# Understanding decisions, intentions and actions a multi-level meta-modal model

Alan Dix, Devina Ramduny-Ellis, Paul Rayson, Ian Sommerville

Computing Department, Lancaster University,
Lancaster, LA1 4YR, United Kingdom
{alan, devina, rayson,is}@comp.lancs.ac.uk,
http://www.hcibook.com/alan/papers/decisions-2004/

**Abstract.** This paper presents a formal model designed to help us understand decisions. We aim to represent the way in which acts affect the world and how decisions create intentions to act. The model is modal in two ways (i) because we are representing the acts and their consequences and (ii) because we wish to represent multiple alternatives for decisions. The model distinguishes different classes of decision: decisions to act, decisions to believe and decisions that create policy for future decisions. In order to create the model we abstract over different kinds of uncertainty, a technique we believe is also novel.

**Keywords:** decisions, formal model, branching time, modelling uncertainty

## 1. Introduction

This paper comes from work on a project, Tracker, which is attempting to understand the way in which past decisions can be captured and tracked and so used to help future decisions [10].

In this paper we will discuss a semantic model of actions and decisions. We are not creating a computational model for embedding within a system, but more a model to help us understand the nature of decisions and how they interact with the world.

The model we will present is a multi-level meta-modal model. It is modal as we need to be able to represent a decision as a choice between alternatives and hence our underlying model of the world must be a modal one. The meeting itself is a discussion of these alternatives and gives rise to intentions, mental states concerning models of the world. That is the decision making process is modal with respect to the modal world model – hence meta-modal. Finally the decision making itself is, of course, an activity within the world and decisions made in one context may well influence the process of a future meeting (for example voting someone onto a company board). Hence it is a reflexive or multi-level meta-modal model.

In the next section we will discuss some of the literature on meetings and decisions. In section 3 will then look in a lttle more detail at the elements of our model as outlined above. This will then be followed, in section 4, by a formalisation of the meta-modal model looking one by one at individual elements including acts in a modal model of the world, intentions as belief states restricting desired evolutions of the world model and finally the modelling of the decision making body itself.

## 2. Background

The meeting has been described as the most pervasive knowledge event in working life [13]. It is a common arena for brainstorming and deliberating complex issues. There has been considerable work on techniques, notations and computer tools to support decision making meetings. These different cognitive and computational tools embody within them models and theories of decisions and decisions making, which in some cases are very explicit, while others require 'digging' beneath the surface.

Typically, the tools with a more explicit model are normative and hence try to impose models and processes upon their users in order to create 'better' decisions. These include the various types of argumentation or design rational including IBIS [4], DRL [5] and QOC [7]. Many of these have had associated graphical tools developed such as gIBIS [2] to support IBIS and SIBYL for DRL [5].

Others tools however take a more open approach by merely facilitating the human interaction in meetings. This includes meeting support systems such as Xerox PARC's early work on Colab [15], which drove much of the early work on understanding groupware. More recently there has been a focus on implicit recording through smart whiteboards, voice and video capture, for example TeamSpace [11].

Yates [16] defines a decision as "the commitment to an action whose aim is to produce (or expected to produce) satisfying outcomes". However, this focuses entirely on decisions that lead to outcomes, ignoring decisions leading to policy or other less tangible results. This tendency in the literature for researchers to focus on very narrow classes of domain and decision is precisely one of the reasons we are seeking a more formal model. However, this class of decisions for action is clearly very important and one we will aim particularly to clarify.

A decision also implies a selection of alternatives. Olson et al [9] found that design teams working in early software design spent 21.5% of design time generating and clarifying alternatives (making solution exploration the largest category of activity) and that for 80% of design issues (problems or questions), two or more alternatives were considered, with an average of 2.5 alternatives per issue. Another study [1] based on the QOC protocol [7] found the mean number of alternatives generated per question to be slightly higher at 2.8. Clearly, the consideration of a single alternative on its own could hardly constitute a decision; it is merely an investigation of an option. There must be at least two or more alternatives, which may or may not be explicitly stated. A decision must therefore include both the investigation and selection of one of the alternatives.

## 3. Understanding decisions

We will now seek to model aspects of decisions and how they interact with acts on the world. We are not going to model the decision making process itself nor discuss the computational models we use to capture decisions in our tools. This is a semantic model intended to help us understand the nature of *what decisions are*.

From the discussion in the previous section, we can see that one of the most frequent outcomes of decisions with meetings are 'Actions'. We can think of the process followed by a decision making group as something like:

(a) group recognises there is a problem or issue to resolve
(b) they consider multiple alternatives
(c) they decide on one alternative
(d) the result of this is one or more 'Actions'
(e) which are carried out and affect the world

This looks simple enough, but in fact there is a big gap between (d) and (e). The 'Actions' in a meeting minutes do not necessarily happen, the individual actioned by the minutes may forget, choose not to follow out the request, be ill or too busy. The 'Action' that is the result of a decision is not an actual act on the world, but an *intention* to act.

This has two implications. First of all one of nomenclature in the following we will use the word *Action* to be the formal resolution of a meeting and *act* to be the event of an agent affecting the world. This is slightly confusing as the word 'action' is used extensively in formal literature, notably in modal action logic, which is closely related to the model of acts we will use here.

Secondly, and more importantly it reminds us that the outcome of a decision is an intention. This intention may be of various forms. It may be an intention to act or may be an establishment of a rule, for example "environmental issues should be taken into account in all future company decisions". The latter is what is called in organisational semiotics a *norm* [6] and in internal project meetings we have sometimes said that the outcome of a meeting is in fact a norm. However, for this paper we will use the more common term intention.

As we have seen for intentions to act, intentions need not be fulfilled and may not even influence the probability of their fulfilment. A child deliberates on whether she would like to be a railway engineer, a ballerina or an intergalactic space pilot when she grows up. After long consideration she decides on the intergalactic space pilot. Now she has made a decision and has a clear intention, but having that intention does not make it any more likely that she will actually become a space pilot.

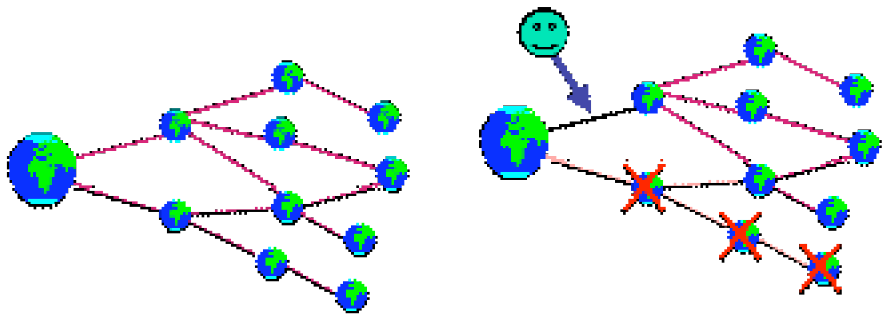

(i) potential futures (ii) acts constrain futures

**Fig. 1.** Modal model of worlds and acts

We said that acts affect the world. This is true in two senses: first the world is typically not the same after an act so it affects the future; and second by performing the act at a particular moment other potential futures never can happen. We can represent this as a branching time model of the world with its current state now and many possible future states. Figure 1(i) illustrates this. Consider time flowing left to right. On the left is the current state of the world (time 0), to its right are two potential next states (time 1), four further states for time 2, etc. The lines represent possible ways in which one state of the world can flow from the previous one. This is a simplification, in the real world time is continuously passing and of course there are many many potential futures!

Suppose our act is to install a Java development kit. This changes the state of the world, there is now (assuming the install works!) a copy of the JDK on our machine. Furthermore, by choosing to install Java at that moment I did not install a C# IDE. Now I may choose to do so later, but if at the moment the JDK is installed I get sent some C# code I cannot run it, this will mean that other things cannot happen etc. My act constrains the possible futures. this is illustrated in Figure 1(ii).

Note in Figure 1(ii), we have shown a small face as we will assume that the act is carried out by some agency: individual, corporate or automatic.

Note too that an act need not be effectual. Imagine you are on a television show. “You have won one million Euros, all you need to do is choose the right button, press NOW”. There are two buttons labelled ‘A’ and ‘B’. You decide to press ‘A’ and do so, but the button does not move. In some ways the world has not changed except you have burnt a tiny bit of energy in pressing. However, the world has changed in the sense that you can no longer press the button ‘B’ at that moment. If you had a chute would have opened and one million Euros would have come pouring around your feet.

Likewise inaction is an act. If you dither too long and do not press either button then similarly your moment has passed and you can never get your one million Euros.

We can see that Figure 1(i) represents a modal model and section 4 will look at a formal semantics for this.

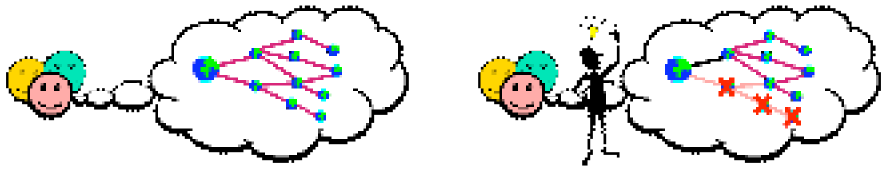

(i) model of possible futures (ii) decisions constrain model

**Fig. 2.** Meta-modal model of decisions

There are many kinds of decisions as we have seen, but for the treatment here we will focus on decisions relating to acts; that is ones where the ultimate outcome is an intention to affect the world.

Figure 2(i) represents the group coming together, perhaps at a meeting, in order to make a decision. Now strictly the making of the decisions, the meeting of the group are all acts in the world. However, in order to make discussions tractable, we will consider the world model illustrated in Figure 1(i) as being the world outside of the decision making group. For example, in the case of a company the model in Figure 1(i) includes the production in the companies factory, the processing of orders and invoices, the delivery to customers, those customers use in the external world, but does not include the activity within the boardroom.

However, inside the blackbox of the board room, the decision makers need to consider the world outside. That is they have some sort of model of the world. However, in order to make decisions they clearly need to have some appreciation that the world can proceed in different ways and has different potential outcomes that they can at least attempt too affect. That is they have some sort of 'mental model' of a modal model of the world. Figure 2(i) shows this as the modal model inside the thought balloon.

Now when we say 'mental model; this does not necessarily mean entirely within any individual's head. In true distributed cognition style, this may be shared between people who understand different aspects of the world, or may be represented in the environment on white boards, computer screens, the outcomes of complex simulations. As a group they have the model.

The model need not be a true reflection of the world. The decision makers' beliefs about the current state of the world may be inaccurate, partial or plain wrong! Even more they may not know how the world can change, what acts are possible and what the effects of those acts will be. Furthermore, they may or may not be aware of the limits of their own knowledge!

Finally, the outcome of the decision, as we noted will not be an act as such, but an intention to act. More generally the outcome (and alternatives considered) is an intention that the actual unfolding of the world will be some constrained form of the potential ones, typically trying to make desired outcomes more likely amongst the potential futures. Figure 2(ii) shows this – just as the act constrains the actual unfolding of the world, the decision constrains the intended unfolding.

Argumentation notations such as QOC, IBIS or DRL are all ways of formalising this modal model of the world [7, 4, 5]. For example, in QOC the options represent different potential futures for a design. Some of the notations are aimed more a recording decisions, but aim to improve decision making. Part of the way they achieve this is by encouraging decision makers to consider a wider range of potential futures that is establish a more accurate modal model of the world.

## 4. Modelling decisions

We will now move on to create a semantic model of some aspects of decisions. It would be possible to construct a modal logic however instead we will build a more constructive model. This is partly a matter of taste, partly because we find the complexity of the meta-modal construction is easier to keep track of with named sets and functions, and partly because anyway it is hard to verify the completeness or consistency of logics without model semantics.

### 4.1 Modelling worlds and their potentiality

In order to simplify the model we will adopt discrete time. This is a reasonable simplification as we are interested in long-term business decisions and so acts take place over hours or days, so a discretisation is unlikely to miss We want a branching temporal version of modal action logic, rather like CTL.

We start with a set of world states, Wrld, agents, Agnt, and acts, Acts. A world state, $w \in$ Wrld, here represents one of the potential states of the world at a particular moment in time.

The world state may contain some trace of the actions that lead to it (like footsteps on the sand), but need not (walking across concrete). However, they do always determine the potentiality for change. This is captured by two functions *active*, which gives the agents potentially active at that moment, and *available*, giving the potential acts that each could perform.

$$\text{active: Wrld} \longrightarrow \text{setof(Agnt)} \quad (1)$$

$$\text{available: Wrld} \longrightarrow \text{Agnt} \nrightarrow \text{setof(Act)} \quad (2)$$

$$\forall\, w \in \text{Wrld} \bullet \text{dom(available(w))} = \text{active(w)} \quad (3)$$

N.B. we are using the Z partial function arrow '⇸', but not other Z symbols.

Note that this definition allows the possibility of having several agents simultaneously acting. In the case of a world model with a single actor this would reduce to simple action semantics. Also we can model turn-taking behaviour by having two agents in Agnt and only having one agent active in each world state.

A particular history is thus an alternating sequence of world states and events:

$$w_0 - e_1 - w_1 - e_2 - w_2 - \ldots\ w_{t-1} - e_t - w_t \quad (4)$$

where $w_i$ in Wrld, $e_i$ in Evnt

Each event in Evnt consists of the actions performed by all the active agents at a particular time:

$$\text{Evnt} = \text{Agnt} \nrightarrow \text{Act} \quad (5)$$

For each transition the event must be consistent with the potentiality of the world state.

$$\forall\ i \in [0,t) \ \bullet \ e \in \text{consistent-event}(w_i) \qquad (6)$$

This relationship consistent is defined in the obvious way:

$$\text{consistent: Wrld} \longrightarrow \text{setof(Evnt)} \qquad (7)$$

$$\forall\ e \in \text{Evnt}, w \in \text{Wrld} \ \bullet \ e \in \text{consistent-event}(w) \ \textbf{iff}$$

$$\text{dom}(e) = \text{active}(w) \qquad (8)$$

$$\wedge \ \forall\ a \in \text{dom}(e) \ \bullet \ e(a) \in \text{available}(w)(a) \qquad (9)$$

Clause (8) says that only and exactly the active agents act and clause (9) says that they can only performs acts that are available to them.

Note that we could have simply defined the function *consistent-event* directly and not bothered with available. The active agents and their available acts would then have been derived from the consistent events. However, this would have allowed semantics where two agents acts are interdependent. By separating these the independence of each agent at each moment is ensured.

Of course one agents acts will affect the subsequent state of the world and hence the available acts of other agents, but this is mediated through their effects on the world not by simultaneous constraints.

### 4.2 Modelling acts in an uncertain world

Of course, the interesting thing is what these effects are ... that is what the acts *do*!

This is modelled by a state update function we will call effect (note not the same as 'effect' in the PIE model [3]):

$$\text{effect: Evnt} \times \text{Wrld} \longrightarrow \text{U(Wrld)} \qquad (10)$$

Two things to note here. First, whilst the *ability* to act is independent for each agent, the effect depends on *all* simultaneous actions. The second things is the notation U(Wrld). This denotes an *uncertainty measure* over World. This abstracts over different kinds of modal models including deterministic, non-deterministic, probabalistic and fuzzy.

Currently we are just using this as a placeholder for a specific uncertainty measures as above, but we intend to more fully develop algebraic structures for uncertainty measures so that it is possible to reason over hybrid probabilistic and non-deterministic situations, as are commonly encountered in statistical reasoning.

The only three properties of the uncertainty we will use for now are that there will be a *support* function giving those elements that have some possibility of occurring.

$$\text{support: U(S):} \longrightarrow \text{setof(S)} \qquad (11)$$

Also some way of combining the uncertainty of several independent things:

$$\text{combine: } U(S_1) \times U(S_2) \times \ldots \times U(S_n) : \longrightarrow U(S_1 \times S_2 \times \ldots \times S_n) \qquad (12)$$

And finally some way of reducing uncertain knowledge about a more complex description of a situation to uncertainty about a reduced representation (e.g. given the probabilities of the way two dice fall we can calculate the probability of the total count):

$$\text{reduce: } U(S_1) \times ( S_1 \longrightarrow U(S_2) ) \longrightarrow U(S_2) \tag{13}$$

In the case of a deterministic modal model U(Wrld) is simply the set Wrld. There is no uncertainty, each combination of acts leads to a unique next world state. The *support* function is then the singleton function that takes an element to the singleton set of itself, the *combine* function is simply the identity on the product set and *reduce* is application of the abstraction function (second argument).

In the case of a non-deterministic model U(Wrld) is the power set, setof(Wrld). That is the effect of a particular combination of agent acts is one of a number of subsequent final world state. In this case *support* function is the identity, *combine* is simply the product of the subsets and reduce is union of the application of the function to the set:

$$\text{combine}_{det}\text{: } ( s_1, s_2, ..., s_n ) = s_1 \times s_2 \times ... \times s_n \tag{14}$$

$$\text{reduce}_{det}\text{: } ( s_1, f ) = \bigcup \{ f(w) \in w \text{ in } s_1 \} \tag{15}$$

In a probabalistic model U(Wrld) is a probability measure over Wrld. That is:

$$U_{prob}(\text{Wrld})\text{: } \text{Wrld} \longrightarrow \text{Real} \tag{14}$$

$$\text{s.t. } \forall\ p \in U_{prob}(\text{Wrld}) \bullet \textstyle\sum_{w \in \text{Wrld}} \bullet\ p(w) = 1.0 \tag{15}$$

In this case the *support* is the set of world states with non zero probability:

$$\text{support}_{prob}\text{: } ( \text{dist} ) = \{ w \in \text{Wrld} \mid p(w) > 0 \} \tag{16}$$

*combine* multiplies probabilities

$$\text{combine}_{prob}\text{: } ( \text{dist}_1, \text{dist}_2, ..., \text{dist}_n ) = \text{dist}$$
$$\text{where dist}( w_1, w_2, ..., w_n ) = \prod \text{dist}_i(w_i) \tag{17}$$

and reduce is the normal Bayesian formula (the second argument is effectively a conditional probability):

$$\text{reduce}_{prob}\text{: } ( \text{dist}_1, f ) = \text{dist}_2$$
$$\text{where dist}_2( w ) = \sum v \in f^{-1}\text{dom}(\text{dist1}) \bullet \text{dist}_1(v) . f(v)(w) \tag{18}$$

Finally, in a fuzzy model we also have U(Wrld) a mapping to [0,1] where the result is interpreted as fuzzy set membership. The *support* function then gives the set of elements with non-zero membership, *combine* uses one of the many fuzzy *and* operators and reduce also uses fuzzy *or*.

We can then look again at a history of alternating world states and events as in (4) and check whether it is self-consistent:

$$\forall\ i \in [0,t) \bullet w_{i+1} \in \text{support}(\text{effect}(e_{i+1}, w_i)) \tag{19}$$

## 4.3 Adding agents behaviour

The *availability* function says how agents are allowed to act at a particular moment and the *effect* function says what the effect of their combined acts will be. The final part of a complete model is that the agents themselves have a behaviour which itself may be understood. This may be deterministic, in the case of a simple automated computational agent, or totally unknown.

If it were deterministic then we would expect a behaviour that is a function of the world state:

$$\text{behaviour}_{det}: \text{Agnt} \times \text{Wrld} \rightarrow\!\!\!\!\!\!+ \ \text{Act} \tag{20}$$

However, we want a more generic expression and so we have instead:

$$\text{behaviour}: \text{Agnt} \times \text{Wrld} \rightarrow\!\!\!\!\!\!+ \ \text{U(Act)} \tag{21}$$

with the obvious side conditions:

$$\forall\, a \in \text{Agnt}, w \in \text{Wrld} \bullet$$
$$\text{dom(behaviour)} \supseteq \{ (a, w) \mid a \in \text{active}(w) \} \tag{21}$$
$$\wedge \ \text{support(behaviour}(a, w)) \subseteq \text{available}(w)(a) \tag{22}$$

The first part (21) says that we must define some behaviour for each agent (which may be simply saying it can non-deterministically do any available action). The second part says that it can only do actions that are available.

A history must of course be consistent with this behaviour too:

$$\forall\, i \in [0,t), a \in \text{active}(w_i) \bullet \quad e_{i+1}(a) \in \text{behaviour}(a, w_i) \tag{23}$$

That is, each agents act in the event must be one that is possible given their behaviour.

Finally for any given uncertainty measure we can work out the likelihood of a particular history by combining the uncertainties of the behaviours at each moment and the uncertainties of the effects.

### 4.4 The meta-modal world

From the above we can see that a complete modal model with all the elements is a tuple:

$$< \text{Wrld, Agnt, Act, active, available, effect, behaviour, U} > \tag{24}$$

We will call such a tuple a world model and call the set of such models WORLD-MODEL. Note we have moved to all upper case to help denote the conceptual higher order nature of this set. However, from a mathematical perspective it is simply the set of all such tuples – just ordinary!

Given this we can look ate relationships between these modal world models. We will look at one example of this.

One difference is that we may take into account the actions of more or less agents specifically or simply regard their behaviour as part of the overall environment of the other agents.

As illustration we can do this for a particular agent A. Given a particular world model as in (24), we get a derived model 'forgetting' agent A:

$$< \text{Wrld, Agnt', Act', active', available', effect', behaviour', U} > \tag{24}$$

These share the same set of world states and have the same model of uncertainty, but differ in other respects. Most of these are simple removals of A from relevant sets.

$$\text{Agnt'} = \text{Agnt} - \{ A \} \tag{25}$$

$$\forall\, w \in \text{Wrld} \bullet \text{active'}(w) = \text{active}(w) - \{ A \} \tag{26}$$

$\forall\, w \in \text{Wrld},\ a \in \text{active'}(w) \bullet$

$$\text{available'}(w)(a) = \text{available}(w)(a) \tag{27}$$

$$\text{behaviour'}(w, a) = \text{behaviour}(w, a) \tag{28}$$

$$\text{Act'} = \bigcup \left\{ \text{available'}(w)(a) \bullet w \in \text{Wrld},\ a \in \text{active'}(w) \right\} \tag{29}$$

The most interesting though is the modified effect. For any world w we have two cases:

**Case (i): a $\notin$ active(w)**

In this case the agent has no effect on the outcome:

$$\forall\, e \in \text{consistent-event'}(w) \bullet \text{effect'}(e, w) = \text{effect}(e, w) \tag{30}$$

Note writing *consistent-event'* for the events allowable in the modified world model. Note that equations (7–9) define this in terms of *active'* and *available'*.

**Case (ii): a $\in$ active(w)**

This case is more complicated as we need to combine the uncertainty of the original effect and that of agent A's acts. In this case, for each event consistent with w in the reduced world, we need to look at all the events including potential acts for agent A and then combine their uncertainties.

To do this we define an extension function that takes an act and extends an event to include agent A doing this act::

$$\text{extend: Evnt'} \times \text{Act} \longrightarrow \text{Evnt} \tag{31}$$

$$\forall\, e \in \text{Evnt'},\ \text{act in Act} \bullet \text{extend}(e,\text{act})(A) = \text{act} \tag{32}$$

$$\forall\, e \in \text{Evnt'},\ \text{ag in Agnt'} \bullet \text{extend}(e,\text{act})(\text{ag}) = e(\text{ag}) \tag{33}$$

We can then use the uncertainty reduce operation to give the uncertainty of any actual event:

$$\text{poss-evnt: Wrld} \times \text{Evnt'} \longrightarrow U(\text{Evnt}) \tag{34}$$

$\forall\, w \in \text{Wrld},\ ev \in \text{Evnt'} \bullet$

$$\text{poss-evnt}(w,ev) = \text{reduce}(B, \text{ext}) \tag{35}$$

where $B = \text{behaviour}(A,w)$ $(\in U(\text{Act}))$

and $\text{ext}(\text{act}) = \text{extend}(ev,\text{act})$ $(\in \text{Act} \longrightarrow \text{Evnt})$

Note, we are using a special case of *reduce* where the second function has no uncertainty.

Finally we use reduce again to produce the final effect:

$\forall\, w \in \text{Wrld}, ,\ ev \in \text{Evnt'} \bullet$

$$\text{effect'}(w,ev) = \text{reduce}(\text{poss-evnt}(w,ev), \text{effect}) \tag{36}$$

## 4.5 Modelling decisions and intentions about the world

Now let's consider a meeting of some formal or informal body. To avoid saying "decision making body" too many times we will simply refer to them as the 'body'.

Let's say that the real world model, not taking into account the decisions to be made, is RealWM. The real world has some agents who are members, employees or otherwise under the behest of the body:

$$\text{ControlledAgnt} \subseteq \text{RealWM.Agnt} \tag{37}$$

Clearly if the body has any ability to influence the world then some of these agents will have possibilities to act in different ways:

$$\exists\ a \in \text{ControlledAgnt},\ w \in \text{RealWM.Wrld} \ \bullet$$

$$a \in \text{active}(w)\ \} \tag{38}$$

$$\wedge\ \ |\ \text{available}(w)(a)\ |\ >\ 1 \tag{39}$$

Furthermore if the agents involved have free will that is influenced by the body then the behaviour of these agents will be non-deterministic with respect to this 'ignoring the decision' model of the real world.

As well as the actual way the world is, the body will also have a model of the world, BodyWM. Of course, as we noted in section 3, this model may or may not agree with the real world.

Now there are many simplifications in expressing the body's model of the world in this way:

- The members of the body are likely to have different models of the world. There must be some agreement in order to sensibly communicate, but they are likely to have different beliefs and knowledge about all sorts of things.
- The models that they do have will certainly not be formal ones as expressed here!
- To the extent that their mental models can be formalised they are more likely to be largely intentional in nature rather than extensional.

However, the important thing we want to discuss here is the nature of the alternatives and the decision between them, so we will use this simplified formal, extensional consensus model for the purposes of this paper. A more complete treatment would need to consider these other factors.

As we noted in section 3 a decision to act is actually establishing an intention to act. There are two aspects to this intention:

(i) a deontic component expressing the will to act

(ii) a descriptive component expressing the nature of the act

We will only deal with the second of these. Now the outcome of a decision to act may be a simple action like "person X will Y", may be a negative action such as "no employee will do Z", or may be more complex "if situation S arises attempt to make T happen". These are of course again typically expressed intentionally, but we can represent the decision outcome as yet another world model from WORLD-MODEL. We will call it OutcomeWM.

Typically OutcomeWM will be some modification of the original BodyWM. Often a filling in of uncertainty, but possibly a change. For example, the RealWM would include the expected behaviours of employees and OutcomeWM may prescribe different acts.

By comparing the OutcomeWM with the original BodyWM, we can differentiate different classes of decision.

The class we identified in section 3 are *intentions to act*. These are characterised by sharing the same 'external' features as the BodyWM, but with the behaviours of the controlled agents changed:

$$\text{IntWM.Wrld} = \text{BodyWM.Wrld} \quad (40a)$$

$$\text{IntWM.active} = \text{BodyWM.active} \quad (40b)$$

etc.

$$\text{IntWM.effect} = \text{BodyWM.effect} \quad (40c)$$

$$\forall\ a \notin \text{ControlledAgnt},\ w \in \text{Wrld} \ \bullet$$
$$\text{IntWM.behaviour(a,w)} = \text{BodyWM.behaviour(a,w)} \quad (41)$$

$$\exists\ a \in \text{ControlledAgnt},\ w \in \text{Wrld} \ \bullet$$
$$\text{IntWM.behaviour(a,w)} \neq \text{BodyWM.behaviour(a,w)} \quad (42)$$

Basically the intention to act does not change knowledge or beliefs about the way the world works (eqn 40a-d). Neither does it change our expectations about the behaviour of other agents (41). However, we are now planning/expecting some controlled agent to act in a way that is different given we have come to the decision.

The opposite case is where the decision does not result in any immediate intention to act, but instead is a decision to believe. This can take many forms.

A marketing meeting may decide that a market growth of between 5 and 10 percent should be used for future planning. This would result in a change in the effect.

$$\text{BeliefWM.effect} \neq \text{BodyWM.effect} \quad (43)$$

They may realise that due to new legislation, it is impossible for they or their competitors to launch a new chocolate bar without a six month food safety trial. That is a change in the active agents and their available acts:

$$\text{BeliefWM.active} \neq \text{BodyWM.active} \quad (44)$$

or:

$$\text{BeliefWM.available} \neq \text{BodyWM.available} \quad (45)$$

They may decide that a competitor is likely to introduce a new ice cream flavour in time for the summer season. That is they change their model of other agents behaviour:

$$\exists\ a \notin \text{ControlledAgnt},\ w \in \text{Wrld} \ \bullet$$
$$\text{BeliefWM.behaviour(a,w)} \neq \text{BodyWM.behaviour(a,w)} \quad (46)$$

Finally, the change in belief may be more fundamental. For example, deciding that the astronomical data suggests that there is a new previously undiscovered planet. That is the actual underlying sets are not preserved:

$$\text{BeliefWM.Wrld} \neq \text{BodyWM.Wrld} \quad (47)$$

Note that this change in belief may be simply a change due to information or may be an active decision of the body.

In practice, decisions about beliefs and decisions about actions are interwoven. The choice of what action to take is intimately connected with the assumptions of what the effects of that action will be. If we believe our competitor is about to release a new ice cream flavour, then we may decide to start a pre-emptive marketing campaign.

### 4.6 Multi-level reflexive meta-modal models

We can apply the world modelling framework in sections 4.1–4.4 to the decision making body themselves. The world states consist of beliefs about the world, options being considered and criteria being applied. The acts in such a model include the consideration of alternative options, the choice of criteria and decisions which change the beliefs/intentions about the world.

We will not fill out the full details of this here, but it will be something like:

$$\text{MetaModalWorld.Wrld} = \langle \text{Wm, Opt, Crtr} \rangle \tag{48}$$

where:

$$\text{Wm} = \text{WORLD-MODEL}$$
$$\text{Opt} = \text{setof( WORLD-MODEL )}$$
$$\text{Crtr} = \text{setof( Hist} \longrightarrow \text{Value )}$$

The first part of this is the body's current model of the world, the second is the alternative options currently being considered and the last is the set of criteria. Note that the typing of the last of these is a little problematic, but it is capturing the fact that criteria are some sort of valuation (not necessarily numeric or monetary) of potential futures.

The decision meeting body is the agent and has acts such as 'make decision':

$$\text{MetaModalWorld.effect( body, (wm,opt,ctr) ) = (wm',opt',ctr)} \tag{49}$$

where:

$$\text{wm}' \in \text{opt}, \qquad \text{opt}' = \{\ \}$$

Of course, this model of the decision making body itself is really not separate from the world that the body are deciding about, but part of it. Furthermore, the body's model of the world will include a model of the body itself. However, we will omit the elaboration of the fixed point formula here!

This is not just a piece of formal nicety, but is actually critical to fully understand decisions such as "when we discuss new packaging we should always consider environmental impact". This is not a decision about acts on the world, or about beliefs about the world, but about how decisions are made. that is meetings often very explicitly talk about themselves.

## 5. Discussion

In this paper we have covered quite a lot of ground.

We began with a more semi-formal look at the nature of actions in the world and contrasting this with the outcomes of decisions which are typically called 'Actions' in meeting minutes, but are really intentions to act. That is the effect of a decision is to change the 'mental state' of a decision making body, in terms of beliefs and intentions, but not to directly affect the external world.

In modelling this complex situation we have inevitably had to make some gross implications. Not least has been the modelling of the decision making body's model of the world as a single extensional model.

The distributed cognition literature has shown that in the complex interplay between people and artefacts there is rarely a single agreed model being used in a 'rational planning' way by all. It is more that the group as a whole may behave as if it had an emergent model, goals and actions. However, this emergent behaviour is not present in any single individual or representation, but about the interaction between the individuals', possible incompatible, models.

Also even these individual models are likely to be fragmentary partially model-based and partially intentional and quite likely again not consistent. People operate with bounded rationality, we do not have the time or reasoning capacity to follow the full implications of our beliefs and so we would find there were *no* extensional models consistent with *all* our beliefs. For the body as a whole with different individuals this is even more extreme.

These limitations, whilst meaning we do not have a 'perfect' model of decision making. However, even this restricted model has allowed us to clarify our informal ideas of decisions and the types of decisions that can be made.

In creating the model we introduced a way of formally dealing with different kinds of uncertainty without having to commit to a deterministic, non-deterministic or probabalistic model. As well as being a convenient abstraction we believe this is an important mechanism for starting to reason about different kinds of uncertainty.

There is considerable studies of human reasoning under different forms of uncertainty [8] and we are known to cope 'badly' with probabalistic situations. This is partly because we are adapted to the real world where we do not have perfect probabalistic knowledge, but a variety of approximate and probabalistic knowledge and complete un-knowledge. In AI there is also a large literature on modelling computationally different kinds of uncertainty [14]. Our goal is to extend the semantic framework to make it possible to *talk about* such mechanisms as well as reason within them.

Similarly, looking back at the reflexive multi-level meta-modal model in this paper; we are using this as a semantic model not a computational one. We have been producing various support tools. One of these uses natural language processing of transcripts of meetings to identify actions and issues using a combination of indicator words, document/font styles and templates of parts-of-speech and semantic tags.. This was built as an extension of the existing requirements identification REVERE toolset [12]. We have also developed simple tablet based meeting management system for a particular organisation. Neither of these includes the meta-modal model within them, but each has their own computational model. For example, the tablet based system only knows about dated actions and their current status. Instead, the conceptual modelling formalised in this paper is more about informing our own thinking and giving a framework in which to understand where our computational models.

## Acknowledgements

This work was funded by the EPSRC funded Tracker and DIRC projects.

## References

1. Ball, L., Ormerod, T.: The Influence of Co-Designers on the Generation and Evaluation of Design Alternatives. In Scrivener, S., Ball, L., Woodcock, A. (eds.): Collaborative Design: Proc. of Co-Designing 2000, Springer, London (2000) 243–252
2. Conklin, J. & Begeman M.: gIBIS: A Tool for All Reasons. Journal of the American Society for Information Science, 40(3), (1989) 200–213.
3. Dix, A.: Formal Methods for Interactive Systems. Academic Press (1991)
4. Kunz, W., Rittel, H.: Issues as elements of Information Systems , Working Paper No. 131, Center for Planning and Development Research, University of California, Berkeley, (1970)
5. Lee, J.: Sibyl: A Tool for Managing Group Decision Rationale. In Proc. of Conf. on Computer Supported Cooperative Work '90, Los Angeles, CA., (1990) 79–92
6. Liu, K.: Semiotics in Information Systems Engineering. Cambridge University Press, Cambridge (2000)
7. MacLean, A., Young, R., Bellotti, V., Moran, T.: Questions, Options, and Criteria: Elements of Design Space Analysis. In Carroll, J., Moran, T. (eds.) Human-Computer Interaction, 6 (3&4), Special Issue on Design Rationale (1991) 201–250
8. Manktelow, K.: Reasoning and Thinking. (Chapter 6: Judging Probability) Psychology Press, Taylor and Francis (2000)
9. Olson, G., Olson, J., Storrøsten, M., Carter, M., Herbsleb, J., Rueter, H.: The Structure of Activity During Design Meetings. In Moran, T., Carroll, J. (eds.): Design Rationale: Concepts, Techniques, and Use, Erlbaum, Mahwah, New Jersey, (1996) 217–239
10. Ramduny-Ellis, D., Dix, A., Rayson, P., Onditi, V., Sommerville, I., Ransom, J.: Artefacts as designed, Artefacts as used: resources for uncovering activity dynamics. Cognition Technology and Work, Springer (in press) (2004)
11. Richter, H., Abowd, G., Geyer, W., Fuchs, L., Daijavad, S., Poltrock, S.: Integrating Meeting Capture within a Collaborative Team Environment. In Proc. of the Intnl Conf. on Ubiquitous Computing, Ubicomp 2001, Atlanta, GA, September, Springer, (2001) 123–138
12. Sawyer, P., Rayson, P., Garside, R.: REVERE: support for requirements synthesis from documents. Information Systems Frontiers Journal. 4(3), (2002) 343–353
13. Selvin, A., Buckingham Shum, S., Sierhuis, M., S., Conklin, J., Zimmermann, B., Palus, C., Drath, W., Horth, D., Domingue, J., Motta, E., Li, G.: Compendium: Making Meetings into Knowledge Events. In Knowledge Technologies 2001, March 4-7, Austin TX (2001)
14. Shafer, G., Pearl, J. (eds): Readings in uncertain reasoning. Los Altos: Morgan Kaufmann (1990)
15. Stefik, M., Bobrow, D., Foster, G.,Lanning, S. Tatar., D.: WYSIWIS revisited: early experiences with multiuser interfaces. ACM Transactions on Office Information Systems, 5(2), (1987) 147–167
16. Yates, F. J. Outsider: Impressions of Naturalistic Decision Making. In Salas E., Klein G. (eds): Linking Expertise and Naturalistic Decision Making, Lawrence Erlbaum Associates (2001) 9–3