Multi-Agent Belief Revision with Linked Plausibilities

Jan van Eijck and Floor Sietsma

July 2, 2008

Abstract

In [11] it is shown how propositional dynamic logic (PDL) can be interpreted as a logic of belief revision that extends the logic of communication and change (LCC) given in [7]. This new version of epistemic/doxastic PDL does not impose any constraints on the basic relations and because of this it does not suffer from the drawback of LCC that these constraints may get lost under updates that are admitted by the system. Here, we will impose one constraint, namely that the agent's plausibility relations are linked. Linkedness is a natural extension of local connectedness to the multi-agent case and it assures that we know the agent's preferences between all relevant alternatives. Since the belief updates that are used in [11] may not preserve linkedness, we limit ourselves to a particular kind of belief change that does preserve it.

Our framework has obvious connections to coalition logic [17] and social choice theory [19]. We show how we can use it to model consensus seeking in plenary Dutch meetings. In Dutch meetings, a belief update is done for all agents in the meeting if a majority beliefs the proposition that is under discussion. A special case of these meetings is judgement aggregation, and we apply our framework to the discursive dilemma in this field [15].

1 Introduction

In [15] a problem in judgement aggregation is considered. This problem is the case of three judges a, b, c with a, b agreeing that p, and b, c agreeing that q, so that both p and q command a majority, but $p \wedge q$ does not. The example shows that majority judgement is not closed under logical consequence. To see the relevance of the example for the practice of law, assume that p expresses that the defendant has done action X, and q expresses that the defendant is under a legal obligation not to do X. Then $p \wedge q$ expresses that the defendant has broken his contract not to do X. This is a standard paradox in judgement aggregation called the discursive dilemma or doctrinal paradox.

The discursive dilemma is an example of a situation where multiple agents have different beliefs. We will present an epistemic/doxastic framework that can be used to model such situations, and present a way to update these frameworks with new beliefs. In the above example, this gives a protocol for judgement aggregation.

The relations in our framework can be seen as relations modeling the beliefs of the agents as is the case in the above example. But ofcourse we could also see them as preference relations to connect the framework to the area of social choice theory.

2 Belief Revision Without Constraints

A preference model $\mathbf{M}$ for set of agents *Ag* and set of basic propositions *Prop* is a tuple (W, P, V) where W is a non-empty set of worlds, P is a function that maps each agent a to a relation P_a (the plausibility or preference relation for a), and V is a map from W to $\mathcal{P}(Prop)$ (a map that assigns to each world a *Prop*-valuation). The tuple (W, P) is called the frame of the model. There are no conditions at all on the P_a. Appropriate conditions will be imposed by constructing the operators for belief and knowledge by means of PDL operations.

We fix a PDL style language for talking about preference (or: plausibility). Assume p ranges over a set of basic propositions *Prop* and a over a set of agents *Ag*. Then the language of epistemic/doxastic PDL is given by:

$$\begin{array}{lcl} \phi & ::= & \top \mid p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid [\pi]\phi \\ \pi & ::= & a \mid a\check{} \mid ?\phi \mid \pi_1;\pi_2 \mid \pi_1 \cup \pi_2 \mid \pi^* \end{array}$$

We use *PROG* for the set of program expressions (expressions of the form π) of this language.

This is to be interpreted in the usual PDL manner, with $[\![\pi]\!]^{\mathbf{M}}$ giving the relation that interprets relational expression π in $\mathbf{M} = (W, V, P)$. $[\pi]\phi$ is true in world w of $\mathbf{M}$ if for all v with $(w, v) \in [\![\pi]\!]^{\mathbf{M}}$ it holds that ϕ is true in v. We adopt the usual abbreviations of $\bot$, $\phi_1 \vee \phi_2$, $\phi_1 \rightarrow \phi_2$, $\phi_1 \leftrightarrow \phi_2$ and $\langle\pi\rangle\phi$. The following additional abbreviations allow us to express knowledge, strong belief, conditional belief and plain belief:

knowledge $\sim_a$ abbreviates $(a \cup a\check{})^*$.

strong belief $\geq_a$ abbreviates a^*.

conditional belief $[\rightarrow_a^\phi]\psi$ abbreviates $\langle\sim_a\rangle\phi \rightarrow \langle\sim_a\rangle(\phi \wedge [\geq_a](\phi \rightarrow \psi))$.

plain belief $[\rightarrow_a]\psi$ abbreviates $[\rightarrow_a^\top]\psi$.

Note that it follows from these conventions that $[\rightarrow_a]\phi$ is equivalent to $\langle\sim_a\rangle[\geq_a]\phi$.

The definition of $\rightarrow_a^\phi$ (conditional belief for a, with condition ϕ) is from Boutillier [9] This definition, also used in [4], states that conditional to ϕ, a believes in ψ if either there are no accessible ϕ worlds, or there is an accessible ϕ world in which the belief in $\phi \rightarrow \psi$ is safe. The definition of $\rightarrow_a^\phi$ matches the well-known accessibility relations $\rightarrow_a^P$ for each subset P of the domain, given by:

$$\rightarrow_a^P := \{(x, y) \mid x \sim_a y \wedge y \in \text{MIN}_{\leq_a} P\},$$

where $\text{MIN}_{\leq_a} P$, the set of minimal elements of P under $\leq_a$, is defined as

$$\{s \in P : \forall s' \in P(s' \leq_a s \Rightarrow s \leq_a s')\}.$$

This logic is completely axiomatized by the standard PDL rules and axioms ([18; 14]) plus the following axioms that describe the relation between the basic programs a and $a\check{}$:

$$\vdash \phi \rightarrow [a]\langle a\check{}\rangle\phi \qquad\qquad \vdash \phi \rightarrow [a\check{}]\langle a\rangle\phi$$

Any preference relation P_a can be turned into a pre-order by taking it reflexive transitive closure ${P_a}^*$. The abbreviation for strong belief introduces the $\geq_a$ as names for these pre-orders. The knowledge abbreviation introduces the $\sim_a$ as names for the equivalences given by $(P_a \cup {P_a}\check{})^*$.

If the P_a are well-founded, $\text{MIN}_{\leq_a} P$ will be non-empty for non-empty P. The canonical model construction for PDL yields finite models; since each relation on a finite model is well-founded, there is no need to impose well-foundedness as a relational condition.

This yields a very expressive complete and decidable PDL logic for belief revision, to which we can add mechanisms for belief update and for belief change.

3 Belief Revision with Linked Preference Relations

The proto-preference relations that serve as the basis for construction of a preference pre-order in Section 2 leave something to be desired. Compare an optometrist who collects answers for a number of lenses she tries out on you: "Better or worse?", (change of lens), "Better or worse?" (change of lens), "Better or worse?".... If you reply "worse" after a change of x to y, and "worse" after a change from y to z, she will most probably not bother to collect your reaction to a change from x to z. But what if you answer "better" after the second swap? Then, if she is reasonable, she will try to find out how x compares to z. It does make sense to impose this as a requirement on preference relations.

There are several ways to do this. Recall that we did not impose a requirement of transitivity on the basic plausibility relations. Here is a definition that does not imply transitivity, but yields that the transitive closures of the basic plausibilities are well-behaved.

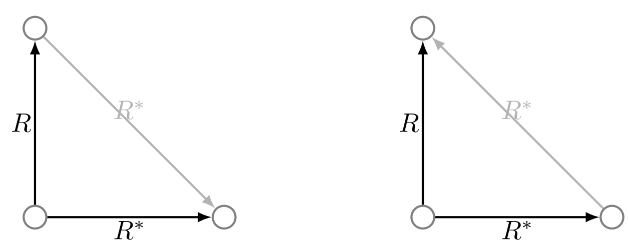

Definition 1. *A binary relation R is* **forward linked** *if the following holds:*

$$\forall x, y, z((xRy \wedge xR^*z) \rightarrow (yR^*z \vee zR^*y)).$$

R is **linked** *if both R and $R\check{}$ are forward linked.*

Note that this is different from the notion of **weak connectedness** (terminology of [12]): a relation R is weakly connected if $\forall x, y, z((xRy \wedge xRz) \rightarrow (yRz \vee y = z \vee zRy))$.

Theorem 1. *R is forward linked iff R^* is weakly connected.*

Proof. The right to left direction is immediate. For the left to right direction, assume R is forward linked. Let wR^*w_1 and wR^*w_2. Then there is an $n \in \mathbb{N}$ with wR^nw_1. We prove the claim by induction on n. If $n = 0$ then $w = w_1$ and $w_1R^*w_2$, and we are done. Otherwise, assume the claim holds for n. We have to show it holds for $n + 1$. Suppose $wR^{n+1}w_1$. Then for some w', $wRw'R^nw_1$. By forward linking of R, either $w'R^*w_2$ or w_2R^*w'. In the first case, use the induction hypothesis to get $w_1R^*w_2$ or $w_2R^*w_1$. In the second case, it follows from w_2R^*w' and $w'R^nw_1$ that $w_2R^*w_1$. □

Starting from relations that are linked, one can upgrade the method from the previous section to construct 'belief revision models' in the style of Grove [13], Board [8], and Baltag and Smets [3; 4] (who call them 'multi-agent plausibility frames').

It is well-known that the following principle characterizes weak connectedness of P_a (cf. [12]):

$$[a]((\phi \wedge [a]\phi) \rightarrow \psi) \vee [a]((\psi \wedge [a]\psi) \rightarrow \phi).$$

The notion of forward linking is characterized by:

$$[a]((\phi \wedge [a^*]\phi) \rightarrow \psi) \vee [a^*]((\psi \wedge [a^*]\psi) \rightarrow \phi). \quad (*)$$

Theorem 2. *Principle (*) holds in a belief revision frame iff P_a is forward linked.*

Proof. Let (W, P) be a frame where P_a is forward linked, and let $\mathbf{M} = (W, P, V)$ be some model based on the frame. We show that (*) holds. Let w be a world in $\mathbf{M}$. Assume $\mathbf{M} \not\models_w [a]((\phi \wedge [a^*]\phi) \rightarrow \psi)$. We have to show: $\mathbf{M} \models_w [a^*]((\psi \wedge [a^*]\psi) \rightarrow \phi)$. From the fact that $\mathbf{M} \not\models_w [a]((\phi \wedge [a^*]\phi) \rightarrow \psi)$, we get that there is a world w_1 with wP_aw_1 and $\mathbf{M} \models_{w_1} \phi \wedge [a^*]\phi \wedge \neg\psi$.

Let w_2 be an arbitrary world with $wP_a^*w_2$. Then by forward linking of P_a, either $w_1P_a^*w_2$ or $w_2P_a^*w_1$. In the first case, it follows from $\mathbf{M} \models_{w_1} [a^*]\phi$ that $\mathbf{M} \models_{w_2} \phi$, and therefore $\mathbf{M} \models_{w_2} (\psi \wedge [a^*]\psi) \rightarrow \phi$. In the second case, it follows from $\mathbf{M} \models_{w_1} \neg\psi$ that $\mathbf{M} \models_{w_2} \neg[a^*]\psi$, and therefore $\mathbf{M} \models_{w_2} (\psi \wedge [a^*]\psi) \rightarrow \phi$. So in both cases, $\mathbf{M} \models_{w_2} (\psi \wedge [a^*]\psi) \rightarrow \phi$, and since w_2 was an arbitrary world with $wP_a^*w_2$, it follows that $\mathbf{M} \models_w [a^*]((\psi \wedge [a^*]\psi) \rightarrow \phi)$.

Next, assume a frame (W, P) where P_a is not forward linked. We will construct a model $\mathbf{M} = (W, P, V)$ and an instance of (*) that does not hold. If P_a is not forward linked, there are w, w_1, w_2 with wP_aw_1, $wP_a^*w_2$, and neither $w_1P_a^*w_2$ nor $w_2P_a^*w_1$. Construct the valuation of $\mathbf{M}$ by setting p true in w_1 and in all worlds w' with $w_1P_a^*w'$, and false everywhere else, and setting q true in w_2 and in all worlds w'' with $w_2P_a^*w''$, and false everywhere else. Note that since not $w_1P_a^*w_2$, p will be false in w_2, and that since not $w_2P_a^*w_1$, q will be false in w_1. So we get $\mathbf{M} \models_{w_1} p \wedge [a^*]p \wedge \neg q$ and $\mathbf{M} \models_{w_2} q \wedge [a^*]q \wedge \neg p$. It follows that

$$\mathbf{M} \models_w \langle a \rangle (p \wedge [a^*]p \wedge \neg q) \wedge \langle a^* \rangle (q \wedge [a^*]q \wedge \neg p),$$

i.e.,

$$\mathbf{M} \not\models_w [a]((p \wedge [a^*]p) \rightarrow q) \vee [a^*]((q \wedge [a^*]q) \rightarrow p),$$

showing that this instance of (*) does not hold in **M**. □

In the multi-agent case there is a further natural constraint. Consider a situation where Alice and Bob have to decide on the chairperson of a program committee. Carol is mediator. Alice says she prefers y to x. Bob counters by saying that he prefers z to x. What should Carol do? Clearly, she should urge *both* of them to compare y and z.

Translating this example to our logic of belief, we want to require that if $x \geq_a y$ and $x \geq_b z$, then either $y \geq_a z$ or $z \geq_a y$ and either $y \geq_b z$ or $z \geq_b y$.

This motivates the following extension of the definition of linkedness to the multi-agent case.

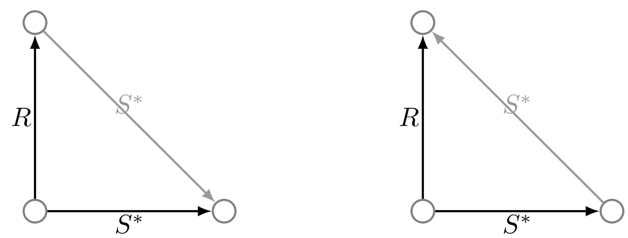

Definition 2. *A set of binary relations* $\mathbf{R}$ *on a domain* W *is* **forward linked** *if for all* R, S *in* $\mathbf{R}$*, if* xRy *and* xS^*z*, then either* yS^*z *or* zS^*y*.* $\mathbf{R}$ *is* **backward linked** *if the set* $\{R^{\smile} \mid R \in \mathbf{R}\}$ *is forward linked.* $\mathbf{R}$ *is* **linked** *if* $\mathbf{R}$ *is both forward and backward linked.*

It follows from Definition 2 that the set $\{R\}$ is forward linked iff R is forward linked according to Definition 1. So Definition 2 gives a natural extension of linking (and of local connectedness) to the multi-agent case.

The following theorem shows that our definition satisfies our motivating requirement that if $x \geq_a y$ and $x \geq_b z$ then either $y \geq_a z$ or $z \geq_a y$:

Theorem 3. *If* R *and* S *are linked then for any* x, y, z*, if* xR^*y *and* xS^*z *then either* yR^*z *or* zR^*y*.*

Proof. Suppose xR^*y and xS^*z. We will prove that for any w on the path from x to z, either wR^*y or yR^*w. This clearly holds for $w = x$. Suppose w is the successor of w' on the path, and the result holds for w'. Suppose $w'R^*y$. Since $w'Sw$ the result holds by forward linking of R and S. Suppose yR^*w'. $w'Sw$ and $w'R^*w'$ so either $w'R^*w$ or wR^*w'. In the first case trivially yR^*w. In the second case the result holds by backward linking of R. □

If R and S are linked relations then common knowledge equals the union of strong common belief and strong reverse common belief:

Theorem 4. *If* R *and* S *are linked, then*

$$(R \cup R^{\smile} \cup S \cup S^{\smile})^* = (R \cup S)^* \cup (R^{\smile} \cup S^{\smile})^*.$$

Proof. The inclusion from right to left is obvious. For the inclusion from left to right, assume $x(R \cup R\breve{} \cup S \cup S\breve{})^* y$. Letting X and Y range over R and S, observe that each $X \circ Y\breve{}^*$ link can be replaced by either a Y^* or a $Y\breve{}^*$ link, and similarly for $X\breve{} \circ Y^*$ links, by linking of R and S. Continuing this process until all one-step links are of the form $R \cup S$ or of the form $R\breve{} \cup S\breve{}$, this yields $x(R \cup S)^* y$ or $x(R\breve{} \cup S\breve{})^* y$. □

The theorem shows that linking of relations simplifies the notion of common knowledge.

The modal characterization of relation linking is given by:

$$[a]((\phi \wedge [b^*]\phi) \rightarrow \psi) \vee [b^*]((\psi \wedge [b^*]\psi) \rightarrow \phi) \qquad \text{(LINK)}$$

Theorem 5. *The set of LINK principles (with a, b ranging over the set of all agents) holds in a belief revision model iff the basic plausibility relations in the model are forward linked.*

Proof. Analogous to the proof of Theorem 2. □

4 Belief Update and Belief Change

For the definition of action models $\mathbf{A}$ and of the update product operation $\otimes$ that combines an epistemic model and an action model, to produce an updated epistemic model, we refer to [2]. In [7] it is shown how extending the PDL language with a extra modality $[\mathbf{A}, e]\phi$ does not change its expressive power. Interpretation of the new modality: $[\mathbf{A}, e]\phi$ is true in w in $\mathbf{M}$ if success of the update of $\mathbf{M}$ with action model $\mathbf{A}$ to $\mathbf{M} \otimes \mathbf{A}$ implies that ϕ is true in (w, e) in $\mathbf{M} \otimes \mathbf{A}$.

Here, we will focus on belief change rather than belief update. In [7], it was proposed to handle factual change by propositional substitution. Relational substitutions were proposed for belief change in [5], and it was shown in [10] that adding relational substitutions for preference change to epistemic PDL makes no difference for expressive power.

A plausibility substitution (or preference substitution) is a map from agents to *PROG* expressions that can be represented by a finite set of bindings

$$\{a_1 \mapsto \pi_1, \ldots, a_n \mapsto \pi_n\}$$

where the a_j are agents, all different, and where the π_i are *PROG* expressions. It is assumed that each a that does not occur in the left hand side of a binding is mapped to a. Call the set $\{a \in Ag \mid \rho(a) \neq a\}$ the *domain* of ρ. If $\mathbf{M} = (W, P, V)$ is a preference model and ρ is a preference substitution, then $\mathbf{M}^\rho$ is the result of changing the preference map P of $\mathbf{M}$ to P^ρ given by:

$$P^\rho(a) \quad := \quad \begin{cases} P_a \text{ for } a \text{ not in the domain of } \rho, \\ [\![\rho(a)]\!]^{\mathbf{M}} \text{ for } a \text{ in the domain of } \rho. \end{cases}$$

Now extend our PDL language with a modality $[\rho]\phi$ for preference change, with the following interpretation:

$$\mathbf{M}, w \models [\rho]\phi \quad :\Longleftrightarrow \quad \mathbf{M}^\rho, w \models \phi.$$

An important thing to note is that since there is there are constraints on the proto-preferences P_a (they are linked), we need to ensure that the belief changing substitutions satisfy these constraints. We leave the general problem of finding a precise characterization of linking-preserving substitutions for future research. Here we focus on a particular kind of belief change. Consider the suggestive upgrade $\sharp_a \phi$ discussed in Van Benthem and Liu [6]:

$$\sharp_a \phi =_{\text{def}} \ ?\phi; a; ?\phi \ \cup ?\neg\phi; a; ?\neg\phi \ \cup ?\neg\phi; a; ?\phi.$$

This is a variation on what is called lexicographic upgrade in the belief revision community (see e.g., [16]). The suggestive upgrade removes all relations from ϕ-worlds to $\neg\phi$-worlds. Clearly, belief revisions with suggestive upgrade do not preserve linking of relations. For consider a case where $wP_a w_1$ and $wP_a^* w_2$ and $w_1 P_a^* w_2$, with ϕ true in w_1 but not in w and w_2. Then suggestive update will remove the a-path from w_1 to w_2:

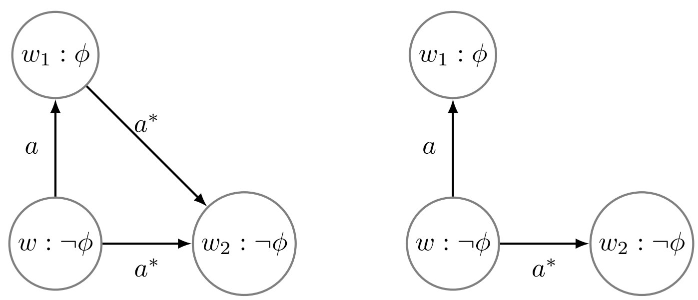

However, if we revise the upgrade procedure so that it adds extra links instead of removing them, as follows, we get a variation that preserves linking:

$$\natural_a \phi =_{\text{def}} \ ?\phi; a^*; ?\phi \ \cup ?\neg\phi; a^*; ?\neg\phi \ \cup ?\neg\phi; (a^* \cup a^{\smile *}); ?\phi.$$

Thus, links from ϕ worlds x to $\neg\phi$ worlds y get reversed, and extra links to x get added to 'support' the new link from y to x. Moreover, ϕ to ϕ links and $\neg\phi$ to $\neg\phi$ links are strengthened to deal with the problem of detours through worlds that assign a different truth value to ϕ.

We can prove the following.

Theorem 6. *If* $\mathbf{M} = (W, P, V)$ *is a belief revision model where* P_a *and* P_b *are linked, and* ϕ *is a PDL formula, then* $[\![\natural_a \phi]\!]^{\mathbf{M}}$ *and* P_b *are also linked.*

Proof. Write a for P_a, b for P_b, and $\natural_a \phi$ for $[\![\natural_a \phi]\!]^{\mathbf{M}}$. First note that for any worlds x and y, if xa^*y then either $x(\natural_a \phi)y$ or $y(\natural_a \phi)x$.

Suppose xby and $x(\natural_a \phi)^* z$. We will show that either wa^*y or ya^*w for all w on the path from x to z. Firstly let $w = x$. Since xby and xa^*x, either xa^*y or ya^*x by linking of a and b. Now let w' be the predecessor of w on the path, so $x(\natural_a \phi)^* w'$ and $w'(\natural_a \phi)w$. Suppose either ya^*w' or $w'a^*y$. Since $w'(\natural_a \phi)w$, either $w'a^*w$ or wa^*w'. If ya^*w' and $w'a^*w$ or wa^*w' and $w'a^*y$, then trivially ya^*w or wa^*y. Suppose $w'a^*y$ and $w'a^*w$. By forward linking of a and Theorem 1, wa^*y or ya^*w. Suppose ya^*w'

and wa^*w'. By backward linking of a and Theorem 1, ya^*w or wa^*y. So then for any w on the path wa^*y or ya^*w, so za^*y or ya^*z, so $z(\natural_a\phi)y$ or $y(\natural_a\phi)z$.

Suppose $x(\natural_a\phi)y$ and xb^*z. Then either xa^*y or ya^*x. In the first case the result follows by Theorem 3. Suppose ya^*x. We will show that for any w on the path from x to z, yb^*w or wb^*y. Firstly let $w = x$. ya^*x and yb^*y so by Theorem 3 the result holds. Suppose w' is the predecessor of w on the path and the result holds for w'. Suppose yb^*w'. Then since $w'bw$, trivially yb^*w. Suppose $w'b^*y$. Then the result holds by linkedness of b. □

Now call a substitution where all bindings are of the form $a \mapsto \natural_a\phi$ a linked substitution. Then we get a complete logic for belief change with linked substitutions, by means of reduction axioms that 'compile out' the belief changes (see [10]):

Theorem 7. *The logic of epistemic preference PDL with belief change modalities for linked substitutions is complete.*

Proof. The preference change effects of $[\rho]$ can be captured by a set of reduction axioms for $[\rho]$ that commute with all sentential language constructs, and that handle formulas of the form $[\rho][\pi]\phi$ by means of reduction axioms of the form

$$[\rho][\pi]\phi \quad\leftrightarrow\quad [F_\rho(\pi)][\rho]\phi,$$

with F_ρ given by:

$$\begin{array}{lll}
F_\rho(a) & := & \begin{cases} \rho(a) & \text{if } a \text{ in the domain of } \rho, \\ a & \text{otherwise,} \end{cases} \\
F_\rho(?\phi) & := & ?[\rho]\phi, \\
F_\rho(\pi_1;\pi_2) & := & F_\rho(\pi_1); F_\rho(\pi_2), \\
F_\rho(\pi_1 \cup \pi_2) & := & F_\rho(\pi_1) \cup F_\rho(\pi_2), \\
F_\rho(\pi^*) & := & (F_\rho(\pi))^*.
\end{array}$$

It is easily checked that these reduction axioms are sound, and that for each formula of the extended language the axioms yield an equivalent formula in which $[\rho]$ occurs with lower complexity, which means that the reduction axioms can be used to translate formulas of the extended language to PDL formulas. Completeness then follows from the completeness of PDL. □

5 Analyzing Plenary Dutch Meetings

A plenary Dutch meeting (Dutch: 'Vergadering') is a simultaneous preference/belief change event where the following happens. Assume an epistemic situation $\mathbf{M}$ with actual world w, and assume proposition ϕ is on the agenda.

- If a majority prefers ϕ to $\neg\phi$, i.e., if

  $$\{i \in Ag \mid \mathbf{M} \models_w [\rightarrow_i]\phi\} > \{i \in Ag \mid \mathbf{M} \models_w [\rightarrow_i]\neg\phi\}$$

  then simultaneous belief change $\{i \mapsto \natural_i\phi \mid i \in Ag\}$ takes place.

- If a majority prefers $\neg\phi$ to ϕ, i.e., if

  $$\{i \in Ag \mid \mathbf{M} \models_w [\rightarrow_i]\phi\} < \{i \in Ag \mid \mathbf{M} \models_w [\rightarrow_i]\neg\phi\}$$

  then simultaneous belief change $\{i \mapsto \natural_i \neg\phi \mid i \in Ag\}$ takes place.

- If there is no majority either way, nothing happens.

In fact, Dutch meetings are procedures for judgement aggregation [15]. Let us return to our example of three judges a, b, c with a, b agreeing that p, and b, c agreeing that q, so that both p and q command a majority, but $p \wedge q$ does not. Using our logic, we can picture the situation as a preference model. We assume that every agent has greater belief in worlds that match her beliefs in more propositions. Then we get the following model:

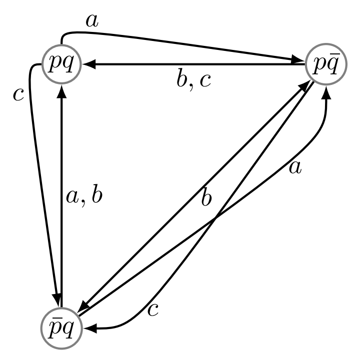

So a has the greatest belief in the world where p and not q hold, but after that she has more belief in a world where p and q both hold than in the world where q and not p hold, because in the first world at least her belief in p is right. Similarly for c. For b, she believes in the world where p and q hold, and values the other worlds equally plausible.

In this model the following formulas hold:

$$[\rightarrow_a]p, [\rightarrow_b]p, [\rightarrow_b]q, [\rightarrow_c]q, [\rightarrow_a]\neg(p \wedge q), [\rightarrow_c]\neg(p \wedge q).$$

This shows that there are majority beliefs in p and in q, but there is also a majority belief in $\neg(p \wedge q)$. If the judges decide to have a Dutch meeting about p, the result will be unanimous belief in p:

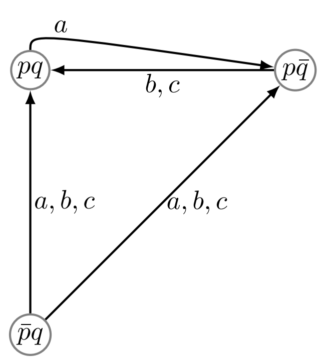

Now if the judges hold a subsequent Dutch meeting about q, the result will be unanimous belief in q:

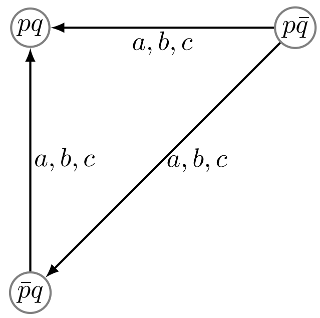

Now the judges unanimously believe in $p \wedge q$, so the defendant will be judged guilty. However, if a Dutch meeting about $p \wedge q$ was held in the first place, the result would be belief in $\neg(p \wedge q)$:

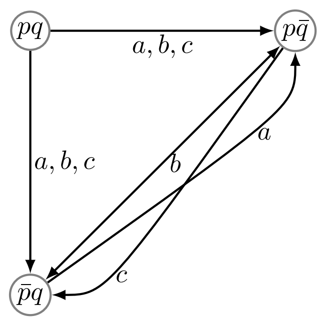

Clearly, in this case the defendant would be acquitted.

Experienced judges are of course familiar with this phenomenon. Procedural discussions about how to decompose a problem, and in which order to discuss the component problems may seem beside the point of a legal issue, but they turn out to be highly relevant for the outcome of the legal deliberations.

6 Further Issues

First and foremost on our research agenda is the problem of finding natural classes of substitutions that preserve relation linking. Next, since our logic provides a general mechanism for simultaneous belief change, it can be used to describe and analyse topics in judgement aggregation, the effects of agenda setting, the effects of subgroups meetings to create general belief, and many further issues of collective rationality. Also, we would like to compare our logic to other proposals to give a modal analysis of judgement aggregation, such as [1]. Putting our logic to work in practical social choice analysis is further work.

References

[1] Thomas Agotnes, Wiebe van der Hoek, and Michael Wooldridge. Reasoning about judgement and preference aggregation. In *The Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 07)*, pages 567–575. IFAAMAS, May 2007.

[2] A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In I. Bilboa, editor, *Proceedings of TARK'98*, pages 43–56, 1998.

[3] A. Baltag and S. Smets. Conditional doxastic models: A qualitative approach to dynamic belief revision. *Electronic Notes in Theoretical Computer Science (ENTCS)*, 165:5–21, 2006.

[4] A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Texts in Logic and Games*. Amsterdam University Press, 2008. To appear.

[5] J. van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 2:129–155, 2007.

[6] J. van Benthem and F. Liu. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logics*, 14(2), 2004.

[7] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.

[8] O. Board. Dynamic interactive epistemology. *Games and Economic Behaviour*, 49:49–80, 2002.

[9] C. Boutilier. Toward a logic of qualitative decision theory. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Proceedings of the 4th International Conference on Principle of Knowledge Representation and Reasoning (KR-94)*, pages 75–86. Morgan Kaufmann, 1994.

[10] Jan van Eijck. Yet more modal logics of preference change and belief revision. In K.R. Apt and R. van Rooij, editors, *New Perspectives on Games and Interaction*, volume 5 of *Texts in Logic and Games*. Amsterdam University Press, 2008.

[11] Jan van Eijck and Yanjing Wang. PDL as a logic of belief revision. CWI, Amsterdam (under submission), 2008.

[12] R. Goldblatt. *Logics of Time and Computation, Second Edition, Revised and Expanded*, volume 7 of *CSLI Lecture Notes*. CSLI, Stanford, 1992 (first edition 1987). Distributed by University of Chicago Press.

[13] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

[14] D. Kozen and R. Parikh. An elementary proof of the completeness of PDL. *Theoretical Computer Science*, 14:113–118, 1981.

[15] C. List and P. Pettit. On the many as one. *Philosophy and Public Affairs*, 33(4):377–390, 2005.

[16] A. C. Nayak. Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 41:353–390, 1994.

[17] Marc Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12:149–166, 2002.

[18] K. Segerberg. A completeness theorem in the modal logic of programs. In T. Traczyck, editor, *Universal Algebra and Applications*, pages 36–46. Polish Science Publications, 1982.

[19] Alan D. Taylor. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press, 2005.