In Praise of Strategies

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

August 2007, *second draft*

Abstract

This note high-lights one major theme in my lecture notes *Logic in Games* (van Benthem 1999 – 2002): the need for explicit logics that define agents' strategies, as the drivers of interaction in games. Our text outlines issues, recalls published results from the last few years, and raises new open problems. Results are mainly quoted, and the mephasis is on new notions and open problems. For more details on the various topics discussed, see the relevant references.

1 Strategies as first-class citizens

Much of game theory is about the question whether strategic equilibria exist. But there are hardly any explicit languages for defining, comparing, or combining *strategies as such* – the way we have them for actions and plans, maybe the closest intuitive analogue to strategies. True, there are many current logics for describing game structure – but these tend to have existential quantifiers saying that "players have a strategy" for achieving some purpose, while descriptions of these strategies themselves are not part of the logical language (cf. Parikh & Pauly 2003, van der Hoek, van Otterloo & Wooldridge 2005). Therefore, I consider strategies 'the unsung heroes of game theory' - and I want to show how the right kind of logic can bring them to the fore. One guide-line of adequacy for doing so, in the fast-growing jungle of 'game logics', is the following: we would like to explicitly represent the elementary reasoning about strategies underlying many basic game-theoretic results. Or in more general terms, we want to explicitly represent agents' reasoning about their plans.

2 Games as models for modal process logics

Van Benthem 2002, mainly an extract from the lecture notes *Logic in Games*, shows how modal logic fits naturally with extensive games, viewed as process models from computer science, viz. labeled transition systems with some special annotation for players' activities.

Basic modal logic Extensive game trees may be viewed as state spaces for some multi-agent process. Labeled modalities $<a>φ$ then express that some move a is available leading to a next node in the game tree satisfying φ. Then modal operator combinations describe potential interaction. For instance, in a self-explanatory notation, the formula

*[move-**A**]<move-**E**>φ*

says that, at the current node of evaluation, player ***E*** has a strategy for responding to ***A***'s initial move which ensures that φ results after two steps of play. Extending this to extensive games up to some finite depth k, and using alternations *[]<>[]<>*... of modal operators up to length *k* to reach the end-points of the game tree, we can express the existence of winning strategies and the like in fixed finite games. Indeed, given this connection, with finite depth, standard logical laws have immediate game-theoretic import. In particular, consider the valid *law of excluded middle* in the following modal form

$$[\,]\langle\rangle[\,]\langle\rangle\ldots\varphi \vee \langle\rangle[\,]\langle\rangle[\,]\ldots\neg\varphi,$$

where the dots indicate the depth of the tree. This expresses the *determinacy* of these games, as stated in Zermelo's Theorem that all finite zero-sum two-player games are determined.

Modal μ-calculus Unfortunately, such modal game-by-game definitions are not 'generic', as they depend on the particular model considered – and Zermelo's inductive argument is rendered much more faithfully by means of just one fixed formula in the *modal μ-calculus* (Bradfield & Stirling 2006). To make the relevant point more generally, let us first define the following 'forcing modality'in games:

> $\boldsymbol{M}, s \models \{i\}\varphi$ iff player *i* has a strategy for the sub-game starting at *s* which guarantees that only nodes will be visited where φ holds, whatever the other does.

Forcing talk is widespread in games, and it is an obvious target for logical formalization. Note that, for convenience, $\{i\}\varphi$ talks about intermediate nodes, not just end nodes of the game. The existence of a winning strategy for player *i* can now be expressed as

$$\{i\}\,(\boldsymbol{end} \rightarrow win_i)$$ [1]

Here is an explicit definition for this assertion in the modal μ-calculus, using some further obvious proposition letters and action symbols for indicating players' turns and moves at nodes of the game tree. In this formula, the symbol *j* is used fort the other player in the game:

$$\{i\}\varphi \;=\; \nu q\bullet\ (\varphi\ \&\ (turn_i\ \&\ \langle move\text{-}i\rangle q) \vee (turn_j\ \&\ [move\text{-}j]q)).$$

This definition is faithful to the obvious recursive meaning of having a strategy for player *i*, regardless of what the others do – and it is generic, since it works in all games viewed as models ***M*** for our language. Incidentally, we use a *greatest fixed-point* operator $\nu q\bullet$ here, rather than a smallest fixed-point operator $\mu q\bullet$, for easier extension to infinite games later on.

Propositional dynamic logic But now to strategies as such! An obvious candidate for defining these is *propositional dynamic logic PDL*, combining propositions about nodes in the game tree with programs defining transition relations between such nodes. Without loss of information, a strategy for player *i* is a binary relation between nodes. At turns for player *i*,

---

[1] Here, ***end*** is a proposition letter for end nodes, or a complex modal formula saying that no move is possible, and win_i is a proposition letter saying that player *i* wins at the current node.

it picks out one transition, at turns for the others, it allows any move. Here is a simple fact for a start. *General PDL-programs now define strategies* σ. Note that, in general, these programs may denote arbitrary transition relations, not just uniquely valued functions. Such relations constrain *i*'s moves at her turns, without necessarily narrowing them down to just a single one. This possible non-determinacy is fine, since it makes eminent sense for plans, and by extension, also for a broader notion of strategies in games.

Now, we can define an explicit version of the earlier forcing modality, indicating the strategy involved – even without recourse to the full power of the modal μ-calculus:

Fact 1 For any program expression σ, *PDL* can define the explicit forcing modality $\{\sigma, i\}\varphi$ stating that σ *is a strategy for player i forcing the game, against any play of the others, to pass only through states satisfying* φ.

The precise definition is an easy exercise in modal logic (cf. van Benthem 2002).

But *PDL* has further uses in this setting. Consider any finite game ***M*** with a strategy σ for player *i*. As a relation, σ is a finite set of ordered pairs *(s, t)*. Thus, it can be defined by enumeration as a program union, if we define these ordered pairs. To do so, assume we have an 'expressive' model ***M***, where states *s* are definable in our modal language by formulas def_s.[2] Then we define transitions *(s, t)* by formulas $def_s; a; def_t$, with *a* the relevant move:

Fact 2 In expressive finite extensive games, all strategies are *PDL*-definable.

Of course, this is a trivial result, but it does suggest that *PDL* is on the right track.

Van Benthem 2002 also discusses further issues about *PDL* as a 'calculus of strategies'. For instance, suppose that player ***E*** plays strategy σ, and at the same time, ***A*** plays strategy τ. What end nodes are reachable in this way? (With standard strategies, a unique outcome will be reached.) This calls for an operation on strategies describing the *joint strategy* of $\{E, A\}$ – and with a little reflection, it is clear that this joint strategy is just the *intersection* $\sigma \cap \tau$ of the relations σ, τ. This operation takes us outside of *PDL* proper, but then, $PDL^{\cap}$ with intersections added is still a reasonably simple modal language.

Thus, propositional dynamic logic does a reasonable job in defining explicit strategies in simple extensive games. In the next sections, we will see whether it can be extended to deal with more realistic game structures, such as preferences and imperfect information. But for here, we end with an open problem concerning the purely modal action part itself.

---

[2] This expressive power can be achieved in several ways, e.g., using backward temporal modalities which can describe the total history leading up to *s*. Cf. Rodenhauser 2001.

Stronger modal logics of strategies? The modal μ-calculus is a natural strengthening of *PDL*, but it has no explicit programs or strategies, as its formulas merely define properties of states. Is there a counterpart to the μ-calculus which also extends *PDL* in terms of defining corresponding transition relations? E.g., a strategy 'keep playing *a*' guarantees infinite a-branches for true greatest fixed-point formulas like $\nu p \bullet \langle a \rangle p$. [3]

3 Preference structure and more realistic games

Real games add preferences for players over outcome states, or utility values beyond 'win' and 'lose'. In this case, defining the so-called Backward Induction procedure for solving extensive games, rather than just Zermelo winning positions, becomes a benchmark for game logics. Here the issue is not *whetherf* this can be defined at all. Any simple game concept can be phrased in some modal-like language with transition relations for moves, provided one adds suitable modalities for the preference order. But can the paraphrase be done in a perspicuous manner, generating some new insight?

Fact 3 The Backward Induction path is definable in modal preference logic.

Solutions have been published by Board, Bonanno, and many others: cf. Harrenstein 2004, De Bruin 2004. Given all this, we do not state an explicit *PDL*-style solution here.

Betterness, preference, and expectation Actually, the situation to be analyzed is somewhat subtle conceptually. A game gives players' direct preferences over outcomes. The Backward Induction algorithm then lifts these to a binary order among intermediate nodes. But as pointed out in van Benthem 2002, 2007D, this order then gets a re-interpretation. It does not just represent what players *prefer*, but what they *expect to happen*, given rationality assumptions about how the other players will proceed. Thus, the resulting binary order is more like the plausibility relations used to interpret conditional beliefs in doxastic logic, generated from a mixture of preference and assumptions about behaviour of other players.

Modal preference languages come in many kinds. Some recent proposals are in Harrenstein 2004, van Otterloo 2005, van Benthem & Liu 2007, van Benthem, Girard & Roy 2007, van Benthem, van Otterloo & Roy 2006. In particular, the latter paper has a new take on what makes Backward Induction tick, using a preference modality

$\langle pref_i \rangle \varphi$: player *i* prefers some node where φ holds to the current one.

It then defines the backward induction path as a unique relation σ, not by a modal formula over models ***M***, but via the following *frame correspondence* on finite structures:

---

[3] Van Benthem 2005A looks at richer fragments than *PDL* with programs as solutions to fixed-point equations of special forms,guaranteeing uniform convergence by stage ω.

Fact 4 The *BI* strategy is definable as the unique relation σ satisfying the following axiom for all propositions *P* – viewed as sets of nodes –, for all players *i*:
$(turn_i \ \& <\sigma^*>(\textbf{end} \ \& \ P)) \rightarrow [move\text{-}i]<\sigma^*>(\textbf{end} \ \& <pref_i>P)$.

Beyond these known things from the current literature, let us now ask a different type of question here, about the right combination of propositional dynamic logic and preferences.

Since *PDL* has both formulas as properties of states and programs as inter-state relations, we can also put preference structure at the same *two levels*. One sort of preference runs between states, interpreting standard modal operators, as in the above preference logics. The other locus places preferences between state transitions (specific 'moves' or 'events'), or global transition relations – as in the dynamic deontic logic of van der Meijden 1996. For this contrast, compare the general distinction made in ethics between 'deontology' and 'consequentialism'. Do we compare the worlds resulting from actions when judging our duties, or do we qualify those actions themselves as 'better' or 'obligatory'?

Problem 1 Design a preference logic for games comparing both worlds and actions.

4 Epistemic logic and extensive games with imperfect information

Next, consider extensive games of imperfect information, which involve 'information sets', or equivalence relations $\sim_i$ between nodes which players *i* cannot distinguish. Van Benthem 2001 points out the obvious, viz. how these games model a combined epistemic modal language including knowledge operators $K_i\varphi$ interpreted in the usual manner as

"φ is true at all nodes $\sim_i$–related to the current one".

This language can make crucial distinctions such as knowing 'de dicto' that one has a move with effect φ, versus having some move of which one knows 'de re' that it yields φ:

$K<a \cup b>\varphi$ versus $K<a>\varphi \vee K<b>\varphi$

Moreover, as epistemic relations describe what agents can observe in the course of a game, this language can define special properties of agents through modal frame correspondences. An example is the following syntactic/semantic analysis of Perfect Recall for a player *i*:

Fact 5 The axiom $K_i[a]\varphi \rightarrow [a]K_i\varphi$ holds for player *i* w.r.t. any proposition φ iff ***M*** satisfies *Confluence*: $\forall xyz: ((xR_a y \ \& \ y\sim_i z) \rightarrow \exists u: ((x\sim_i u \ \& \ uR_a z)$.

Similar analyses work for other memory assumptions, and other types of observational powers for agents. For instance, looking in the opposite direction of Perfect Recall, agents satisfy the principle of 'No Miracles' when their epistemic uncertainty relations can only disappear when they observe two subsequent events which they can distinguish. [4]

Now once again for explicit strategies! As before, we can add *PDL*-style programs here to define players' strategies under the new circumstances. But there is a twist. Especially relevant then are the '*knowledge programs*' of Fagin et al. 1995, whose only test conditions for actions are knowledge statements for agents. In such programs, the actions prescribed for an agent can only be guarded by conditions which the agent knows to be true or false. It is easy to see that knowledge programs can only define *uniform strategies*, i.e., transition relations where a player always chooses the same move at any two game nodes which she cannot distinguish epistemically. A converse also holds, modulo some assumptions on expressiveness of the game language defining nodes in the game tree (van Benthem 2001):

Fact 6 On expressive finite games of imperfect information, the uniform strategies are precisely those definable by knowledge programs in epistemic *PDL*.

But there is still more to games of imperfect information. As with adding preferences, there are *two levels* for making our base logic *PDL* epistemic. One can connect worlds, as with the above language with standard epistemic modalities K_i. But one can also place epistemic structure on the moves themselves, as in *dynamic epistemic logic*. This raises the issue what sort of games correspond to models for dynamic-epistemic logic. We refer to van Benthem & Liu 2004, van Benthem, Gerbrandy & Pacuit 2007 for details of definitions and proofs:

Fact 7 An extensive game is isomorphic to a repeated product update model *Tree(**M**, **E**)* over some epistemic event model ***E*** iff it satisfies, for all players: (a) Perfect Recall, (b) (uniform) No Miracles, and (c) Bisimulation Invariance for domains of moves. [5]

But again, there are other issues. Epistemized *PDL*, either way, is also a good setting for pursuing the famous distinction between "knowing that" versus "knowing how" (Gochet 2006). Strategies are ways of achieving goals, and hence they represent *procedural know-how*. Many authors have proposed the latter as a challenge when epistemic logic comes in:

Problem 2 Define what it means to 'know a strategy' in epistemic *PDL*, and then develop a version of *DEL* with explicit strategies (but see Section 9 below).

This question comes from van Benthem 2006, Section 8, which discusses ways of making standard algorithmic tasks epistemic. In particular, it states the desideratum that strategies □ should be *epistemically transparent*, in the sense that when a part of □ has been played, players know at all intermediate stages that playing the rest will achieve the intended result.

---

[4] Our analysis restates that of Halpern & Vardi on epistemic-temporal logic (Fagin et al. 1995).

[5] I.e., two epistemically bisimilar nodes in the tree have the same moves possible at them.

When players are not assumed to have Perfect Recall, defining this in a generic manner seems non-trivial. But of course, there is much more to analyzing epistemic know-how.

5 Ignorance about the future, beliefs and expectations

The phrase 'imperfect information' covers two intuitively different senses of knowledge, which are sometimes confused. One is *observation uncertainty*: players may not have been able to observe all events so far, and so they do not know just where they are in the game tree. This 'past-oriented' knowledge and ignorance is found in *DEL* or epistemic temporal logics (van Benthem & Pacuit 2006). But there is also 'future-oriented' *expectation uncertainty*: players may not know where the game is heading since they do not know what others, or they themselves, are going to do. Modeling the latter type of knowledge and ignorance is not necessarily reducible to uncertainty between local nodes: it may involve current uncertainty between whole future histories, or between players' strategies (i.e., whole ways in which the game might evolve). Here are a few relevant observations from the class notes van Benthem 2004 on describing information update and belief revision along the branches of a game tree.

Branching epistemic temporal models The following structure is common to many fields in computer science and philosophy (cf. the surveys van Benthem & Pacuit 2006, van Benthem, Gerbrandy & Pacuit 2007). In tree-like models for branching time, 'legal histories' *h* represent possible evolutions of a given game. At each stage of the game, players are in a node *s* on some actual history whose past they know, either completely or partially, but whose future is yet to be fully revealed:

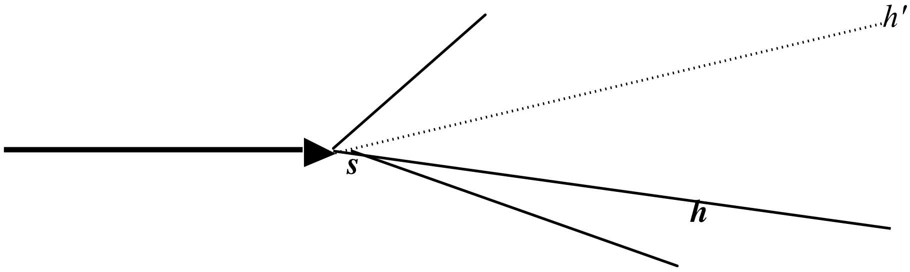

This can be described in an action language with knowledge, belief, and added temporal operators. We first describe games of perfect information (about the past, that is):

(a) $\boldsymbol{M}, h, s \models F_a\varphi$ iff $s \frown \langle a \rangle$ lies on h and $\boldsymbol{M}, h, s \frown \langle a \rangle \models \varphi$

(b) $\boldsymbol{M}, h, s \models P_a\varphi$ iff $s = s' \frown \langle a \rangle$, and $\boldsymbol{M}, h, s' \models \varphi$

(c) $\boldsymbol{M}, h, s \models \langle\rangle_i\, \varphi$ iff $\boldsymbol{M}, h', s \models \varphi$ for some h' equal for i to h up to stage s.

Now, as moves are played publicly, players receive 'public announcements' of these:

Fact 8 The following valid principle is the temporal equivalent of the key *DEL* reduction axiom for public announcement: $F_a\langle\rangle\varphi \leftrightarrow (F_aT\ \&\ \langle\rangle F_a\varphi)$.

As in the earlier modal setting, commutation of a temporal and an epistemic operator implies a form of *Perfect Recall*: agents' present uncertainties are always inherited from past ones.

Adding beliefs Next, if players i also have beliefs about the course of the game, we add binary relations $\leq_i$ of *relative plausibility* between histories, and we add a doxastic modality:

(d) $\boldsymbol{M}, h, s \models \langle B, i\rangle\varphi$ iff $\boldsymbol{M}, h', s \models \varphi$ for some history h' coinciding with h up to stage s and most plausible for i according to the given relation $\leq_i$.

Beliefs may change gently or drastically, and this matters to how our models should behave. Let B_s be the most plausible histories for i at $h, s,$ while s' is some stage later than s. First suppose some histories in B_s agree with h' up to s'. Then *coherence* says the most plausible histories for i at h', s' are the *intersection* of B_t with all continuations of s'. But when some unexpected move a is played at s, the most plausible histories at $s^\frown\langle a\rangle$ may be wholly disjoint from those at s. Here are axioms for both scenarios, with the second one involving an operator of conditional belief, which is needed for present 'pre-encoding' of later plausibility changes that may take place (van Benthem 2007A): [6]

Fact 9 The following temporal principles are valid for belief revision along a game tree:

$$\langle B, i\rangle F_a T \rightarrow (F_a\langle B, i\rangle\varphi \leftrightarrow (F_a T \ \& \ \langle B, i\rangle F_a\varphi))$$
$$F_a\langle B, i\rangle\varphi \leftrightarrow (F_a T \ \& \ \langle B, i\rangle(F_a T, F_a\varphi))$$

Richer models Uncertainty between histories is not sufficient for modeling 'higher' hypotheses about the future, e.g., about players' strategies ('am I playing a simple automaton, or a sophisticated learner?'). To model these, one needs full-fledged epistemic *game models* with worlds including whole strategy profiles. [7]

But more to the point here, there is the issue of what strategies players follow on the branching temporal 'playgrounds'. The extended version of van Benthem & Pacuit 2006 has a 'logic of protocols' for this purpose, but it is not yet a definitive answer to the next

Problem 3 Develop doxastic-epistemic temporal logics of explicit strategies. [8]

6 Public announcement and changing games

Next, we consider other questions about logics of strategies, moving away from fixed games represented in single models $\boldsymbol{M}$. What about dynamic settings where games can *change*?

---

[6] Similar principles have been rediscovered by Bonanno 2007 which formalizes AGM theory.

[7] For further variations on the above scenarios, including ways of 'localizing' global temporal update to DEL-style scenarios, or vice versa, ways of replacing step-by-step local update by public update on richer game models, cf. van Benthem 2004, van Benthem & Pacuit 2007.

[8] For an epistemic temporal logic with explicit strategies, cf. van Benthem & Pacuit 2006.

Promises and intentions Following van Benthem 2007D, one can break the impasse of a bad Backward Induction solution by changing the game through making suitable *promises*. E.g., in the following game, the bad equilibrium *(1, 0)* can be avoided by ***E***'s promise that she will not go left – and the new equilibrium *(99, 99)* results, making both players better off by restricting the freedom of one of them!

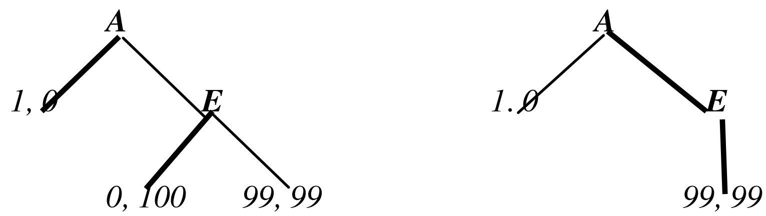

But one can also add new moves to a game. [9] Van Otterloo 2005 has a logic of strategic enforceability plus preference, where models change by announcing players' intentions.

'Rational Dynamics' In the global setting of strategic game forms, van Benthem 2003 uses public announcement to analyze other solution procedures. Strategic games induce epistemic models ***M*** of strategy profiles with preferences and uncertainty relations for players who know their own strategy, but not that of the others. Then a combined modal-preference language can formulate statements of Weak Rationality ("no player chooses a move which she knows to be worse than some other available one") and Strong Rationality ("every player chooses a move which she thinks may be the best possible one"). When announced, these eliminate worlds, and iterating this, one finds a smallest sub-model where announcements of WR and SR have no further effect: *WR* or *SR* are now common knowledge.

Fact 10 The result of iterated announcement of *WR* is the usual solution concept of Iterated Removal of Strictly Dominated Strategies; and it is definable inside ***M*** by means of a formula of a modal μ–calculus with inflationary fixed-points. The same for iterated announcement of *SR* and game-theoretic Rationalizability. [10]

In this scenario of 'internal deliberation' players keep recalling their rationality. But one can announce many further types of statement. A similar analysis applies to extensive games:

Fact 11 The Backward Induction solution for extensive games is obtained through repeated announcement of the assertion "no player chooses a move all of whose further histories end worse than all histories after some other available move".

---

[9] Yes, one could code up all such changes beforehand in one grand initial 'Super Game', but that would lose all the flavour of understanding what happens in a stepwise manner.

[10] If the iterated assertion *A* has so-called 'existential-positive' syntactic form (for instance, *SR* does), then the relevant definition can even be formulated in the standard μ–calculus.

Here is how this works out for a Centipede game, with three turns as indicated, branches indicated by name, and pay-offs indicated for ***A***, ***E*** in that order:

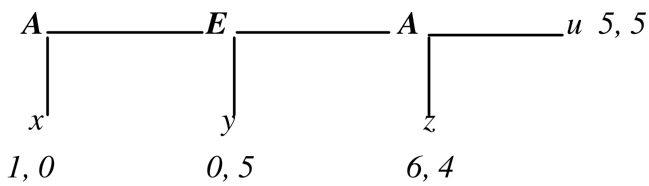

Stage 0 of the announcement procedure rules out branch *u*,
Stage 1 then rules out *z*, while Stage 2 finally rules out *y*.

Again, this iterated announcement procedure for extensive form games (or alternatively, for the temporal models of the preceding section) ends in largest sub-models in which players have common belief of rationality, or other relevant assertions that have been made.

Problem 4 Give an epistemic-doxastic temporal preference language where the final sub-models can be defined. Does it need fixed-point operators, as for strategic games?

Again, there are many options here than the ruthless egotism of Backward Induction. Van Benthem 2003, 2004 discuss history-oriented announcements, where players steer their future actions by reminding themselves of the *legitimate rights of other players*, because of 'past favours received'. Likewise, current type space models for games (Brandenburger 2007) would allow for much greater freedom in making assumptions about other players.

But now for the strategies! How can we talk about explicit strategies in all of this? If we use *PDL* for strategies and moves in games, as suggested earlier, this leads to the obvious logic *PDL+PAL* adding public announcements *[!A]*. It is easy to show that *PDL* is closed under relativization to definable sub-models, both in its propositional and its program parts, using a recursive operation $\sigma|A$ for programs σ which surrounds every atomic move with tests *?A*.

Fact 12 *PDL+ PAL* is axiomatized by merging their separate laws while adding the following reduction axiom: $[!A]\{\sigma\}\varphi \leftrightarrow (A \rightarrow \{\sigma|A\}[!A]\varphi)$. [11] [12]

---

[11] For versions of *PDL* plus full *DEL*-style product update, cf. van Benthem, van Eijck & Kooi 2006, and the subsequent μ–calculus-based analysis in van Benthem & Ikegami 2007.

[12] *The wrong way around?* Maybe this reduction axiom misses the real issue. It explains what old plan I should have had in order to make some new plan work in the changed game model. But usually, I already have a plan σ for playing game *G* to obtain a certain effect φ. Now *G changes to G':* my machine lost some functionality, my game got some extra moves, etc. How should I *revise that current plan* σ to get some intended effect φ in the new game *G'*? This may be as hard as still open model-theoretic problems like finding a good Łos-Tarski theorem syntactically characterizing those *PDL* formulas which are preserved under sub-models.

But of course, we also want to know about versions with epistemic preference languages – and hence there are many further questions following up on this initial observation.

7 A more global level: logics of powers

Here is another approach to our topic: first making us see what we are missing. One can also see what good strategies do by first considering a game setting where they are conspicuously *lacking* – and then trying to introduce them after all!

Powers over outcomes A more global description of games uses a player *i*'s *powers*, being the sets of outcomes *X* for which *i* has a strategy forcing the game to end inside *X*, whatever the other players do. This is the level of the initial forcing modality in Section 1, and it supports its own style of modal logic over modal 'neighbourhood models' (van Benthem & Blackburn 2006), including its own bisimulations, etc. Here is also where Parikh's *dynamic game logic* lives (*DGL*, Parikh & Pauly 2003). Its models ***M*** encode forcing relations $\rho^{G,\,i}\, sX$ for players *i* in atomic games *G*, played from some initial sate *s*. [13]This notion is then lifted inductively to the right powers for both players in arbitrary games *G* formed using further operations of dual, choice, sequential composition, and iteration. One sets

$$\boldsymbol{M}, s \models \{G, i\}\varphi \text{ iff } \exists X: \rho^{G,\,i}\, sX \;\&\; \forall x \in X: \boldsymbol{M}, x \models \varphi$$

The axioms for this logic reflect some basic properties of powers, with two players ***A***, ***E***:

C1	if $\rho_{\boldsymbol{G}}{}^{\boldsymbol{i}}\, sY$ and $Y \subseteq Z$, then $\rho_{\boldsymbol{G}}{}^{\boldsymbol{i}}\, sZ$		*Monotonicity*
C2	if $\rho_{\boldsymbol{G}}{}^{\boldsymbol{A}}\, sY$ and $\rho_{\boldsymbol{G}}{}^{\boldsymbol{E}}\, sZ$, then Y, Z overlap		*Consistency*

In the usual version od dynamic game logic version, they also satisfied *Determinacy*:

C3 if not $\rho_{\boldsymbol{G}}{}^{\boldsymbol{A}}\, sY$, then $\rho_{\boldsymbol{G}}{}^{\boldsymbol{E}}\, sS{-}Y$; and the same for ***E*** versus ***A***

Conversely, these three conditions are also *all* that must hold:

Fact 13 Families F_1, F_2 of subsets of some set *S* satisfying *C1, C2, C3* are the powers for ***A***, ***E*** at the root of some determined two-step game.

The result also extends to games with imperfect information (van Benthem 2001):

Fact 14 Families F_1, F_2 of subsets of some set *S* satisfying just *C1, C2* are powers for ***A***, ***E*** at the root of some two-step game with imperfect information.

The power analysis is natural, and it fits with the 'Thompson Transformations' in game theory (Osborne & Rubinstein 1994), which say when two games of imperfect information

---

[13] It is crucial that the same game can be played from different starting states on board models.

are the same (for players having Perfect Recall). It is also the natural level at which logicians view 'logic games' of evaluation, model comparison, or proof (van Benthem 2007C).

For concreteness, we conclude with some crucial compositional axioms of *DGL* (cf. Pauly 2001 for details on the model theory and proof theory of this system). These concern the sequential operations of game composition ***;***, choice $\square$ for the distinguished player, and game dual $^{\boldsymbol{d}}$ (role switch between the two players), respectively:

Fact 15 The following principles are valid in modal forcing semantics:

$$\{G_1 ; G_2\}\square\square\ \{G_1\}\{G_2\}\square \qquad \{G_1\square\, G_2\}\square\square\ \{G_1\}\square\ \ \{G_2\}\square$$

$$\{G^d\}A\ \square\ \neg\{G\}\neg A \qquad \{?P\}\square\ \square\ (P\ \square\ \square)$$

The soundness arguments for these axioms involve some basic reasoning with strategies. Pauly 2001 points out that this operational view gets closer to describing extensive games, but let us now address that issue more directly. Let us formalize the strategic reasoning!

8 'Strategizing' logics of powers

DGL 'dynamifies' the compositional structure of the games – at least as far as sequential operations are concerned (parallel game composition is another story: see below). But one can even try to 'do a *double PDL*' here, by also bringing in a description of strategies:

Problem 5 Add strategies to *DGL*, with forcing modalities $\{\square, i, G\}\square$ meaning: "strategy $\square$ for player *i* forces outcome $\square$ in game *G*"; and determine the logic.

This issue is significant, and there are grounds for optimism. In particular, we already know that making witness objects for existential quantifiers explicit can be done with beneficial effects in the study of proof and computation. Artemov 1998's 'logic of proofs' did this for modal provability logic, replacing boxes *[]* for 'provability' by boxes *[t]* containing terms *t* for proofs or pieces of evidence – and it has been extended to deal with epistemic evidence in a more general sense since by the 'New York School'. [14]

But strategizing *DGL* also involves some tricky issues about the appropriate setting. In some sense, the issue seems crystal-clear. The reasoning behind the validity of the *DGL*-axioms is strategy-laden, and here is a telling illustration. Viewing a strategy as something which we can unpack into a 'head' (the first move to be played) and a 'tail' (the rest), we have that

$$\{\square, i, G\square H\}\square\square\ \ \{tail(\square), i, G\}\square\ \ \{tail(\square), i, H\}\square$$ [15]

---

[14] Renne 2006 proposes explicit strategies here, linking up with dynamic epistemic logic.

[15] Incidentally, there is an interesting sense of directionality here. *DGL* analyzes strategies *top-down* into sub-strategies in component games. But what about matching *bottom-up* principles for strategy construction – such as: $\{\square, i, G\}\square\square\ \ \{\langle LEFT; \square\rangle, i, G\square H\}\square$?

What is the right semantic setting to make precise sense of this? *DGL* expressions are interpreted on abstract *board models*, while strategies seem to live inside concrete games. Van Benthem 1999 (Chapter 5) proposed *merging boards and games* for this purpose – but no standard way of doing this has emerged so far. Also, Parikh's games are generic: they can start at arbitrary board states, so what about generic strategies? For proper balance, a solution to Problem 5 may even have to work with *two languages*, one game-external over board models, and one-game internal, referring to turns, moves, and other procedural features.[16]

Moving on, there are further natural game constructions involving strategies. *DGL* is only about sequential game operations. What about concurrency, and operations for *parallel* play – surely, a natural scenario in practice? Van Benthem, Ghosh & Liu 2007 propose a *concurrent dynamic game logic CDGL*, bringing in ideas from concurrent dynamic logic.[17] Its crucial axiom is this simple reduction for game products $G_1 \times G_2$ where players pay simultaneously without any intermediate communication:

$$\{G_1 \times G_2, i\}\square\square \ \{G_1, i\}\square\square\{G_2, i\}\square$$

But again, explicit strategies are missing here…

Problem 5, continued Strategize concurrent dynamic game logic.

Concurrent game semantics also occurs in the area of linear logic, which also considers product operations that do allow for communication between subgames: see Section 11.

Finally, opening up existential modalities of abilities for agents also seems natural in temporal logics. Consider the logic of agency *STIT* (Belnap et al. 2001), which describes agents' powers for 'seeing to it that' certain states of the world are realized without mentioning explicit actions. The intuition underlying *STIT* is that of some agents acting simultaneously, knowing only their own action, but not that of the others – something reminiscent of games with a mild form of imperfect information. This seems close to our earlier framework of epistemic *PDL*, perhaps enriched with an intersection operation for parallel action.

Problem 6 Do an explicit action/strategy-based version of *STIT*,
analyzing its time steps as parallel actions for a bunch of players.

I conclude with a few brief illustrations of further directions.

9 *Digression A* Coalition logics and group strategies

Most game logics deal with two players, often for convenience, sometimes since they must. Now consider the logic of agents' powers in Pauly 2001. Here $\{G\}\square$ means that the agents

[16] Some concrete proposals are found in van Benthem, Dégrémont, Ghosh & Liu 2007.

[17] *CDGL* works with *non-determined* games, as parallelism leaves the realm of determinacy.

in the group G can achieve a set of outcomes all of which satisfy the proposition □. In such a language we can talk about group powers, coalitions in games, and so on – and interesting principles emerge. For instance, if the groups G, H are disjoint, the following principle describes their cooperation:

{G}□ & {G}□ □ {G□H} □&□

No explicit actions or strategies are provided. But what we really want to say here is this:

Fact 17 The principle *{G, □}□ & {H, □}□ □ {G□H, □#□} □&□* is valid with □#□ intersection of strategies (when these are viewed as relations).

There are many further interesting laws to be had here. For instance, one can also state laws to the effect that powers of a coalition must be combinations of powers for subgroups.

Problem 7 Give an explicit action/strategy version of Coalition Logic.

Coalition logic becomes more powerful when explicit preferences of the players are added in the form of suitable modal operators. Based on joint work with Michael Wooldridge, van der Hoek 2007 shows what significant collective scenarios can be analyzed then.

10 *Digression B* Infinite games, and linear logic

From finite to infinite games Many of the motivating examples in game logics seem finite. But infinite games, too, can be analyzed with explicit logics of strategies, taking whole histories as outcomes. But then we seem to need temporal logics as in several earlier sections – and some of the points in Section 5 apply. I do not have a crisp set of open problems here; but many of the above ones seem to still make sense in this setting. Here is one illustration.

Following the methodology that game logics should be able to formalize basic game-theoretic arguments, one obvious candidate beyond Zermelo's Theorem (Section 2) is the *Gale-Stewart Theorem*. It says that all infinite two-player games of perfect information, viewed as topological trees, where the winning condition for one of the players is an open set, are determined. Now the key point in the proof of this result is the following universally valid temporal principle of 'Weak Determinacy':

> "A player either has a winning strategy, or the other player *j*
> has a strategy making sure that the first player *i* never reaches
> a stage in the game where she has a winning strategy."

This cries out for a temporal logic describing these powers – and here is a natural candidate, using our earlier forcing modalities, plus a temporal modality *G* ('always in the future'). Weak Determinacy then becomes the following formal observation:

Fact 18 *{i}□ {j} G ¬ {i}□* is a valid law for infinite games.

Problem 8 Find a complete temporal logic of players' powers over infinite game trees. And 'strategize' it using some version of *PDL*, or other logics of protocols.

11 *Digression C* Linear logic

Game semantics for linear logic Finally, there is another area of game logic where strategies are crucial, viz. the game semantics of Abramsky 1997 for linear logic and concurrency in general. Again, it is striking that reasoning about strategies is crucial to understanding the soundness of the axioms of linear logic. The protagonist is an evergreen from game theory, viz. 'Copy Cat' (or 'Tit for Tat') copying moves from one parallel game to another. But again, important as they are, these strategies are not explicitly represented in the formal language of linear logic, which merely describes generic game forms. For example, a basic formula of linear logic like $A + A^d$, with + parallel disjunction, is interpreted as follows:

In each concrete game of this form, the distinguished player ***P*** has a winning strategy.

This is even less informative than the global power formulas of Parikh's *DGL*, which were able to talk about different propositions being true after the game. And so we ask:

Problem 10 Develop a logic of explicit strategies which can formalize the soundness arguments for linear logic, and make further statements about the course of games. [18]

Strategies and proofs This illustration also raises a general issue, which surfaced briefly before, in our discussion of Artemov's Logic of Proofs. Linear logic arose originally from mathematical proof theory. And, as has been known ever since Lorenzen's dialogue games (cf. van Benthem 2007C), a proof is like a strategy for winning an interactive debate, or performing some computation. This analogy is well-known in proof theory and its category-theoretic versions. Thus, versions of proof theory or type theory with explicit proof terms *are* an explicit calculus of strategies! But in this paper, we have mainly followed a model-theoretic take on games and strategies, as definable subrelations, leading to modal logics of actions and processes, and strategies in the style of dynamic logic. How these broad styles of thinking, proof-theoretic and model-theoretic, are related in general remains a vexed issue. [19]

12 Conclusion

Explicit logics of strategies are feasible, useful, and fun to explore. *Let's strategize!*

PS Something to think about. Strategies in games can be made explicit using formulas in logical languages. Logical formulas can be interpreted in evaluation games. Their truth then amounts to stating that 'there exists' a winning strategy for the Verifier. And so on...

[18] I plan to work on this with Samson Abramsky, tying in 'linear' and 'modal' traditions.

[19] In another guise, this is the fundamental issue of semantic versus deductive views of information in logic. Cf. van Benthem & Martinez 2007 for a state of the art.

Acknowledgment I thank the students of our 2007 ILLC graduate seminar in logic of games for their responses, as well as Rineke Verbrugge for her editorial comments.

References

S. Abramsky, 1997, ‘Semantics of Interaction: an Introduction to Game Semantics’, in *Proceedings of the 1996 CLiCS Summer School*, Isaac Newton Institute, P. Dybjer and A. Pitts, eds. (Cambridge University Press) 1997, 1 – 31.

S. Artemov. 1998, `Logic of Proofs: a unified semantics for modality and lambda-terms,' Technical Report CFIS 98-06, Cornell University.

N. Belnap, M. Perloff, & M. Xu, 2001, *Facing the Future: Agents and Choices in Our Indeterminist World*, Oxford University Press, Oxford.

J. van Benthem, 1999 – 2003, *Logic in Games*, lecture notes, Institute for Logic, Language and Computation, University of Amsterdam.

2001, 'Games in Dynamic Epistemic Logic', *Bulletin of Economic Research* 53:4, 219–248.

2002, ‘Extensive Games as Process Models’, *Journal of Logic, Language and Information* 11, 289–313.

2003, 'Rational Dynamics', Department of Political Economy, University of Siena. In *International Game Theory Review* 9:1, 2007, 13 – 45. Erratum reprint, Volume 9:2, 377 – 409.

2004, ‘Update and Revision in Games’, manuscript, ILLC, University of Amsterdam.

2005A, ‘A Fixed-Point Look at Update Logic’, manuscript, ILLC, University of Amsterdam.

2005B, 'Open Problems in Logic and Games', in S. Artemov, H. Barringer, A. d'Avila Garcez, L. Lamb & J. Woods, eds., *Essays in Honour of Dov Gabbay*, King's College Publications, London, 229–264.

2006, 'Computation as Conversation', *Computation in Europe*, to appear.

2007A, ‘Dynamic Logic of Belief Revision’, *Journal of Applied Non-Classical Logics*, to appear.

2007B, Interview, *Five Questions in Game Theory*, V. Hendricks, ed., Automated Press.

2007C, 'Logic Games: from tools to models of interaction', in J. van Benthem, A. Gupta & R. Parikh, eds., *Proceedings First Winter School on Logic and its Relations to Other Disciplines*, IIT Bombay, Mumbai.

2007D, 'Rationalizations and Promises in Games', *Philosophical Trends*, special issue.

J. van Benthem & P. Blackburn, 2006, ‘Modal logic, a Semantic Perspective’, in P. Blackburn, J. van Benthem & F. Wolter, eds., *Handbook of Modal Logic*, Elsevier Science Publishers, Amsterdam, 1 – 84.

J. van Benthem, C. Dégrémont, S. Ghosh & F. Liu, 2007, ‘Straegizing Game Logics’, manuscript, ILLC, University of Amsterdam.

J. van Benthem, J. van Eijck & B. Kooi, 2006, ‘Logics of Communication and Change’. *Information and Computation* 204(11), 1620 – 1662

J. van Benthem, J. Gerbrandy & E. Pacuit, 2007, ‘Merging Frameworks for Interaction: DEL and ETL’, *Proceedings TARK 2007*, Namur.

J. van Benthem, P. Girard & O. Roy, 2007, ‘A Logic for Ceteris Paribus Preferences’, ILLC, University of Amsterdam.

J. van Benthem & D. Ikegami, 2007, ‘Modal Fixed-Point Logic and Changing Models’, manuscript, ILLC, University of Amsterdam.

J. van Benthem & F. Liu, 2004, ‘Diversity of Logical Agents in Games’, *Philosophia Scientiae 8:2,* 163 – 178.

J. van Benthem & F. Liu, 2007, ‘Dynamic Logic of Preference Upgrade’, *Journal of Applied Non-Classical Logics*, to appear.

J. van Benthem, S. Ghosh & F. Liu, 2007, ‘Modelling Simultaneous Games with Concurrent Dynamic Logic’, ILLC Preprint 2007–26, University of Amsterdam. Also in J. van Benthem, S. Ju, & F. Veltman, eds., *A Meeting of the Minds*, Proceedings LORI Beijing 2007, Colle Publications, London.

J. van Benthem & M. Martinez, 2007, ‘Logic and Information’, in P. Adriaans & J. van Benthem, eds., *Handbook of the Philosophy of Information*, Elsevier, Amsterdam, to appear.

J. van Benthem, S. van Otterloo & O. Roy, 2006, ‘Preference Logic, Conditionals, and Solution Concepts in Games’, in H. Lagerlund, S. Lindström & R. Sliwinski, eds., *Modality Matters*, University of Uppsala, 61 - 76.

J. van Benthem & E. Pacuit, 2006, 'The Tree of Knowledge in Action', *Proceedings AiML IV,* Melbourne.

J. van Benthem & E. Pacuit, 2007, ‘Modeling Protocols and Transforming Temporal Models’, manuscript, ILLC, University of Amsterdam,

G. Bonanno, 2007, ‘Axiomatic Characterization of the AGM Theory of Belief Revision in a Temporal Logic’, *Artificial Intelligence* 171: 2-3, 144-160

J. Bradfield & C. Stirling, 2006, ‘Modal □–Calculi’, in P. Blackburn, J. van Benthem & F. Wolter, eds., *Handbook of Modal Logic*, Elsevier Science Publishers, Amsterdam, 721 – 756.

A. Brandenburger, 2007, *Tutorial on Game Theory*, First Amsterdam–Lausanne –London Graduate Workshop, London School of Economics.

B. de Bruin, 2004, *Explaining Games*, Dissertation 2004–03, ILLC, University of Amsterdam.

P. Gochet, 2006, ‘Un Problème Ouvert en Epistémologie: la formalisation du savoir-faire’, Institut de Philosophie, University de Liege.

P. Harrenstein, 2004, *Logic in Conflict*, Dissertation, Department of Computer Science, Utrecht University.

W. van der Hoek, 2007, 'Extended Modal Logic for Social Software', invited lecture, LORI Workshop, Beijing.

W. van der Hoek, S. van Otterloo & M. Wooldridge, 2005, 'A Logic for Strategic Reasoning', *Proceedings of the fourth international joint conference on Autonomous Agents and Multiagent Systems AAMAS*, Utrecht, 157 – 164.

R. van der Meyden, 1996, 'The Dynamic Logic of Permission', *Journal of Logic and Computation* 6:3, 465 - 479.

M. Osborne & A. Rubinstein, 1994, *A Course in Game Theory*, The MIT Press, Cambridge, Mass.

S. van Otterloo, 2005, *A Security Analysis of Multi-Agent Protocols*, Dissertation, ILLC, University of Amsterdam, and Department of Computing, University of Liverpool.

R. Parikh & M. Pauly, 2003, 'Game Logic - An Overview', *Studia Logica* 75(2): 165-182.

M. Pauly, 2001, *Logic for Social Software*, Dissertation, ILLC, University of Amsterdam.

B. Renne, 2006, 'Propositional Games with Explicit Strategies', in G. Mints and R. de Queiroz (Eds.), *Proceedings of the 13th Workshop on Logic, Language, and*(WoLLIC), Stanford, CA.

B. Rodenhauser, 2001, *Updating Epistemic Uncertainty*, Master Thesis MoL-2001-07, ILLC, University of Amsterdam.