Exploring a Theory of Play

Johan van Benthem
Institute for Logic, Language and Computation, Universiteit van Amsterdam
Department of Philosophy, Stanford University
Johan.vanBenthem@uva.nl

ABSTRACT

We explore some recent directions for the logical foundations of social action that emerge from contacts between logic, game theory, philosophy, and computer science.[1]

Categories and Subject Descriptors

F.4.1 [**Mathematical Logic**]: Modal Logic

General Terms

Theory

Keywords

backward induction, dynamic epistemic logic, theory of play

1. LOGIC OF SOCIAL ACTION

The traditional focus of logic has been on activities of single agents, but social action is now widely seen as essential. Rational - or reasonable - agency involves interaction over time in a balance between information and evaluation of states of the world. There are many strands to this that invite logical analysis, and we may be just at the start.

2. LOGIC AND GAME THEORY

Many issues in the logic of social action become sharper when thinking about games. This lecture will take its examples mainly from that interface, though we make no claims about how much good that will do for game theory (or for logic) per se. Moreover, no exhaustive survey is intended: cf. [8, 32] for more.

[1]What follows is an extended abstract for the TARK lecture, not a full paper.

3. BACKWARD INDUCTION REVISITED

Like the Muddy Children puzzle, Backward Induction keeps suggesting new logical perspectives, being a miniature of non-trivial rational agency:

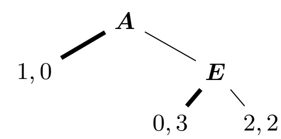

Are the bold-face moves the 'rational outcome'? Reasoning underpinning this involves action, preference, knowledge, belief, conditionals (all of philosophical logic in one tree).

4. DELIBERATION CREATES BELIEFS

In the spirit of logical dynamics [27], we put the spotlight on the *BI* algorithm rather than its outcomes, as a deliberation procedure producing initial expectations when a game starts. These expectations are encoded in a plausibility ordering of the endpoints of the game, which gets updated step by step in a systematic way. These updates may be viewed as steps of belief revision during deliberation:

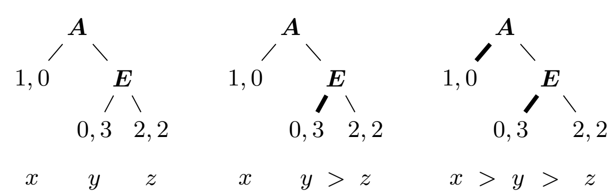

Definition 1. Move x *dominates* sibling move y *in beliefs* if the *most plausible* end nodes reachable after x along any path in the game tree are all better for the active player than all *most plausible* end nodes reachable after y. The assertion of *rationality-in-beliefs* (**rat**) says at a node that no player has played a move in the past that was dominated in beliefs.

Definition 2. Given a proposition P, the operation of *radical upgrade* $\Uparrow P$ changes a current plausibility model $\boldsymbol{M}$ to $\boldsymbol{M}\Uparrow P$: all P-worlds are now better than all $\neg P$-worlds; within zones, the old order remains.

THEOREM 1. *On finite trees, the Backwards Induction strategy is encoded in the limit plausibility order for leaves created by iterated upgrade* $\Uparrow\boldsymbol{rat}$ *with rationality-in-belief.*

In the limit of this procedure, players have acquired *common belief in rationality.*

Encoding strategies as plausibility relations [5] Each subrelation R of the *move* relation induces a total plausibility order $ord(R)$ on leaves x, y of the game tree: $x\, ord(R)\, y$ iff, 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. Conversely, every such 'tree-compatible' total order $\leq$ on leaves of the game also induces a subrelation $rel(\leq)$ of the move relation via an obvious stipulation.

5. EQUILIBRIA AND FIXED-POINT LOGIC

The outcome of this dynamic analysis is a relational strategy σ that can be defined in standard logical systems, namely first-order fixed-point logic:

THEOREM 2. *BI is the largest subrelation σ of the* move *relation in a finite game tree satisfying*

(a) the relation has a successor at each intermediate node,

(b) a Confluence property for all players i at all nodes:

$$\begin{aligned} CF: \textstyle\bigwedge_i \forall x \forall y \Big(& (Turn_i(x) \wedge x\sigma y) \\ & \rightarrow \big(x\; move\; y \wedge \\ & \quad \forall z (x\; move\; z \rightarrow \\ & \qquad \exists u \exists v (end(u) \;\wedge end(v) \wedge \\ & \qquad\quad y\sigma^* u \;\wedge z\sigma^* v \;\wedge v \leq_i u)))\Big) \end{aligned}$$

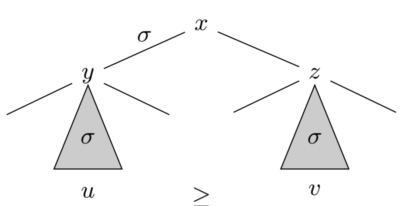

Inspecting the syntax of *CF*, we can use the language of *first-order fixed-point logic LFP(FO)*:

THEOREM 3. *The BI relation is definable in LFP(FO).*

The same analysis works for other variants of *BI*, with definability in inflationary fixed-point logic (for details, see [29]). This is just the start of exploring general connections between game-theoretic equilibrium and fixed-point logics.

6. ZOOMING IN OR ZOOMING OUT: MODAL LOGIC OF BEST ACTION

Fixed-point logics describe detailed mechanics of game solution. Practical reasoning can also zoom out, hiding details:

Can we axiomatize the modal logic of finite game trees with a *move* relation (plus $move^*$), turns and preference for players, and a new relation *best* computed by Backward Induction? We need a preference modality $\langle pref_i \rangle \varphi$: i prefers some node with φ to the current one.

FACT 1 ([31]). *Confluence corresponds to the following modal axiom, for all propositions p - viewed as sets of nodes - and for all players i:*

$$\begin{aligned} \big(turn_i \wedge \langle best \rangle\, [best^*]\, (\boldsymbol{end} \rightarrow p)\big) \\ \rightarrow [move_{-i}]\, \langle best^* \rangle\, (\boldsymbol{end} \wedge \langle pref_i \rangle\, p) \end{aligned}$$

Does Rationality, meant to make behavior predictable, actually make game logic complex? In modal logics of *action* and *knowledge*, Perfect Recall can cause Π^1_1-completeness by grid encoding [15]. Rationality, too, forces grid-like patterns, cf. the picture for CF.[2]

7. LOGIC OF LIMIT PHENOMENA

The above scenario is driven by iterated upgrade with one particular formula φ that can be true or false at nodes of a game tree. Let us look further.

Iterated public announcement This drives Muddy Children puzzles, or game solution procedures that announce rationality, pruning the initial game until a first fixed-point [26].

Definition 3. The *update limit* $(\varphi, \boldsymbol{M})^\sharp$ is the first model reached by iterated announcements $!\varphi$ in $\boldsymbol{M}$ that no longer changes. If this model is non-empty, φ holds in all nodes: common knowledge results (self-fulfilling) - if empty, $\neg\varphi$ was common knowledge (*self-refuting*). Rationality assertions $\boldsymbol{rat}$ are self-fulfilling, the ignorance statement for the Muddy Children is self-refuting.[3]

FACT 2. *Limit update models for 'positive-existential' modal formulas φ are definable in the modal μ-calculus. Arbitrary formulas require inflationary μ-calculus.*

Why is rationality self-fulfilling? And why is disagreement in beliefs self-refuting [11]?

OPEN PROBLEM 1. *Characterize the self-fulfilling and self-refuting formulas syntactically.*[4]

THEOREM 4 ([20]). *PAL plus iteration is Π^1_1-complete.*

All these questions also make sense for first-order logics of game trees, as used above.

Iterated upgrade of plausibility order More complex [4]: cycles, new logics. Links with learning theory: [13].

General logic of protocols For a more general program in the background, see [12, 28, 33].

8. PARADOX OF BACKWARD INDUCTION

Is the analysis stable under inversion [6]? After deliberation, we observe the *actual play of the game.* Do players now get cold feet?

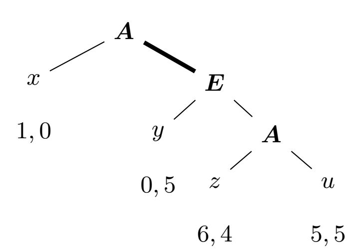

[2]Complexity of a logic is not task complexity for agents: [9].
[3]This is the global version, we discuss local versions with an actual world in the full paper.
[4]For many relevant results about the non-limiting case, see several recent papers by W. Holliday, T. Hoshi & Th. Icard, Stanford Logical Dynamics Lab: `http://stanford.edu/~thoshi/ldl/Home.html`

Backward Induction says that $\boldsymbol{A}$ will go left at the start, on the basis of logical reasoning available to both players. But if $\boldsymbol{A}$ plays *right*, what should $\boldsymbol{E}$ conclude? Perhaps $\boldsymbol{A}$ is not following the *BI* reasoning, and all bets are off as to what he will do later on - especially in long games?

Many theorems characterizing *BI* assume common knowledge [2] or true common belief in rationality. A richer analysis would take in *revision* policies by players learning a fact contradicting their beliefs in the course of a game [23]. These may stay close to rationality, or reflect other hypotheses. $\boldsymbol{E}$ might think: (a) '$\boldsymbol{A}$ is telling me that he is willing to take risks', (b) '$\boldsymbol{A}$ is an automaton with a general rightward tendency', and so on.

Conclusion: One should not just analyze games, but also the styles of the agents playing them.

9. RATIONALIZATION

Rationality is a 'bridge law' relating observable and postulated theoretical properties like Newton's laws in mechanics. It is interesting to hold on to it for a while:

> **Rationalize** By playing a move, a player gives information about her beliefs. These beliefs are such as to rule out that her actual move is strictly dominated-in-beliefs.[5]

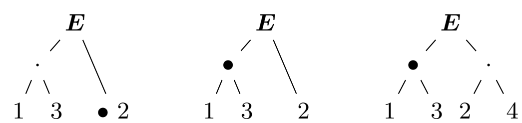

(a) (b) (c)

The play in Game (a) is rational by ascribing a belief to $\boldsymbol{E}$ that choosing left would result in outcome 1. Game (b) may be rationalized by ascribing a belief to $\boldsymbol{E}$ that the game will now reach 3. Game (c) suggests that $\boldsymbol{E}$ thinks she will reach 3, while she would have reached 2 if she had gone right.[6]

The point is not that one rule now replaces Backward Induction by 'forward induction' [7]. It is rather that *the past is informative*, telling us which choices players made in coming here. Here is an example adapted from [22]:

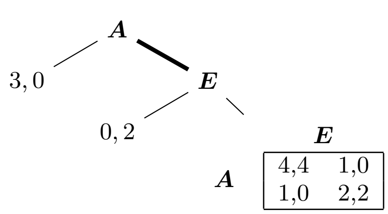

The most rational thing to do for $\boldsymbol{E}$ when $\boldsymbol{A}$ plays *Right* is to go *Right* as well. Perea's more general algorithm raises interesting logical issues, for which we refer to our full paper.

[5]Stronger rationality assertions will be discussed in the full paper. Also, the scenarios to follow assume that players see one unique most plausible history: but this can be relaxed.

[6]Beliefs of a player $\boldsymbol{E}$ do double duty here. Connected to a turn for the other player $\boldsymbol{A}$, they correspond to *expectations* about what $\boldsymbol{A}$ will do. But with a turn for $\boldsymbol{E}$ herself, they are more like *intentions*.

10. NEW LOGICAL ISSUES

Games must now have a distinguished point s, encoding the position where the actual play stands:

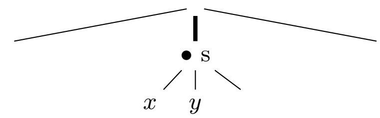

Players can let their choice depend on (a) the remaining game, (b) what players did so far in the larger game. While *BI* created one uniform plausibility relation $x \leq y$ among histories x, y, we now get a *ternary plausibility relation* $\leq_s xy$. This allows for differences between what players expect hypothetically if another move had been played than the actual one (say, that would have been a 'stupid move') and how they would feel if that move were actually played.

Thus, rationalization algorithms need not produce uniform expectations, not even going down one history.[7] To describe this in logic, we may need to 'lift' game trees to more complex standard models for games, since beliefs need not have a simple encoding any more as with *BI*.[8] In dynamic logics for plausibility upgrade, we then need to allow world-dependent relations (like in conditional logic). Behind any specific rationalization algorithms, there is a general issue here of developing the dynamic (limit) upgrade logic for ternary plausibility models.

11. PREFERENCE CHANGE AND GAME CHANGE

Folklore results make sense of almost any observed behavior - if we can construct preferences on the fly. Any strategy against the strategy of another player with known preferences can be rationalized by assigning suitable preferences.

The background to such algorithms are dynamic logics of *preference change* [18, 14] describing evaluation changes triggered by observing moves of a game.

Players need not even know precisely which game they are playing: a reality of social life. And even if they do know the game, they may want to change it. Dynamic-epistemic logics of *game change* have been proposed e.g. for modeling promises that change strategic equilibria [25] or for adding players.[9] But we need more 'cross-game logic'.[10]

12. THEORY OF PLAY

There is no unique way of defining 'best action'. The missing ingredient is information about the types of agent we are interacting with. The structure of a game by itself does not provide this, unless we make strong uniformity assumptions. We need more input.

The term coined for this perspective in [30] is *Theory of Play*. To make sense of what happens in a game, we must

[7]Backward Induction might be the *only* uniform and monotonic algorithm creating expectations.

[8]This need for 'lifting' is known for reasoning about strategies, cf. Ch. 10 of [27].

[9][21] discusses agents manipulating knowledge during play.

[10]Recall our analogy with mechanics. Why is postulating a force function behind observed particle behavior more than an ad-hoc device? This function still makes sense when we change the physical situation, adding or removing objects. Getting to grips with such uniformities is a major challenge.

combine information about game structure plus the agents in play. Game theory allows each player her own preferences, but the Backward Induction algorithm assumes uniformity on how players think and act, witness the symmetries. But we need much more variety: in computational limitations, belief revision policies, etc. [11]

At present, there is no Theory of Play in my sense of the term: only interesting bits and pieces that might help us create one. Here is what I see as some relevant tasks.

Taxonomy of players. In principle, there can be huge spaces of possible hypotheses concerning players. We need to constrain these to small sets of relevant options - and much literature has relevant proposals. These options seem to come in several different kinds. One is *processing properties* of agents: what are their powers of memory, observation, or even of inference? Another is *update* policies of agents: how will they revise their beliefs, or more generally, what learning methods do they follow? And a third dimension might be called *balance types* between information and evaluation: agents can be more optimistic or pessimistic in pursuing their goals, and so on.

Where to locate the diversity. One way of implementing such a taxonomy would be in an explicit model of agents, say using a class of automata endowed with beliefs and preferences. But diversity also lives elsewhere. Dynamic-epistemic logics of knowledge or belief change have no explicit agents, but they highlight diversity of observational access or plausibility shift in different signals (technically, 'event models') and the updates produced by these.[12] Of course, in doing so, they may still have hidden presuppositions that can be brought out, and then varied. Here is a result from [28]:

THEOREM 5. *An imperfect information game arises from iterated epistemic DEL update iff players have Perfect Memory and No Miracles (all learning is by observation).*[13]

As to possible variations, there are also natural *DEL*-style update rules for memory-bounded agents.

Objection: messiness. Theory of Play comes at the cost of a large space of hypotheses about agents, with models that can be much more complex than game trees.[14] We need to find simple taxonomies. Also, logical systems acknowledging variety tend to get complex. But this may be a matter of choosing the right architecture. Consider belief revision. Prima facie, it dissolves into many policies for relational update, with complete dynamic logics for each [24]. But [3] lets event models encode the variety, leaving only one rule of Priority Update with simple axioms. The challenge for a Theory of Play is acknowledging diversity, while letting logic do its usual job of abstraction and idealization.

[11] [17] has a suggestive map of agent diversity from the standpoint of dynamic-epistemic logic.

[12] E.g., it is not the agent that is 'radical', but a current way of taking an input signal may be radical.

[13] The cited paper suggests that synchronicity is also built in, but cf. [10] for an alternative analysis.

[14] 'Worlds' might be nodes in game trees, histories in game trees, or even thicker possible worlds that encode games, strategy profiles, and other features. Theory of Play seems to need all three levels.

13. REPERCUSSIONS

Bringing in agent diversity and theory of play is something that happens in many disciplines. Consider results on game play in computer science in terms of 'positional strategies' scenarios where simple memory-free agents can do an optimal job [1]. Or consider empirical results on actual behavior in auctions [19]. These illustrate the earlier mismatch between deliberative rationality and actual play, where preferences may change in the heat of battle. Theory of play may even affect philosophy. What is 'fair play' in ethics given the undeniable diversity of agents? Are uniformity assumptions the greatest justice, or the greatest injustice? There seem to be no easy answers.[15]

Theory of Play might even reach logic itself. What about a Theory of Inference describing human or computational agents engaging in deduction and other activities, and their different styles of doing so? Say, finite automata doing first-order proof or competing in logic games? Can logic get closer to actual reasoning if we relax its standard uniformity assumptions?

14. REFERENCES

[1] K. Apt and E. Grädel, editors. *Lectures in Game Theory for Computer Scientists*. Cambridge University Press, 2011.

[2] R. J. Aumann. Backward induction and common knowledge of rationality. *Games and Economic Behavior*, 8(1):6–19, 1995.

[3] A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory (LOFT7)*, volume 3 of *Texts in Logic and Games*, pages 13–60. Amsterdam University Press, Amsterdam, The Netherlands, 2008.

[4] A. Baltag and S. Smets. Group belief dynamics under iterated revision: fixed points and cycles of joint upgrades. In Heifetz [16], pages 41–50.

[5] A. Baltag, S. Smets, and J. A. Zvesper. Keep 'hoping' for rationality: a solution to the backward induction paradox. *Synthese*, 169(2):301–333, 2009.

[6] C. Bicchieri. Common knowledge and backward induction: A solution to the paradox. In M. Y. Vardi, editor, *TARK*, pages 381–393. Morgan Kaufmann, 1988.

[7] A. Brandenburger. Forward induction. Stern School of Business, 2007.

[8] B. de Bruin. *Explaining Games: The Epistemic Programme in Game Theory*. Springer, 2010.

[9] C. Dégremont, L. Kurzen, and J. Szymanik. Cognitive plausibility of epistemic models: Exploring tractability borders in epistemic tasks. ILLC Amsterdam and IAI Groningen, 2011.

[10] C. Dégremont, B. Löwe, and A. Witzel. The synchronicity of dynamic epistemic logic. In *Proceedings TARK*, 2011.

[11] C. Dégremont and O. Roy. Agreement theorems in dynamic-epistemic logic. In Heifetz [16], pages 91–98.

[15] Should 'fair' exams be individualized to the intelligence level of individual students?

[12] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about knowledge*. The MIT Press, Cambridge, Mass., 1995.

[13] N. Gierasimczuk. *Knowing One's Limits. Logical Analysis of Inductive Inference*. PhD thesis, Institute for Logic, Language and Computation (ILLC), Universiteit van Amsterdam (UvA), Amsterdam, The Netherlands, 2010. ILLC Dissertation series DS-2010-11.

[14] P. Girard. *Modal Logic for Belief and Preference Change*. PhD thesis, Department of Philosophy, Stanford University, Stanford, CA, USA, Feb. 2008. ILLC Dissertation Series DS-2008-04.

[15] J. Y. Halpern and M. Y. Vardi. The complexity of reasoning about knowledge and time. I. lower bounds. *Journal of Computer and System Sciences*, 38(1):195–237, Feb. 1989.

[16] A. Heifetz, editor. *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2009), Stanford, CA, USA, July 6-8, 2009*, 2009.

[17] F. Liu. Diversity of agents and their interaction. *Journal of Logic, Language and Information*, 18(1):23–53, 2009.

[18] F. Liu. *Reasoning About Preference Dynamics.* Synthese Library. Springer Science Publisher, 2011.

[19] S. McClure. Decision making, 2011. Lecture slides SS100, Stanford University.

[20] J. S. Miller and L. S. Moss. The undecidability of iterated modal relativization. *Studia Logica*, 79(3):373–407, 2005.

[21] R. Parikh, C. Tasdemir, and A. Witzel. The power of knowledge in games, 2011. Working paper, CUNY Graduate Center & New York University.

[22] A. Perea. Belief in the opponents' future rationality, 2011. Working paper, Epicenter, Department of Quantitative Economics, University of Maastricht.

[23] R. Stalnaker. Extensive and strategic form: Games and models for games. *Research in Economics*, 53:293–291, 1999.

[24] J. van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 17(2):129–155, 2007.

[25] J. van Benthem. In praise of strategies. In J. van Eijck and R. Verbrugge, editors, *Foundations of Social Software*. College Publications, London, 2007.

[26] J. van Benthem. Rational dynamics. *International Game Theory Review*, 9(1):13–45, 2007. Errratum reprint: 9(2), 377-409.

[27] J. van Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, 2010.

[28] J. van Benthem, J. Gerbrandy, T. Hoshi, and E. Pacuit. Merging frameworks for interaction. *Journal of Philosophical Logic*, 38(5):491–526, 2009.

[29] J. van Benthem and A. Gheerbrant. Game solution, epistemic dynamics and fixed-point logics. *Fundamenta Informaticae*, 100(1-4):19–41, 2010.

[30] J. van Benthem, E. Pacuit, and O. Roy. Toward a theory of play: A logical perspective on games and interaction. *Games*, 2(1):52–86, 2011.

[31] J. van Benthem, S. van Otterloo, and O. Roy. Preference logic, conditionals and solution concepts in games. In H. Lagerlund, S. Lindström, and R. Sliwinski, editors, *Modality Matters: Twenty-Five Essays in Honour of Krister Segerberg*, number 53 in Uppsala Philosophical Studies, pages 61–76. University of Uppsala, Upsala, 2006.

[32] W. van der Hoek and M. Pauly. Modal logic for games and information. In P. Blackburn, J. van Benthem, and F. Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical Reasoning*, pages 1077–1148. Elsevier Science Inc., Amsterdam, 2006.

[33] Y. Wang. *Epistemic Modelling and Protocol Dynamics*. PhD thesis, Institute for Logic, Language and Computation (ILLC), Universiteit van Amsterdam (UvA), Amsterdam, The Netherlands, Sept. 2010. ILLC Dissertation series DS-2010-06.