

A MINI-GUIDE TO LOGIC IN ACTION

Johan van Benthem, Amsterdam & Stanford, December 2003

Abstract The *Dynamic Turn* in logic makes actions of communication and general information update into explicit objects of investigation. This paper is a brief tour of this research program in a new version, bringing together ideas from logic, philosophy, computer science, and game theory. In particular, we discuss epistemic, dynamic, and temporal logics, belief revision, and game logics.

1 The dynamic turn

Classical logic is about propositions which we can know or believe, and unchanging inferential relationships between them. But inference is first and foremost an *activity*, for which propositions are merely the input, and the result. In recent years, there has been a growing awareness that various activities of reasoning, evaluation, belief revision, or communication, are themselves typical themes for logical investigation, and that their dynamic structure can be studied explicitly by logical means.¹ For instance, it seems strange to study only the statics of what it means to 'know' a proposition, when knowledge usually results from basic actions of learning that we perform all the time, such as asking a question and getting an answer. Indeed, asking questions and giving answers are just as much logical core activities as drawing conclusions! This line can be extended: the natural dynamic counterpart of static epistemic logic is the theory of arbitrary individual or social *learning mechanisms*. Similar trajectories from static to dynamic arise when we look at inference in such stages, first as a zero-agent mathematical relationship between static propositions, then as a one-agent activity of drawing conclusions, and finally as a many-agent interactive process of argumentation. This broadening of perspective, sometimes called the 'Dynamic Turn', started around 1980 with work on interpretation procedures for natural language, as well as belief revision in artificial intelligence. But how should logic incorporate actions as first-class citizens into its scope? Plausible formal frameworks to this effect come from the philosophy of action, temporal logic, and systems for analyzing programs in computer science, such as dynamic logic.

¹ The same 'static'/'dynamic' distinction makes sense when we extend our notion of classical logic, e.g. by including definitions and expressive power of languages. Expressive power has to do with *activities* of evaluation of statements, making distinctions between given situations, and so on.

Moreover, further influences have come from process theories in computer science, as well as game theory, and this contact between disciplines is still continuing. This paper sketches one trajectory of 'dynamification', reflecting my personal interests. A much broader survey is given in van Benthem 1996.

2 From epistemic logic to communication

Questions and answers A typical illustration of the Dynamic Turn arises in epistemic logic. Let us move beyond the usual concerns that most of us were raised with, such as 'is knowledge true justified belief?', or 'which modal axioms should we choose for the epistemic operator K ?'. Instead, consider the most basic episode of communication. I ask you a simple *YES/NO* question

"Is Amsterdam at the same latitude as Peking?",

and you answer me truly. By the way, the actual answer is

"No" ²

Now much more information flows in this simple question-answer episode than meets the eye. Under normal circumstances, my question is only felicitous when certain *preconditions* are satisfied. First, I indicate to you that I do not know the answer. But there is more. The fact that I am asking *you* indicates that I think it is at least possible that you know the answer. ³ Now to the effects of the answer. By telling me, you make me learn the relevant fact P . But more is true afterwards. You know that I know, I know that you know that I know, and so on to any depth of iteration. We achieve what is called *common knowledge* in the philosophical and logical literature. These are the so-called *postconditions* of a truthful answer. ⁴

Incidentally, most preconditions and postconditions noted here involve knowledge about other people's knowledge. This may seem a somewhat redundant social side-effect of communication. But in reality, such iterated knowledge levels are often crucial to effective physical action. Suppose that I know that you have stolen my watch and are now wearing it, but also I know that you do not know that I know it. Then I will try to quickly grab it back. But if I think that you may know that I know

² At least, according to the little globe standing on my desk as I write this.

³ These are normal cooperative questions. Neither condition holds when a teacher asks a didactical question to students in class – or in *games*, where questions may serve to mislead an opponent.

⁴ 'Preconditions' and 'postconditions' are notions from program analysis in computer science.

you have it (note that this involves 3 iterations!), I must retrieve my stolen watch in some more sophisticated manner. Thus both communication and genuine physical action involve careful handling of knowledge assertions of various shades.

Epistemic logic The preconditions and postconditions of the preceding episode can be written in standard epistemic logic, which is an excellent formalism for displaying knowledge about facts and about other people's information. For instance, the question indicated that the following was true, with K for the epistemic modal operator "knows that" and $\langle \rangle$ for the dual modality "holds it possible that":

$$\neg K_I P \ \& \ \neg K_I \neg P \\ \langle I \rangle (K_{you} P \vee K_{you} \neg P)$$

Moreover, after the answer has been given, the following are true:


$$K_{you} P, K_I P, K_{you} K_I P, \text{ etc.} \\ \text{all the way to common knowledge } C_{\{you, I\}} P.$$

More precisely, these epistemic formulas refer to the usual semantic models

$$M = (W, \{\sim_j\}_j, V)$$

for epistemic logic, consisting of a set W of possible worlds (the ways the actual world might be), accessibility relations \sim_j for each agent j , and a valuation function V giving each proposition letter a truth value in each world. A formula $K_j P$ is then true at a world s if P is true in all worlds t with $s \sim_j t$. The much stronger formula $C_G P$ is true if P holds at all worlds that are reachable from s by any finite chain $s \sim_j t \sim_k \dots$ where the relations may be for arbitrary agents. For convenience, one often assumes that the \sim_j are equivalence relations, making the logic a poly-modal $S5$ system in a language with a common knowledge operator. But similar ideas will work for much weaker logics, modeling agents' belief instead of knowledge.


Dynamics: changing information states But there is more to be done. An explicit account of what happens in a question-answer episode does not just record statements that are true before and after. It will also model the change of information state directly, in terms of transitions between states in some information space:




To make this precise, we need to 'dynamify' traditional epistemic logic. First, for the successive information states in a conversation, we can take epistemic models (M, s)

as above with a designated actual world s for the real state of affairs. These models describe 'snapshots' of the current information available to the agents. Normally, we keep one such model M fixed, and evaluate formulas ϕ as to their truth or falsity in some world. But now, we must look at sequences of such models, because speech acts of assertion change them according to some update rule.

E.g., in a simple question/answer scenario, the initial model might be as follows, indicating that Questioner (Q) does not know whether P , but Answerer (A) does:




The black dot stands for the actual world. (In this particular model, by the rules of epistemic logic, Q even knows that A knows the answer – though this is not strictly required for asking a genuine question.) Next, A 's answer triggers an *update* of this information model, eliminating the option *not-P*, to yield the one-point diagram




At this stage, P has become common knowledge between Q and A .

The general dynamics here is as follows. *Public announcement* $\phi!$ of a true proposition ϕ eliminates all those worlds from the current model which fail to satisfy ϕ :



With larger epistemic models, world elimination can acquire striking effects.


Games Card games are nice examples, with non-trivial information flow even in simple cases. Let three players $1, 2, 3$ draw a card from 'red', 'white', 'blue', with an actual distribution rbw . Each sees only his own card. The epistemic model is this



The diagram says the following. Though they are in *rw**b***, no player knows this. As they ponder their group situation, they must take into account all 6 worlds. Now

1 says truly: “*I do not have the blue card*”.

What do players know about the cards after this event? Solving this in words can be complicated, but here is the correct update, removing the two worlds starting with *b* :



This shows at once that 2 knows the distribution, 3 knows that 2 knows, and 1 knows only that 2 or 3 knows. But, e.g., it is not common knowledge that 2 knows! For, 1 thinks it possible that 2 has the blue card, in which case the first assertion would not have helped her. The diagram also predicts the effects of further assertions. E.g., if 3 now were to say truly “*I still don't know*”, only the left-most worlds would remain, and 2 would find out the correct distribution.

More general update Models like this clarify, e.g., the famous Muddy Children puzzle and other scenarios, as shown in Fagin, Halpern, Moses & Vardi 1995. A simple exposition of the ideas and resulting general questions is found in van Benthem 2001. Such scenarios have been the starting point for a whole line of research on update mechanisms for more sophisticated forms of communication, including hiding, forgetting, or cheating. These may mix public and private information (as happens with security protocols on the Internet), where agents may even become systematically misinformed. The best current system is *product update* for states with actions: see Baltag-Moss-Solecki 1998, van Ditmarsch 2000. ⁵

3 Epistemic process logics

Dynamic logic The preceding dynamification still has no explicit calculus for defining update actions and reasoning about them. A truly two-level static-dynamic system implementing the Dynamic Turn imports an idea from computer science, viz. the coexistence of propositions and action expressions in so-called *dynamic logics*.

These can describe conditions true in states resulting from performing actions:

⁵ In general product update, epistemic models may also grow in size, as a conversation or a game proceeds, and there may be no straightforward descent to common knowledge of the actual world!

$[a]\phi$ ϕ holds after every successful execution of action a

In the same vein, we can now state epistemic effects of communication, such as

$[A!]K_j\phi$ after a true public announcement of A , j knows that ϕ

This combined language mixes modalities from dynamic logic with epistemic modalities. Their order records the interaction of preconditions and postconditions. For instance, here is a simple statement, that may seem obvious:

$[A!]C_G A$ public announcement leads to common knowledge

We will see later how plausible this is as a general logical law of communication. As another illustration, here is a valid principle in the obvious semantics relating knowledge achieved after a public announcement to what agents know beforehand:

$[A!]K_j\phi \leftrightarrow (A \rightarrow K_j(A \rightarrow [A!]\phi))$

This says that knowledge of ϕ afterwards corresponds to knowledge of a suitably relativized version of ϕ beforehand. This is just one law for reasoning about communication in a complete system of dynamic-epistemic logic for public announcement which is known to be axiomatizable and decidable. This seems the simplest logical calculus of communication.⁶ More sophisticated systems exist for more complex product updates. Thus, dynamic-epistemic logic promises a more systematic logical taxonomy and understanding of general communication.

Analyzing speech acts All this takes a new look at old issues in philosophy. Consider *complete epistemic specifications* in speech act theory. Say, what do we learn from a public announcement of ϕ ? The above 'learning principle' suggested it always produces *common knowledge* of ϕ . But this is false! E.g., if A had answered

ϕ "You don't know it, but P ",

this would have been true, the same update would have occurred, but the assertion ϕ would become false by the very update, because Q now knows that P ! Philosophers will recognize Moore's Paradox here, now as an issue in dynamic epistemic logic.⁷

⁶ This calculus consist of basic epistemic logic plus simple reduction axioms decomposing postconditions recursively. But there are subtleties, as the reduction axiom for common knowledge after an announcement requires enriching the static base language to a richer epistemic logic with an operator $C_G(A, \phi)$ of *relativized common knowledge* within the set of worlds satisfying A . Thus, a dynamic superstructure may also suggest modifications of its static base structure.

⁷ The technical question which forms of epistemic assertion do produce common knowledge when announced is still open. A connection with the 'Fitch paradox' is explored in van Benthem 2003.

Thus, update logics in the Dynamic Turn take up old issues with new techniques. Indeed, even a simple formula like $[A!]K_j\phi$ encodes ideas from linguistic speech acts, philosophical epistemology, and program logics from computer science.

Program structure But the analogy between communicative actions and programs goes still further. Computer programs are typically constructed from basic actions hardwired in a computer using software *constructions*, such as

| | |
|--------------------|-----------------------|
| composition | $S;T,$ |
| conditional choice | $IF P THEN S ELSE T,$ |
| guarded iteration | $WHILE P DO S.$ |

Especially, the latter structure is typical for computation, where we may not be able to tell beforehand how often the computer has to repeat some instruction. But this analogy persists for communication. A public announcement is a basic instruction, which modifies an information state in a way that is hardwired into our social conventions, or even our brains. But on top of that, there is 'communicative software'. We can give people complex instructions like

"First ask how she is doing, and then state your request",
or "If the teacher asks A , then say B , else say C ".

And even iterations occur. Thus we can think of conversation as a sort of imperative programming, where the 'machines' are the social settings that we influence.

A nice concrete example of iteration occurs in the following well-known puzzle :

Muddy Children *After playing outside, two of three children have mud on their foreheads. They all see the others, but not themselves, so they do not know their own status. Now their Father comes and says: "At least one of you is dirty". He then asks: "Does anyone know if he is dirty?" The children answer truthfully. As this question–answer episode repeats, what will happen?*

Nobody knows in the first round. But upon seeing this, the muddy children will both know in the second round, as each of them can argue as follows:

"If I were clean, the one dirty child I see would have seen only clean children around her, and so she would have known that she was dirty at once. But she did not. So I must be dirty, too!" This reasoning is symmetric for both muddy children – so both know in the second round. The third child knows it is clean one round later, after they announced that.

The puzzle is easily generalized to other numbers of clean and dirty children. It involves an iteration "keep stating your ignorance until you know", which may be repeated any number of times, depending on the composition of the group.

To analyze this puzzle completely, we need a dynamic-epistemic logic which allows for complex actions π in assertions $[\pi]\phi$. Axioms for such constructions are known from computer science, such as the program reduction law $[S;T]\phi \leftrightarrow [S][T]\phi$.

General logic of communication There is much more to logic of communication. Van Benthem 2002 explores new sorts of issue, such as "Tell All":

How to describe the best possible outcome that can be achieved
by a group of agents that are out to inform each other optimally?

A very rich source is Baltag-Moss-Solecki 2003. E.g., it contains the result that the dynamic-epistemic logic of public update with Kleene iterations^{*} of assertions added is *undecidable*. Thus, the background logic of puzzles like Muddy Children⁸ is rich enough to encode significant mathematical problems! This is one of many 'complexity thresholds' in the spectrum of human communicative activities.

4 Revising beliefs and expectations

From update to revision Information update is just one cognitive activity that we engage in. Another key source for the Dynamic Turn is the theory of belief revision (Gärdenfors & Rott 1995), which highlights the interplay of three processes:

- (a) information *update*, adding certain propositions
- (b) information *contraction* leaving out certain propositions
- (c) *belief revision* changing prior beliefs to accommodate new ones.

Belief revision theory proposes representations of information states plus an account of the revision process via basic postulates, and optional ones reflecting more conservative or more radical policies for changing one's beliefs. Moreover, there is not just transformation of propositional information. One can also change agents' plausibility orderings between worlds, or their preferences, or indeed any parameter in logical semantics that admits of meaningful variation over time.⁹

It is still an open issue how to best combine these ideas with epistemic update logics as proposed above. One way of doing this works by dynamifying *conditional logic*, the study of implications $A \Rightarrow B$ interpreted as saying that

B is true in all most preferred or most plausible worlds satisfying A .

⁸ Or those late-night alcoholic conversations where we tend to repeat ourselves...


⁹ Even the language itself, encoding the conceptual framework, may be subject to revision.

Dynamic actions then involve changes in plausibility orderings, in addition to just removing worlds or uncertainty links. Some relevant publications are Veltman 1996, Aucher 2003. Eventually, something like must be done, even when modelling very simple scenarios in understanding conversation, as we shall see with games below.

Learning theory Evidently, people have various strategies for revising theories, or just their ordinary opinions. Belief revision theory is not out-and-out dynamics yet, as those processes themselves are not manipulated as first-class citizens in the calculus. An example of the latter move is the modern theory of *learning mechanisms*, merging ideas from the philosophy of science, mathematical topology, and computer science. Hendricks 2002 makes an extensive plea for the broad epistemological relevance of this move. Update, revision, and learning form a coherent family of issues, going upward in complexity and range from short-term to long-term cognitive behaviour. Van Benthem 2003A discusses the whole picture in some more detail.

5 Goals, strategies, and games

The broader setting of communication Public announcements are building blocks for arguments or conversations. But in those larger settings, we do not just ask *what* people are telling us, but also *why*. What are my partners trying to achieve, and for that matter, what are my own goals in choosing what to say or ask? E.g., consider the following scenario. *A* has to choose an action, and then the turn passes to *E*, who can choose an ending from *x*, *y* or from *z*, *u*. The first to know where the story ends wins a prize. Players have made up their mind what to do in each case.



Now suppose *A* asks *E*:


"What are you going to play on the left? "

This is a genuine question, as *A* does not know, and *A* even knows that *E* knows the answer. But there seems to be more information in this question than just these preconditions from the earlier epistemic update logics. For, *why* would *A* ask this? It only seems to make sense to know this if he is going to play 'left'. But the latter information would tell *E* exactly what is going to happen, since she already knows her own move, and so she can win the prize even before answering the question. So, *is* it justified for *E* conclude this? It depends on what sort of conversational partner she takes *A* to be: rational, stupid, etc. Moreover, pay-offs matter. Suppose that

announcing the wrong solution makes the prize go the other player. Then *A* might have asked the question in order to fool *E* into making a wrong announcement...

The considerations in this simple example all point toward strategic interaction in rounds of conversation, and planning for various future contingencies. A good paradigm extending update logics for this broader purpose is found in *game theory*. Games are a model for a group of agents trying to achieve certain goals through interaction. They involve two new notions compared with what we had before: agents' *preferences* among possible outcome states, and their longer-term *strategies* providing successive responses to the others' actions over time. In particular, strategies take us from the micro-level to a description of longer-term behaviour.

Conversation games Consider two people who are not equally informed. I do not know if we are in Holland (P) or not ($\neg P$), and if the year is 2004 (T) or not ($\neg T$). You know that I do not know the place, but think that I might know about the time. But I do know whether we are together for a good reason (R), whereas you don't. In fact, we are in Holland in 2004, and indeed for a good reason. Here is a concrete epistemic model for this situation, with the black dot indicating the actual world:



Now we want to discover the true situation – and the one who finds out first wins. I can ask you a question first, and it needs to be genuine: in particular, I do not know its answer. Then you can ask, and so on. At each stage, someone who knows the precise facts can announce this, and wins. (There might be a draw if both announce simultaneously). Now I can clearly ask better or worse questions.


Suppose I ask you about the time. Then you learn that I do not know if T holds, which eliminates the two bottommost worlds. But then you know the facts (as we are really in the black world with P, T, R , and there are no uncertainty lines from there left for you), and so you win at once. Therefore, I should rather ask about the place (P). This gives away no information which you don't already have, because it is compatible with all four worlds. But your positive answer eliminates the two right-most worlds, after which I know the facts and you still do not know about R .

This choice between better and worse questions (or things to say in general) is the beginning of a game dynamics of conversation, where players must select questions so as to profit most while leaving their opponents in the dark as much as possible.

Whether this can be done depends not just on the epistemic model, but also on the schedule of questions and answers. (Clearly, you could win the above game if you could start.) But matters of timing, too, are very much a feature of real games.¹⁰

Game theory and logic Game theory studies sets of strategies that reflect optimal long-term behaviour for players, according to Nash equilibria or other plausible notions of game solution, where players do not gain by deviating from their strategy given what the others have chosen. These notions apply to concrete games of any sort (economics, war, amusement), but also to generic games for social activities of language use or logical reasoning. Much of the mathematics of the field is about finding equilibria and their properties, for players having more or less information at their disposal. There are many techniques for this, from leaf-to-root analysis of game trees to much more complex results (Osborne & Rubinstein 1994). Despite obvious differences in scope and aims, game theory and logic also have natural connections. Van Benthem 1999–2003 presents a panorama of games inside logic that model semantic evaluation, argumentation and other activities. This idea of *logic games* may be extended to uses of games as a model for interactive computation. The result is a merge of logical calculi for programs and logical calculi for defining games and studying their computational properties (Parikh 1985, Abramsky 1998).

Game logics The other side of the contact between logic and game theory are logical investigations of deliberation, decision and action by players. For general games, this involves an abstraction step as compared with the earlier update logics. We have a complete game tree of all possible moves, with players' turns indicated at the nodes, and we wish to analyze which particular sequence(s) of actions will be taken by agents who can reflect on their strategies. For a simple example, consider the following three game trees, with respective values for *A*, *E* indicated at the end:



Each of these games is a model for a modal logic of its basic actions – in this case, 'left', 'right'. Game structure and strategies may then be formulated in standard terms.

¹⁰ There is much more to the issue of asking best questions in a conversational setting, and real conversation games might easily involve more probabilistic considerations. Cf. van Rooy 2003.

E.g., out of her 4 possible strategies (maps from turns to moves), the best strategy for *E* in the first game (a) is to do the opposite of what *A* has done:

"if he has gone left, go right – if he has gone right, go left" #

This strategy is a simple program that can be studied in a standard dynamic logic (van Benthem 2001B). Interpreting the value '1' as 'winning', we see that this is a *winning strategy* for player *E*: by following it, she wins no matter what *A* plays. Most logic games go no further than this notion. But in the middle game (b), with finer preferences among outcomes, better predictions can be made. Again *E* will play strategy # at her two turns, assuming she is rational. But given that, *A* will choose left, as it will give him 1/2, as opposed to the 0 on the right. This predicts the unique 'subgame-perfect' Nash equilibrium of this game, which lets *E* play her winning strategy, while *A* plays 'left'. In logical terms, an argument like this involves expressions for values of nodes, perhaps even a full-fledged preference logic.

Finally, the game (c) introduces a new feature, viz. *imperfect information*. At her turn, *E* does not know what move was played by *A*, as indicated by dotted line between the two nodes in the middle. Imperfect information arises happens in many games, e.g. because of restricted powers of observation – as in card games. Such games are models for a joint dynamic-epistemic language with basic actions *a*, *b*, ... corresponding to the moves, and epistemic operators K_j standing for players' knowledge. This language can express special information patterns in games, such as the fact that player *E* does not know which move will make her win:

$$\neg K_E \langle \text{right} \rangle \text{win}_E \wedge \neg K_E \langle \text{left} \rangle \text{win}_E$$

It can also express general laws describing special types of agent. Van Benthem 2001A has typical illustrations of this interface between logical and game-theoretic notions. For instance, players *j* with *perfect recall* of the past history of the game will have an ignorance pattern satisfying this knowledge-action interchange axiom:

$$K_j[a] \phi \rightarrow [a]K_j \phi$$

This is like an earlier axiom for actions of public announcement, which basically related $[A!]K_j\phi$ to $K_j[A!]\phi$.¹¹ This assumed perfect recall for all agents involved. By

¹¹ The axiom for update logic has an *equivalence* between the two operator orders for $[]$ and K . The extra implication reflects a further condition that players never lose ignorance 'spontaneously'.

contrast, players with some finite *bounded memory* will only remember the past up to some fixed 'window', and their behaviour will satisfy different logical laws.¹²

Information update in games In this game setting, the earlier update logic still makes sense. Intuitively, players move through a game tree, as moves are played. That is, at every stage, they learn more about events that took place, while their horizon of possible future developments decreases. First, consider the former.

Imperfect information games as described here encode structural uncertainty about the game, which gets modified systematically by observing moves. For this purpose, one can use the earlier epistemic update mechanisms, as a means of explaining how the dotted uncertainty lines arose in the above pictures. One starts at the root, perhaps with some initial epistemic model M . In general, players have only partial powers of observation for moves as the game unfolds. This may be encoded in an

epistemic *action model* \mathbf{A} consisting of concrete events,
with uncertainty relations between them indicated as for worlds.

E.g., I may observe that you are drawing a card, but for all I know you are either drawing the Queen of Hearts or the King of Spades. Both actions will occur in \mathbf{A} , but there will be an uncertainty line for me between them. Now, successive layers of the game tree arise by computing successive update products in the sense of Section 2:

$M, M \times \mathbf{A}, (M \times \mathbf{A}) \times \mathbf{A}, \text{ etc.}$

Given this special update mechanism, their pattern of dotted lines for the complete game tree will satisfy special requirements (one of them is the above perfect recall), which can be determined precisely. The full story is in van Benthem 2001A.

Managing expectations But information update by observed events is only half of the story of reasoning in games. Even when they know the whole game structure perfectly well, including all past moves, players still play a game with expectations about their own future behaviour and that of others – and that anticipation is also the essence of all human activity. Stable predictions of this sort are indeed the point of the game-theoretic notion of a strategic equilibrium. But expectations can really be of any sort. Perhaps, you suspect that I have a one-bit memory, remembering only the last

¹² For further topics at the interface of logic and game theory: cf. van Benthem 1999-2003, 2003B, on powers of players, structural notions of game equivalence, operations constructing new games and their algebra, and analysis of game-theoretic equilibrium concepts in fixed-point logics. Many other interesting strands are found in Stalnaker 1999, Pauly 2001.

move that was played, so that my behaviour only depends on what you did just before. Now, as moves are played, some of those expectations may be refuted. Say, *E* was expecting *A* to start by playing 'left' in game tree (b), but instead, *A* plays 'right'. In this case, expectations about the other player need to be revised, and we enter the area of *belief revision*, as briefly considered in Section 4. A proper account of the two sorts of mechanism combined: information update and expectation management, seems just around the corner in current logical studies of games. ¹³


6 Temporal evolution

We started with the logic of single steps in communication, and the corresponding updates of information states for groups of agents. Then we moved to longer-term behaviour in games, where players want to achieve goals through finite sequences of actions, responding to what others do. This requires stronger logics, including reasoning about strategies. But eventually, communication and games lie embedded in an even larger temporal setting of human practices over time. We briefly consider some aspects of this more general perspective here.

Finite versus infinite Games seem finite terminating activities, like proofs or talks. But computer science also studies useful infinite processes, like the running of an operating system allowing many special-purpose programs to perform finite tasks. The same dichotomy occurs with cognitive processes in the Dynamic Turn. Some activities are meant to terminate, others provide the operating system for these. Examples of the latter are logical proof systems, or Grice's well-known maxims in running conversation. Likewise, game theory also studies infinite games and players' behaviour in them, such as repeated Prisoner's Dilemma in social co-operation.

Temporal logic To study these phenomena, the above logical systems need to be embedded in a temporal system, allowing for discussion of epistemic multi-agent *protocols* over time, and other long-run notions. E.g., a protocol may encode general regularities relevant to communication, like my knowing that you speak the truth only half of the time. The usual picture here is the familiar Tree of forking paths:

¹³ Abstract games and update by observing moves still relate to concrete conversation in many ways. Suppose that players have already chosen their strategies in a game tree, but the art is now to find out where the game will end. The player to know this first gets a prize. This is again an imperfect information game where information can be revealed through statements and questions. In particular, just failure to claim the prize, implying ignorance of where the game will end, can convey useful information, as it may rule out certain moves. See van Benthem 2004 for details.



This temporal universe, with epistemic structure added, seems the right stage for putting together single update steps, finite game-like activities, and relevant infinite processes running in the background.^{14 15}

Dynamic logic and dynamical systems In modern game theory, this infinite setting leads to *evolutionary* considerations, where social behaviour is analyzed in terms of equilibrium features of infinite dynamical systems, often with a state-transition function of some biological sort (cf. Osborne & Rubinstein 1994). This is a very different mathematical style of thinking about long-term behaviour, where stable structures emerge as statistical properties of populations. Nevertheless, there is an interesting challenge how to interface this with the logical approach of this paper .

7 From analysis to synthesis

The final relevant aspect of the Dynamic Turn that we wish to advertize lies on a different dimension. Most of logic is about analyzing and understanding *given behaviour*, of language users or reasoners. But of equal interest is the undeniable fact that logical investigations also create *new ways* of expressing ourselves, reasoning, and computation. Well-known examples in computer science are formal specification languages, or logic programs. But the same move from analysis to *design* makes sense in general cognition. For instance, any working voting procedure is a designed piece of 'social software' (Parikh 2002), where we create a new pattern of behaviour

¹⁴ Cf. the computer run model of Fagin et al. 1995, the infinite games of Abramsky 1996, the protocol model for messages in Parikh & Ramanujam 2003, the universe for learning mechanisms in Kelly 1996, or the philosophical theory of deliberation and action in Belnap et al. 2001.

¹⁵ Uncertainty between finite sequences of actions in these models naturally generalizes earlier notions from dynamic epistemic logic. E.g., in the Tree setting, epistemic product update says that two sequences X, Y are indistinguishable if they are of equal length, and all their matching members X_i, Y_i are indistinguishable. By contrast, systems based on finite automata for their memory will only require indistinguishability up to some fixed finite set of preceding positions.

for beneficial purposes. Analyzing these may be hard by itself¹⁶, but designing better ones is even more of a challenge! And the same is true for the stream of new games that appear in this world, and which are assimilated into our repertoire of human activities.¹⁷ The systems of this paper can also be used in this more 'activist' mode, as a way of designing behaviour, and changing the world. An example from the original update logics is the 'Moscow Puzzle' (van Ditmarsch 2002):

"A gets 1 card, B and C get 3 cards each. What should B, C tell each other, in A's hearing so that they find out the distribution, while A does not?"

Going beyond such puzzles, one might even think about creating new games, and other practices, using dynamic logics as a means of suggesting possibilities, and as a way of keeping our thinking straight about the intended effects.

8 Conclusion

This paper has sketched a broad view of logic in a setting of communication, computation, and cognition. This merges the traditional analysis of reasoning and definition with that of revising beliefs, planning actions, playing games, and their embedding in longer-term patterns of social behaviour. We gave some examples of how this might be done – but admittedly, most of this is still wishful thinking, rather than solid experience. But then, experience does tell us that wishes may come true...

9 References

- S. Abramsky, 1998, 'From Computation to Interaction, towards a science of information', BCS/IEE Turing Lecture.
- G. Aucher, 2003, *A Joint System of Update Logic and Belief Revision*, Master of Logic Thesis, ILLC University of 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.
- Updated version, 2003, department of cognitive science, Indiana University, Bloomington, and department of computing, Oxford University.

¹⁶ A nice example is the impenetrable selection procedure for the Doges adopted in Venice in 1268, with its vast array of safeguards against family influence and patronage, including many stages of voting plus drawing by lots. Norwich's *History of Venice* (Vintage Books, New York, 1989) says it "must surely rank among the most complicated ever instituted by a civilized state".

¹⁷ Cf. also the study of 'mechanism design' in modern game theory.

- N. Belnap, M. Perloff & M. Xu, 2001, *Facing the Future*, Oxford University Press, Oxford.
- J. van Benthem, 1996, *Exploring Logical Dynamics*, CSLI Publications, Stanford.
- J. van Benthem, 1999 – 2003, *Logic in Games*, electronic lecture notes, ILLC Amsterdam & philosophy Stanford (occasional paper versions).
- J. van Benthem, 2001A, 'Dynamic Epistemic Logic', *Bulletin of Economic Research* 53:4, 219–248 (Proceedings LOFT-4, Torino).
- J. van Benthem, 2001B, 'Extensive Games as Process Models', *Journal of Logic, Language and Information* 11, 289–313.
- J. van Benthem, 2001C, 'Logics for Information Update', *Proceedings TARK VIII (Siena 2001)*, Morgan Kaufmann Publishers, Los Altos, 51–88.
- J. van Benthem, 2002, 'One is a Lonely Number: on the logic of communication', Tech Report PP-2002-27, ILLC Amsterdam. To appear in P. Koepke et al. eds., *Colloquium Logicum, Muenster 2001*.
- J. van Benthem, 2003A, 'Logic and the Dynamics of Information', in L. Floridi, ed., *Minds and Machines* 13:4, 503–519.
- J. van Benthem, 2003B, 'Rational Dynamics and Epistemic Logic in Games', in S. Vannucci, ed., *Logic, Game Theory and Social Choice III*, University of Siena, department of political economy, 19–23.
- J. van Benthem, 2003C, 'What One may Come to Know', Tech Report PP-2003-22, ILLC Amsterdam. To appear in *Analysis*, 2004.
- J. van Benthem, 2004, 'Representing Arbitrary Games through Conversation', manuscript, ILLC Amsterdam.
- H. van Ditmarsch, 2000, *Knowledge Games*, dissertation DS-2000-06, ILLC Amsterdam and University of Groningen.
- H. van Ditmarsch, 2002, 'Keeping Secrets with Public Communication', department of computer science, University of Otago.
- R. Fagin, J. Halpern, Y. Moses & M. Vardi, 1995, *Reasoning about Knowledge*, The MIT Press, Cambridge (Mass.).
- P. Gärdenfors & H. Rott, 1995, 'Belief Revision', in D. M. Gabbay, C. J. Hogger & J. A. Robinson, eds., *Handbook of Logic in Artificial Intelligence and Logic Programming* 4, Oxford University Press, Oxford 1995.
- V. Hendricks, 2002, 'Active Agents', PHILOG Newsletter, Roskilde. In J. van Benthem & R. van Rooij, eds., special issue on Information Theories, *Journal of Logic, Language and Information* 12:4, 469–495.
- M. Osborne & A. Rubinstein, 1994, *A Course in Game Theory*, The MIT Press, Cambridge (Mass.).

- R. Parikh, 1985, 'The Logic of Games and its Applications',
Annals of Discrete Mathematics 24, 111–140.
- R. Parikh, 2002, 'Social Software', *Synthese* 132, 187–211.
- R. Parikh & R. Ramanujam, 2003, 'A Knowledge Based Semantics of Messages', CUNY New York & Chennai, India. In J. van Benthem & R. van Rooy, eds., special issue on Information Theories, *Journal of Logic, Language and Information* 12:4, 453–467.
- M. Pauly, 2001, *Logic for Social Software*, dissertation DS-2001-10, Institute for Logic, Language and Computation, University of Amsterdam.
- R. van Rooy, 2003, "Quality and Quantity of Information Exchange. In J. van Benthem & R. van Rooy, eds., special issue on Information Theories, *Journal of Logic, Language and Information* 12:4, 423–451.
- R. Stalnaker, 1999, 'Extensive and Strategic Form: Games and Models for Games', *Research in Economics* 53:2, 93-291.
- F. Veltman, 1996, 'Defaults in Update Semantics', *Journal of Philosophical Logic* 25, 221–261.