Logic*; The MIT Press: Cambridge, UK, 2000.
28. Harrenstein, B.; van der Hoek, W.; Meyer, J.J.; Witteveen, C. A modal characterization of Nash Equilibrium. *Fund. Inform.* **2003**, *57*, 281–321.
29. Van Benthem, J. In Praise of Strategies. Technical report, PP-2008-03; Institute for Logic, Language and Computation: Amsterdam, The Netherlands, 2008.
30. Van Benthem, J.; van Otterloo, S.; Roy, O. Preference logic, conditionals, and solution concepts in games. In *Modality Matters: Twenty-Five Essays in Honour of Krister Segerberg*; Lagerlund, H., Lindström, S., Sliwinski, R., Eds.; Uppsala University Press: Uppsala, Sweden, 2006; Vol. 53.
31. Ebbinghaus, H.D.; Flum, J. *Finite Model Theory*; Springer: Berlin, Germany, 1995.
32. Gheerbrant, A. Fixed-Point Logics on Trees. Ph.D Thesis, Dissertation Series DS-2010-08; Institute for Logic, Language and Information, University of Amsterdam: Amsterdam, The Netherlands, 2010.
33. Van Benthem, J. *Logical Dynamics of Information and Interaction*; Cambridge University Press: Cambridge, UK, 2010.
34. Van Benthem, J. Rational Animals: What is ‘KRA’? Presented at ESSLLI Summer School, Malaga, Spain, 2006.
35. Baltag, A.; Smets, S. Conditional Doxastic Models: A Qualitative Approach to Dynamic Belief Revision. Proceedings of the 13th Workshop on Logic, Language, Information and Computation (WoLLIC 2006), Stanford University, CA, USA, 18–21 July 2006; In *Electronic Notes in Theoretical Computer*; Mints, G., de Queiroz, R., Eds.; Elsevier: Amsterdam, The Netherlands, 2006; Vol. 165, pp. 5–21.
36. Schwitzgebel, E. Belief. In *The Stanford Encyclopedia of Philosophy*, Fall 2008 ed.; Zalta, E.N., Ed.; Stanford University: Stanford, CA, USA, 2008.
37. Plaza, J. Logics of public communications. Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems, Charlotte, NC, USA, 1989; Emrich, M.L., Pfeifer, M.S., Hadzikadic, M., Ras, Z., Eds.; pp. 201–216.
38. Gerbrandy, J. Bisimulations on Planet Kripke. Ph.D Thesis, (DS-1999-01), Institute for Logic, Language and Computation, University of Amsterdam: Amsterdam, The Netherlands, 1999.
39. Fagin, R.; Halpern, J.; Moses, Y.; Vardi, M. *Reasoning about Knowledge*; The MIT Press: Cambridge, MA, USA, 1995.
40. Vanderschraaf, P.; Sillari, G. Common knowledge. In *The Stanford Encyclopedia of Philosophy*, Spring 2009 ed.; Zalta, E.N., Ed.; Stanford University: Stanford, CA, USA, 2009.
41. Lewis, D. *Convention*; Harvard University Press: Cambridge, MA, USA, 1969.
42. Cubitt, R.P.; Sugden, R. Common knowledge, salience and convention: A reconstruction of David Lewis’ game theory. *Econ. Phil.* **2003**, *19*, 175–210.
43. Van Benthem, J. Modal frame correspondences and fixed-points. *Studia Logica* **2006**, *83*, 133–155.
44. Baltag, A.; Moss, L.; Solecki, S. The Logic of Common Knowledge, Public Announcements and Private Suspicions. Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 98), Evanston, IL, USA, 1998; Gilboa, I., Ed.; Morgan Kaufmann Publishers: San Francisco, CA, USA, 1998; pp. 43–56.

45. Van Benthem, J.; van Eijck, J.; Kooi, B. Logics of communication and change. *Inform. Comput.* **2006**, *204*, 1620–1662.
46. Alchourrón, C.E.; Gärdenfors, P.; Makinson, D. On the logic of theory change: Partial meet contraction and revision functions. *J. Symbolic Logic* **1985**, *50*, 510–530.
47. Van Benthem, J. Dynamic logic for belief revision. *J. Appl. Non-Classical Logics* **2004**, *14*, 129–155.
48. van Ditmarsch, H. Prolegomena to dynamic logic for belief revision. *Synthese: Knowl. Ration. Action* **2005**, *147*, 229–275.
49. Board, O. Dynamic interactive epistemology. *Games Econ. Behav.* **2004**, *49*, 49–80.
50. Boutilier, C. Conditional Logics for Default Reasoning and Belief Revision. Ph.D Thesis, University of Toronto: Toronto, Canada, 1992.
51. Lamarre, P.; Shoham, Y. Knowledge, Certainty, Belief and Conditionalisation. Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning (KR'94), Bonn, Germany, 24–27 May; Doyle, J., Sandewall, E., Torasso, E., Eds.; Morgan Kaufmann: San Francisco, CA, USA, 1994; pp. 415–424.
52. Shoham, Y.; Leyton-Brown, K. *Multiagent Systems*; Cambridge University Press: Cambridge, MA, USA, 2008.
53. Stalnaker, R. Iterated belief revision. *Erkenntnis* **2009**, *70*, 189–209.
54. Battigalli, P.; Siniscalchi, M. Strong belief and forward induction reasoning. *J. Econ. Theory* **2002**, *105*, 356–391.
55. Baltag, A.; Smets, S. ESSLLI 2009 course: Dynamic logics for interactive belief revision, Bordeaux, France, 2009; Available online: <http://alexandru.tiddlyspot.com/> (accessed on 15 February 2011).
56. Baltag, A.; Smets, S. A qualitative theory of dynamic interactive belief revision. In *Logic and the Foundation of Game and Decision Theory (LOFT7)*; Bonanno, G., van der Hoek, W., Wooldridge, M., Eds.; Amsterdam University Press: Amsterdam, The Netherlands, 2008; Vol. 3, pp. 13–60.
57. Rott, H. Shifting priorities: Simple representations for 27 iterated theory change operators. In *Modality Matters: Twenty-Five Essays in Honour of Krister Segerberg*; Lagerlund, H., Lindström, S., Sliwinski, R., Eds.; Uppsala University Press: Uppsala, Sweden, 2006; Vol. 53, pp. 359–384.
58. Van Benthem, J. Dynamic logic for belief revision. *J. Appl. Non-Classical Logics* **2007**, *17*, 129–155.
59. Van Benthem, J.; Gerbrandy, J.; Hoshi, T.; Pacuit, E. Merging frameworks for interaction. *J. Phil. Logic* **2009**, *38*, 491–526.
60. Aumann, R. Agreeing to disagree. *Ann. Stat.* **1976**, *4*, 1236–1239.
61. Monderer, D.; Samet, D. Approximating common knowledge with common beliefs. *Games Econ. Behav.* **1989**, *1*, 170–190.
62. Samet, D. Agreeing to disagree: The non-probabilistic case. *Games Econ. Behav.* **2010**, *69*, 169–174.
63. Geanakoplos, J.; Polemarchakis, H. We can't disagree forever. *J. Econ. Theory* **1982**, *28*, 192–200.

64. Bonanno, G.; Nehring, K. Agreeing to Disagree: A Survey. Document prepared of an invited lecture at the Workshop on Bounded Rationality and Economic Modeling, Firenze, Italy, 1–2 July 1997.
65. Dégrement, C.; Roy, O. Agreement theorems in dynamic-epistemic logic. Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge, Stanford, CA, USA, 2009; Heifetz, A., Ed.; ACM Digital Library: New York, NY, USA, 2009; pp. 91–98.
66. Bacharach, M. Some extensions of a claim of Aumann in an axiomatic model of knowledge. *J. Econ. Theory*, **1985**, *37*, 167–190.
67. Baltag, A.; Smets, S. Talking Your Way into Agreement: Belief Merge by Persuasive Communication. Proceedings of the Second Multi-Agent Logics, Languages, and Organisations Federated Workshops, Turin, Italy, 7–10 September, 2009; Baldoni, M., Eds.; CEUR-WS.org: Tilburg, The Netherlands, 2009; Vol. 494, pp. 129–141.
68. Holliday, W.; Icard, T. Moorean phenomena in epistemic logic. Proceedings of Advances in Modal Logic, Moscow, Russia, 2010; Beklemishev, L., Goranko, V., Shehtman, V., Eds.; College Publications: London, UK, 2010; Vol. 8, pp. 178–199.
69. Van Benthem, J. One is a lonely number: On the logic of communication. Proceedings of the Annual European Summer Meeting of the Association for Symbolic Logic And Colloquium Logicum, Logic Colloquium '02, Munster, Germany, 3–11 August 2006; Chatzidakis, Z., Koepke, P., Pohlers, W., Eds.; A K Peters: Natick, MA, USA, 2006; Vol. 27, pp. 96–129.
70. Roelofsen, F. Exploring Logical Perspectives on Distributed Information and Its Dynamics. Master Thesis, ILLC (LDC 2005-05), University of Amsterdam: Amsterdam, The Netherlands, 2005.
71. Baltag, A.; Smets, S. Group Belief Dynamics under Iterated Revision: Fixed Points and Cycles of Joint Upgrades. Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge, Stanford, CA, USA; Heifetz, A., Ed.; ACM Digital Library: New York, NY, USA, 2009; pp. 41–50.
72. Dégrement, C. The Temporal Mind. Observations on the Logic of Belief Change in Interactive Systems. Ph.D Thesis, Institute for Logic, Language and Computation (DS-2010-03), University of Amsterdam: Amsterdam, The Netherlands, 2010.
73. Gupta, A. Truth and Paradox. *J. Phil. Logic* **1982**, *11*, 1–60.
74. Herzberger, H. Notes on naive semantics. *J. Phil. Logic* **1982**, *11*, 61–102.
75. Visser, A. Semantics and the Liar Paradox. In *Handbook of Philosophical Logic*; Gabbay, D., Guenther, F., Eds.; Kluwer Academic Publishers: Dordrecht, The Netherlands, 2004; Vol. 11, pp. 149–240.
76. Kelly, K. *The Logic of Reliable Inquiry*; Oxford University Press: Oxford, UK, 1996.
77. Osherson, D.; Stob, M.; Weinstein, S. *Systems that Learn*; The MIT Press: Cambridge, MA, USA, 1986.
78. Schulte, O. Formal Learning Theory. In *The Stanford Encyclopedia of Philosophy*, Winter 2008 ed.; Zalta, E.N., Ed.; Standford University: Stanford, CA, USA, 2008.
79. Gold, E.M. Language identification in the limit. *Inform. Control*. **1967**, *10*, 447–474.
80. Dégrement, C.; Gierasimczuk, N. Can Doxastic Agents Learn? On the Temporal Structure of Learning. Proceedings of LORI-II: Workshop on Logic, Rationality and Interaction, Chongqing,

- China, 6–11 October 2009; He, X., Horty, J., Pacuit, E., Eds.; Springer-Verlag: Berlin, Germany, 2009; pp. 90–104.
81. Gierasimczuk, N. *Knowing Ones Limits: Logical Analysis of Inductive Inference*. Ph.D Thesis, Institute for Logic, Language and Computation, University of Amsterdam: Amsterdam, The Netherlands, 2011.
 82. Hendricks, V. *Mainstream and Formal Epistemology*; Cambridge University Press: Cambridge, UK, 2005.
 83. Baltag, A.; Gierasimczuk, N.; Smets, S. Truth-tracking by belief revision. 2010 (Unpublished).
 84. Kalai, E.; Lehrer, E. Rational learning leads to Nash Equilibrium. *Econometrica* **1993**, *61*, 1019–1045.
 85. Van Benthem, J.; Gheerbrant, A. Game solution, epistemic dynamics and fixed-point logics. *Fund. Inform.* **2010**, *100*, 1–23.
 86. Baltag, A.; Smets, S.; Zvesper, J. Keep ‘hoping’ for rationality: A solution to the backward induction paradox. *Synthese* **2009**, *169*, 301–333.
 87. Zvesper, J. *Playing with Information*. Ph.D Thesis, (DS-2010-02), Institute for Logic, Language and Computation, University of Amsterdam: Amsterdam, The Netherlands, 2010.
 88. Bicchieri, C. Self-refuting theories of strategic interaction: A paradox of common knowledge. *Erkenntnis* **1989**, *30*, 69–85.
 89. Aumann, R. Backward induction and common knowledge of rationality. *Games Econ. Behav.* **1995**, *8*, 6–19.
 90. Samet, D. Hypothetical knowledge and games with perfect information. *Games Econ. Behav.* **1996**, *17*, 230–251.
 91. Stalnaker, R. Knowledge, belief and counterfactual reasoning in games. *Econ. Philos.* **1996**, *12*, 133–163.
 92. Stalnaker, R. Substantive rationality and backward induction. *Games Econ. Behav.* **2001**, *37*, 425–435.
 93. Velázquez-Quesada, F.R. *Small Steps in the Dynamics of Information*. Ph.D Thesis, Institute for Logic, Language and Computation, University of Amsterdam: Amsterdam, The Netherlands, 2011.
 94. Plaza, J. Logics of public communications. *Synthese: Knowl. Ration. Action* **2007**, *158*, 165–179.