

Article

Toward a Theory of Play: A Logical Perspective on Games and Interaction

Johan van Benthem^{1,*}, Eric Pacuit¹ and Olivier Roy²

¹ TiLPS, Tilburg University, Tilburg, The Netherlands

² Faculty of Philosophy, University of Groningen, The Netherlands

* Author to whom correspondence should be addressed; J.vanBenthem@uva.nl

Version February 9, 2011 submitted to *Games*. Typeset by \LaTeX using class file *mdpi.cls*

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 $i \in N$:

- $\sigma \succeq_i \sigma'$ 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,

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.

- $\sigma \approx_i \sigma'$ 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

$$\mathcal{M} = \langle S, \{\sim_i\}_{i \in N}, \{\approx_i\}_{i \in N}, \{\succeq_i\}_{i \in N} \rangle$$

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 [cf. 12,14]. A modal *betterness model* for a set N of players is a tuple $\mathcal{M} = \langle W, \{\succeq_i\}_{i \in N}, V \rangle$ where W is a nonempty set of states, for each $i \in N$, $\succeq_i \subseteq W \times W$ 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 \succeq_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 \succeq_i , the corresponding *strict* preference ordering is written $>_i$.

A modal language to describe betterness models uses modalities $\langle \succeq_i \rangle \varphi$ saying that “there is a world at least as good as the current world satisfying φ ”, and likewise for strict preference:

- $\mathcal{M}, w \models \langle \succeq_i \rangle \varphi$ iff there is a v with $v \succeq_i w$ and $\mathcal{M}, v \models \varphi$
- $\mathcal{M}, w \models \langle >_i \rangle \varphi$ iff there is a v with $v \succeq_i w$, $w \not\succeq_i v$, and $\mathcal{M}, v \models \varphi$

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.

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.

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.

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

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.

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.

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 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

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.

¹²For further illustrations of logics on strategic games, cf. [23–26].

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:

- $R_{\alpha \cup \beta} := R_\alpha \cup R_\beta$
- $R_{\alpha; \beta} := R_\alpha \circ R_\beta$
- $R_{\alpha^*} := \cup_{n \geq 0} R_\alpha^n$. $R_\alpha^0 = \text{Id}$ (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 [cf. 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

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”.

¹³In what follows, we shall mainly work with *finite games*, though current dynamic and temporal logics can also deal with infinite games.

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 :*

$$(\text{turn}_i \wedge \langle \sigma^* \rangle (\text{end} \wedge p)) \rightarrow [\text{move}_i] \langle \sigma^* \rangle (\text{end} \wedge \langle \geq_i \rangle p)$$


Here $\text{move}_i = \bigcup_{a \text{ is an } i\text{-move}} a$, 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 $(\text{turn}_i \wedge [\sigma^*](\text{end} \rightarrow p)) \rightarrow [\text{move}_i] \langle \sigma^* \rangle (\text{end} \wedge \langle \text{pref}_i \rangle p)$ 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:



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.

¹⁴“Game frames” here are extensive games extended with one more binary relation σ .

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 $LFP(FO)$:

$$BI(x, y) = \nu S.xy \cdot x \text{ move } y \wedge \bigwedge_i (Turn_i(x) \rightarrow \forall z(x \text{ move } z \rightarrow \exists u \exists v (end(u) \wedge end(v) \wedge S.yv \wedge S.zu \wedge u \leq_i v)))$$

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.

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*.

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].

¹⁶Note 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.

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 At 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:

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. ×

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 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: $C_G\varphi \rightarrow E_GC_G\varphi$
- 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.

¹⁷[39] and [40] provide an extensive discussion.

¹⁸Cf. [42] for an alternative reconstruction.

¹⁹Cf. [43] for an easy way of seeing why the next principles do the job.

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]:

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 $v \in Min_{\preceq_i}([w]_i)$, $\mathcal{M}, v \models \varphi$.

This is the usual notion of belief which satisfies standard properties,

²⁰Well-foundedness is only needed to ensure that for any set X , $Min_{\preceq_i}(X)$ 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 $v \preceq_i w$.

- *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 [cf. 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., $B_i^\psi \perp$ 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} .

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:



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.


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 (eg., $w_1 \models K_2(p \wedge q)$), 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 : $w_1 \models B_1^p B_2q$. 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 $B_i^p(p \wedge \neg K_i p)$ is satisfiable but $[!p]B_i(p \wedge \neg K_i p)$ 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 [cf. 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 states in such a way that φ is believed, and perhaps other desiderata are met. There are many "policies" for doing this [cf. 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?

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.

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 [cf. 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 $\llbracket \varphi \rrbracket_i^w = \{x \mid \mathcal{M}, x \models \varphi\} \cap [w]_i$ (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 $x \in \llbracket \varphi \rrbracket_i^w$ and $y \in \llbracket \neg\varphi \rrbracket_i^w$, set $x \prec_i^{\uparrow\varphi} y$,
- for all $x, y \in \llbracket \varphi \rrbracket_i^w$, set $x \preceq_i^{\uparrow\varphi} y$ iff $x \preceq_i y$, and
- for all $x, y \in \llbracket \neg\varphi \rrbracket_i^w$, 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:

$$\mathcal{M}, w \models [\uparrow_i\varphi]\psi \text{ iff } \mathcal{M}^{\uparrow\varphi}, w \models \psi.^{26}$$

Here is how belief revision under soft information can be treated:

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.

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 $best_i(\varphi, w) := Min_{\preceq_i}([w]_i \cap \{x \mid \mathcal{M}, x \models \varphi\})$.

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 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.

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".

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

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.

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:*

$$\mathcal{M}, w \models C_I \left(\bigwedge_{i \in I} \mathbb{B}_{n,i} \varphi \right)$$

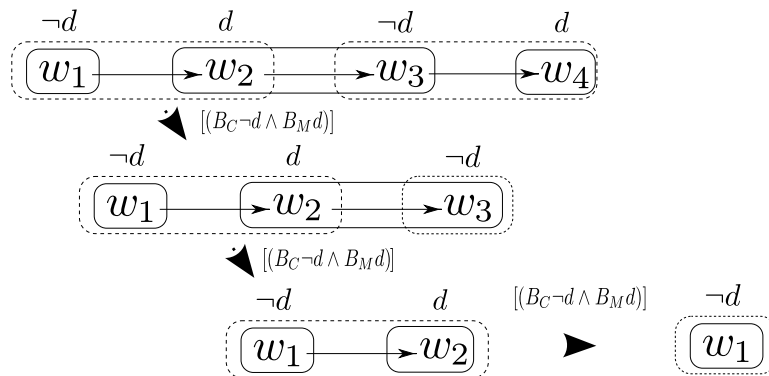
The reader accustomed to static agreement theorems will see that we are now only a small step away from concluding that sequences of simultaneous conditionings lead to agreements, as it is indeed the case in our example. Since common prior and common belief of posteriors suffice for agreement, we get:

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.

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,

³⁰See the remarks on page 15 contrasting public announcement and belief conditioning.

³¹Our analysis also applies to infinite models: see the cited papers.

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 [cf. 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 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 $\lim(\mathcal{M}, \varphi)$ where φ has either become true throughout (it

³²Thanks to Alexandru Baltag for pointing out this feature to us.

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, 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 $[\varphi]^{\mathcal{M}_1} = \{w_3, w_1\}$, we have $\mathcal{M}_1^{\uparrow\varphi} = \mathcal{M}_2$. Furthermore, $[\varphi]^{\mathcal{M}_2} = \{w_2, w_1\}$, 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:

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 [cf. 68].



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.

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

³⁵Infinite iteration of plausibility reordering is in general a non-monotonic process closer to philosophical theories of truth revision in the philosophical literature [cf. 73,74]. The technical theory developed on the latter topic in the 1980s may be relevant to our concerns here [cf. 75].

³⁶Van Benthem [33, Chapter 12], elaborates this connection in more technical detail.

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.

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.

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.

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.

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:



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:

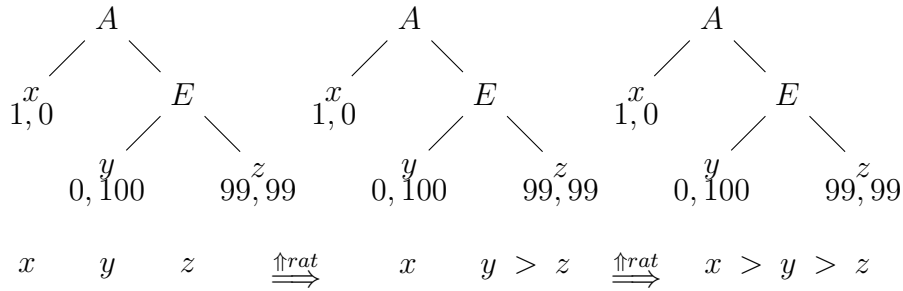
A move x for player i dominates its sibling y *in beliefs* if the *most plausible* end nodes reachable after x along any path in the whole game tree are all better for the active player than all the *most plausible* end nodes reachable in the game after y .

³⁹Other versions of our scenario would rather make them equi-plausible.

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 [cf. 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 .

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.

⁴⁰We refer to [85] for technical details.

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:

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 [cf. 15,18]. But now consider the non-epistemic property of rationality studied above, that mixes

⁴¹ See the dissertation [32] for details, linking up with computational logic.

⁴² Recall our earlier remarks in Section 1 on the complexity of strategic games.

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

Example [The ‘Paradox of Backwards Induction’.] Recall the style of reasoning toward a Backward Induction solution, as in our earlier simple scenario:



⁴³There is a large literature focused on this “paradox” of backwards induction which we do not discuss here. See, for example, [88].

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. 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

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.

⁴⁶One 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].

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

1. Bacharach, M.; Gerard-Varet, L.; Mongin, P.; Shin, H., Eds. *Epistemic Logic and the Theory of Games and Decisions*; Kluwer Academic Publishers, 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, 2006; Vol. 3, *Studies in Logic*, pp. 1077 – 1148.
3. Bonanno, G. Modal logic and game theory: two alternative approaches. *Risk, Decision and Policy* **2002**, 7, 309–324.
4. van Benthem, J. Games in Dynamic Epistemic Logic. *Bulletin of Econ. Research* **2001**, 53, 219 – 248.
5. Brandenburger, A. The power of paradox: some recent developments in interactive epistemology. *International Journal of Game Theory* **2007**, 35, 465–492.
6. Aumann, R. Interactive Epistemology I: Knowledge. *International Journal of Game Theory* **1999**, 28, 263–300.
7. Seligman, J. Hybrid Logic for Analyzing Games. Lecture at the Workshop *Door to Logic*, Beijing May 2010, 2010.
8. van Benthem, J. *Modal Logic for Open Minds*; CSLI Publications, 2010.
9. Blackburn, P.; de Rijke, M.; Venema, Y. *Modal Logic*; Cambridge University Press, 2002.
10. von Wright, G.H. *The Logic of Preference*; Edinburgh, 1963.
11. Hansson, S.O.; Grüne-Yanoff, T. Preferences. In *The Stanford Encyclopedia of Philosophy*, Spring 2009 ed.; Zalta, E.N., Ed.; 2009.
12. Liu, F. Changing for the Better: Preference Dynamics and Agent Diversity. PhD thesis, Institute for Logic, Language and Computation (DS-2008-02), 2008.
13. van Benthem, J.; Girard, P.; Roy, O. Everything Else Being Equal: A Modal Logic for *Ceteris Paribus* Preferences. *Journal of Philosophical Logic* **2009**, 38, 83–125.
14. Girard, P. Modal Logic for Belief and Preference Change. PhD thesis, Stanford University, 2008.

15. Halpern, J.; Vardi, M. The complexity of reasoning about knowledge and time. *Journal of Computer and System Sciences* **1989**, *38*, 195 – 237.
16. Marx, M. Complexity of Modal Logic. Handbook of Modal Logic; Blackburn, P.; van Benthem, J.; Wolter, F., Eds., 2007, pp. 139 – 189.
17. Gabbay, D.; Kurucz, A.; Wolter, F.; Zakharyashev, M. *Many-Dimensional Modal Logics: Theory and Applications*; Elsevier, 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*; Governatori, G.; Hodkinson, I.; Venema, Y., Eds.; King's College Press, 2006; pp. 87 – 106.
19. van Benthem, J. *Exploring Logical Dynamics*; CSLI Press, 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. *Journal of Economic Theory* **2008**, *141*.
22. van Benthem, J. Rational Dynamics and Epistemic Logic in Games. *International Journal of Game Theory Review* **2007**, *9*, 13 – 45.
23. de Bruin, B. *Explaining Games: The Epistemic Programme in Game Theory*; Springer, 2010.
24. Bonanno, G. A syntactic approach to rationality in games with ordinal payoffs in: (eds.), Logic and the Foundations of Game and Decision Theory; Bonanno, G.; van der Hoek, W.; Wooldridge, M., Eds. Amsterdam University Press, 2008, Texts in Logic and Games Series, pp. 59 – 86.
25. van Benthem, J. Open Problems in Logic and Games. 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, 2005, Vol. 1.
26. Lorini, E.; Schwarzentruher, F.; Herzig, A. Epistemic Games in Modal Logic: Joint Actions, Knowledge and Preferences all together. LORI-II Workshop on Logic, Rationality and Interaction, Chongqing, China; He, X.; Horty, J.; Pacuit, E., Eds. Springer-Verlag, 2009, pp. 212–226.
27. Harel, D.; Kozen, D.; Tiuryn, J. *Dynamic Logic*; The MIT Press, 2000.
28. Harrenstein, B.; van der Hoek, W.; Meyer, J.J.; Witteveen, C. A Modal Characterization of Nash Equilibrium. *Fundamenta Informaticae* **2003**, *57*, 281 – 321.
29. van Benthem, J. In Praise of Strategies. Technical report, ILLC Technical Reports, 2008.
30. van Benthem, J.; van Otterloo, S.; Roy, O. Preference Logic, Conditionals, and Solution Concepts in Games. Festschrift for Krister Segerberg. University of Uppsala, 2005.
31. Ebbinghaus, H.D.; Flum, J. *Finite Model Theory*; Springer, 1995.
32. Gheerbrant, A. Fixed-Point Logics on Trees. PhD thesis, Institute for Logic, Language and Information Dissertation Series DS-2010-08, 2010.
33. van Benthem, J. *Logical Dynamics of Information and Interaction*; Cambridge University Press, 2010.
34. van Benthem, J. Rational Animals: What is 'KRA'? invited lecture Malaga ESSLLI Summer School 2006, 2005.

35. Baltag, A.; Smets, S. Conditional Doxastic Models: A Qualitative Approach to Dynamic Belief Revision. Proceedings of WOLLIC 2006, Electronic Notes in Theoretical Computer Science; Mints, G.; de Queiroz, R., Eds., 2006, Vol. 165, pp. 5 – 21.
36. Schwitzgebel, E. Belief. In *The Stanford Encyclopedia of Philosophy*, Fall 2008 ed.; Zalta, E.N., Ed.; 2008.
37. Plaza, J. Logics of public communications. Proceedings, 4th International Symposium on Methodologies for Intelligent Systems; Emrich, M.L.; Pfeifer, M.S.; Hadzikadic, M.; Ras, Z., Eds., 1989, pp. 201–216 (republished as [94]).
38. Gerbrandy, J. Bisimulations on Planet Kripke. PhD thesis, Institute for Logic, Language and Computation (DS-1999-01), 1999.
39. Fagin, R.; Halpern, J.; Moses, Y.; Vardi, M. *Reasoning about Knowledge*; The MIT Press, 1995.
40. Vanderschraaf, P.; Sillari, G. Common Knowledge. In *The Stanford Encyclopedia of Philosophy*, Spring 2009 ed.; Zalta, E.N., Ed.; 2009.
41. Lewis, D. *Convention*; Harvard University Press, 1969.
42. Cubitt, R.P.; Sugden, R. Common Knowledge, Salience and Convention: A Reconstruction of David Lewis' Game Theory. *Economics and Philosophy* **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); Gilboa, I., Ed., 1998, pp. 43 – 56.
45. van Benthem, J.; van Eijck, J.; Kooi, B. Logics of Communication and Change. *Information and Computation* **2006**, *204*, 1620 – 1662.
46. Alchourrón, C.E.; Gärdenfors, P.; Makinson, D. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic* **1985**, *50*, 510 – 530.
47. van Benthem, J. Dynamic Logic for Belief Revision. *Journal of Applied Non-Classical Logics* **2004**, *14*, 129 – 155.
48. van Ditmarsch, H. Prolegomena to Dynamic Logic for Belief Revision. *Synthese: Knowledge, Rationality, and Action* **2005**, *147*, 229 – 275.
49. Board, O. Dynamic Interactive Epistemology. *Games and Economic Behavior* **2004**, *49*, 49 – 80.
50. Boutilier, C. Conditional Logics for Default Reasoning and Belief Revision. PhD thesis, University of Toronto, 1992.
51. Lamarre, P.; Shoham, Y. Knowledge, Certainty, Belief and Conditionalisation. Proceedings of the International Conference on Knowledge Representation and Reasoning, 1994, pp. 415 – 424.
52. Shoham, Y.; Leyton-Brown, K. *Multiagent Systems*; Cambridge University Press, 2008.
53. Stalnaker, R. Iterated Belief Revision. *Erkenntnis* **2009**, *70*.
54. Battigalli, P.; Siniscalchi, M. Strong belief and forward induction reasoning. *Journal of Economic Theory* **2002**, *105*, 356 – 391.
55. Baltag, A.; Smets, S. ESSLLI 2009 course: Dynamic Logics for Interactive Belief Revision. Slides available online at <http://alexandru.tiddlyspot.com/#%5B%5BESSLLI09%20COURSE%5D%5D>, 2009.

56. Baltag, A.; Smets, S. A Qualitative Theory of Dynamic Interactive Belief Revision. Logic and the Foundation of Game and Decision Theory (LOFT7); Bonanno, G.; van der Hoek, W.; Wooldridge, M., Eds. Amsterdam University Press, 2008, Vol. 3, *Texts in Logic and Games*, pp. 13–60.
57. Rott, H. Shifting Priorities: Simple Representations for 27 Iterated Theory Change Operators. Modality Matters: Twenty-Five Essays in Honour of Krister Segerberg; Lagerlund, H.; Lindström, S.; Sliwinski, R., Eds., 2006, Vol. 53, *Uppsala Philosophical Studies*, pp. 359 – 384.
58. van Benthem, J. Dynamic Logic for Belief Revision. *Journal of Applied Non-Classical Logics* **2007**, *17*, 129 – 155.
59. van Benthem, J.; Gerbrandy, J.; Hoshi, T.; Pacuit, E. Merging Frameworks for Interaction. *Journal of Philosophical Logic* **2009**, *38*, 491–526.
60. Aumann, R. Agreeing to Disagree. *Annals of Statistics* **1976**, *4*, 1236 — 1239.
61. Monderer, D.; Samet, D. Approximating Common Knowledge with Common Beliefs. *Games and Economic Behavior* **1989**, *1*, 170 – 190.
62. Samet, D. Agreeing to Disagree: The Non-Probabilistic Case. *Games and Economic Behavior* **2010**, *69*, 169 – 174.
63. Geanakoplos, J.; Polemarchakis, H. We Can't Disagree Forever. *Journal of Economic 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 Modelling, 1997.
65. Dégrement, C.; Roy, O. Agreement Theorems in Dynamic-Epistemic Logic. Proceedings of TARK'09 (Theoretical Aspects of Rationality and Knowledge); Heifetz, A., Ed. ACM Digital Library, 2009.
66. Bacharach, M. Some extensions of a claim of Aumann in an axiomatic model of knowledge. *Journal of Economic 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, 2009, Vol. 494, pp. 129 – 141.
68. Holliday, W.; Icard, T. Moorean Phenomena in Epistemic Logic. Advances in Modal Logic; Beklemishev, L.; Goranko, V.; Shehtman, V., Eds., 2010, Vol. 8pp., College Publications., pp. 178 – 199.
69. van Benthem, J. One is a lonely number: on the logic of communication. In *Logic Colloquium '02*; Chatzidakis, Z.; Koepke, P.; Pohlers, W., Eds.; ASL & A.K. Peters, 2006; Vol. 27, *Lecture Notes in Logic*, pp. 96 – 129.
70. Roelofsen, F. Exploring Logical Perspectives on Distributed Information and Its Dynamics. Master's thesis, ILLC University of Amsterdam (LDC 2005-05), 2005.
71. Baltag, A.; Smets, S. Group Belief Dynamics under Iterated Revision: Fixed Points and Cycles of Joint Upgrades. Proceedings of Theoretical Aspects of Rationality and Knowledge, 2009.
72. Dégrement, C. The Temporal Mind. Observations on the Logic of Belief Change in Interactive Systems. PhD thesis, ILLC University of Amsterdam Dissertation Series DS-2010-03, 2010.
73. Gupta, A. Truth and Paradox. *Journal of Philosophical Logic* **1982**, *11*, 1 – 60.
74. Herzberger, H. Notes on Naive Semantics. *Journal of Philosophical Logic* **1982**, *11*, 61 – 102.

75. Visser, A. Semantics and the Liar Paradox. *Handbook of Philosophical Logic*; Gabbay, D.; Guenther, F., Eds. Springer, 2004, Vol. 11.
76. Kelly, K. *The Logic of Reliable Inquiry*; Oxford University Press, 1996.
77. Osherson, D.; Stob, M.; Weinstein, S. *Systems that Learn*; The MIT Press, 1986.
78. Schulte, O. Formal Learning Theory. In *The Stanford Encyclopedia of Philosophy*, Winter 2008 ed.; Zalta, E.N., Ed.; 2008.
79. Gold, E.M. Language identification in the limit. *Information and Control* **1967**, *10*, 447–474.
80. Dégremon, C.; Gierasimczuk, N. Can Doxastic Agents Learn? On the Temporal Structure of Learning. *Proceedings of LORI-II*; He, X.; Horty, J.; Pacuit, E., Eds., 2009, pp. 90 – 104.
81. Gierasimczuk, N. Knowing Ones Limits: Logical Analysis of Inductive Inference. PhD thesis, ILLC University of Amsterdam, 2011.
82. Hendricks, V. *Mainstream and Formal Epistemology*; Cambridge University Press, 2005.
83. Baltag, A.; Gierasimczuk, N.; Smets, S. Truth-tracking by belief revision. Unpublished manuscript, 2010.
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. *Fundamenta Informaticae* **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. PhD thesis, ILLC University of Amsterdam Dissertation Series DS-2010-02, 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 and Economic Behavior* **1995**, *8*, 6 – 19.
90. Samet, D. Hypothetical Knowledge and Games with Perfect Information. *Games and Economic Behavior* **1996**, *17*, 230 – 251.
91. Stalnaker, R. Knowledge, Belief and Counterfactual Reasoning in Games. *Economics and Philosophy* **1996**, *12*, 133 – 163.
92. Stalnaker, R. Substantive Rationality and Backward Induction. *Games and Economic Behavior* **2001**, *37*, 425 – 435.
93. Velázquez-Quesada, F.R. Small Steps in the Dynamics of Information. PhD thesis, ILLC University of Amsterdam, 2011.
94. Plaza, J. Logics of Public Communications. *Synthese: Knowledge, Rationality, and Action* **2007**, *158*, 165 – 179.