Reasoning About Strategies

Johan van Benthem[1]

University of Amsterdam
and Stanford University

Abstract. Samson Abramsky has placed landmarks in the world of logic and games that I have long admired. In this little piece, I discuss one theme in the overlap of our interests, namely, logical systems for reasoning with strategies - in gentle exploratory mode. [1]

1 Reasoning about strategies, a priori analysis or rather logical fieldwork?

The notion of a strategy as a plan for interactive behavior is of crucial importance at the interface of logic and games. Truth or validity of formulas corresponds to existence of appropriate strategies in systems of game semantics, and in game theory, it is strategies that describe multi-agent behavior interlocked in equilibria. But strategies themselves are often implicit in logical systems, remaining "unsung heroes" in the meta-language (5). To put them at centre stage, two approaches suggest themselves. One is to assimilate strategies with existing objects whose theory we know, such as proofs or programs. This is the main line in my new book (6). However, one can also drop all preconceptions and follow a "quasi-empirical approach". A traditional core business of logic is analyzing a given reasoning practice to find striking patterns, as has happened with great success in constructive mathematics or in formal semantics of natural language. In this piece, I will analyze a few set pieces of strategic reasoning in basic results about games, and just see where they lead. I restrict attention to two-player games (players will be called ***i, j***), and usually, games of winning and losing only. Also, given the limitations of size for this paper, I will just presuppose many standard notions.

2 The Gale-Stewart theorem and its underlying temporal logic of forcing

Two basic theorems Consider determined games, where one of the players has a winning strategy. This is the area where basic mathematical results about games and strategies started:

[1] I thank the two readers of this paper, and also Chanjuan Liu and Prakash Panangaden for their generous practical help.

Theorem 1 (Zermelo's Theorem) *Games with finite depth are determined.*

Proof . The proof is essentially an algorithm computing positions where players have winning strategies, a precursor to the game-theoretic method of Backward Induction (16). Its key recursion defines predicates WIN_i ("player $\boldsymbol{i}$ has a winning strategy from now on") at nodes of the game tree in terms of auxiliary predicates $\boldsymbol{end}$ ("endpoint"), $\boldsymbol{turn}_i$ ("it is player $\boldsymbol{i}$'s turn to move"), $move_i$ ("the union of all currently available moves for $\boldsymbol{i}$"), and $\boldsymbol{win}_i$ ("player $\boldsymbol{i}$ wins at this node"):

$$WIN_i \leftrightarrow ((\boldsymbol{end} \wedge \boldsymbol{win}_i) \vee (\boldsymbol{turn}_i \wedge \langle move_i \rangle WIN_i) \vee (\boldsymbol{turn}_j \wedge [move_i] WIN_i))$$

∎

Notice the different existential and universal modalities in the two cases.[2]

Now we move to infinite games. An *open winning condition* is a set X of histories h with $h \in X$ iff some initial segment of h has all its extensions in X. Call a game "open" where at least one of the players has such a winning condition. Here is another classical result:

Theorem 2 (Gale-Stewart Theorem) *Open infinite games are determined.*

Proof . The proof revolves around this property of all infinite games:

Weak Determinacy
Either player $\boldsymbol{i}$ has a winning strategy, or player $\boldsymbol{j}$ has a strategy ensuring that player $\boldsymbol{i}$ never reaches a position in the game where $\boldsymbol{i}$ has a winning strategy.

If $\boldsymbol{i}$ has no winning strategy, then $\boldsymbol{j}$ has a "nuisance strategy" by Zermelo reasoning. At $\boldsymbol{i}$'s turns, no move for her can guarantee a win, and so $\boldsymbol{j}$ can "wait and see". If $\boldsymbol{j}$ is to move, there must be at least one successor state where $\boldsymbol{i}$ has no winning strategy: otherwise, $\boldsymbol{i}$ has a winning strategy after all. Continuing this way, $\boldsymbol{j}$ produces runs as described.

Next, without loss of generality, let $\boldsymbol{i}$ be the player with the open winning condition. Then the nuisance strategy is winning for $\boldsymbol{j}$. Consider any history r that it produces. If r were winning for $\boldsymbol{i}$, some initial segment $r(n)$ would have all its continuations winning. But then "play whatever" would be a winning strategy for $\boldsymbol{i}$ at $r(n)$: quod non. ∎

A temporal logic of forcing powers Now we introduce some minimal machinery formalizing these arguments. Extensive games may be viewed as branching tree models $\boldsymbol{M}$ for time, with histories as complete branches h, and stages s as points on these histories:

[2] A correctness proof for the algorithm is essentially "excluded middle writ-large": either player $\boldsymbol{i}$ has a response to every move by $\boldsymbol{j}$ yielding φ, or player $\boldsymbol{j}$ has a move such that each follow-up by $\boldsymbol{i}$ yields $\neg\varphi$.

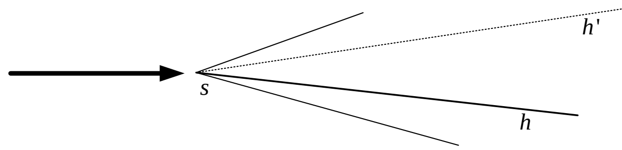

The bold-face line is the actual history, only known up to stage s so far. Points can have local properties encoded, while total histories can also have global properties such as Gale-Stewart winning conditions, or the total discounted pay-offs used in evolutionary games.

Such structures, assuming discrete time, interpret a standard branching temporal language ((10) has a survey of flavours), in the format

$\mathbf{M}, h, s \models \varphi$ formula φ is true at stage s on history h

with formulas φ constructed using proposition letters, Boolean connectives, existential and universal temporal operators *F, G, H, P, O* (future and past on branches, with O for "at the next moment"), as well as existential and universal modalities $\Diamond$, $\Box$ over all branches at the current stage. Here are the truth conditions for some major operators:

$\mathbf{M}, h, s \models F\varphi$ iff $\mathbf{M}, h, t \models \varphi$ for some point $t \geq s$,

$\mathbf{M}, h, s \models O\varphi$ iff $\mathbf{M}, h, s+1 \models \varphi$ with $s+1$ the immediate successor of s on h,

$\mathbf{M}, h, s \models \Diamond\varphi$ iff $\mathbf{M}, h', s \models \varphi$ for some h' equal to h up to stage s.

To this description of the basic structure of the model, we now add a *strategic forcing modality* $\{\boldsymbol{i}\}\varphi$ describing the powers of player $\boldsymbol{i}$ at the current stage of the game:

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

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: any successful strategy may have to deviate from the current "road to perdition".

As an illustration of the perspicuity of this language, Weak Determinacy becomes the following simple formula:

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

Valid principles Some obvious laws of reasoning for the resulting *temporal forcing logic* are a combination of some well-known components:

Fact 3 *The following principles are valid in temporal forcing logic:*

(a) *the standard laws of* branching temporal logic*,*

(b) *the standard logic of a* monotonic neighborhood modality *for* $\{\boldsymbol{i}\}\varphi$, *plus one for its strongly modalized character:* $\{\boldsymbol{i}\}\varphi \rightarrow \Box\{\boldsymbol{i}\}\varphi$,

(c) *three more specifically game-oriented principles:*

$$(c_1)\ \{\boldsymbol{i}\}\varphi \leftrightarrow ((\boldsymbol{end} \wedge \varphi) \vee (\boldsymbol{turn}_i \wedge \Diamond O\{\boldsymbol{i}\}\varphi) \vee (\boldsymbol{turn}_j \wedge \Box O\{\boldsymbol{i}\}\varphi))$$

$$(c_2)\ \alpha \wedge \Box G((\boldsymbol{turn}_i \wedge \alpha) \rightarrow \Diamond O\alpha) \wedge ((\boldsymbol{turn}_j \wedge \alpha) \rightarrow \Box O\alpha))) \rightarrow \{\boldsymbol{i}\}\alpha$$

$$(c_3)\ (\{\boldsymbol{i}\}\varphi \wedge \{\boldsymbol{j}\}\}\varphi) \rightarrow \Diamond(\varphi \wedge \psi)$$

For the list of principles meant under (a), see (10). For those under (b), see (14). The first law of (c) is the fixed-point recursion in the Zermelo argument, and the second an introduction law reminiscent of the axiom for the universal iteration modality in propositional dynamic logic.[3] The third principle is a simple form of independence of strategy choices for the two players that occurs in many logics of simultaneous action.

Proving our basic results formally These laws allow us to derive our earlier results. Here are the essential steps in the proof of Weak Determinacy:

- $(\mathbf{turn_i} \wedge \neg\{\boldsymbol{i}\}\varphi) \rightarrow \Box O\neg\{\boldsymbol{i}\}\varphi$ from (c_1)
- $(\mathbf{turn_j} \wedge \neg\{\boldsymbol{i}\}\varphi) \rightarrow \Diamond O\neg\{\boldsymbol{i}\}\varphi$ from (c_1)
- $\neg\{\boldsymbol{i}\}\varphi \rightarrow \{\boldsymbol{j}\}\neg\{\boldsymbol{i}\}\varphi)$ from (c_2)

Now we can also derive the Gale-Stewart Theorem formally. Suppose that φ is an open condition, i.e.:

$$\varphi \rightarrow F\Box G\varphi$$

Then it is easy to derive formally that $\{\boldsymbol{j}\}\neg\{\boldsymbol{i}\}\varphi \rightarrow \{\boldsymbol{j}\}\neg\varphi$, and combined with Weak Determinacy, this makes the game determined:

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

Zermelo's Theorem follows as well, since "having an endpoint" is an open property of branches, satisfying the implication

$$F\boldsymbol{end} \rightarrow F\Box GF\boldsymbol{end}.$$

Temporal forcing logic Viewed as a system, temporal forcing logic on our tree models has some familiar laws:

Fact 4 *The modal* ***K4****-axiom* $\{\boldsymbol{i}\}\alpha \rightarrow \{\boldsymbol{i}\}\{\boldsymbol{i}\}\alpha$ *is valid in temporal forcing logic.*

[3] Note that the principle stated here is less strong than it may seem: to see this, just apply it to a global winning condition.

This is not so much the usual "introspection" for knowledge-like modalities, but a sort of "safety": following a winning strategy never takes one outside of the area where one has a winning strategy. But it is also interesting to look at non-validities of the system:

Example 1. Some informative non-validities:

(a) The modal T-axiom $\{\boldsymbol{i}\}\alpha \rightarrow \alpha$ fails since the current history need not be the one recommended by $\boldsymbol{i}$'s strategy forcing α. [4] (b) Also invalid is the implication $G\{\boldsymbol{i}\}\alpha \rightarrow F\alpha$, that might look plausible as a principle of eventual success. However, it fails anywhere on the infinite $\neg\alpha$ branch in the following model, viewed as a one-person game:

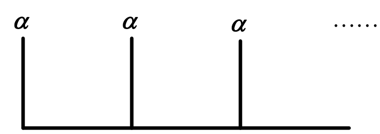

Even though we do not know a complete axiomatization for temporal forcing logic, we do have the wind in our sails:

Fact 5 *Temporal forcing logic is decidable.*

Proof . All temporal modalities, but also the forcing modality, can be defined in monadic second-order logic *MSOL* on trees with successor relations. Histories are maximal linearly ordered sets of nodes, and strategies can be identified with subsets of the tree as well, in a manner shown in (7). Then Rabin's Theorem on decidability of *MSOL* tree logic applies.[5] ∎

Remark A short piece like this cannot do justice to links with existing temporal logics for games. Classics such as (2) come to mind as obvious comparisons. (6) explores further connections between our forcing-based logic of strategies with various game-related systems in computational logic.

3 Nondeterminacy, strategy stealing, and temporal forcing logics of special games

Within our general logic of strategies, further properties come to light in special models. Going beyond the Gale-Stewart Theorem, consider a standard nondetermined game.

[4] But valid again in temporal forcing logic is the special instance $\{\boldsymbol{i}\}\{\boldsymbol{i}\}\alpha \rightarrow \{\boldsymbol{i}\}\alpha$.

[5] Many strategy-related modalities on trees are even bisimulation-invariant, so by the main theorem in (15), they are also definable in the modal μ-calculus.

Example 2. The 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 a sequence like this:

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

Player $\boldsymbol{i}$ wins if the union of all intervals chosen by her is in U- otherwise, $\boldsymbol{j}$ wins. Winning sets are not open, as sets in U are not determined by finite initial segments. This interval game is not determined, by a so-called "strategy stealing" argument:

Lemma 1. *Player $\boldsymbol{i}$ has no winning strategy.*

Proof . Suppose that player $\boldsymbol{i}$ had a winning strategy, then $\boldsymbol{j}$ could actually use it with a delay of one step to copy $\boldsymbol{i}$'s responses to her own moves, now 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, since U is free. Player $\boldsymbol{j}$ has no winning strategy for similar reasons. ∎

Analyzing this proof in detail reveals interesting logical structure. Let $\boldsymbol{i}$ start, the other case is similar. 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. Now crucially, in the interval game, 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. What this presupposes is the following special, but natural property of a game:

Composition Closure: Any player can play any concatenation of available successive moves as one single move.[6]

Now the game tree has the following property. The two stages described here start the same subgames in terms of available moves, but with all turn markings interchanged. Thus, one subgame is a "dual" of the other.[7] The core of $\boldsymbol{j}$'s strategy is now that he uses $\boldsymbol{i}$'s strategy in the other game to produce *identical runs* in both subgames, except for the inverted turn marking. This leads to a contradiction via the following logical *Copy Law*:

Fact 6 *In games with composition closure, the following formula is valid:*

$\{\boldsymbol{i}\}\varphi \to \Diamond OO\{\boldsymbol{j}\}\varphi_d$, *where φ_d is the formula φ with all turn occurrences for players $\boldsymbol{i}, \boldsymbol{j}$ interchanged.*

[6] One could define this property formally in a modal-temporal action language suitably extending our earlier formalism.

[7] This is not the standard game-theoretic dual, since we do not interchange winning conditions. See (6) for more discussion of different dualizations in games.

Many further questions make sense about powers of players in games with special structure, but here, we only conclude that both general and special temporal forcing logics have an interest of their own.[8]

4 Explicit logics of strategies as programs

Forcing modalities profess a general love for strategies without an interest in any specific one. We now go one step further in our logical analysis, introducing *terms that define strategies*, thus enabling us to reason explicitly about strategies themselves. A wide array of motivations for taking this step can be found in (5). Suitable languages can take various forms, but one obvious candidate is *propositional dynamic logic*.

Transition relations and programs Strategies are functions defined on players' turns, with typical instructions like "if she plays this, then I play that". Plans like this may allow more than one "best move", so general relations make sense as well, providing at least one move per turn. Thus, strategies are additional relations on a game tree that can be defined by programs. Since we need one-step actions only, normally, *flat programs* suffice using only atomic actions, tests, sequence; and choice $\cup$ - often just unions of guarded actions of the form

$$?\varphi; \alpha; ?\psi(9).^{9}$$

However, consecutive moves become important when we think of forcing outcomes. Using *PDL* programs, we now introduce a new forcing language with a key 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.

While this notion is natural, it still has an explicit definition in more familiar terms, viz. program modalities:

Fact 7 *For any game program expression σ, PDL can define $\{\sigma, \boldsymbol{i}\}\varphi$.*

Proof . The formula $[((?\mathbf{turn}_i; \sigma) \bigcup (?\mathbf{turn}_j; move_j))^*]\varphi$ is the required equivalent, as is easy to see from its truth conditions. [10] ∎

[8] Yet further questions would arise if we also introduce "intermediate" forcing modalities $\{\sigma\}^*\varphi$ saying that *partial* strategy σ guarantees reaching a *barrier* of intermediate positions in the game satisfying φ. This would connect with current modal logics of barriers and "cut-sets".

[9] It is easy to see that, on "expressive" finite game trees (each node is uniquely definable), each strategy is definable by our simplest flat *PDL* programs. But, if definitions are to be uniform across models, fixed-point languages are needed (7).

[10] In the same style, properties of the outcome of running joint strategies σ, τ, too, can be described in *PDL*.

Still, working with an explicit forcing modality $\{\sigma, i\}\varphi$ provides a natural notation for strategic behavior, and it fits well with actual examples of reasoning about games and interaction.

Remark *PDL* programs can even do a lot more, since they also model *partial strategies* that can be combined. See (4), (11) for recent work on on propositional dynamic logics of strategy combination, where the key operation is *intersection* of relations. Laws of such systems mix our earlier forcing modalities with program terms, as in the following implication:

$$(\{\sigma\}\varphi \wedge \{\tau\}\psi) \rightarrow \{\sigma \cap \tau\}(\varphi \wedge \psi)$$

Further benchmarks Our earlier "quasi-empirical" approach would now compile a repertoire of ubiquitous strategies, and formalize basic reasoning about their properties. We will not do so here. Also, *PDL* programs are geared toward finite termination, whereas we also want to look at natural non-terminating strategies such as "keep moving" – but we omit this extension as well.

5 Zoom, levels, invariants, and definability

Zooming in and out It now looks as if we have two competing approaches to logics of strategies, one with existentially quantified forcing modalities, and one with explicit program terms that define strategies. But in practice, both options are natural. The fact of the matter is that logic provides different levels of "zoom" on reasoning practices. Sometimes, we want to see underlying details, sometimes we want the broad picture. That is precisely why logical languages come in hierarchies of expressive power.

In the case of games, it may even be useful to *combine* our two formats. It might look as if explicit forcing modalities $\{\sigma, \boldsymbol{i}\}\varphi$ are just more informative than implicit $\{\boldsymbol{i}\}\varphi$. But this is misleading. If we want to say that a player *lacks a strategy* for achieving some purpose, then we need expressions $\neg\{\boldsymbol{i}\}\varphi$, and no natural explicit equivalent will do.

Even so, this combined language of forcing also has some surprises in store. Here is a "triviality result" saying that implicit can always become explicit by means of a strategy "be successful":

$$\sigma_{\varphi,i} = ?\mathbf{turn}_i; move_i, ?\{\boldsymbol{i}\}\varphi$$

Fact 8 *The following equivalence is valid:*

$$\{\boldsymbol{i}\}\varphi \leftrightarrow \{\sigma_{\varphi,i}, \boldsymbol{i}\}\varphi$$

The proof is easy and follows the earlier-mentioned valid recursion principles that govern temporal forcing.

Definitions for strategies Here is how we view the preceding observation. Most strategies have bite since they employ *restricted tests* on local assertions about the present or the past of the current node, but not about the *future* (like the above program did with the forward-looking test $?\{\boldsymbol{i}\}\varphi$).

This fine-structure suggests a study of *formats for definability* of strategies in temporal tree models beyond what we have done in the above with our simple *PDL* approach. Key strategies with great power are often defined by finite automata, with Samson's beloved Copy-Cat as a pet example. As a still more special case, *memory-free strategies* have turned out important in game semantics (1), in the field of logics, games and tree automata (13), and interestingly also, in the guise of *Tit-for-Tat*, in evolutionary game theory (16).

Two-level views and invariants Our view of what is going on here combines levels. Often we want two views together. *Games* have moves and internal properties, such as marking of nodes as turns or wins. But there is also an external *game board* recording observable or other relevant behavior. An example are the ubiquitous "graph games" of computational logic where the graph is the board (19). Usually, there is an obvious *reduction map* ρ sending game states to matching states on the game board satisfying a certain amount of back-and-forth simulation ((6) has many examples). Now, strategies in a game often consist in maintaining some *invariant* at the level of its board. Defining strategies then has to do with defining such invariants. In fact, the forcing modalities in the above triviality result may be seen as, somewhat bleak, invariants.

Excursion This perspective suggests interesting questions. One of the crucial results about graph games is the Positional Determinacy Theorem (12) saying that graph games with parity winning conditions are determined with positional strategies whose moves depend only on the graph component of the current game state of play. What this suggests is that the set of winning positions projects via the reduction map to a set of board positions that can be definable. A logical explanation of positional determinacy would then be the existence of a *translation* from modal forcing statements in the game to equivalent modal fixed-point assertions about associated graph states.

6 Strategy logics with operations on games

Finally, moving closer to Samson's trademark compositional methodology, we can go yet one step further in our formalizations. So far, we had forcing modalities $\{\boldsymbol{i}\}\varphi$, and when needed, we put in explicit terms for strategies $\{\sigma, \boldsymbol{i}\}\varphi$. But all this still takes place inside the setting of some game that is just given. However, it also makes sense to add *explicit descriptions of games* to the logical language, to obtain a notation, say,

$\{\sigma, \boldsymbol{i}, G\}\varphi$, with a game term G, saying that following the strategy σ forces φ-outcomes only for player $\boldsymbol{i}$ in game G.

Now we can reason about strategies in different games, and how they can be combined. There are in fact several logical systems in the literature that treat relevant operations on games that make sense here– such as choice, sequence, dual, and parallel composition.

Dynamic game logic One available line is the dynamic logic of games in (17) that extends our forcing modalities with game terms, where the formulas are now interpreted, not inside games, but on their associated game boards. The resulting system is a two-agent *PDL* on neighbourhood models, with typical decomposition axioms such as the one for "choice games":

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

whose validity can be established by an elementary soundness argument. Other axioms proceed on analogies with *PDL* as well, except that for the game dual.

Such soundness arguments provide nice material for the logical fieldwork of this paper, since we can tease out something that was left implicit in Parikh's notation: the underlying calculus of strategies.

Example 3. Strategizing power logic.

Consider the above axiom for choice. 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 $G \cup H$.

Under the surface, a general strategy calculus is at work here. Our first argument involved two operations: $head(\sigma)$ gives the first move of strategy σ, and $tail(\sigma)$ the remaining strategy, in a way that validates

$$\sigma = (head(\sigma), tail(\sigma))$$

The second part of the argument prefixed an action a to a given strategy σ, yielding $\alpha; \sigma$ satisfying obvious laws like

$$head(\alpha; \sigma) = \alpha, tail(\alpha; \sigma) = \sigma.$$

Dynamic game logic encodes a natural notion of game equivalence based on equal powers for players across games, and it has a literature of its own.[11]

Still, it is clear that the strategy calculus we just elicited does not look like our earlier *PDL* programs. The basic operations of *head* and *tail* rather suggests a

[11] The system has been extended to some kinds of parallel games in (8).

co-algebraic perspective of observing and then looking at the rest of the strategy. This brings us to another line in logics with explicit game terms, namely, the "game semantics" of Samson himself. It would be tedious to explain this extensive research program in a brief paper like this, and so I will just make a few points connecting with the above.

Linear game logic In this case, the logical formulas are just game terms, and systems of linear logic encode game equivalence or inclusion. There is no explicit forcing modality–though one might say that the precise notion of validity associates statements about winning powers with game terms. Still, game semantics takes place in the same temporal models that we have used so far, so it can be analyzed by earlier techniques. In particular, we could add a description language for what goes on inside Samson's games, with forcing modalities and names of specific strategies. I have ideas on how to do such a two-level logic, but these would transcend the boundary of this paper.

A concrete "quasi-empirical" challenge for such systems is similar to what we suggested for dynamic game logic. Begin with the absolute basics, look at the soundness arguments for linear logic in game semantics, and extract the minimum needed to make its reasoning about game constructions work. This reasoning will be more sophisticated than what we have considered before. In particular, parallel games involve "shadow arguments" (say for the soundness of the Cut Rule) about what can take place in subgames, and I am not sure how to represent these minimally.[12] [13]

7 Knowledge, preference, and game theory

Many topics in the above are reminiscent of real game theory. Strategy stealing proofs and copy-cat behavior are reminiscent of the central role in game theory for simple strategies like *Tit-for-Tat* in infinite evolutionary games (3). I end with mentioning just two points about new structure that should enter if we want to engage with real games.

Knowledge In the background of many arguments about strategies is what players *know*. I can hardly "copy" or "steal" a strategy if I do not know what it looks like. Now in many standard arguments for existence of strategies, the talk of knowledge is just didactical wrapping. But it is of interest to take it seriously, merging strategy logics with epistemic logics or other ways of representing information. Next on this road are *imperfect information games*, where

[12] More sophisticated arguments about "shadow matches", copying strategies in games, and representing parallel by sequential play, occur in the theory of graph games (19).

[13] But we could also start our fieldwork in this area with minimal logical specification calculi for effects of basic strategies, such as *Copy Cat*.

players need not know exactly where they are in the game tree. Such games, even when finite, are notoriously non-determined, and analyzing them might throw new light on game logics. Finally, strategies in this case will typically have knowledge-dependent instructions, and what also becomes essential is the informational nature of players: endowed with perfect memory, observation-driven, or yet otherwise. Even their beliefs and policies for belief revision become important in the usual foundations of strategic behavior in game theory. (6) explores this area in detail, but at the end of it all, an overall strategy calculus remains to be found.

Preference Another obvious feature of real games is the much more sophisticated dynamics of evaluation that drives behavior and mathematical equilibrium theory. The balance of available moves, beliefs, and *preference* is what drives *rational play* in the usual sense. Players can have any preferences between outcomes of a game (whether endpoints or infinite histories), and again this structure requires extending our logics of strategies. Issues this time include new notions of game equivalence, perhaps dependent on rationality types of players, but also just the analysis of basic game-theoretic arguments about solution methods. In particular, (6) has an extensive study of the typical algorithm of Backward Induction that already poses many challenges to the above. For one, while it does have a natural definition in the first-order fixed-point logic *LFP(FO)*, it does not seem to have an obvious program definition in the above *PDL* terms. For another, the current game-theoretic discussion between Backward Induction, a purely future-looking reasoning style, and "Forward Induction", a way of factoring in the past of the game so far (see (18)), seems to connect with choices in logical modeling at many points.

I believe that merging the best of computational logics of games and of game theory has a great future, but as will be amply clear, a lot remains to be done.

8 Conclusion

Logical analysis of strategic reasoning is a rich topic that unifies across the study of computation and social interaction. I have looked at a number of ways of pursuing this, in consecutive steps of explicitly defining forcing, strategies, and games. I believe that my interests in doing so are close to Samson's, but there is a caveat. Samson is a type theorist or category theorist at heart, while I am a model theorist. We may be looking at the very same things, and Samson sees a rabbit, while I see a deer. Proof theory versus model theory is a major divide in logic, but it is also a constructive case of complementarity, as has been shown again and again. This mixture of shared interests and different inclinations leads me to a conclusion whose phrasing I borrow from Immanuel Kant: I *can know* that Samson and I are allies, but I *may hope* that we are friends.

References

1. Abramsky, S.: Information, processes and games. In: van Benthem, J., Adriaans, P. (eds.) Handbook of the Philosophy of Information, pp. 483–549. Elsevier Science Publishers, Amsterdam (2008)
2. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. J. ACM 49(5), 672–713 (2002)
3. Axelrod, R.: The Evolution of Cooperation. Basic Books, New York (1984)
4. van Benthem, J.: Extensive games as process models. J. of Logic, Lang. and Inf. 11(3), 289–313 (2002)
5. van Benthem, J.: In praise of strategies. In: Eijck, J., Verbrugge, R. (eds.) Games, Actions and Social Software, pp. 96–116. Lecture notes in computer science 7010, Springer, Heidelberg (2012)
6. van Benthem, J.: Logic in Games. The MIT Press, Cambridge (2013)
7. van Benthem, J., Gheerbrant, A.: Game solution, epistemic dynamics and fixed-point logics. Fundam. Inf. 100(1-4), 19–41 (2010)
8. van Benthem, J., Ghosh, S., Liu, F.: Modelling simultaneous games in dynamic logic. Synthese 165(2), 247–268 (2008)
9. van Benthem, J., Liu, F.: Dynamic logic of preference upgrade. Journal of Applied Non-Classical Logics 17(2), 157–182 (2007)
10. van Benthem, J., Pacuit, E.: The tree of knowledge in action: Towards a common perspective. In: Advances in Modal Logic. pp. 87–106 (2006)
11. Eijck, J.v.: PDL as a multi-agent strategy logic. In: Schipper, B.C. (ed.) TARK 2013 – Theoretical Aspects of Reasoning About Knowledge, Proceedings of the 14th Conference – Chennai, India. pp. 206–215 (2013)
12. Emerson, E.A., Jutla, C.S.: Tree automata, mu-calculus and determinacy. In: Proceedings of the 32nd Annual Symposium on Foundations of Computer Science. pp. 368–377. IEEE Computer Society, Washington, DC, USA (1991)
13. Grädel, E., Thomas, W., Wilke, T. (eds.): Automata, Logics, and Infinite Games: A Guide to Current Research, Lecture Notes in Computer Science, vol. 2500. Springer, Heideberg (2002)
14. Hansen, H.H., Kupke, C., Pacuit, E.: Neighbourhood structures: Bisimilarity and basic model theory. Logical Methods in Computer Science 5(2), 1–38 (2009)
15. Janin, D., Walukiewicz, I.: On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic. In: Montanari, U., Sassone, V. (eds.) CONCUR. Lecture Notes in Computer Science, vol. 1119, pp. 263–277. Springer (1996)
16. Osborne, M., Rubinstein, A.: A Course in Game Theory. MIT Press, Cambridge (1994)
17. Parikh, R.: The logic of games and its applications. In: Annals of Discrete Mathematics. vol. 24, pp. 111–140. Elsevier (1985)
18. Perea, A.: Epistemic Game Theory. Cambridge University Press (2012)
19. Venema, Y.: Lectures on the modal mu-calculus. Tech. rep., Institute for Logic, Language and Computation, University of Amsterdam (2007)