LOGICS OF POWERS AND STRATEGIES

Johan van Benthem, Amsterdam & Stanford, http://staff.science.uva.nl/~johan/

Workshop *Social Dynamics of Information Change*, Amsterdam, 3 Dec 2013

Abstract We discuss reasoning about strategies. We propose logics that analyze elementary game-theoretic reasoning, and identify a few challenges for new research, some of them close to home.

1 Reasoning about strategies, a priori or fieldwork? *Approach 1*: identify strategies with objects we know already: programs (*PDL*), proofs (lambda calculus). *Approach 2*: analyze basic proofs in game theory or computer science about strategies, without preconceptions.

2 Classical theorems live in temporal forcing logic Zermelo's Theorem: 2-player zero-sum games with finite depth for histories are determined. Heart of the proof is a modal recursion:

$$WIN_i \leftrightarrow ((\mathbf{end} \wedge \mathbf{win}_i) \vee (\mathbf{turn}_i \wedge < move_i > WIN_i) \vee (\mathbf{turn}_j \wedge [move_i]\ WIN_i))$$

Open set X of histories h: $h \in X$ iff some initial segment of h has all its extensions in X.

Gale-Stewart Theorem Infinite games with an open winning condition are determined.

The proof revolves around a general principle that follows the above recursion:

Weak Determinacy Either player ***i*** has a winning strategy, or player ***j*** has a strategy which ensures that ***i*** never reaches a position where ***i*** has a winning strategy.

Temporal logic of forcing powers Extensive games may be viewed as branching tree models ***M*** for time, with histories as branches h, and stages s as points on these:

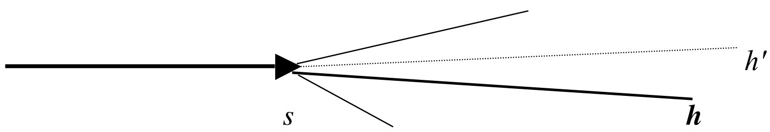

Standard modal temporal language $\boldsymbol{M}, h, s \models \varphi$ plus a *strategic forcing modality* $\{i\}\varphi$:

$\boldsymbol{M}, h, s \models \{i\}\varphi$ player i has a strategy from s onward which ensures that only histories h' result for which, at each stage $t \geq s$, $\boldsymbol{M}, h', t \models \varphi$

While this looks local to stages s, φ can also be a global stage-independent property of the histories h'. Note that the condition does not imply that the actual history h satisfies φ.

Fact Weak Determinacy is defined by the formula $\{i\}\varphi \vee \{j\}\neg\{i\}\varphi$

Valid principles The following principles are valid in temporal forcing logic:

(a) standard laws of *branching temporal logic*
(b) standard logic of a *monotonic neighborhood modality* for $\{i\}\varphi$, plus $\{i\}\varphi \rightarrow \Box\{i\}\varphi$
(c1) $\{i\}\varphi \leftrightarrow ((end \wedge \varphi) \vee (turn_i \wedge \Diamond O(\{i\}\varphi) \vee (turn_j \wedge \Box O(\{i\}\varphi)))$
(c2) $(\alpha \wedge \Box G ((turn_i \wedge \alpha) \rightarrow \Diamond O\alpha) \wedge ((turn_j \wedge \alpha) \rightarrow \Box O\alpha))) \rightarrow \{i\}\alpha$
(c3) $(\{i\}\varphi \wedge \{j\}\psi) \rightarrow \Diamond(\varphi \wedge \psi)$

Proving our basic results formally The essential steps in the proof of Weak Determinacy:

$(turn_i \wedge \neg\{i\}\varphi) \rightarrow \Box O\neg\{i\}\varphi)$	from (c1)
$(turn_j \wedge \neg\{i\}\varphi) \rightarrow \Diamond O\neg\{i\}\varphi)$	from (c1)
$\neg\{i\}\varphi \rightarrow \{j\}\neg\{i\}\varphi$	from (c2)

We can also derive the Gale-Stewart Theorem for open φ satisfying $\varphi \to F\Box G\varphi$:

$$\{i\}\varphi \vee \{j\}\neg\varphi$$

Zermelo's Theorem follows, 'having an endpoint' is open: $F\ \boldsymbol{end} \to F\Box G\ F\ \boldsymbol{end}$.

Temporal forcing logic The modal *K4*-axiom $\{i\}\alpha \to \{i\}\{i\}\alpha$ is valid: it expresses 'safety'. Non-valid principles also of interest: why $G\{i\}\alpha \to F\alpha$ fails on the 'infinite comb'.

Fact Temporal forcing logic is decidable.

Proof All temporal modalities plus the forcing modality can be defined in monadic second-order logic *MSOL* on trees with successor relations. Then Rabin's Theorem applies. ■

3 Nondeterminacy, strategy stealing, and logics of special games Interval selection game: Take any free ultrafilter U on the natural numbers N. Two players pick successive closed initial segments of N of arbitrary finite lengths, producing an infinite sequence like this:

$\boldsymbol{i}$: $[0, n_1]$, with $n_1 > 0$, $\boldsymbol{j}$: $[n_1{+}1, n_2]$, with $n_2 > n_1{+}1$, etc.

Player $\boldsymbol{i}$ wins if the union of all intervals chosen is in U – otherwise, $\boldsymbol{j}$ wins. Winning sets not open. Game non-determined, by "strategy stealing" argument: Player $\boldsymbol{i}$ has no winning strategy. For, if he had one, $\boldsymbol{j}$ could use it with a delay of one step to *copy* $\boldsymbol{i}$'s responses to her own moves disguised as $\boldsymbol{j}$-moves. Both resulting sets of intervals (disjoint up to some finite initial segment) would have their unions in U: which cannot be, as U is free. Likewise, player $\boldsymbol{j}$ has no winning strategy.

Analyzing this proof. The strategy σ gives $\boldsymbol{i}$ a first move $\sigma(-)$. Now let $\boldsymbol{j}$ play any move e. $\boldsymbol{i}$'s response is $\sigma(\sigma(-), e)$, after which it is $\boldsymbol{j}$'s turn again. Crucially, the same sequence of events can be viewed differently, as a move $\sigma(-)$ played by $\boldsymbol{i}$, followed by a move $e\ ;\ \sigma(\sigma(-), e)$ played *by* $\boldsymbol{j}$, after which it is $\boldsymbol{i}$'s turn. Presupposes a special, but natural property of a game:

Composition closure Any player can play any concatenation of available successive moves.

The two stages described start the same subgame but for turn markings interchanged (one of the 'duals' of Venema 2007). $\boldsymbol{j}$'s strategy uses $\boldsymbol{i}$'s strategy in the other game to produce *identical runs* in both subgames, except for the turn switch. Contradiction via a *Copy Law*:

Fact In games that satisfy composition closure, the following modal formula is valid:
$\{i\}\varphi \to \Diamond OO\{j\}\varphi_d$, where φ_d is φ with all turn occurrences of i, j interchanged.

4 Logics of strategies as programs We can define explicit strategies as transition relations, say in *PDL*. As we work move by move, *flat programs* suffice using only atomic actions, tests, sequence $;$, and choice $\cup$ – and often even just unions of guarded actions of the form $?\varphi\ ;\ a\ ;\ ?\psi$. Consecutive moves become important when we think of forcing outcomes:

Fact For any game program expression σ, *PDL* can define an explicit forcing modality $\{\sigma, \boldsymbol{i}\}\varphi$ stating that σ is a strategy for player $\boldsymbol{i}$ forcing the game, against any play of the others, to pass only through states satisfying φ.

Proof The *PDL*-formula $[((?\boldsymbol{turn}_i\ ;\ \sigma) \cup (?\boldsymbol{turn}_j\ ;\ move_j))^*]\varphi$ is equivalent to $\{\sigma, \boldsymbol{i}\}\varphi$. ■

Extra *PDL* programs also describe strategy combination via intersection (van Eijck 2012).
Issue Strategies in infinite games need not involve terminating programs (*WHILE* vs. *PDL*).

5 Zoom, combination, invariants Logic: different levels of 'zoom' on reasoning practices. Sometimes, we want to see details, sometimes the broad picture. Often both are needed:

Fact 'Lacking a strategy' needs the non-explicit earlier modality: $\neg\{\boldsymbol{i}\}\varphi$.

Combined language has surprises. The strategy 'be successful': $\sigma_{\varphi, i} = ?\boldsymbol{turn}_i\ ;\ move_i,\ ?\{\boldsymbol{i}\}\varphi$

Fact The following equivalence is valid: $\{i\}\varphi \leftrightarrow \{i, \sigma_{\varphi, i}\}\varphi$

Definition levels for strategies Explanation: need levels of definability. E.g., finite memory, copying (Abramsky 2008, Graedel, Thomas & Wilke 2002, Osborne & Rubinstein 1994).

Two-level views and invariants *Games* have moves and internal properties, marking nodes as turns or wins. There is also an external *game board* recording observable or other relevant behavior. Examples: `graph games', logic games. Often a reduction map ρ sends game states to corresponding states on the game board satisfying a certain amount of back-and-forth simulation. Strategies in a game maintain some *invariant* at the level of its board. In fact, the forcing modalities in the above triviality result may be seen as, somewhat bleak, invariants.

Games in computational logic Fast-growing theory. E.g., Positional Determinacy Theorem (Emerson & Jutla 1991, Mostowski 1991) says that graph games with parity conditions are determined with 'positional strategies', whose moves depend only on the graph component of the current state of play. Thus, winning positions project via ρ to a set of board positions. Logical explanation of positional determinacy might be: *translation* from modal forcing statements in the game to equivalent modal fixed-point assertions about matching graph states.

6 Strategy logics with operations on games It also makes sense to add terms for games to the logic, with relevant operations such as choice, sequence, dual, and parallel composition.

Dynamic game logic Extend forcing modalities with game terms, with formulas interpreted, not inside games, but on their game boards. The result is a two-agent *PDL* on neighbourhood models, with typical axioms such as the one for `choice games':

$$\{G \cup H, \boldsymbol{i}\}\varphi \leftrightarrow \{G, \boldsymbol{i}\}\varphi \vee \{H, \boldsymbol{i}\}\varphi$$

The elementary soundness argument can help tease out the underlying calculus of strategies:

Player $\boldsymbol{i}$ starts a game $G \cup H$ by choosing to play either G or H. If $\boldsymbol{i}$ has a strategy σ forcing φ-outcomes in $G \cup H$, its first step describes her choice, *left* or *right*, and the rest forces φ-outcomes in the chosen game. Vice versa, if she has a strategy σ forcing φ in game G, prefixing it with a move *left* gives her a strategy forcing φ in $G \cup H$. General operations involved: $head(\sigma)$ gives the first move of the strategy σ, and $tail(\sigma)$ the remain-der, in a way that validates $\sigma = (head(\sigma), tail(\sigma))$. The second part prefixed an action a to a strategy σ, yielding $a\ ;\ \sigma$ with laws like $head(a\ ;\ \sigma) = a$, $tail(a\ ;\ \sigma) = \sigma$. ■

This strategy calculus does not look like *PDL*. It rather suggests a *co-algebraic* perspective of 'observing' and then looking at the rest of the strategy. This brings us to a next approach.

Linear game logic Logical formulas are game terms, *linear logic* encodes game equivalence or inclusion. Key strategy: Copy-Cat for 'parallel excluded middle'. Games live in the above temporal models, so branching temporal logic can still analyze basic arguments, such as soundness of the Cut Rule that turns on the properties of Copy-Cat plus wining conditions. Reasoning in parallel games involves 'shadow arguments' about non-actual subgames, reminiscent of proofs about games in computational fixed-point logics.

7 Game theory, knowledge, preference Adding notions from philosophical logic, agent systems and game theory, real game theory. What players prefer, know, believe.

Knowledge I can hardly 'copy' or 'steal' a strategy if I do not know it. Merge strategy logics with epistemic logics or other logics of information. Next come *imperfect information games* where strategies can have knowledge-dependent instructions, and an essential informational nature of players, endowed with perfect memory, observation-driven, or otherwise.

Preference The balance of moves, beliefs, and preference drives rational play in the usual sense. Challenge: extend ideas from computational logic and philosophical logic to deal with strategic reasoning in this case. Van Benthem & Gheerbrant 2011 has an extensive study of Backward Induction in the first-order fixed-point logic *LFP(FO),* but without an explicit program definition in the above terms. Other issues: game algebra, game equivalence. And: preference can run between infinite histories – what does this do to our *MSOL* analysis?

Pang of conscience 1 Our forcing modality only depends on histories reachable by further play with moves of my strategy and your arbitrary moves. It does *not* say what might happen counterfactually if I had played differently (Nash equilibrium!). But strategies as usually defined give 'counterfactual' moves for me at nodes I would not even reach when playing the strategy from the root. Is there a technical difference between the 'actual' and the full 'counterfactual' version of the strategy logics considered so far? Should we add models for games with whole strategy profiles after all?

8 Simultaneous moves, logics of joint action Tree models as above, now with simultaneous actions, arise in philosophy (*STIT*), and (evolutionary) game theory. Additional modality

$<i>\varphi$: *i* can choose an action such that whatever the others do, the next state satisfies φ

Add $<i, j>\varphi$: what can be made true by simultaneous action of both players: *STIT*, *ATL*, coalition logic. We then define again the earlier strategy modalities $\{i\}\varphi$ referring to what is true throughout any history where player *i* plays according to her strategy.

Fact Laws of temporal forcing logic simplify: e.g., $\{i\}\varphi \leftrightarrow (\varphi \,\&\, <i>\{i\}\varphi)$ for both players.

Determinacy goes away: only monotone and consistent powers of players in simultaneous games. Connect with powers in repeated temporal *STIT* rounds. Is some analogue of Weak Determinacy left? What about a connection with *Blackwell Determinacy* (Vervoort, p.c.)?

Conjecture The temporal logic of simultaneous forcing powers logic is still embeddable into monadic second-order *MSOL* over trees with a finite set of successor relations.

With *WD*-ish reasoning, a difference emerges between strong and weak powers. Either *i* has a strong power forcing a winning strategy at the next stage ($<i>\{i\}\varphi$ ($<i\text{'s actions}>[j\text{'s actions}]\{i\}\varphi$), or *j* has a weak power, some pattern of behavior for any action of i's that blocks the assertion: $[i\text{'s actions}]<j\text{'s actions}>\{i\}\varphi$. A weak power is not something you can choose to do deliberately, but it is a possibility about your behavior that others cannot rule out. Weak powers may have nice logical laws, and we may want to use ideas from *STIT* logic. (Alternative: probabilistic features now become essential?)

Issue: how to define strategies explicitly? Will a version of *PDL* do the job? What makes specific strategies special? Say, characterize the ubiquitous *Copy-Cat* as the unique strategy satisfying natural conditions? A relative in game theory is *Tit-for-Tat*, but not quite the same. *Copy-Cat* creates a history whose projections for both players are the same, so that one of the players wins one of these if his opponent is his own dual. But the pay-off for player *i* in an infinite history of iterated one-shot Prisoner's Dilemma games cannot be read off from the subhistory of *i*'s moves only: the context matters, viz. what the other player *j* did.

Pang of conscience 2 Two views of strategies: *concrete* or *generic*. Do we enumerate moves, or give a generic description? Tricky: "do what the other does" is a best equilibrium for Prisoner's Dilemma.

Likewise, the benchmark proof of a simple fact such as *TfT* being in Nash equilibrium with itself is less straightforward than the usual facts about *Copy-Cat*, mixing preference and actions. What base logic would naturally represent the basic proofs about Nash equilibria in classical and evolutionary game theory (cf. Binmore, Axelrod, or Maynard-Smith)?

9 Challenges/connections to be made *Computational logic*: sophisticated results on infinite games, automata, and fixed-point recursion (Venema 2007). *Philosophical logic*: powers in *STIT* and other logics of agency. *Descriptive set theory*: determinacy results for Blackwell games. *Evolutionary game theory*: logics for limit behavior in dynamical systems.

10 References (a) J. van Benthem, 2014, *Logic in Games*, The MIT Press, Cambridge MA. (b) JvB & E. Pacuit, 2014, 'Logics of Choice and Change', in Th. Moeller, ed., *Nuel Belnap on Indeterminism and Free Choice*, Springer, Dordrecht. (c) JvB, S. Ghosh & R. Verbrugge, eds., 2014, *Reasoning about Strategies*, Springer Lecture Notes in AI, FoLLI Series. (d) M. Vervoort, 2000, *Games, Walks and Grammars*, Ph.D. thesis, ILLC. (e) Y. Venema, 2007, *Lectures on the Modal μ–Calculus*, ILLC. (b) B. Loewe, 'Set Theory of Infinite Imperfect Information Games', in A. Andretta, ed., *Set Theory: Recent Trends and Applications*, Napoli 2005 [Quaderni di Matematica 17], 137–181.