Rational Dynamics and Epistemic Logic in Games

Johan van Benthem, 1 december 2002

Abstract I propose a bare-bones look at epistemic models for games, with a focus on update procedures for reaching equilibrium 'zones'. Connections are given with standard update and fixed-point logics. This is just a 'methods' paper, as readers will want to play with models, uncertainty relations, and announcements different from those used here for the purposes of illustration.

1 Reaching equilibrium

Nash equilibrium describes optimal strategy profiles for players of a game. But it has some well-known problems, where predicted equilibria do not match most plausible behaviour. Examples arise both in strategic games such as Prisoner's Dilemma, and extensive ones such as the Centipede (cf. Binmore 1992, Osborne & Rubinstein 1994). One important line of analysis tries to analyze this situation by identifying suitable epistemic assertions about players' behaviour which imply the usual equilibrium solutions. For this purpose, game theorists have introduced *epistemic models* for a given game, representing what players know and do not know about possible strategy profiles – plus perhaps other features of the game, its players, and its environment. One such type of behaviour is *rationality*: choosing one's best available response to the action profile of the other players. And *common knowledge* of rationality is taken to imply the Nash equilibrium solution concept for strategic games, and suitably formulated, also the backward induction solution for extensive games. This Note takes a closer look at the logic of the implication from rationality to equilibrium. In this first part, we focus on strategic games, where players choose one action each, simultaneously – as in the familiar 2-player game matrices. There may be a second part later.

The main thrust of what follows is a familiar point of 'logical dynamics'. In addition to the standard *static* notion of equilibrium, game-theoretic intuitions seem to be just as much about *dynamic* procedures that produce equilibria. E.g., in Iterated Removal of Dominated Strategies for strategic games, or Backward Induction for extensive games, equilibrium arises as a fixed-point of some process of deliberation. Dynamic intuitions concerning activities of deliberation and communication lie behind much of standard epistemic logic and related themes in game theory – but they are usually left implicit. To make the equilibrium dynamics explicit, we apply ideas from current update logics enriching static epistemic logic with a dynamic logic describing changes in epistemic models.

2 Strategic game forms as epistemic models

2.1 Full game models Consider any strategic game G with a set of players I, and sets of actions A_i for each player $j \in I$. A tuple of actions for each player is a *strategy profile*, which determines a unique outcome, and each player has his own *preference relation* among these outcomes. Now for the epistemic superstructure. The *full model over G* is a multi-S5 epistemic model whose states consist of all strategy profiles, and accessibility $\sim_j$ for player j is defined as agreement of profiles in the j'th coordinate. Thus, a game matrix is such a model

	E	*a*	*b*	*c*
A	*d*	*2, 1*	*1, 3*	*0, 1*
	e	*2, 0*	*1, 2*	*0, 1*
	f	*-1, 0*	*1, 1*	*0, 2*

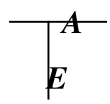

In this case, ***E***'s uncertainty relation runs along columns, and ***A***'s among rows. Thus, we assume that players know their own actions. This assumption may be somewhat curious, and it can be relaxed, but we will retain it in this Note. Also, models for a game often allow more complex 'worlds', with strategy profiles as just one ingredient. We stick with the simplest case to make our main points.

> *Note* We will concentrate on two-player games henceforth, looking at the case of more players only occasionally.

2.2 Epistemic language As usual, these models support assertions in an epistemic language with knowledge modalities for single agents and for groups:

Individual knowledge

$\boldsymbol{M}, s \models K_j\, \phi$ if ϕ is true in all j's epistemic alternatives at s

Common knowledge in a group

$\boldsymbol{M}, s \models C_G\phi$ if ϕ is true at all states that are reachable from s via some finite sequence of uncertainty relations for players in G

The general logic for this language over arbitrary models with equivalence relations $\sim_j$ is multi-S5. But special model classes may validate special principles. E.g., in full 2-player game models, any point can reach any other via a composition of the relations $\sim_A$ and $\sim_E$. Thus, both K_AK_E and K_EK_A both express the universal modality, and the following commutation principle is valid:

$$K_AK_E\, \phi \rightarrow K_EK_A\, \phi$$

2.3 General game models In addition to full game models, there are other models over a game. Certain strategy profiles may be left out from the full space, resulting in submodels that we shall call *general game models*. These represent common knowledge that the missing profiles will never occur. In addition, uncertainty relations might vary, too. E.g., players might not know their own actions, which requires bigger uncertainty relations representing some initial state of deliberation about which action to choose. But in this Note, 'general game models' will be submodels of full models, with the original uncertainty relations:

for any two strategy profiles X, Y, $X \sim_j Y$ iff $X_j = Y_j$

In general game models, the above commutation principle is no longer valid. Indeed, we have total generality – and the logic drops to multi-S5:

Theorem Every multi-S5 model has a bisimulation with a general game model.

Proof See van Benthem, 1996, p. 193. The idea is essentially to unravel arbitrary multi-S5 models into bisimilar tree-like models, and then turn the latter into the required special models whose states are vectors.

3 Rationality

3.1 Best response and Nash equilibrium Rationality is usually taken to involve players' best response to the actions of the others. For each world ω, let $\omega(j)$ be the action performed in that world by player j. Then, the world $\omega(j/a)$ is the strategy profile ω with the component $\omega(j)$ replaced by action a. Now we set:

$$\boldsymbol{M}, \omega \models RAT_j \qquad \textit{iff} \qquad \&_{\{a \in Aj \mid a \neq \omega(j)\}} \ \omega(j/a) \leq_j \omega$$

This is the usual notion of best response. We can now define

Nash Equilibrium $\quad NE \Leftrightarrow \&_j RAT_j \quad ('RAT_G')$

3.2 Rationality, absolutely and relatively best The strength of the preceding best-response assertions RAT_j depends on the comparison set of actions over which we take the conjunction $\&_{\{a \in Aj \mid a \neq \omega(j)\}} \ \omega(j/a) \leq_j \omega$ to the right:

absolute version all actions A_j available to j in the whole game,
relative version all actions that player j still considers possible.

To achieve the latter effect, one considers only those modified strategy profiles $\omega(j/a)$ in the comparison that are *present* in the domain of the epistemic model. This version, too, has its plausibility. Why bother about strategy profiles of which it is common knowledge that they will not occur? Of course, in full models of a game, the absolute and relative versions agree. But in the relative version, a one-world game model with just one profile represented will be 'rational' for all players involved, as there are no relevant alternatives for them to compare with.

Aside An even more radical relative version of 'rationality' would only look at actions which j still considers possible from the current world ω. In the present setting, this would trivialize by our working assumption that j knows his strategy.

In this discussion, the term 'rationality' may be misleading. What RAT_j describes is j's current action *being in fact* a best response at ω. But j need not know that, as she need not know what the other players are doing. In other words,

RAT_j does not imply $K_j\, RAT_j$,

But 'rationality' as the notion is normally used seems to be a feature of agents which they know by introspection. It seems to mean something like

doing the best you can *given what you know*.

Thus, *'best response'* would be a better term for the above definition of RAT_j,, as the latter clearly does not carry connotations about what j knows. Incidentally, for exactly the same reasons, being in a Nash equilibrium of a game model does not imply that players *know* that they are in such a state. In Section 4, we return to the issue of what players do know in such an equilibrium state.

It will be clear where our sympathies lie between the two notions of best response. Nevertheless, in what follows, we mainly look at absolutely best responses. In the latter case, the assertion RAT_j is true or false absolutely at states in a game model, and we may write it as an atomic *proposition letter* B_j. E.g., here is how these proposition letters would occur in our original example:

	E	*a*	*b*	*c*
A	*d*	B_E, B_A	B_A	-
	e	B_E	-	-
	f	-	B_E	B_A

3.3 Epistemic versions of best response The relative version of best response may be described in epistemic terms, though not referring to the player making the response, but to the other(s). E.g., in the 2-player case, it really says that

the other player knows that j's current action
is at most as good for j as j's action at ω

This is because running over all possible actions for j yields precisely the information set for the other player. More generally, consider a total group of players G. Then the above best response for j amounts to the following assertion

> "j's current action is at most as good for j as j's action at ω"
> is *implicit knowledge* at ω *for the rest of the group* $G–\{j\}$.

Implicit knowledge refers to the *intersection* of all uncertainty relations for group members. Thus $I_{G–\{j\}}$ only looks at variants of a world ω where all players other than j play the same action they do in ω. Intuitively, this refers to what the others might come to know by 'pooling' their information. One typical theme in update logics is how groups of agents can convert their implicit knowledge into *common knowledge* through internal communication. This theme will return below.

4 Technical excursion 1: extended modal languages

The basic epistemic language is too poor to express all assertions made in the above, and we need some extensions to formalize what we are doing.

4.1 Hybrid languages Players' preferences make *cross-world comparisons* through worlds while referring back to some fixed world. E.g., one may want to refer to j's action at ω in worlds different from ω. This requires additional operators as in *hybrid modal languages*. E.g., to express that an agent prefers the actual world among all his alternatives, we might want to say something like this:

$$K_j\, P_j\, @$$

with a *preference modality* $P_j\phi$: ϕ holds at some world no less j–preferred than the current one, and the *nominal* $@$ a proposition letter only true at the actual world. These extensions all stay in the decidable 2-variable fragment of first-order logic.

4.2 Polyadic quantifiers But also purely modally, the natural cross-world comparisons in defining best response seem to involve two accessibility relations:

$X \sim_j Y$	$X_j = Y_j$, while all other coordinates may differ
$X \sim^j Y$	$X_k = Y_k$ for all $k \in G$, $k \neq j$, while j's value may differ

Both relations occur in epistemic modelling, and they both validate multi-S5. But it also makes sense to have them *together* in one language, plus modalities

> $<j>$ for the former relation and (in view of the close analogy with the existential quantifier in first-order logic) $\exists j$ for the second.

More generally, instead of the $\exists j$ version, one can introduce accessibility relations *for (sub-)groups G* keeping all coordinates for members of that group fixed – or achieve the same effect by means of *intersections* of identity relations in single coordinates. Technically, the resulting modalities on game models are exactly like single and *polyadic* quantifiers in so-called 'generalized assignment models' for the first-order language: cf. van Benthem 1996, Chapter 9.

5 Common knowledge of rationality and removing strategies

Let us now look at solution procedures for games. The well-known textbook Binmore 1992 justifies the most basic procedures on game matrices as follows. We successively remove dominated strategies, until we reach a fixed-point – and there, we may encounter a Nash equilibrium, or at least some more 'stable' subspace of the total game model. Consider again the above matrix

	E	*a*	*b*	*c*
A	*d*	*2, 1*	*1, 3*	*0, 1*
	e	*2, 0*	*1, 2*	*0, 1*
	f	*-1, 0*	*1, 1*	*0, 2*

Here, one can first remove the dominated right-hand column (***E***'s action *c*), and after that, the bottom row for ***A***'s action *f*, and after that ***E***'s action *b*, and after that, ***A***'s action *e*, arriving at the unique Nash equilibrium *(a, d)* of this game. These successive eliminations are justified by appealing to levels of iterated knowledge of rationality. E.g., ***A*** can be sure that ***E*** will disregard the right-hand column, as ***E*** is rational. And ***E*** can later remove the second column because she knows that ***A*** knows her rationality, and will therefore remove the bottom-most row. And so on.

But this argument seems somehow flawed. If a statement of common knowledge of rationality is true in the model, it is true *at every world* in the model: that is the semantics of a *C*– operator. Therefore, the assertion that players are rational in the sense of playing their objectively best response will also be true everywhere. But obviously, the latter assertion is *false* in the right-most column. Similar objections vitiate appeals to more modest iterations of knowledge about each other. Indeed, in the full game model, as pointed out earlier, composing the uncertainty relations for the two players, in either order, gives the universal relation: so two levels of iterated knowledge is the same as common knowledge.

This is not to deny that the above intuition behind iterated removal of dominated strategies makes good epistemic sense. But its justification must be different from the above. In what follows, we shall look at two possible analyses that fare better.

6 Static version: find the right submodels

Common knowledge of rationality will not normally hold in full game models. But it may hold in *submodels* of these – and indeed, we can determine exactly which general game models satisfy it. Recall that rationality of all agents in a group just said the current world is a Nash equilibrium. Suppose also that values for outcomes are unique: no two outcomes are entirely the same for players.

Fact A submodel of a game model satisfies $C_G\, RAT_G$ if and only if it consists of a single profile which is a Nash equilibrium.

Proof Since RAT_G only holds at Nash equilibrium profiles, the formula can only hold in models each of whose worlds is such an equilibrium. But then, assuming that different worlds have different utility vectors, there can be only one.

This observation is essentially trivial. Moreover, it does not capture the intuition that players might arrive at equilibria through reflection and deliberation. To get at the latter, we need a more *dynamic* view of what went on in the above procedure. The static nature of the Fact shows in the epistemic logic of common knowledge:

Any valid implication $C_G\, RAT_G \rightarrow NE$ implies its modalization
$C_G\, C_G\, RAT_G \rightarrow\; C_G\, NE,$ while also $C_G\, RAT_G \rightarrow\; C_G\, C_G\, RAT_G,$
– and, therefore, we must also have $C_G\, RAT_G \rightarrow\; C_G\, NE.$

That is, if we had that common knowledge everywhere, we also have Nash equilibrium everywhere: nothing needs to be removed – let alone 'iteratively'.

As readers of a first draft have pointed out, the above might reflect no more than the poverty of bare epistemic models. More sophisticated versions arise when additional structure is taken into account – such as players' *expectations*. Also, it might be possible to change the result to a more interesting one involving *minimal modifications* of given game models satisfying common knowledge of rationality. Finally, another line might be this. The common knowledge behind removal of strictly dominated strategies might not concern rationality of things which players *know*, but of actions: things which players *do*. This take on the algorithm is much closer to the dynamic analysis which forms the major theme of this Note.

7 Dynamic version: announcing what one knows about rationality

7.1 Repeating assertions in public Successive elimination of parts of an epistemic model suggests an analogy with *model update* by *public assertions*. Every time a player announces something non-trivial, the common epistemic model shrinks, so that new assertions can be made, and so on. This is the way in which the children arrive at common knowledge of their status in the well-known puzzle of the *Muddy Children*, by repeatedly announcing their ignorance publicly, until no further shrinking of the epistemic model takes place. This amounts to reaching a *fixed-point* of the semantical update procedure triggered by the announcement – a technical theme which will return in Section 9. Now, in the Muddy Children scenario, children are set on their course by a deus ex machina, viz. their Father, who made a helpful initial announcement. But sometimes, just sharing information available inside a group gets things going, and we will consider the additional game import of adding initial announcements later. But for the moment, let's concentrate on the stabilizing effects of saying the same thing over and over again – the way (some) people keep saying "I love you".

What informative things can players tell each other in a game model, if we assume – as is usual in update logics – that they may only say things which they know to be true? One option would be as follows. Since we assume that players know *their own action*, they could announce *that*, and communication ends at once in a single profile. But this is boring, and it would be like a card game ending at once because all players reveal their cards. Let's hasten on to more interesting phenomena, having to do with generic types of assertion.

So, let us look at the dynamics of repeated *rationality assertions*. Can people just tell each other that 'their action is the best possible'? Recall that, even at a Nash equilibrium profile, in the full game model, players need not know that their own action is a best response to the actions of the others – as they do not know what these actions are. Here is a concrete counterexample:

	E	*a*	*b*
A	*c*	*2, 2*	*1, 0*
	d	*0, 1*	*2, 2*

E.g., at the top left state *(c, a),* player ***E*** does not know that her action *a* is best, as she also holds it possible that ***A*** plays *d*, for which *b* would be the better response. Hence, we cannot let players announce that they know their own action to be best:

$K_j B_j$ is not a suitable rationality assertion.

But we can formulate at least two other versions which players *will* know in this diagram, as they hold throughout their information sets, viz. their column or row.

7.2 Announcing weak rationality The first rationality assertion is suggested by removal of strictly dominated rows or columns in a strategic game. It is a weak modalized variant expressing 'possible best response', and lets *j* state that

Weak Rationality (WR) For each alternative action of mine, I think that the current one may be at least as good.

In other words, this says about player j that

there is no alternative action which j knows to be better.

This assertion obviously fails at strictly dominated rows or columns in a game model. Weak Rationality can be formulated in a logical language in the form

$\&_{\{a \mid a \neq \omega(j)\ \&\ \omega(j/a) \in M\}} \langle j \rangle$ 'j's current action is at least as good for j as a'

For instance, in a two-agent matrix, *WR* says for a column x that for each other column y, there is at least one row where $\boldsymbol{E}$'s value in x is at least as good as that in y. Such columns always exist, and generalizing, we have this observation:

Fact Every general game model has worlds where *WR* holds.

In models for multi-S5, true assertions of the logical form of *WR* are known by j, and hence they can be communicated successfully. The main formal features of *WR* are two dependencies: (a) the epistemic possibility is inside the scope of the universal quantifier over actions, and (b) the set of relevant actions for comparison is restricted by the profiles available in the current model $\boldsymbol{M}$. These features make successive announcements of weak rationality just the right driving force for removal of currently (!) strongly dominated rows and columns:

Fact Repeated announcement of Weak Rationality leads to precisely the same stages as iterated removal of dominated strategies.

Proof This is immediate from the definition of *WR* and its explanation.

To formalize *WR* in a suitable formalism, we need a hybrid modal language as in Section 4. E.g., the modal pattern in *WR*-$\boldsymbol{E}$ would be of the technical form

$K_A \downarrow \langle \boldsymbol{E} \rangle\ act_E \geq_E \downarrow act_E$

where act_E refers to ***E***'s action in the current world, while the 'down arrow' $\downarrow$ takes evaluation back to the world where $\downarrow$ was first introduced. Such an enriched formal language seems more germane to game models than just the basic epistemic one. It would be of independent interest to describe its complete logic.

Once again, *WR* is relative to alternative actions in the current model. By contrast, the next, stronger type of rationality assertion can be taken to be absolute.

7.3 ***Strong rationality*** Something quite a bit stronger than *WR* holds in full game models, even those without any Nash equilibria at all. In any 2x2 matrix, each column *x* – standing for some action of ***E*** – has an entry *(a, x)* whose value for ***A*** is highest along the column, where *a* is an action for ***A*** which is that best response to *x*. Throughout the row *a* crossing this maximal entry *(a, x)*:

A holds it possible that his action is a best response to what ***E*** does.

Of course, the same observations may be made for ***E***. This inverts the above order of conjunction and epistemic possibility in Weak Rationality, and we get

Strong Rationality for player j (SR-j)
I think that my current action may be at least as good as all others"
$\langle j \rangle$ **&** $_{a \neq \omega(j)}$ '*j*'s current action is at least as good for *j* as *a*'

Moreover, our argument about matrices has 'best' as an absolute comparison with all alternative actions, whether or not represented in the current model. The two features are independent, and we could also formulate an intermediate notion, say *WSR,* taking best in the relative sense of the current model employed with *WR*.

Fact *SR* implies *WSR*, *WSR* implies *WR*, and no implication can be converted.

Proof The implications are obvious. As to the non-implications, in the following matrix of ***E***-values, the right-most column satisfies *WR*-***E***, but not *SR*-***E***:

1	*3*	*2*
3	*1*	*2*

As to the other non-implication, *WSR*-***E*** holds in the right-hand column viewed as a game model on its own, whereas *SR*-***E*** still would not hold in that model.

We will leave out the variant *WSR* in what follows – but it is one more example of the sort of assertion covered by our general analysis.

Using a notation from Section 3.2, we can view *SR-j* as a formula of the form

$$<j>B_j$$

where the proposition letter B_j expresses at any world ω that j's action at ω is an absolutely best response to the action profile of the others at ω. Once again, this assertion is true or false absolutely in general game models. Finally,

> *Strong Rationality* in general is the conjunction of all
> separate Strong Rationality assertions for each player j.

Strong Rationality need not be satisfied anywhere in general game models. It is easy to draw counterexamples. But we do have the following initial situation.

Fact Each finite full game model has worlds where Strong Rationality holds, and indeed, there are *SR–loops* of the form $s_1 \sim_A s_2 \sim_E \ldots \sim_A s_n \sim_E s_1$ with $s_1 |{=} B_E$, $s_2 |{=} B_A$, $s_3 |{=} B_E$, ...

Proof By an earlier observation about maxima on columns and rows, the following statements hold everywhere in a full game model:

$$<\boldsymbol{E}>B_A, \qquad <\boldsymbol{A}>B_E$$

Therefore, there must be sequences of worlds $B_E \sim_E B_A \sim_A B_E \sim_E B_A \ldots$ – which loop back on themselves as the model is finite. But now, looking *backwards* along such a loop, and using the symmetry of the uncertainty relations, we have the required pattern with $<\boldsymbol{E}>B_E \wedge <\boldsymbol{A}>B_A$ described above.

7.4 Excursion: more than two players The same analysis works for *SR* announcements of rationality by more than two players – including the fixed-point characterization still to come. Here is how the loop argument can be extended.

Fact Each finite full game model has 3-player *SR* loops.

Proof In such models, by the earlier observation, the following assertions are true everywhere. Here modalities $<\boldsymbol{i}, \boldsymbol{j}>$ refer to the accessibility relation $\sim_{\{i, j\}}$ keeping the coordinates for both $\boldsymbol{i}$ and $\boldsymbol{j}$ the same – i.e., the intersection of $\sim_i$ and $\sim_j$:

$$<\boldsymbol{B}, \boldsymbol{C}>B_A, \qquad <\boldsymbol{A}, \boldsymbol{C}>B_B, \qquad <\boldsymbol{A}, \boldsymbol{B}>B_C$$

But then, again, repeating this, we must have loops of the form

$$B_A \sim_{\{A, C\}} B_B \sim_{\{A, B\}} B_C \sim_{\{B, C\}} B_A \sim_{\{A, C\}} \ldots$$

returning to the initial world with B_A. Now, any world in such a loop satisfies

$$<\boldsymbol{A}>B_A \wedge <\boldsymbol{B}>B_B \wedge <\boldsymbol{C}>B_C$$

E.g., if the world itself has B_A, by reflexivity, it satisfies $<\boldsymbol{A}>B_A$. Looking back at its mother B_C via $\sim_{\{B, C\}}$, by symmetry, it satisfies $<\boldsymbol{C}>B_C$. And looking at its grandmother B_B via $\sim_{\{B, C\}}$ and $\sim_{\{A, B\}}$, by transitivity, it also satisfies $<\boldsymbol{B}>B_B$.

7.5 Repeated announcement of strong rationality Announcements of Strong Rationality by all players, either one by one, or simultaneously, drive a dynamic elimination procedure. This will rule out dominated rows and columns in the current model, and sometimes even more. Here is one *SR*-update sequence:

$$\boxed{\begin{array}{ccc} B_E & B_A & - \\ - & - & B_E, B_A \\ B_A & - & B_E \end{array}} \quad \boxed{\begin{array}{cc} B_E & - \\ - & B_E, B_A \\ B_A & B_E \end{array}} \quad \boxed{\begin{array}{cc} - & B_E, B_A \\ B_A & B_E \end{array}} \quad \boxed{\begin{array}{c} B_E, B_A \\ B_E \end{array}} \quad \boxed{\begin{array}{c} B_E, B_A \end{array}}$$

In this particular sequence, a one-world Nash equilibrium model is reached at the end, where further update with *SR* has no effect. But in the following sequence, successive announcement of *SR* gets stuck in a larger zone with a '4-cycle':

$$\boxed{\begin{array}{ccc} B_E & - & B_A \\ - & B_A & B_E \\ B_A & - & B_E \end{array}} \qquad \boxed{\begin{array}{cc} B_E & B_A \\ - & B_E \\ B_A & B_E \end{array}} \qquad \boxed{\begin{array}{cc} B_E & B_A \\ & \\ B_A & B_E \end{array}} \qquad \boxed{\begin{array}{cc} B_E & B_A \\ & \\ B_A & B_E \end{array}}$$

The fact that repeated announcements of *SR* may have new update effects has the following explanation. Public announcements of epistemic formulas of the form $<j>B_j$, either sequential or simultaneous, may have self-invalidating effects known from the literature on update logics. By itself, a separate statement announcing that *j* is strongly rational leaves only states where *SR-j* is true: i.e., *j*'s rationality becomes *common knowledge*. But this announcement may eliminate worlds from the model, thereby invalidating the analogous *SR* statement for other players, as their existential modalities now lack a witness. This is precisely what happened in the above update sequences, including that for the original 3x3 matrix. For the same reason, simultaneous announcements of everyone's rationality need not result in common knowledge of this joint rationality.

On the other hand, in any model during the update procedure, worlds in an *SR loop* keep satisfying Strong Rationality, and will never be eliminated. Thus, in finite models, the procedure stops when all worlds lie on such loops, and Strong Rationality for the whole group has become common knowledge. In particular, Nash equilibria, if present in the game, survive into the fixed-point, being *SR*-loops of length 0. But the above showed a non-Nash example which survives:

$$\begin{tikzcd} B_A \arrow[r, dash, "\boldsymbol{A}"] \arrow[d, dash, "\boldsymbol{E}"'] & B_E \arrow[d, dash, "\boldsymbol{E}"] \\ B_E \arrow[r, dash, "\boldsymbol{A}"'] & B_A \end{tikzcd}$$

Even when Nash equilibria exist, we cannot hope to get just these, though, as the above description also allows for states that have enough links to Nash equilibria, as in the earlier example of 'Battle of the Sexes', which is also an SR-loop:

	E	*a*	*b*
A	*c*	*2, 2*	*1, 0*
	d	*0, 1*	*2,2*

More generally, we say that a statement

> ϕ is *self-fulfilling* on a model $\boldsymbol{M}$ if ϕ has become common knowledge at the first fixed-point of its repeated update sequence computed on $\boldsymbol{M}$.

Fact Strong Rationality is self-fulfilling on finite full game models.

Proof All finite full game models had *SR*-loops, which survive all updates.

On infinite game models, things are more complicated, as there need not be best profiles on rows or columns. And even if there are, the greatest fixed-point may be empty. E.g., suppose that the best-response pattern runs diagonally as follows:

B_E	–	–	–
B_A	B_E	–	–
–	B_A	B_E	–
–	–		

Then, every sequence $B_E \sim_A B_A \sim_E B_E \sim_A B_A \ldots$ must break off at the top.

Next, we can be more precise about the technical mechanism of the *SR* procedure!

8 Elimination algorithms and greatest fixed-points

We can also analyze iterated removal of dominated rows or columns as just another algorithm. What makes it a bit unusual, perhaps, is that a solution is not computed bottom-up from safe instances, but rather top-down, by first overshooting the mark, and then successively discarding options. But even so, this looks very much like a procedure for *computing a fixed-point*. As we are not adding but removing, this must be a *greatest fixed-point*, with an associated top-down procedure – rather than a smallest fixed-point corresponding to a bottom up procedure. Indeed, here is the correct description for Strong Rationality:

Fact The stable set of worlds reached via repeated announcement of *SR* is defined inside the original full game model by the greatest fixed-point formula $\nu p \bullet (\langle \boldsymbol{E} \rangle (B_E \wedge p) \wedge \langle \boldsymbol{A} \rangle (B_A \wedge p))$.

Proof Both inclusions are easy to show. First, the set of non-eliminated worlds in the *SR* procedure itself has the required closure properties, and hence it is included in the greatest fixed-point. But also conversely, it is easy to see that each world in that fixed-point can never be eliminated.

More precisely, the top-down approximation sequence for such a formula $\phi(p)$ looks like this – starting from a formula true everywhere in the model:

$$T, \qquad \phi(T), \qquad \phi(\phi(T)), \qquad \dots$$

taking intersections at limit ordinals. It is easy to see the precise correspondence between these successive stages and elimination stages in game matrices.

This is just one instance of a general style of analyzing repeated updates. E.g., iterated announcement of Weak Rationality can be analyzed in the same fashion. But recall that *WR* involved best response assertions with a 'relative' feature, viz. using a conjunction depending on the available actions in the model.

Greatest fixed-point formulas like that for *SR* can be defined in the so-called modal *μ–calculus*. But there is more to iterated announcement than this...

9 Technical excursion 2: update logic and greatest fixed-points

9.1 The announcement limit Consider any statement ϕ in our language. For any epistemic model $\boldsymbol{M}$, we can keep announcing ϕ, retaining just those states where ϕ holds. In any model – after some finite or infinite number of stages – this process must reach a fixed-point, perhaps the empty set. We call this submodel

$$\#(\phi, \boldsymbol{M})$$

Of course, various things may be true in this final stage. The above *self-fulfilling* formulas in $\boldsymbol{M}$ had the property that ϕ becomes common knowledge in $\#(\phi, \boldsymbol{M})$. This feature also holds in one step for all *persistent* formulas ϕ which have the property that, written in the notation of update logic:

$$\phi \to [\phi!]\, C_G\phi$$

More generally, self-fulfilment involves formulas satisfying

$$\phi \to (C_G\phi)^{\#(\phi, \boldsymbol{M})}$$

By contrast, the repeated ignorance assertions of Muddy Children were 'self-defeating', leading to common knowledge of their negation.

9.2 Inflationary fixed-points It is of interest to see which technical fixed-point procedure is at work here. The announcement limit arises by continued application of the following function, taking intersections at limit ordinals:

$$F_{\boldsymbol{M},\phi}\ (X) \;=\; \{s \in X \mid \boldsymbol{M}|X, s \models \phi\}$$

with $\boldsymbol{M}|X$ the restriction of $\boldsymbol{M}$ to the set X

In general, this function F is *not monotone*, and we cannot apply the usual theory of fixed-points. The reason is the well-known fact from update logic that statements ϕ may change their truth value as we pass from $\boldsymbol{M}$ to submodels $\boldsymbol{M} \mid X$. In particular, we do not recompute stages inside one unchanging model, as with formulas $\nu p \bullet \phi(p)$, but in ever smaller models, thereby changing the range of the modal operators or quantifiers all the time. Thus

we mix *fixed-point computation* with *relativization*.

Aside This is connected to the two senses of an ordinary *restriction operator* for quantifiers $QA\ \phi$: do we mean that all A-objects satisfy ϕ *in the whole model*, or that ϕ holds for objects in the submodel consisting of just the A-points?

Despite the non-monotonicity of its update function, iterated announcements form a fixed-point procedure of sorts, as we can analyze it as follows:

Fact The iterated announcement limit is an *inflationary* fixed point.

Proof Take any ϕ, and relativize it to a fresh proposition letter p, yielding

$$(\phi)^p$$

In the latter formula, p need not occur positively (it becomes negative when relativizing positive box modalities or universal quantifiers), and hence an ordinary fixed-point operator is forbidden. For instance, consider

$$(<>[]q)^p \quad = \quad <>(p \wedge [](p \rightarrow q))$$

Now the usual Relativization Lemma says that

$$\boldsymbol{M}, s \models (\phi)^p \quad \text{iff} \quad \boldsymbol{M}/X, s \models \phi$$

Therefore, the above definition of $F_{M,\phi}(X)$ as $\{s \in X \mid \boldsymbol{M}|X, s \models \phi\}$ equals

$$\{s \in \boldsymbol{M} \mid \boldsymbol{M}, s \models (\phi)^p\} \cap X$$

But this computes a greatest *inflationary fixed-point* (Ebbinghaus & Flum 1995) for arbitrary formulas $\psi(p)$, without monotonicity requirements. ■

The fact that the announcement limit cannot be defined in ordinary fixed-point logic recalls Baltag–Solecki–Moss 2002's observation that update logic with iterated announcements, as in $<(A!)^*>C_GA$, is not part of the modal μ–calculus. Perhaps, the best logic here is some inflationary version of the latter system.

9.3 Monotone fixed points after all But then, why did iterated announcement of Strong Rationality work via an ordinary greatest fixed-point operator? The reason is that the above update function $F_{M,\phi}(X)$ *is* indeed monotone for set inclusion with special sorts of formulas.

Fact $F_{M,\phi}(X)$ is monotone for *existential modal formulas*.

Proof For formulas of this special sort, the above relativization $(\phi)^p$ introduces only positive occurrences of p, so that F must be monotone – and therefore, we can do an ordinary greatest fixed $(\phi)^p$-point computation.

Note that this is also the format of the repeated announcements in Muddy Children. We can do a similar analysis for the relative version of Rationality, where the alternative actions considered only ran over those represented in the model via strategy profiles. This amounted to announcing assertions of the form

$$<j> I_{G-\{j\}} \&_k p_k$$

As these *universal* modalities inside are preserved under passing to submodels, the preceding procedure does not change essentially.

9.4 Digression: comparing update sequences The general proposal in this Note is the exploration of different epistemic assertions that can be announced iteratively, and the investigation of their greatest fixed-points as possible 'solution zones' for strategic games. But then, we have the following natural question (Boudewijn de Bruin). What about comparing our variants *WR* and *SR*? In particular,

> Will the end result of repeated announcement of *SR* be a *submodel*
> of the greatest fixed-point for repeated announcement of *WR*,
> given that the latter statement is weaker than the former?

The answer is positive in this case: repeated announcement of *SR* eventually lands in a subzone of iterated removal of strictly dominated strategies.

Fact $\#(SR, \boldsymbol{M}) \subseteq \#(WR, \boldsymbol{M})$

Proof Clearly, by their definitions, *SR* implies *WR* at any world in any general game model. Next, it suffices to compare the approximation stages of the fixed-point computation. We always have that

$$T^{\alpha}{}_{SR}\,(\boldsymbol{M}) \subseteq T^{\alpha}{}_{WR}\,(\boldsymbol{M})$$

The reason for this is the following inclusion

if $X \subseteq Y$, then $F_{M,\,SR}\;(X) \subseteq F_{M,\,WR}\;(Y)$

This is not a general fact, but a consequence of the special form of our assertions. If $\boldsymbol{M}|X,\ s \models SR$ and $s \in X$, then $s \in Y$ and also $\boldsymbol{M}|Y,\ s \models SR$ – by the modal *existential* definition of *SR*, which makes it preserved under model extensions. But then $\boldsymbol{M}|Y,\ s \models WR$, by our valid implication.

But the general situation is more complex. Even when ϕ implies ψ in a model $\boldsymbol{M}$, $\#(\phi, \boldsymbol{M})$ need not be included in $\#(\psi, \boldsymbol{M})$. A counter-example are the formulas

$$\phi = p \wedge (\langle\rangle\neg p \rightarrow \langle\rangle q) \qquad\qquad \psi = \phi \vee (\neg p \wedge \neg q)$$

In the following epistemic model, the update sequence for ϕ stops in one step with the world ***1***, whereas that for ψ runs as follows: $\{\boldsymbol{1}, \boldsymbol{2}, \boldsymbol{3}\}$, $\{\boldsymbol{1}, \boldsymbol{2}\}$, $\{\boldsymbol{2}\}$.

$$\begin{array}{ccccc} \boldsymbol{1} & \text{———} & \boldsymbol{2} & \text{———} & \boldsymbol{3} \\ p, \neg q & & \neg p, \neg q & & \neg p, q \end{array}$$

10 Greatest fixed points in game generally

The above suggests rethinking fixed-points in game analysis. Initially, *smallest fixed-points* seem most appropriate, as these can be used to analyse the mother of all solution procedures, the bottom-up Zermelo colouring algorithm, step by step. But here is another look at that procedure.

Fact The Zermelo algorithm also works as a *greatest fixed point*.

Proof Using the formula

$$\nu p \bullet \; (end \,\&\, win_E) \; \vee (turn_E \wedge \langle \boldsymbol{E} \rangle p) \vee (turn_A \wedge [\boldsymbol{A}]p)$$

will colour every node as a win for ***E*** first – but then, using the universal set as a first approximation, stage by stage, the right colours for ***A*** will appear. ■

More generally, strategies seem like never-ending resources like our doctors which can be tapped in case of need, and then return to their original state. This fits well with greatest fixed-points, and also with treatments of strategies in *co-algebra*. Finally, the latter thinking extends more smoothly to infinite games, where we often give the win to the player being able to 'keep the ball rolling' indefinitely – as in graph games, or Ehrenfeucht games.

11 The moral of Muddy Children in games

Iterated announcement is high-lighted in epistemic logic by the Muddy Children puzzle. That puzzle still has some further suggestions for game analysis. In fact, its typical cube models looks like full 3-agent game models with 2 actions each:

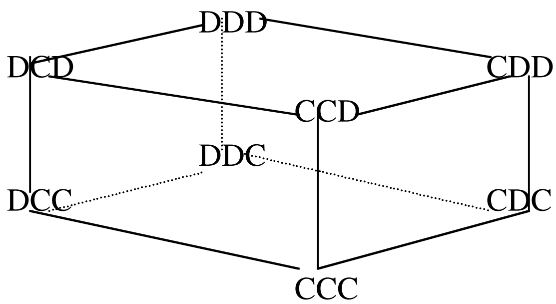

11.1 Initial announcement as an enabling action First, the procedure needs a 'jump-start', viz. Father's initial announcement. Internal communication only reaches the desired goal of common knowledge of the situation after some external help. This may also be the case in games, where some equilibria are only reached after external sources have removed some strategy profiles. Think of the Father saying "You will not be bad boys" – removing some troublesome equilibrium at the start. Introducing such 'asymmetries' seems the only way of breaking into overly large fixed-points in game models. Suitable initial announcements remove worlds, and prune *SR*-loops. But of course, cutting too many worlds can also overshoot the mark, and make *SR* self-defeating in the resulting general game model – in that its iterated fixed-point becomes empty!

[To be added: some explicit examples, including cases with 'relative' rationality assertions.]

11.2 The use of self-defeating announcements Next, notice that the ignorance assertions in Muddy Children are *self-defeating* instead of self-fulfilling: repeated announcements of ignorance eventually results in knowledge. This suggests alternative communication procedures in games than announcing rationality, reaching solution zones by processes such as

> repeatedly announcing players' *worries* that
> non-optimal outcomes are still a live option.

In fact, one can reinterpret the story of the Muddy Children as it stands as such a scenario, with actions 'dirty', 'clean'. Children keep saying "my action might turn out well, or badly" – until the first time they know what is in fact the case.

12 Refinements of the model

One can question the plausibility of the preceding models. E.g., why assume that players know their own action - as this is precisely what they are trying to find out through deliberation? We may need more structure in the models, such as *beliefs* versus knowledge, or *current intentions* which can be modified by the iteration procedure. This is a static point of optimal representation. By contrast, the main point of this Note is a dynamic one of procedure. There is an orthogonal semantic degree of freedom, even on simple epistemic models, namely, the choice of statements to be announced whose updates can be used in pruning these models.

13 Two kinds of update in extensive games

Extensive games involve a division of labour. The above fixed-point procedure had updates triggered by internal communication, with no external information coming in. But in real games, actions take place, and hence there are also updates of players' information through observation of actions. Thus, there are

(a) internal updates by reflection on one's information,

(b) external updates because of observations

The two need to be integrated to get a true account of Backward Induction or other solution concepts. E.g., when some move is played, this should obviously eliminate all strategy profiles from the current model which exclude that move – and even more sophisticated belief revision will be required in general. It seems feasible to have a 'product' of the above elimination procedure for strategy profiles with update by observation. But this is beyond the scope of the present Note.

14 References

A. Baltag, L. Moss & S. Solecki, 2002, 'The Logic of Public Announcements, Common Knowledge and Private Suspicions', updated version, department of cognitive science, Indiana University, Bloomington, and department of computing, Oxford University, Oxford.

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

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

B. de Bruin, 2000 – onward, *Rationality Concepts in Games*, dissertation project, ILLC, University of Amsterdam.

K. Binmore, 1992, *Fun and Game*s, Heath & Company, Lexington (Mass.).

H-D Ebbinghaus & J. Flum, 1995, *Finite Model Theory*, Springer, Berlin.

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

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

R. Stalnaker, 1999, ‘Extensive and Strategic Form: Games and Models for Games’, Research in Economics, 53:293-291.

APPENDIX *Belief update, conditionals, and public announcement*

Here is one standard mechanism for adding belief revision to knowledge update.

Adding belief to epistemic models In the simplest semantics, epistemic models ***M*** get absolute orderings $\leq_j$ of *relative plausibility* among the worlds which agent j cannot epistemically distinguish. Belief is truth in all most plausible alternatives:

$$\boldsymbol{M}, s \models B_j\ \phi \qquad \text{iff} \qquad \text{for all } \leq_j\text{-}best \text{ worlds in } \{t \mid t \sim_j s\}: \boldsymbol{M}, t \models \phi$$

This is less demanding than knowledge: in particular, agents' beliefs can be false when the actual world is not among the $\leq_j$-best ones. The combined logic validates laws like (a) knowledge implies belief, (b) if j believes that ϕ, then j knows that j believes ϕ, and (c) if j does not believe that ϕ, j knows that j does not believe ϕ.

Public announcement and changing belief Now we can look at the systematic effect of announcement – i.e., learning some true fact – on belief statements. This time, the right reduction axiom involves an agent-dependent *conditional* $\Rightarrow_j$:

$$[A!]B_j\ \phi \ \leftrightarrow\ (A\ \Rightarrow_j\ [A!]\ \phi)$$

This assumes that a conditional $A\ \Rightarrow_j\ B$ is read as follows:

> *B* holds in all most *j*-plausible worlds where *A* holds.

On finite models, this is the standard truth condition. Thus, the conditional is a static counterpart of belief update, encoding what agents would believe when updated, through information about their plausibility orderings. This reflects thewell-known Ramsey Test for conditionals. Much ordinary update logic generalizes in this way.

But the richer setting also invites a more complex story of information update. Public announcement of what a truthful speaker knows eliminates worlds as before, perhaps overriding earlier beliefs. But update gets trickier when agents tell each other things they merely believe to be true – as in ordinary conversation. Then, agents may decide whether or not to accept the information as true. Things get even more complex when we add agents' preferences to the models. They might then even decide that a statement is a lie, or at least misleading, made merely to lure us to outcome states which the speaker prefers. Cf. Clark 1996, van Rooy 2002, van Benthem 2003, for an upward ladder of ever more complex ways of taking updates.

Richer structure, even for strategic games Many further features of modality, knowledge and belief can be built into these models. We list a few:

- Conditionals are often used to express *counterfactuals*: "if A had been the case, B would have been the case". This needs extra structure, because it may refer to things we know to be false – requiring worlds outside of our epistemic model.

- In general conditional semantics, plausibility is not a binary but a *ternary* relation: j finds x more plausible than y from the standpoint of s. Compare a ternary geometrical predicate like 'x lies closer to s than y does', as in Lewis spheres.

- Updates might do more than just restrict the ordering $\leq_j$ of $\boldsymbol{M}$ to the set $\boldsymbol{M}|A$. There can be explicit actions *changing plausibility relations*. E.g., dynamic semantics for default conditionals (Spohn, Veltman): announcement of $A \to B$ does not eliminate worlds, but makes $A\&B$-worlds more plausible than $A\&\neg B$-ones.

- Plausibility is often run together with *preference* (also in the AI-literature), but we need a bridge to get from one to the other: e.g., rationality assumptions!

- Analyze interesting philosophical thoughts in a formal manner on these models, such as the Dretske/Nozick formula $K_j\phi \leftrightarrow (\phi \to_j B_j\phi) \wedge (\neg\phi \to_j B_j\neg\phi)$.

Additional references

J. van Benthem, 2003, 'Logic and the Dynamics of Information', in L. Floridi, ed., special issue of *Minds and Machines* on philosophy of information.

H. Clark, 1996, *Using Language*, Cambridge University Press, Cambridge.

R. van Rooy, 2002, 'Quality and Quantity of Information Exchange'. ILLC Amsterdam. To appear in *Journal of Logic, Language and Information*, special issue on formal theories of information.