The epistemic logic of *IF* games

Johan van Benthem, Amsterdam & Stanford

December 2002

To appear in the Hintikka Volume, *Library of Living Philosophers.*

Abstract

We analyze *IF*/hyper-classical games by bringing together two viewpoints from Jaakko Hintikka's work: game semantics, and epistemic logic. In the process, we link up between logic and game theory.

1 Logic meets games

Game theory and logic met in the 1950s – and Jaakko Hintikka has been a pioneer ever since in introducing game-theoretic viewpoints into logic, from his early basic evaluation games for predicate logic to his more recent ‘information-friendly’ logic based on extended games that go far beyond classical systems. The grand philosophical program behind these technical efforts is found in his books “Logic, Language Games and Information” (1973), “The Game of Language” (1985), the Handbook of Logic & Language chapter with Gabriel Sandu on ‘Game-Theoretical Semantics’ (1997), and many recent papers and manifestoes (cf. Hintikka 2002). Connections between logic and games are attracting attention these days, ranging from special-purpose ‘logic games’ to 'game logics' analyzing general game structure (cf. the general program in van Benthem 1999–2002). *IF* logic is intriguing in this respect, as it sits at the interface of ordinary logic games, whose players have perfect information about their position during play, and general game theory, where players may typically have imperfect information of various sorts. My aim in this paper is to explore the game content of Hintikka’s systems using tools from epistemic logic, and more generally, clarify their thrust at the interface of logic and game theory.

Exegetically, however, this is a somewhat tricky business. There is much less game content to Hintikka’s systems than one might expect. His true interest is closer to the classical logical agenda of meaning and expressive power, mainly for quantifier expressions, viz. the notion of *(in-)dependence.* Despite occasional declarations of love for games as such as the basis of rational enquiry, they remain mostly a didactic device for studying dependence in quantification – and a way of drawing battle-lines in that well-trodden war zone of compositionality. By contrast, I myself am a post-Hintikkean radical, whatever the original motivations. Games are important *per se* as models for action and information flow, and the interface of logic and game theory has a logical agenda of its own, which may make the classical one less urgent. So, admittedly, by taking the games too seriously in this essay, there is a grave risk of missing the point of Hintikka's work by pursuing a shallow dispensable metaphor. But I will cheerfully accept that stigma, provided – fair is fair – I can take the credit for all the pleasant new views that arise by setting off resolutely on my shallow path.

But what vistas can there be? One look at a game theory book shows that the field is driven by concerns far removed from Hintikka's evaluation games, or other logic games due to Lorenzen, Ehrenfeucht, and more recent authors such as Hodges, Blass, or Abramsky. Game theorists look at such issues as players' preferences, strategic equilibria, imperfect information, uncertainty and probability, bounded rationality, repeated behaviour, or the powers of coalitions. The intersection between logic and game theory may be as thin as just the shared notion of a *strategy*. Well, let us see. In this paper, I will first analyze Hintikka's original first-order games (Section 2), usually thought rather trivial, and uncover lots of general game-theoretic structure. Then I analyze the more mysterious *IF* games (Section 3) as imperfect information versions of the original games, using a mix of the game theoretic notions just found and *epistemic logic*. Both analyses broaden the bridge between logic games and general game theory, and show the contours of a new game logic. In Section 4, I discuss *IF* logic once more, but now from a general game-theoretic perspective. Finally, I state my conclusions and suggestions in Section 5. The tools for all this are two: (a) some unbiased reflection on the role of games in logic, and (b) the use of an explicit epistemic language of actions and knowledge. Both are things we have learnt from Jaakko Hintikka, and thus, the title of this essay has been explained.

2 First-order evaluation games

We start with the simplest games which Hintikka proposed back in the 1960s, taking off with some well-known facts, and becoming airborne in a few pages.

2.1 Evaluation games, truth and winning strategies

Let two parties disagree about a statement in some model ***M*** under discussion: *Verifier* ***V*** claims it is true, *Falsifier* ***F*** that it is false. Evaluation games describe their moves of defense and attack – with a schedule of turns driven by the statement:

atoms	*test* to determine who wins
disjunction $A \vee B$	***V*** *chooses* which disjunct to play
conjunction $A \wedge B$	***F*** *chooses* which conjunct to play
negation $\neg A$	*role switch* between the two players, play continues with respect to A
existential quantifiers $\exists x\, A(x)$	***V*** *picks an object d*, after which play continues with respect to $A(d)$
universal quantifiers $\forall x\, A(x)$	likewise, but now for ***F***

E.g., consider the first-order formula

$$\forall x\, \exists y\, x \neq y$$

on a model with two objects *s, t*. The game may be pictured as a tree of possible moves of object picking and fact testing, with the schedule read from top to bottom:

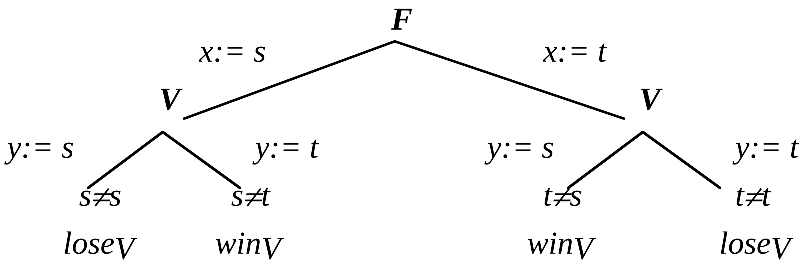

Falsifier starts, Verifier must respond. There are four possible runs of the game, with two wins for each player. Games like this are easy to play in class, and they sharpen the students' sense of first-order expressive power and model checking complexity. A bit more precisely (though one can go even further), think of the states as *pairs*

$$\langle s, \psi \rangle$$

where s is an assignment of objects in ***M*** to the variables in the original formula ϕ, and ψ is a subformula of ϕ. In particular, the game must start from some initial assignment, which can be modified by quantifier moves, and whose descendants eventually serve to identify the relevant atomic fact to be tested.

In the preceding game, players are not evenly matched. For, ***V*** has a *winning strategy*, a map from her turns to available moves that guarantees a winning outcome against every play by the opponent: she just needs to play the object different from the one picked by ***V***. This makes sense, as she has Truth on her side. This illustrates a general connection between evaluation games and standard first-order semantics:

<u>Proposition</u> The following two assertions are equivalent:

(a) Formula ϕ is *true* in model ***M*** under assignment *s*

(b) Verifier has a *winning strategy* for ϕ's evaluation game played in ***M*** starting from the initial state $<s, \phi>$

This equivalence seems at best a Pyrrhic victory for game-theoretical semantics. It says that the game-theoretic analysis amounts to a notion that we knew already. So, it yields nothing new, except for a pleasant didactic tool for feeding our students the Tarskian fare we had decided they should eat anyway. But the result has many interesting features, and it is worth-while to take our time, and think about these.

2.2 Exegetic intermezzo: the importance of strategies

Strategies First, the Proposition highlights the role of winning strategies, or generally, strategies. In particular, it suggests a new semantic notion. Verifier may have more than one winning strategy in the game for a given formula. E.g., for a disjunction with both disjuncts true, there are two winning strategies 'choose left', 'choose right'. (The number of strategies can be computed for any formula and model.) Thus, winning strategies are *a more fine-grained semantic object* than the usual denotations (truth values, predicates): say, patterns of verification, or reasons for truth. Classical model theory does not deal with these as such, unless in the auxiliary guise of Skolem functions, but they have lots of nice features.

In fact, as most logic games capture basic notions by winning strategies for some player (Proponent, Duplicator, etc.), a general *calculus of strategies* is a mechanism underlying much of logic. For instance, take a classically valid inference like

$$A\&(B\vee C) \quad \models \quad (A\&B)\vee C$$

At the finer-grained level of semantic reasons, this says that any winning strategy σ for Verifier in an $A\&(B\vee C)$–game can be transformed explicitly into one in an $(A\&B)\vee C$–game. In the latter, ***V*** makes the same choice at the start that σ prescribes in the premise game if ***F*** were to play “right” . After that she can sit back and wait…

Powers Strategies do not just serve to win. Any strategy gives a player a *power*, a certain control over the outcomes of the game, no matter what the other player does. The Proposition says that a winning power for ***V*** amounts to truth. But ***F*** may still have powers, too. E.g., in the above game, even though ***V*** can always win, it is up to ***F*** to decide where that winning takes place. This, too, may be a crucial feature of a game. Think of my having a strategy ensuring I will defeat you, but either in some boring meadow, or a picturesque location. If it is up to you to decide, you will go for the Last Stand at Thermopylae, as bards will sing about your defeat for centuries… To get its full impact, this story needs finer preferences for players than just zeroes and ones, but nothing prevents us from introducing these, and liven up logic games. But bare powers are of interest by themselves, and we will pursue them later on.

Games and boards Another striking feature of the Proposition is the juxtaposition of *two* relevant objects: an external *game board* – here, a model ***M*** plus all variable assignments over it – and a *game tree* with internal states for the game played over this board, generated by the formula ϕ. One board can accommodate many games. The Proposition says that some game-internal property, the existence of a winning strategy for player ***V***, *reduces* to an external first-order property of the game board.

Activities versus assertions But the Proposition contains one more juxtaposition! It distinguishes *games* as dynamic activities from *assertions* about games. This is just as in dynamic logics of programs, which have two kinds of expression on a par: terms denoting actions and formulas denoting propositions. In the present setting, the distinction is easy to overlook, since the same letter 'ϕ' denotes a game in clause (b) and a standard proposition in (a). In fact, most of the literature on game-theoretic semantics wavers on this issue, using 'ϕ' both for the game and the assertion that ***V*** has a winning strategy in it, or the assertion expressed by ϕ without any games at all. This may reflect the earlier point that people are not really interested in the games, but in their good old logical propositions. I will try to be explicit about the difference where it matters. Indeed, when all is said and done, the Proposition makes a plea for having *three* kinds of entity on a par: games, assertions, and strategies.

Even this discussion has just skimmed the surface of the Proposition! For a more elaborate analysis of Adequacy Theorems for logic games, cf. van Benthem 2002C.

2.3 Logical laws, players' powers, and game equivalence

The Proposition is a bridge between logic and game theory. Let's take a walk on it. For a start, logical laws now acquire game-theoretic import.

Determinacy Consider the classical law of *excluded middle* $A \vee \neg A$. That Verifier has a winning strategy for it in every model means she can choose to play either A as Verifier, or $\neg A$ as Verifier, i.e., A as Falsifier, and still have a winning strategy for the remainder. But this just expresses a well-known notion from game theory:

<u>Fact</u> All evaluation games are *determined*: one player has a winning strategy.

This is true for a very general game-theoretic reason:

<u>Theorem</u> (Zermelo 1913) All two-player games with perfect information that are *zero-sum* and have *finite branch depth* are determined.

Determinacy is important in descriptive set theory and foundations of mathematics: Zermelo's theorem started a long line of results on classes of determined games. Nevertheless, this first link between logic and game theory may be misleading. Not all games are determined, and excluded middle is not the most significant logical law from a game-theoretic viewpoint. We will do better in a moment.

Powers once more Determinacy emphasizes powers of one player only. Indeed, Hintikka's work has a bias towards Verifier. But a more general description of games must state what *both* players can achieve – especially in non-determined settings such as the *IF* games of Section 3. Here is a somewhat more formal definition.

<u>Definition</u> A player's *powers* in a game are all sets of outcomes X for which the player has a strategy in the game which ensures that all its outcomes, regardless of the opponent's moves, lie inside X.

Consider the following abstract version of our earlier game:

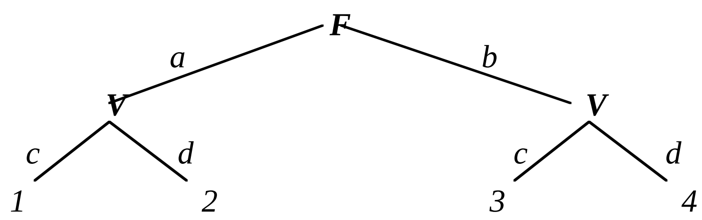

Here is the complete description of the power structure in this game.

> ***F*** has two strategies: 'left' with power *{1, 2}*, and 'right' with *{3, 4}*,
> ***V*** has four strategies: 'left, left' with power *{1, 3}*, 'left, right' with *{1, 4}*, 'right, left' with power *{2, 3}*, and 'right, right' with *{2, 4}*.

This tells us much more about the interaction encoded by the game. More generally, players' powers satisfy some general conditions which together are necessary and sufficient for representability in a determined game (van Benthem 2001A):

Monotonicity	If ***j*** has power X and $X \subseteq Y$, then ***j*** has power Y
Consistency	If ***V*** has power X and ***F*** has power Y, then X, Y overlap
Determinacy	If ***V (F)*** lacks power X, then ***F (V)*** has power $-X$

Now, here is the deeper connection with logical laws. Many of these have the form of equivalences. Now consider a propositional tautology like *distribution*:

$$p \wedge (q \vee r) \leftrightarrow (p \wedge q) \vee (p \wedge r)$$

Here are the two games corresponding to the formulas on the left and the right:

Fact Both players have the same powers in both games.

On the left, ***F*** has strategies 'left' and 'right' yielding powers *{p}*, *{q, r}*, while **V** has 'left', 'right' yielding *{p, q}*, *{p, r}*. On the right, ***V*** has two strategies yielding again *{p, q}*, *{p, r}*, while ***F*** has *four*, yielding *{p}*, *{p, r}*, *{q, p}*, *{q, r}*. But as supersets represent weaker powers, two are redundant, and ***F*** has really powers *{p}*, *{q, r}*. Power equivalence is an excellent notion of game equivalence overall, and we have:

<u>Theorem</u> All valid equivalences of predicate logic, with their formulas interpreted as evaluation games, give players equal powers on both sides.

From logic games to game logics There is much more to this style of analysis. One can design richer modal power languages (Parikh 1985, Pauly 2001) with operators *{G, **j**}*ϕ expressing that player ***j*** has the power to enforce proposition ϕ by the end of game G. Such a language for describing games expresses many further properties preserved under power equivalence. Cf. also van Benthem 2002A for links with modal languages in computational process theories. Actually, this is a momentous move – even though we will downplay it in this paper, to keep the focus on Hintikka games. Logic now plays two roles. We started with logic games: very specific games for analyzing logical formalisms. But now we also have game logics, formalisms that describe properties of games in general. And then the mill starts turning: there are also logic games for analyzing game logics, and so on: mind-boggling, but useful!

2.4 Compositionality and operations on games

Perhaps the most lively discussion concerning Hintikka games has been the issue of their *compositionality*. I, too, would love to write on this fascinating subject, but must honour the ten-year moratorium on the subject imposed at Amsterdam.

General operations in evaluation games Instead, let me point out a related, and equally interesting aspect of the above evaluation games, viz. the completely general *game-forming operations* embodied in them:

(a) offer a *choice* between two games G, H to one of the players:
a disjunction $\vee$ gave this to ***V***, and a conjunction **&** to ***F***

(b) negation *switches the roles* in G to get the dual game G^d

(c) *compose* two games $G;H$, playing one after the other

The latter operation occurs in a quantified formula like $\exists x\ Px$, where ***V*** first picks some object for x, and then an atomic test is played. Properly understood, first-order evaluation games are operational compounds of two sorts of semantic base game:

(i)	*object picking*	(single quantifiers)
(ii)	*fact testing*	(atomic formulas)

But there are other natural operations on games, which are less *sequential* and more *parallel*, such as playing two games interleaved ("having a family breakfast" while "reading one's paper"). The latter are more prominent in that other grand tradition of logic games, running from the pioneering work of Paul Lorenzen to modern game semantics for linear logic (Blass 1992, Abramsky 1996).

Game algebra Where there are operations, there must be algebra. The above set of {choices, switch, composition} support a natural abstract game algebra. Its criterion for validity of an identity $G=H$ is – as above – that, when interpreted on any game board, the two expressions G,H define games in which both players have the same powers. The earlier distribution law is generally valid in this sense:

$$p\wedge(q\vee r) \leftrightarrow (p\wedge q)\vee(p\wedge r)$$

Game-theoretically, it says that one can reverse the scheduling order of players without affecting their powers. Many other equivalence laws of first-order logic are game-valid, too. But game algebra also includes some further principles, such as the following laws for composition that go beyond first-order syntax:

$(G \vee H) ; K$	$=$	$(G ; K) \vee (H ; K)$	left-distribution
$(G ; H)^d$	$=$	$G^d ; H^d$	dualization

Typically non-valid, however, would be *right*-distribution

$$G ; (H \vee K) = (G ; H) \vee (G ; K)$$

To refute this, set $G = \forall x$, $H = Px$, $K = Qx$. Basic game algebra is decidable and axiomatizable, cf. Goranko 2000. Moreover, it tells us something new about first-order logic from a game-theoretic perspective. The corresponding set of valid equivalences $\phi\leftrightarrow\psi$ may be viewed as a new *decidable sublogic* of first-order logic. The above criterion of algebraic validity then amounts to the following:

> The equivalence between two formulas should hold no matter what formulas we substitute for their atomic predicates, and also no matter what quantifiers (or general game expressions) we substitute for their quantifier occurrences.

The latter clause explains why right-distribution fails, even though predicate logic validates $\exists x(Px\vee Qx) \leftrightarrow \exists xPx\vee\exists xQx$ for the special case of the existential quantifier. One more result of interest here is that each non-validity of general game algebra can be refuted by such predicate-logical equivalences. In that precise sense, logic games are complete for game logics (cf. the representation theorem in van Benthem 2002B).

2.5 Finer levels of game structure: extensive games and modal logic

Choosing an invariance What we have so far suffices for analyzing logic games as usually understood. But from a game-theoretic viewpoint, we have still missed an important issue. In any field, a crucial test on understanding its structures is asking when two presentations are the same. In the philosophers' terms, we need a *criterion of identity*. Now, the above power equivalence is one answer, but it seems rather coarse and global, disregarding details of players' turns and moves. Indeed, asking when two games are equivalent is an excellent test on one’s understanding of any game-semantics. So, let's go back once more to our distribution example:

In terms of powers, these two games were the same. In game theory, this corresponds to looking at *strategic forms* of games, which only care about input-output relations.

But game theory also studies *extensive games*, the full trees of what can happen (cf. Osborne & Rubinstein 1994 for this, and other game-theoretic points). And then, the two games have important differences of detail. Their scheduling of turns clearly differs, and also the intermediate powers. E.g., ***V*** might get a choice between *q* and *r* on the left, but this will never happen in the game on the right.

Modal logic Extensive games are like process graphs in computer science or Kripke models which can be studied using modal and dynamic logics. Typically, modal logic allows us to express the key difference between the two games:

$<>(<>q \ \& <>r)$ is true in the root on the left, but not on the right

Also, ***V***’s having a winning strategy in the game of Section 2.1 is expressed by a modal-dynamic formula with assignment actions and choices $\cup$ inside the boxes:

$$[x{:=}s \ \cup \ x{:=}t] <y{:=}s \ \cup y{:=}t> win_V$$

A more complex example is the earlier Zermelo Theorem on determinacy, whose proof involves this modal inductive clause for computing winning positions of player ***E***, with $\boldsymbol{\mathcal{E}}$ the union of all moves available to her, and $\boldsymbol{\mathcal{A}}$ the same for player ***A***:

$$\mathrm{WIN}_{\mathbf{E}} \leftrightarrow (end \ \& \ win_{\mathbf{E}}) \vee (turn_{\mathbf{E}} \ \& <\boldsymbol{\mathcal{E}}>\mathrm{WIN}_{\mathbf{E}}) \vee (turn_{\mathbf{A}} \ \& \ [\boldsymbol{\mathcal{A}}]\mathrm{WIN}_{\mathbf{E}})$$

Specialized to first-order evaluation games, this schema may be seen as an alternative formalization of the recursive mechanics of the truth definition. For more on this fine-grained analysis using modal logic as the game logic, cf. van Benthem 2001A, 2002A. For the moment, we just remark that game equivalence at this level would be more like modal *bisimulation*, a much finer sieve than power equivalence. What level of detail wants depends very much on the intended application.

2.6 A first summary

We have not yet reached *IF* games! And we have already found the beginnings of a research program about the connections between evaluation games and game theory. Moreover, this perspective passes one test: it tells us things about first-order logic that we did not know before. One striking example was the discovery of a decidable game algebra lying underneath its surface. But one can find such things basically anywhere. For instance, take our final excursion into modal logic, at the game level which 'did not fit' first-order equivalence. Actually, the issue of finer levels than standard equivalence at which to identify *logical propositions* has a long history, going back at least to Russell. Bisimulation of evaluation games provides one such answer to this, and more generally, different levels of game representation might provide different accounts of logical *propositions*: some more ‘extensional’, some more ‘intensional’ (cf. Moschovakis 1994). Thus there is much more game structure to evaluation games than you’d think, once you stop waving the classical tourbook.

3 *IF* games and imperfect information

3.1 *IF* logic in a nutshell

It is high time to turn to Hintikka’s more spectacular proposals, changing standard evaluation games into an engine for general *information-friendly logic* ('*IF* logic'). In what follows, we presuppose familiarity with this system on the part of the reader. Here is the program in a nutshell. Standard first-order logic imposes a linear operator order, which introduces hosts of dependencies, since 'later means under'. There are many reasons for breaking away from this – in logic, philosophy, linguistics, computer science, or even physics. We wish to allow for more complex non-linear constellations of quantifiers with only partial dependencies. In terms of the above evaluation games, perfect information meant that players have access to all previous moves by their opponent and themselves. Breaking with this constraint requires new games where players may have to make their choices of objects independently from what the other player has done before. One typical way of achieving this involves *imperfect information*, where players need not know where they are in the game tree. This is the typical situation in card games, where we do not know each other's hands – and indeed, in game theory, imperfect information is a well-established subject (cf. Osborne & Rubinstein 1994). In logic, however, it is a major innovation, whose repercussions are still widely debated (Hodges 1997, Janssen 2002).

In this Section, we make just one major logical point. Imperfect information means that players cannot distinguish between different states of a game. This is *precisely* the standard semantics of epistemic logic, and hence we can introduce explicit epistemic knowledge operators to formalize various aspects of *IF* games. The benefits of such a move are the same as those of epistemic logic generally: clarity of analysis, and suggestiveness for further topics. Using this tool, we will look at *IF* games more or less as in Section 2, at various levels. Viewed as theatres for players operating under ignorance, we analyze them using a dynamic-epistemic language. Viewing them as just outcome-producing 'machines', we extend the earlier 'power equivalence' to deal with *IF* equivalence, and relate the result to known game theory.

3.2 Getting acquainted

Many discussions of *IF* games start with perplexities, and attempts at formulating the design intuitions behind the system. Here, we will just make a brief tour of issues.

Slash syntax and nondetermined games *IF* logic has a lush syntax of slashes, indicating that quantifiers are independent, or that players may have imperfect information about previous moves. In what follows, we forego a priori limitations: players may be uncertain about their own, or the opponent's moves. There are some syntax restrictions in *IF* logic, such as ***F***'s never being uncertain about ***V*** – but these seem mainly remnants of a statement focus, making Verifier the prima donna. Having grasped the general scene, we will discuss systematic restrictions later. As an example, consider the earlier game with a 2-object domain, but now for the game

$$\forall x\ \exists y/\mathbf{x}\ x \neq y$$

where the slash indicates that Verifier no longer has access to the first object mentioned by Falsifier. She may have forgotten, the object may have been presented in a sealed envelope, etc. Intuitively, this game has the following tree:

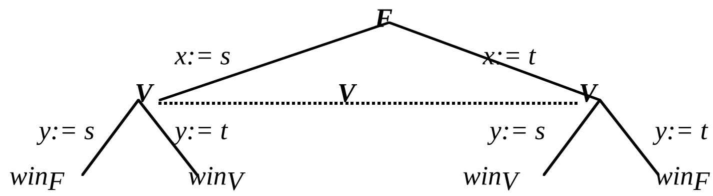

Here the dotted line is a standard game-theorist's device indicating Verifier's natural equivalence relation of indistinguishability between the two game states in the middle. Equivalence classes of this relation are called players' 'information sets' in game theory – a deviant terminology going back to an independent rediscovery of Hintikka's epistemic logic (Hintikka 1962) by game theorists in the 1970s.

Crucially, the new game is *non-determined*, in a sense appropriate to the extended setting. Verifier still has her old winning strategy, but it is not useable. What she needs is a *uniform winning strategy*, whose prescribed actions are the same across all game states that are indistinguishable to her. Game theorists would even call this the only strategies for ***V*** in this game, as strategies assign moves to information sets. But no such uniform strategy exists in this game. The only two candidates ('choose object *s*', 'choose object *t*') cannot guarantee a win. But neither does Falsifier have a winning strategy: anything he does might be countered by a move for Verifier.

Skolem forms and complexity In logical terms, the statement that Verifier has a winning strategy corresponds closely to normal forms using *Skolem functions*. E.g., a standard first-order formula $\forall x\ \exists y\ R(x, y)$ is equivalent to

$$\exists f\ \forall x\ R(x, fx)$$

Likewise, the statement that ***V*** has a uniform winning strategy in the above game for $\forall x\ \exists y/\mathbf{x}\ x \neq y$ can be written as follows, dropping one variable dependency:

$$\exists f\ \forall x\ R(x, f)$$

This gets more exciting in more complex examples. As an illustration consider

$$\forall x\ \exists y\ \forall z\ \exists u/\mathbf{x}\ \ R(x, y, z, u)$$

Here the statement that ***V*** has a winning strategy amounts to saying that

$$\exists f\ \exists g\ \ \forall x\ \ \forall z\ R(x, f(x), z, g(z))$$

There is a body of technical theory on this (cf. Sandu & Väänänen 1992), showing that the expressive power of *IF* logic goes up to fragments of second-order logic. That is, the statement that Verifier has a uniform winning strategy in an *IF* game can leading to branching non-first-order quantification patterns over Skolem functions.

Of course, in terms of Section 2.2, this says something about the complexity of some statements about *IF* games. It does not tell us much about the games themselves. Imperfect information is all around us: in card games, or in parlour games, with sometimes quite sophisticated mechanisms of information hiding. The logic of those mechanisms is an exciting ongoing story (cf. Baltag, Moss & Solecki 1999, 2002, van Ditmarsch 2000, van Benthem 2001B), but it has taught us at least this. That some technical statements *about* imperfect information games need high complexity is orthogonal to the issue whether the games themselves, as activities, are easy or hard to play. Some might even be *easier* to play than their perfect information counterparts, as there may be fewer things to keep in mind in small memories.

Can IF games be played at all? Even so, all this does not address the question *how, or even whether, one can play IF games*. *IF* syntax allows arbitrary slashing of quantifiers and connectives, suppressing dependencies on any earlier operators. Does this correspond to realistic settings where players find themselves in such circumstances? Hintikka and Sandu never provide a definition of *IF* games. We are not given the game trees, let alone specific mechanisms that would make arbitrary *IF* games playeable. Parts of the syntax suggest imperfect information about moves (as in the above game), others memory loss, perhaps even just intermittent:

$$\forall x\ \exists y\ \forall z\ \exists u/\mathbf{x}\ \forall v\ \exists s\ R(x, y, z, u, v, s)$$

One interpretation offered in the folklore is that all slashes make sense when we assume that ***V*** and ***F*** are really *teams* whose members work in parallel. This would be like the typical game-theoretic notion of a *coalition* (cf. Section 4) – but no precise interpretation of this form has been specified so far by Hintikka or his critics.

My own view is the following. *IF* syntax is a specification for patterns of knowledge and ignorance. It does not address the issue of designing actual games that meet these specifications. Also, it ignores finer distinctions. Some ignorance is public, and part of the legitimate design of a game. Examples are putting moves in envelopes, shuffling cards, or dealing hands to players. Such games can be played by ideal players without limitations on their capacities for reasoning and observation. Another, quite different source of ignorance are players' limitations: they may not pay attention, have bounded memory, cheat, and so on. This might even happen with games of perfect information. These different sources of ignorance are run together in *IF* syntax, so that discussion is bound to remain confused.

Way-out! At this stage, I should offer the reader an escape hatch. All these worries only hurt if one takes the games seriously. On the other exegetic hypothesis, *IF* logic is just about *(in-)dependence*, and the game metaphor can be thrown away as soon as it becomes a nuisance. I intend to pursue the games, but one does not have to follow. To focus what follows here are a few more concrete test questions to play with.

Test problems Here is a first example. Many people claim that the above game

$\forall x\ \exists y/\mathbf{x}\ x \neq y$

is not yet an issue, because it is 'really' equivalent to the first-order formula

$\exists y\ \forall y\ x \neq y$

For, in order to win the first game, Verifier must put up an object that works against anything that Falsifier may have mentioned. But upon reflection, this story is strange! The first game is *non-determined*, the second game has perfect information. So they differ in significant properties – and one would expect them to come out as being different. (They are.) Equivalence judgments are a nice test for understanding any proposed semantics. Typically, when quizzed on equivalence of slash formulas, people will either quote Hintikka, or try to look mysterious and appeal to private semantic intuitions. We will analyze what goes on in neutral game-theoretic terms, leading to a different outcome – which is actually nicer purely logically.

A second example are the beautiful signalling phenomena found in Hodges 1997. Consider a slight modification of $\forall x\ \exists y/\mathbf{x}\ Rxy$, with a vacuous quantifier inserted:

$\forall x\ \exists z\ \exists y/\mathbf{x}\ x \neq y$

Some people's intuitions tell them a vacuous quantifier never makes a difference 'since it is redundant in standard logic'. This reasoning is hard to understand, since one of the purposes of *IF* logic was to extend standard logic, so that intuitions from that original area need to be sifted as to what should generalize and what should not. Indeed, vacuous quantifiers are additional moves, which do matter in game theory. This time, Verifier does have a uniform winning strategy:

"use your *z*-move to copy ***F***'s first move,
then copy that for your own *y*-move".

This is admissible. If we want to prevent ***V*** from using her earlier *z*-response, we should rather consider a different *IF* game, with some obvious extended slash syntax:

$\forall x\ \exists z\ \exists y/\mathbf{\{x, z\}}\ x{\neq}y$

There is nothing mysterious here, and we will provide precise game-theoretic details.

I will now show how to treat at least some *IF* games as imperfect information game trees, and we can then embark on the program outlined at the end of Section 3.1.

3.3 *IF* games as imperfect information games

Intriguing examples In Section 2, it was easy to define an extensive game tree *game(ϕ, **M**, s)* for any first-order formula *ϕ*, model ***M***, and variable assignment *s*. We have not really specified the last details of this, but it can be done, given enough industry. Can we do the same when *ϕ* is a slash formula from *IF* syntax? We already did the example of $\forall x\ \exists y/\mathbf{x}\ x{\neq}y$ on a two-object domain, whose form will also be clear for arbitrary models ***M***. The underlying game tree was the ordinary one for $\forall x\ \exists y\ x{\neq}y$, while the slash told us where to put dotted lines in that tree for players' uncertainties. Next, consider the two (non-)signalling examples from Section 3.2, again for convenience over a two-object domain. Lots of things will emerge.

An imperfect information game for $\boldsymbol{\forall x\ \exists z\ \exists y\ /x\ x{\neq}y}$ It may not be immediate from this first slash formula how to draw dotted lines for ***V*** in the underlying game tree. But a rather simple algorithm does exist – and it is implicit in the next picture:

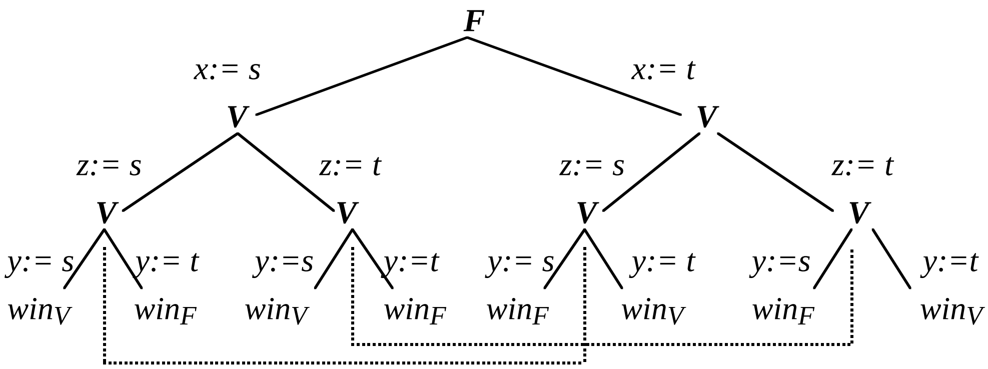

The dotted lines represent ***V***'s uncertainty in the third round about ***F***'s first move. But at the same time, in that round, the lines show that she knows her own move in the second round. We indicate a uniform winning strategy for her with bold-face arrows:

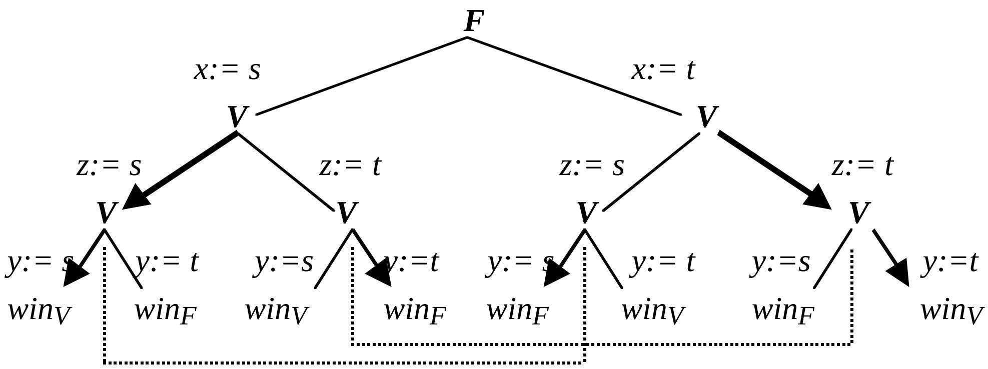

There are some subtleties here in interpretation! ***V***'s strategy is indeed uniform by definition, as it assigns the same move to states that she cannot distinguish. Moreover, it is a winning strategy, in that, if she follows it, she will in fact end in a winning state, whatever ***F*** does. But when the third round has come, ***V*** will not *know* that her strategy is winning, as she considers it possible that ***F*** played another move, so that her prescribed move will make her lose. In other words, one can have a uniform winning strategy without knowing at each stage of following it that playing the rest of the strategy is in fact winning. This is like following a guide through a bog, having forgotten the reasons that convinced us that the guide was going to get us across. Some people find such subtleties annoying: I myself find them delightful.

In game-theoretic terms, notice that the above is a game *without Perfect Recall*. In the third round, ***V*** has forgotten information which she did have in the second round. Such games are notoriously harder to interpret than games where players cope with uncertainty without memory failures, like expert card players.

Knowledge about strategies The interpretation of what happens under various scenarios in imperfect information games remains a contested issue, even in game theory. Incidentally, these difficulties reflect those of interpreting *IF* syntax, and so they strengthen, rather than weaken the connection that we are making. For instance, consider this. Since ***V*** is just as rational as you and me, she can see that the above strategy must make her win. Will not this knowledge assure her in the third round that she must win? Well, for that to happen, she must *remember* her strategy. But modeling the latter knowledge goes beyond knowing where one is in the game tree: it presupposes a richer representation, including information about possible strategies. Such 'meta-models' of games have existed in game theory since the 1970s in discussions of rational behaviour (cf. Osborne & Rubinstein 1994, Stalnaker 1999), but they would take us too far here. In particular, if ***V*** remembers her strategy throughout the game, her information should only contain the game played according to that strategy. But then, the above picture changes, and we get a 'cut-off version':

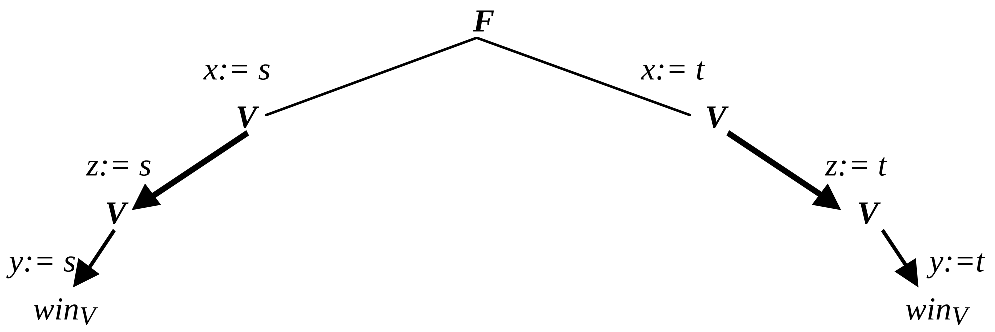

An imperfect information game for ***∀x ∃z ∃y*** */{x, z}* ***x≠y*** The preceding game may be contrasted with the next, where only the dot pattern for ***V*** changes:

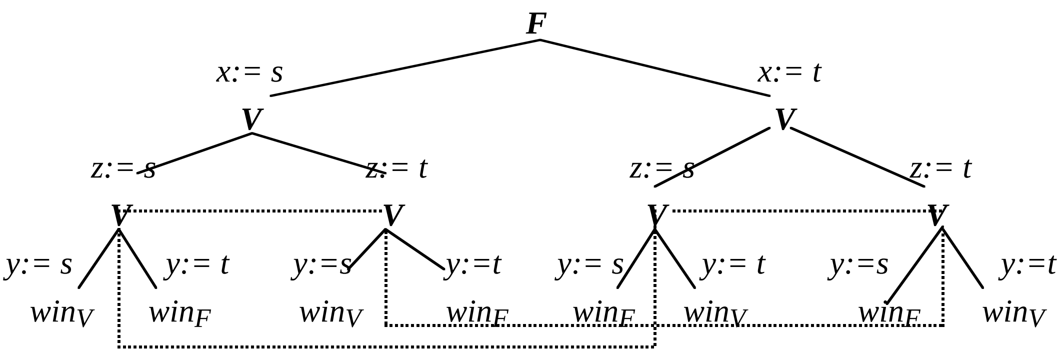

With the uncertainty lines in this game, ***V*** has no strategy which she knows to work.

The phenomenon encountered with *∀x ∃z ∃y/***x*** x≠y* is called *self-signalling*. Players may be able to derive officially unavailable information by a roundabout route. To work well, signalling arguments depend on epistemic assumptions about knowledge, such as *players know their available moves, how many moves have been played,* etc. In general games of imperfect information, players need not know how many moves were played. But *IF* syntax always seems to assume at least this much: ***V*** may know nothing about the object which ***F*** chose, but she does know that he made a choice.

I have not formulated a precise algorithm for drawing game trees for *IF* formulas, but the general method should be clear from these examples. One first draws the slash-free game tree, and then, at the level corresponding to an operator, one connects all histories for ***E*** which differ only in positions on which her choice should not depend.

3.4 Powers, game equivalence, and game algebra

Hintikka & Sandu sometimes call *IF* games *three-valued*: either ***V*** has a winning strategy (this is 'truth'), or ***F*** has a winning strategy ('falsity'), or neither ('third truth value'). This is about the most niggardly way of giving imperfect information games some additional structure beyond that of perfect information games.

Uniform powers Instead, let us look at the power analysis of Section 2.3. This extends immediately to imperfect information games, but this time considering only *uniform* strategies. Thus, in the game of Section 3.2, ***F*** retains the powers he had in the perfect information version, but ***V*** loses two former powers, retaining only those for her remaining strategies "left, left" and "right, right":

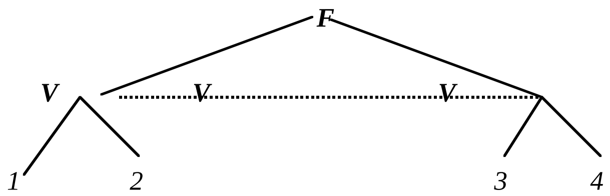

powers of ***F*** $\{1, 2\}, \{3, 4\}$
powers of ***V*** $\{1, 3\}, \{2, 4\}$

This list seems poorer than for the perfect information version. But one can also see it as a more subtle form of power sharing where ***V*** and ***F*** have become more equal. In fact, imperfect information is often *needed* in designing organizations giving members just the right amount of influence. This time, the only general conditions that hold are Monotonicity and Consistency. Van Benthem 2001A shows that these suffice to represent any power list for two players by an imperfect information game.

Power equivalence Game equivalence may again be analyzed in terms of powers in this new sense. An interesting check is that there already exists a calculus to this effect in game theory, the 'Thompson transformations' (Osborne & Rubinstein 1994, Chapter 11). These match the predictions of power equivalence precisely, at least for games with Perfect Recall. We are now in a position to answer an earlier question:

What is the correct game equivalent for $\forall x\ \exists y/\mathbf{x}\ x{\neq}y$?

The answer is, not $\exists y\ \forall x\ x{\neq}y$, but the much more symmetric formula:

$\exists y\ \forall x/\mathbf{y}\ x{\neq}y$

This corresponds to the above game tree, with turns and outcomes interchanged:

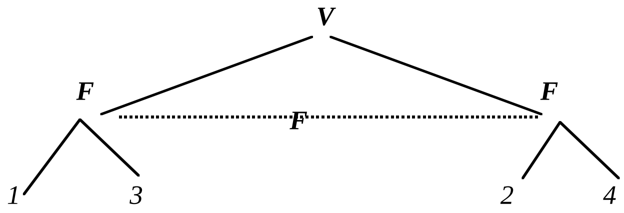

Clearly, players' uniform powers are exactly the same here as in the above game. This scheduling equivalence is about the most basic Thompson transformation.

A similar analysis of the Hodges example $\forall x\ \exists z\ \exists y/\mathbf{x}\ x{\neq}y$ shows that it is equivalent to the formula $\forall x\ \exists y\ x{\neq}y$, which is slash-free and determined! On the other hand, this is not a Thompson transformation, as we are not assuming Perfect Recall.

IF logic as game calculus These observations raises some interesting logical issues. We can see the equivalential part of *IF*-logic as a calculus for game equivalence, just as first-order logic encoded such a calculus for perfect information games in Section 2.4. For instance, the equivalence between $\forall x\ \exists y/\mathbf{x}\ \phi$ and $\exists y\ \forall x/\mathbf{y}\phi$ is a valid distribution law of sorts. It also has propositional equivalents, such as

$$(A \boldsymbol{\vee}/\wedge\ B) \wedge (C \boldsymbol{\vee}/\wedge\ D) \quad \leftrightarrow \quad (A \boldsymbol{\wedge} C/\vee) \vee (B \boldsymbol{\wedge}/\vee D)$$

One interesting question is this:

> Does *IF* logic have a simple subsystem of operator equivalences
> which axiomatizes uniform power equivalence over general games?

Operations and game algebra But there are also pitfalls in extending the account of Section 2. Hintikka's well-known quarrels with compositionality reflect the game-theoretic difficulty that imperfect information games 'have no good notion of a subgame'. Their dotted lines mess up the compositional structure of the underlying game tree. So, are there natural operations at all on imperfect information games?

Parallel products Perhaps a shift in perspective is needed. Abramsky 2000 embeds some *IF* games in linear game semantics, using parallel composition to achieve imperfect information. This would embed part of *IF* logic into *linear logic*, although the sense in which is a bit unclear, given the different complexities. Netchitajlov 2000 proposes further parallel products, allowing for interleaved play.

Here is a simpler observation from van Benthem 2002B with a similar point. A basic structure in game theory are strategy matrices. Two players move in parallel, with four possible outcomes. A parallel phenomenon in logic is 'branching quantification':

$$\left.\begin{matrix}\forall x\ \ \exists y \\ \\ \forall z\ \ \exists u\end{matrix}\right> Rxyzu$$

This lets choices for prefixes take place independently – bringing them together at the end to evaluate the matrix assertion *Rxyzu.* Such games involve a mild form of imperfect information: ignorance of others' moves played at the same time. We can define the corresponding game operation more generally as

product $G \times H$

whose runs are pairs of separate runs for *G, H* with the product of their end states as the total end state. In terms of players' powers, this works out as follows:

$$\rho^{i}_{GxH}\ (s, t), X \quad \text{iff} \quad \exists U\!:\ \rho^{i}_{G}\ s, U\ ,\ \exists V\!:\ \rho^{i}_{H}\ t, V : UxV \subseteq X$$

Players' powers in such games are no longer determined, but they still satisfy Monotonicity and Consistency, and there is an analogue of the above representation.

Fact The following identities of game algebra hold for product games:

$A\ x\ (B \cup C)$	=	$(AxB) \cup (AxC)$
$(A \cup B)\ x\ C$	=	$(AxC) \cup (BxC)$
$(AxB)^d$	=	$A^d\ x\ B^d$
$G\ x\ H$	=	$H\ x\ G$

This may be proved by straightforward analysis of players' powers. The fourth line also assumes that the component order in product states *(s, t)* is immaterial.

But now back to *IF* games. What would it mean to play an evaluation game $\phi x \psi$? Consider the above branching quantifier. Here is a corresponding slash formula:

$$\forall x\ \exists y\ \forall z/\{\mathbf{x}, \mathbf{y}\}\ \exists u/\{\mathbf{x}, \mathbf{y}\}\ Rxyzu$$

which suppresses all information flow between the two prefixes. In game-algebraic terms, this would be written as follows, with a 'test game' at the end:

$((\forall x \,\mathbf{;}\, \exists y)\ \mathbf{x}\ (\forall z \,\mathbf{;}\, \exists u))\ \mathbf{;}\ Rxyzu\mathbf{?}$

Thus, again, at least this fragment of *IF* logic seems a mixture of general game algebra and special facts about first-order semantic procedures. Game-algebraic laws now have *IF*-instances that allow one to manipulate quantifier prefixes, such as

$$(\forall x \,;\, \exists y)\ x\ ((\forall z \,;\, \forall u) \cup (\exists v \,;\, \exists u)) \ =$$
$$((\forall x \,;\, \exists y)\ x\ (\forall z \,;\, \forall u)) \ \cup\ ((\forall x \,;\, \exists y)\ x\ (\exists v \,;\, \exists u))$$

Also, valid principles of *IF* logic show up as algebraic validities. E.g., the above

$$\forall x\ \exists y/\mathbf{x}\ Rxy \quad \leftrightarrow \quad \exists y\ \forall x/\mathbf{y}\ Rxy$$

says in game-algebraic terms that

$$(G\ x\ H)\ ;\ K \ = \ (H\ x\ G)\ ;\ K$$

This principle follows easily from the above game algebra. But *IF* logic can also detect invalid algebraic principles. Here is an example of the latter:

$$(A\ x\ B)\ ;\ C \ = \ (A\ ;\ C)\ x\ (B\ ;\ C)$$

An *IF*-counterexample is the slash formula $\exists x\ \forall y/\mathbf{x}\ Rxy$, whose evaluation game is not equivalent to that for the perfect information game for $\exists x\ Rxy\ \mathbf{;}\ \forall y\ Rxy$. Given these observations, can we extend the representation theorem for perfect information games via evaluation games in van Benthem 2002B to *IF* logic after all?

3.4 Dynamic-epistemic logic of actions and knowledge

Games as dynamic-epistemic models Inside information games many interesting phenomena occur as players move through a game. To bring this out, we must move to the action level of Section 2.5, with a formalism to describe what players *know*. Games of imperfect information have states, moves, and epistemic equivalence relations $\sim_i$ for players ***i*** between states. The resulting models look like this:

$$\mathbf{M} = (S, \{R_a \mid a\in A\}, \{\sim_i \mid i\in I\}, V)$$

In principle, any uncertainty pattern might occur. Players need not know what the opponent has played, or they themselves, they need not know whether it is their turn, if the game has ended, and so on. Game theorists sometimes impose restrictions like common knowledge of the current turn or moves – which can be special axioms.

Dynamic-epistemic logic These models support a standard combined dynamic-epistemic language, with action modalities $[a]$, $\langle a \rangle$ for moves as in Section 2.6, and knowledge operators for each player:

$$\boldsymbol{M}, s \models K_i\phi \quad \text{iff} \quad \boldsymbol{M}, t \models \phi \ \ \textit{for all } t \textit{ s.t. } s \sim_i t$$

Now we can talk about knowledge and ignorance of players when the game has reached a certain state. This illuminates the situation depicted in Section 3.1. At the intermediate states, Verifier's knowledge may be described as follows:

$$K_V(\langle y{:=}s\rangle win_V \vee \langle y{:=}t\rangle win_V)$$

$$\neg K_V\langle y{:=}s\rangle\, win_V \;\&\; \neg K_V\langle y{:=}t\rangle\, win_V$$

This is a familiar distinction from intensional logic. ***V*** knows *de dicto* that she has a winning move, but she lacks a *de re* version: there is no particular move which she knows to be winning. You may know the ideal partner is walking around in this dark rain-swept town without knowing of any passer-by whether (s)he is that partner... Players can also have *iterated* knowledge about others' knowledge and ignorance via formulas like $K_iK_j\phi$, $K_i\neg K_j\phi$, which may be crucial to understanding the course of a game. Also players may achieve *common knowledge* about certain facts:

$C_{\{1, 2\}}\,\phi$ — ϕ is true in all those states that can be reached from the current one in a finite number of $\sim_1$ and $\sim_2$ steps

E.g., in the above game, ***E***'s plight is common knowledge between the players.

As for systematic reasoning about players' actions, knowledge, and ignorance in such models, the complete set of axioms for validity in dynamic-epistemic logic is

(a) the *minimal dynamic logic* for the modal operators $[a]$

(b) *epistemic S5* for each knowledge operator K_i

With a common knowledge operator added, we also get the minimal logic of that (cf. Fagin et al. 1995). There are no further axioms in general – but see below.

Defining uniform strategies In these game models, we can also define players' strategies as in Chapter 2. Recall that the relevant strategies now are the *uniform* ones, which have to prescribe the same moves at indistinguishable nodes for a player where it is her turn. Speaking generally, this restricts the possible behaviours. The above examples suggest that the uniform strategies are the ones of which a player *knows* that they lead to the desired result. This will show in available strategy definitions of this language, which may contain instructions like

"IF $K_{you}P$ THEN *do a* ELSE *do b*"

It was suggested in Section 3.3 that uniform strategies are precisely those that force a set of outcomes such that their owners *know* at each stage of using them that they will produce that set. But this is still imprecise, and not always true – e.g. with 'self-signalling' examples like the above $\forall x\ \exists z\ \exists y/\mathbf{x}\ x{\neq}y$. A more precise description is found in van Benthem 2001A, which shows that for players with Perfect Recall, uniform strategies and 'fully predictive' strategies of this epistemic kind indeed coincide. This may be seen as a kind of epistemic analysis of Skolem functions.

Varieties of imperfection Within the total universe, specific imperfect information games may validate additional epistemic-dynamic axioms, such as the game-theoretic assumptions mentioned above, which also hold for *IF* games:

(a) The fact who is to move is common knowledge between players
(b) All indistinguishable nodes have the same possible actions

More generally, in this way, we can do an epistemic analysis of types of imperfect information game, distinguishing different strands inside full *IF* syntax. In particular, the cited reference shows how one can distinguish *ways of playing games*, and describe their effects. For instance, consider players who have the above-mentioned feature of *Perfect Recall*, operating with perfect memory amidst the structural uncertainties introduced by the game itself. In particular, the latter may arise through defective observation of other players' moves. This restricts the pattern of dotted uncertainty lines in ways expressed by two additional principles:

(a) $turn_i\ \&\ K_i[a]\phi\ \rightarrow\ [a]K_i\phi$
(b) $\neg turn_i\ \&\ K_i[A]\phi\ \rightarrow\ [\boldsymbol{A}]K_i\phi$

with $\boldsymbol{A}$ the union of all actions available to the other player

These axioms say that moves of player ***i*** in this game *commute* with her knowledge. This commutation fails in general dynamic-epistemic logic, since my normal actions can have epistemic side-effects. I may know that "having a beer" will lead to my "being a bore", without knowing I am being a bore once I have drunk the beer. The resulting restriction on games is a commutative diagram. E.g., for (a) we get:

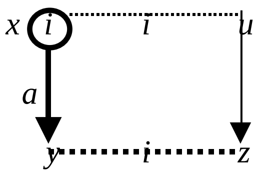

$$\forall xyz\colon TURN_i x \ \&\ R_a xy \ \&\ y \sim_i z$$
$$\to\ \exists u\colon x \sim_i u \ \&\ R_a uz$$

One can take this restriction to *IF* games, and ask just which syntactic slash patterns obey this commutative condition. In particular, the preceding diagram says that ***V***'s slashes at some level in a quantifier prefix must have ancestors at the preceding prefix position, if the latter is an existential quantifier. The result of the restricted syntax might be a simpler sublogic for game equivalence.

At an opposite extreme to Perfect Recall, players have *Bounded Memory*, allowing them to remember only the last *k* moves played for some *k*. This, too, can be expressed in dynamic-epistemic logic. E.g., with $k=1$, we get a characteristic axiom

$$E(\langle a\rangle T \ \&\ \phi) \to\ U[a^{\cup}]\neg K_i \neg \phi$$

where *E (U)* is an *existential* (*universal*) modality. This, too, is a pattern of *IF* syntax: slashes should start appearing beyond a certain distance in the quantifier prefix. More generally then, it would be of interest to look at the fine-structure of *IF* games, and characterize those fragments of *IF* syntax which model natural ways of playing games, with their a corresponding axioms in our logic. The dynamic-epistemic language will allow us to reason about ***V*** and ***F***'s interaction in these games.

3.5 A second summary

We have shown how *IF* games may be seen as perfectly ordinary games of imperfect information. The fact that there are some difficulties of making intuitive sense of them merely reflects intriguing similar subtleties in game theory. At the level of players' powers, on can do an analysis of equivalence in terms of uniform strategies, and even an incipient game algebra. This led to open questions about *IF* logic serving as a complete algebra for varieties of imperfect information games. But perhaps the most interesting perspective is the more detailed action level. There we can use explicit epistemic logic to describe players' progress in a game, and define interesting types of special behaviour. These correspond to *IF* sublanguages, with perhaps better-behaved logics. Thus, we get a handle on the fine-structure of *IF* logic.

4 Discussion

Finally, we point at some further issues concerning *IF* logic which we had to forego here. After that, we state our main claims, and draw our general conclusions.

Logical aspects There are many further logical aspects to *IF* logic in the above.

First, it would be good to add an explicit component of *strategy calculus,* dealing, amongst others, with Skolem functions. But the latter would have to be generalized, as we are dealing not just with ***V*** but also with independent powers of ***F***. Instead of embedding *IF* logic into second-order logic as it is now, this might enrich second-order logic to a system with a duality between Skolem functions for two players.

Another aspect is the question whether the first-order language, even when slashed, is really the right formalism for the enterprise. With many other deviant logics, one is led to introduce *new logical operators* reflecting the new setting. Examples are the product operations of linear logic, which made their appearance in Section 3.3. Other examples might be polyadic quantifiers, letting players pick bunches of objects at the same time. Redesign of the *IF* syntax is also taken up in Hintikka 2002.

Finally, a widely noted desideratum, one would like to have a more general account of '*IF*-ing', which can also be applied to study imperfect information versions of other logic games, such as proofs where participants do not know some earlier moves, or model comparisons where players must make do with a finite memory.

Game-theoretic aspects *IF* games introduce one type of more realistic structure over traditional logic games of semantic evaluation, viz. imperfect information. But there are other candidates. For instance, real game theory is about games where players have finer *preferences* than the 'win'/'lose' of logic games. This can also be done in logic games, giving ***V***, ***F*** independent evaluations of atomic facts in a model. This fits with modern default logics using preferences for players in their models. An example was the non zero-sum battle-field game of Section 2.2:

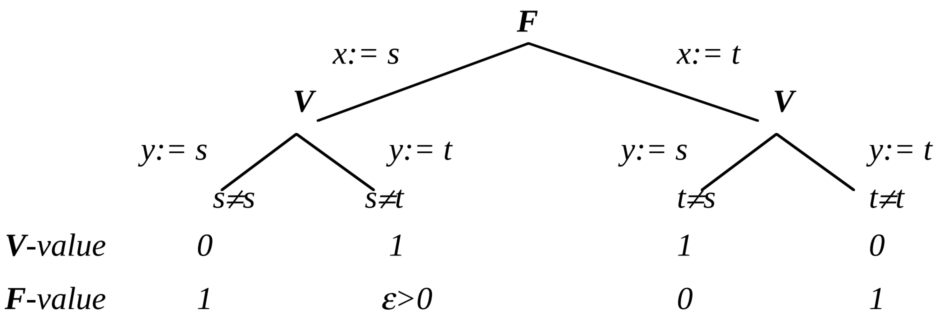

Now the deeper issue is this. In game theory, preferences lead to Nash's notion of *strategic equilibrium*, and finer predictions of behaviour. E.g., the original game of Section 2.2 has two Nash equilibria: ***V*** plays her winning strategy, ***F*** any strategy:

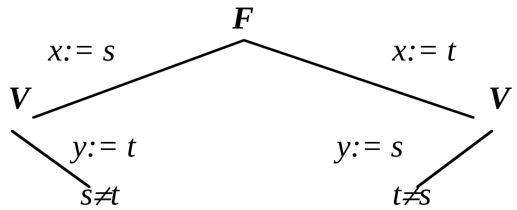

The game with the *ε*-preference has only one equilibrium, resulting in

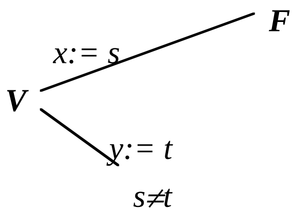

The set of Nash equilibria of its evaluation games might be a good candidate for a more radical game-theoretic denotation of a logical formula!

Another realistic game-theoretic feature are *coalitions*. Perhaps the most significant move in epistemic logic after Hintikka's pioneering work has been the introduction of operators that are typical to groups of agents, such as *common knowledge* (Fagin et al. 1995, Osborne & Rubinstein 1994). It suggests a similar extension of *IF*-logic with joint actions for groups of agents. In particular, the players ***V***, ***F*** themselves might be teams. A first logical analysis of coalitions is found in Pauly 2001, but epistemic and dynamic logic have not yet taken in this notion in its full generality.

***Probability**?* But things are even more intriguing. The basic insight in game theory has been that strategic equilibria may only exist in a game when we move from pure to *mixed strategies*, using probabilistic mixtures of pure strategies. This will not arise in the standard logic games of Section 2 with preferences added, as we can always find Nash equilibria there using the well-known algorithm of 'backward induction', a numerical version of the proof of the Zermelo Theorem on determinacy.

Probabilistic solutions do arise in some *IF* games! The little game in Section 3.2

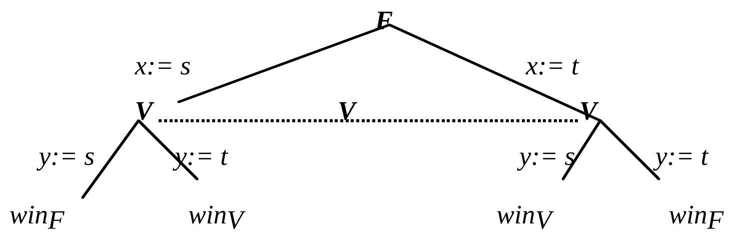

is just the game-theoretic classic of ‘Matching Pennies’. This has an optimal value *(1/2, 1/2)*, achieved by players using their uniform strategies with probability *1/2*.

I find this observation extremely intriguing from a logical point of view. We know that probability sometimes emerges naturally in pure logic, telling us something about long-term behaviour. An example are the *Zero-One laws* of first-order logic, which state that with increasing finite domain size, the probability that any given first-order formula is true goes to either *1* or *0*. Could it be that *IF* games also involve an essential probabilistic feature, which we just have not been able to identify yet?

Architecture of intended applications The original grand motivations of game-theoretical semantics (Hintikka & Sandu 1997) had to do with describing large-scale cognitive systems such as natural language. Evaluation games are just a small part of this story, and more can be said. Natural language, or ordinary reasoning, involve many different games. There are terminating finite-depth games for short-term tasks, such as evaluation or proof. Nice recent examples on very different principles from *IF* games are the interpretation games of van Rooy 2001 with speaker/hearer preferences from linguistic optimality theory, and the argumentation games of Rubinstein 2000 analyzing Gricean pragmatics. But there are also infinite games providing the hopefully never-ending 'operating system' for these short-term tasks, such as the procedural rules of civilized conversation or debate. Another missing feature then is an account of architecture: how do different games fit together into one coherent system? How can information be passed from one game to another?

5 Conclusion

What has been shown in this paper is that logical evaluation games, either Hintikka's original ones or their *IF* versions, can be linked systematically with game-theoretic themes, some of them present in existing game theory, and some of them new. In particular, this has led to some insights and new questions about game-theoretical semantics, summarized in Sections 2.6 and 3.5, which need not be repeated here. Whatever its merits as an exegesis of Hintikka's intentions, I hope this has been a convincing sample of the lively current interface between logic, computer science, and game theory, which naturally covers many more topics (cf. van Benthem 1999–2002, and dedicated conferences such as TARK, LOFT, and GAMES).

More specifically toward *IF* games, we propose viewing them in a systematic game-theoretic light, which suggests a host of new perspectives and questions. In doing so, one encounters Hintikka's pioneering work once more, since the tool of choice in this area is his very own epistemic logic in its original form. This is somewhat surprising, since many of the informal explanations behind the Hintikka & Sandu approach involve players' knowledge, but they are left implicit. Admittedly, there are published *IF* versions of epistemic languages, but my point is rather that one can illuminate the workings of any *IF* by means of standard static epistemic operators. Of course, one can then play a carrousel research game of systems *EL(IF(EL(IF*…

Finally, let us briefly reconsider the central *IF* motivation of possible independence between quantifiers. As said before, Hintikka's writings leave open an interpretation where *this* is the central topic, and the games just a discardable wrapping. This is a crucial decision point. For, *if* we take the games seriously, then eventually, dependence and independence will *not* be the central notions. They are rather derivatives from something still more central, namely, *interaction* between players. I will make no further defense of the latter here, but the issue cannot be evaded.

Nevertheless, I totally agree that independence is a crucial logical topic. But, then we must be radical, and account for the fact that it is quite diverse, with intuitively different sources. Hintikka describes one of these: *procedural dependencies* which arise in a process of evaluating assertions. These are absolutely important – but there are also *objective dependencies* in the nature of things, lying encoded in models whether or not we interpret anything at all. Objective dependencies have been studied by van Lambalgen 1995 on the logic of independent events in probabilistic reasoning. Another example, with more published results, is 'generalized assignment semantics' for first-order logic (Németi 1985, van Benthem 1996, 1997), which drops the assumption of the standard Tarskian models that values for variables can be modified completely independently of what happens to other variables. This, too, generates a new base logic of dependent and independent quantifiers different from standard first-order logic, while supporting a richer logical vocabulary of polyadic quantifiers and substitutions. But the intriguing difference with *IF* logic is that in this case, the true first-order basic logic without any built-in objective independence assumptions has lower, rather than higher complexity: it becomes *decidable*! Were I to write a paper on independence – which I have not done – a comparison between these varieties of independence logics would be the first order of business.

Despite possible divergences in interests, one only writes a long paper like this if the subject seems worth-while. *IF* logic offers an attractive laboratory for studying the logic of imperfect information and information flow generally. Jaakko's broad intellectual vision and challenging ideas over the years have been remarkable – and like many colleagues, I am still happy to experience his continued influence.

6 References

S. Abramsky, 1996, Semantics of Interaction: an Introduction to Game Semantics, in P. Dybjer & A. Pitts, eds., Proceedings 1996 CLiCS Summer School, Isaac Newton Institute, Cambridge University Press, Cambridge, 1–31.

S. Abramsky, 2000, 'Games in Computer Science and Logical Dynamics', invited lecture, 12th Amsterdam Colloquium, ILLC Amsterdam.

A. Baltag, L. Moss & S. Solecki, 1998, 'The Logic of Public Announcements, Common Knowledge and Private Suspicions, *Proceedings TARK 1998*, 43–56, Morgan Kaufmann Publishers, Los Altos. Update in 2002: department of cognitive science, Indiana University, Bloomington, and department of computing, Oxford University.

J. van Benthem, 1996, *Exploring Logical Dynamics*, CSLI Publications, Stanford.

J. van Benthem, 1997, 'Modal Foundations for Predicate Logic', *Bulletin of the IGPL* 5:2, 259–286, London and Saarbruecken. (R. de Queiroz, ed., *Proceedings WoLLIC, Recife 1995*).

J. van Benthem, 1999 – 2002, *Logic in Games*, lecture notes, ILLC Amsterdam & Stanford University.

J. van Benthem, 2001A, 'Games in Dynamic Epistemic Logic', in G. Bonanno & W. van der Hoek, eds., *Bulletin of Economic Research* 53:4, 219–248.

J. van Benthem, 2001B, 'Logics for Information Update', *Proceedings TARK VIII*, Morgan Kaufmann, Los Altos, 51–88.

J. van Benthem, 2002A, 'Extensive Games as Process Models', *Journal of Logic, Language and Information* 11, 289–313.

J. van Benthem, 2002B, 'Logic Games are Complete for Game Logics', to appear in *Studia Logica.*

J. van Benthem, 2002C, 'What Logic Games are Trying to Tell Us', to appear in R. Downey, ed., *Proceedings 7th and 8th Asian Logic Conferences*, World Scientific Press, Singapore.

A. Blass, 1992, 'A Game Semantics for Linear Logic', *Annals of Pure and Applied Logic* 56, 183 – 220.

H. van Ditmarsch, 2000, *Knowledge Games*, dissertation DS-2000-06, Institute for Logic, Language and Computation, University of Amsterdam.

R. Fagin, J. Halpern, Y. Moses & M. Vardi, 1995, *Reasoning about Knowledge*, MIT Press, Cambridge (Mass.).

V. Goranko, 2000, 'The Basic Algebra of Game Equivalences', Preprint 2000-12, Institute for Logic, Language and Computation, University of Amsterdam.

J. Hintikka, 1962, *Knowledge and Belief*, Cornell University Press, Ithaca.

J. Hintikka, 1973, *Logic, Language Games and Information*, Clarendon Press, Oxford.

J. Hintikka, 2002, 'Hyperclassical Logic (a.a.a. IF Logic) and its Implications for Logical Theory', *Bulletin of Symbolic Logic* 8:3, 404–423.

J. Hintikka & J. Kulas, 1985, *The Game of Language*, Kluwer, Dordrecht.

J. Hintikka & G. Sandu, 1997, ‘Game-Theoretical Semantics’, in J. van Benthem & A. ter Meulen, eds., *Handbook of Logic and Language*, Elsevier, Amsterdam, 361–410.

W. Hodges, 1997, 'Compositional Semantics for a Language of Imperfect Information', *Journal of the IGPL* 5, 539–563.

Th. Janssen, 2002, 'Independent Choices and the Interpretation of *IF* Logic', *Journal of Logic, Language and Information* 11, 367–387.

M. van Lambalgen, 1995, 'Natural Deduction for Generalized Quantifiers'. In J. van der Does & J. van Eyck, eds., *Quantifiers, Logic and Language*, CSLI Lecture Notes, vol. 54, Stanford University, 225–236.

Y. Moschovakis, 1994, 'Sense and Denotation as Algorithm and Value', *Lecture Notes in Logic* #2, J. Oikkonen and J. Väänänen, eds., 210-249.

I. Németi, 1985, 'The Equational Theory of Cylindric Relativized Set Algebras is Decidable', Preprint No 63/85, Mathematical Institute, Hungarian Academy of Sciences, Budapest, 38 pp.

Y. Netchitajlov, 2000, *An Extension of Game Logic with Parallel Operators*, Master of Logic thesis, Institute for Logic, Language and Computation, University of Amsterdam.

M. Osborne and A. Rubinstein, 1994, *A Course in Game Theory*, MIT Press, Cambridge (Mass.).

Rohit Parikh, 1985, ‘The Logic of Games and its Applications’, *Annals of Discrete Mathematics* 24, 111–140.

M. Pauly, 2001, *Logic for Social Software*, dissertation DS-2001-10, Institute for Logic, Language and Computation, University of Amsterdam.

R. van Rooy, 2001, Signalling Games Select Horn Strategies', manuscript, Institute for Logic, Language and Computation, University of Amsterdam. To appear in *Linguistics and Philosophy*.

A. Rubinstein, 2000, *Economics and Language*, Cambridge University Press, Cambridge.

Gabriel Sandu & Jouko Väänänen, 1992, 'Partially Ordered Connectives', *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik* 38, 361 – 372.

R. Stalnaker, 1999, ‘Extensive and Strategic Form: Games and Models for Games’, Research in Economics 53:2, 93-291.

E. Zermelo, 1913, 'Ueber eine Anwendung der Mengenlehre auf die Theorie des Schachspiels', In E. Hobson & A. Love, eds., *Proceedings of the Fifth International Congress of Mathematicians*, Cambridge University Press, Cambridge, 501–504.