Preference Change and Information Processing

Fenrong Liu*

June 15, 2006

1 Introduction

Incoming information not only changes our knowledge but also our preferences. Decisions are made according to the preferences, which are eventually based on our evaluations of the options. In this paper, we will explore the ways new information affects our evaluations to see how this results in a preference change. A *qualitative* investigation was undertaken in [BL06] in which the preference relation in the initial model is manipulated according to incoming information. Here we will take a more *quantitative* approach by introducing an evaluation function. Interestingly, in this manner it becomes possible to consider the subtlety of information processing.

As an example, suppose that you plan to buy an apartment. There are two candidate apartments d_1 and d_2 available, located in different places. You have your own judgement based on your current knowledge: they could be equally preferable, or one is more preferable than the other. To mark your evaluation difference, you assign two numbers to d_1 and d_2, respectively. A newspaper article that "the government is planning to build a park near d_1" may *increase* your value for d_1. In contrast, getting to know that the criminal rate is going up in the neighborhood of d_1 may *decrease* your value for d_1. The idea is: you start off with the initial values of the options, and keep *scoring* in accordance with the new information, either adding points if the information has a positive influence on the option, or dropping points in case it has a negative effect, the number zero is added when it does not have any effect or is irrelevant. Altogether this brings about an evaluation change from which the preference change can be induced.

2 An evaluation language and model

Following [Spo88] and [Auc03], a language of graded preference modalities is introduced to indicate the strength of preference. Here we take a simple design ([Liu04]), which is more workable and perspicuous.

Definition 2.1 (Language) *Let a finite set of proposition variables Φ and a finite set of agents G be given. The epistemic evaluation language $\mathcal{L}$ is defined by the rule*

$$\varphi := \top \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid q_a^m \mid K_a\varphi$$

where $p \in \Phi$, $a \in G$, and $m \in \mathbb{Z}$.

A propositional constant q_a^m is added to the language for each agent $a \in G$ and each value $m \in \mathbb{Z}$. The intended interpretation of the formula q_a^m is 'the agent a assigns the state where she stands the value at most m', and the intended interpretation of the formula $K_a\varphi$ is 'the agent a knows that φ'. We will see that the language of [Auc03], $\mathcal{L}_A$, can be simplified with this language.

Definition 2.2 (Evaluation models) *An evaluation model for the epistemic evaluation language is a tuple $\mathcal{M} = (S, \{\sim_a \mid a \in G\}, \{v_a \mid a \in G\}, V)$* [1] *such that S is a non-empty set of states, $\sim_a$ is an epistemic*

*University of Amsterdam & Chinese Academy of Social Sciences

[1] I will sloppily write it as $\mathcal{M} = (S, \sim_a, v_a, V)$ when G is clear from the context.

equivalence relation on S, v_a is an evaluation function assigning each state an element from $\{-\infty\} \cup \mathbb{Z} \cup \{\infty\}$[2], and V is a function assigning to each proposition variable p in Φ a subset $V(p)$ of S.

Evaluation functions induce a total ordering in the obvious way, namely, from $v_a(s) \leq v_a(t)$ we can obtain $s \preceq_a t$. In this way, we are making use of the information about the qualitative ordering encoded in the evaluation functions. However, we will see that the quantitative information part will play a big role in many situations in the following sections. For instance, considering information about the intensity of preference will lead to a new definition of bisimulation.

Definition 2.3 (Truth conditions) *Suppose s is a state in a model $\mathcal{M} = (S, \sim_a, v_a, V)$. Then we inductively define the notion of a formula φ being true in $\mathcal{M}$ at state s as follows:*

$\mathcal{M}, s \vDash \top$

$\mathcal{M}, s \vDash p$ *iff* $s \in V(p)$, *where* $p \in \Phi$

$\mathcal{M}, s \vDash \neg\varphi$ *iff not* $\mathcal{M}, s \vDash \varphi$

$\mathcal{M}, s \vDash \varphi \wedge \psi$ *iff* $\mathcal{M}, s \vDash \varphi$ *and* $\mathcal{M}, s \vDash \psi$

$\mathcal{M}, s \vDash K_a\varphi$ *iff for all* $t \in S$ *such that* $s \sim_a t$ *and* $\mathcal{M}, t \vDash \varphi$

$\mathcal{M}, s \vDash q_a^m$ *iff* $v_a(s) \leq m$, *where* $m \in \mathbb{Z}$.

For the sake of comparison, we give the definition for $B_a^m\varphi$ in [Auc03] as follows,

$\mathcal{M}, s \vDash B_a^m\varphi$ iff for all $t \in S$ such that $s \sim_a t$ and $v_a(t) \leq m$, $\mathcal{M}, t \vDash \varphi$.

Theorem 2.4 (Soundness) *Epistemic Evaluation Logic (EEL) consists of the following axioms and derivation rules. Furthermore, it is sound with respect to evaluation models.*

1. *All propositional tautologies*
2. $K_a(\varphi \to \psi) \to (K_a\varphi \to K_a\psi)$
3. $K_a\varphi \to \varphi$
4. $K_a\varphi \to K_aK_a\varphi$
5. $\neg K_a\varphi \to K_a\neg K_a\varphi$
6. $q_a^m \to q_a^n$ *for all* $m \leq n \in \mathbb{Z}$
7. *From* $\vdash \varphi$ *and* $\vdash \varphi \to \psi$ *infer* $\vdash \psi$
8. *From* $\vdash \varphi$ *infer* $\vdash K_a\varphi$.

We take the standard notion of proof. In case a formula φ is provable in EEL, we write $\vdash_{EEL} \varphi$.

Theorem 2.5 (Completeness) *The logic EEL is complete with respect to evaluation models.*

Proof. The proof is standard. First we define the canonical model as follows: $\mathcal{M}^c = (S^c, \sim_a, v_a, V)$

- S^c = $\{s_S$: S maximal EEL-consistent set$\}$
- $\sim_a$ = $\{(s_S, s_T)$: $S/K_a \subseteq T\}$ where S/K_a = $\{\varphi$: $K_a\varphi \in S\}$
- $v_a(s_S)$= min$\{m : q_a^m \in S\}$ (∞ if $\{m : q_a^m \in S\}$ is empty, $-\infty$ if $\{m : q_a^m \in S\} = \mathbb{Z}$.)
- $s_S \in V(p)$ iff $p \in S$.

[2] In [Auc03] the range is natural numbers up to a maximal element (Max). The values are normalized to Max. For me the distance between the numbers seems essential, so normalization is not an option. Similarly I like to be able to subtract unrestrictedly.

We need to show that

$$\varphi \in T \text{ iff } \mathcal{M}^c, s_T \models \varphi.$$

By induction on the structure of the formula φ. We only consider the case of the constant q_a^m:

(⇒) Assume $q_a^m \in T$. We have $v_a(s_T) \leq m$. Then by Definition 2.3, we get $M^c, s_T \models q_a^m$.

(⇐) Assume $M^c, s_T \models q_a^m$. We know $q_a^{v_a(s_T)} \in T$ and $v_a(s_T) \leq m$. By axiom 6, $q_a^{v_a(s_T)} \to q_a^m$. So, we get $q_a^m \in T$. This is to say that we have proved that

> Every EEL-consistent set Γ of formulas is satisfiable in some epistemic model.

The completeness result follows. ■

To conclude this section we look at the relation between $\mathcal{L}_A$ and $\mathcal{L}$. From $\mathcal{L}_A$ to $\mathcal{L}$, we can define a translation: a formula of the form $B_a^m \varphi$ is translated into $K_a(q_a^m \to \varphi)$. This is to say that in the language $\mathcal{L}$, we can express the same notions as [Auc03] without introducing additional epistemic operators. This advantage leads to the much simpler completeness proof we have just seen. It becomes even more prominent when constructing reduction axioms for dynamics in the later sections. On the other hand, we can easily translate $\mathcal{L}$ back into $\mathcal{L}_A$: q_a^m will be $\neg B_a^m \bot$, which means that $\mathcal{L}_A$ and $\mathcal{L}$ are equivalent.
Having set up the base language for evaluation models, we now proceed to the dynamic superstructure that we have in mind.

3 Finer modelling of evaluation changes

3.1 Preliminaries: product update

To model knowledge change due to incoming information, the most powerful mechanism is dynamic epistemic logic, which has been developed intensively by [Pla89], [Ben96], [BMS98], [Ger99], [DHK06], etc. Here we briefly recall the basic ideas and techniques.

Definition 3.1 (Event models) *An event model is a tuple $\mathcal{E} = (E, \sim_a, PRE)$ such that E is a non-empty set of events, $\sim_a$ is a binary epistemic relation on E, PRE is a function from E to the collection of all epistemic propositions.*

The intuition behind the function PRE is that it gives the *preconditions* for an action: an event e can be performed at world s only if the world s fulfills the precondition $PRE(e)$.

Definition 3.2 (Product update) *Let an epistemic model $\mathcal{M} = (S, \sim_a, V)$ and an event model $\mathcal{E} = (E, \sim_a, PRE)$ be given, the product update model is defined to be the model $\mathcal{M} \otimes \mathcal{E} = (S \otimes E, \sim'_a, V')$ such as*

- $S \otimes E = \{(s, e) \in S \times E : (\mathcal{M}, s) \models PRE(e)\}$
- *(s, e)* $\sim'_a$ *(t, f) iff both* $s \sim_a t$ *and* $e \sim_a f$
- $V'(p) = \{(s, e) \in \mathcal{M} \otimes \mathcal{E} : s \in V(p)\}$.

The above notions suggests an extension of the epistemic language.

Definition 3.3 (Dynamic epistemic language) *Let a finite set of proposition variables Φ, a finite set of agents G, a finite set of events E be given. The dynamic epistemic language is defined by the rule*

$$\varphi := \top \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_a\varphi \mid [e]\varphi$$

where $p \in \Phi$, $a \in G$, and $e \in E$.

We could also add the usual action operations of composition, choice, and iteration from propositional dynamic logic to the event vocabulary - but in this paper, we will have no special use for these. The language has new dynamic modalities $[e]$ referring to epistemic events, and these are interpreted in the product update model as follows:

$$\mathcal{M}, s \models [e]\varphi \text{ iff } \mathcal{M} \otimes \mathcal{E}, (s, e) \models \varphi.$$

Reduction axioms in dynamic epistemic logic play an important role to encode the changes when the events take place. For example, the following axiom concerns agents' knowledge change.

$$[e]K_a\varphi \leftrightarrow PRE(e) \rightarrow \bigwedge_{f \in E}\{K_a[f]\varphi : e \sim_a f\}.$$

Intuitively, after an event e takes place the agent a knows φ, is equivalent to saying that if the event e can take place, a knows beforehand that after e (or any other event f which a can not distinguish from e) happens φ would hold.
The above update setting can be extended to preference upgrade[3] over evaluation models. We will make this precise below.

3.2 Evaluation product upgrade

We have defined evaluation models in section 2. Now we need to do the same thing to event models.

Definition 3.4 (Evaluation event model) *A evaluation event model is a tuple $\mathcal{E} = (E, \sim_a, v_a, PRE)$ such that E is a non-empty set of events, $\sim_a$ is a binary epistemic relation on E, v_a is an evaluation function assigning each action an element from $\mathbb{Z}$, PRE is a function from E to the collection of all epistemic propositions.*

Based on the values they assign to events, the evaluation functions v_a indicate which events agents prefer. Note that this is a major change as compared with standard uses of evaluation: we do not just evaluate static states of affairs, but also actions or events!

Definition 3.5 (Evaluation product upgrade) *Let an evaluation model $\mathcal{M} = (S, \sim_a, v_a, V)$ and an evaluation event model $\mathcal{E} = (E, \sim_a, v_a, PRE)$ be given, the evaluation product upgrade model is defined to be the model $\mathcal{M} \otimes \mathcal{E} = (S \otimes E, \sim'_a, v'_a, V')$ such that*

- $S \otimes E = \{(s, e) \in S \times E\}$
- *(s, e) $\sim'_a$ (t, f) iff both $s \sim_a t$ and $e \sim_a f$*
- $v'_a(s, e) = v_a(s) + v_a(e)$ *(Addition rule)*
- $V'(p) = \{(s, e) \in \mathcal{M} \otimes \mathcal{E} : s \in V(p)\}$.

Note that we keep all world/event pairs (s, e) represented, as these are the non-realized options that we can still have regrets about. For the evaluation upgrade, we simply take the *sum* of the value for the previous state and that for the event. The Addition rule is best understood by looking at the example in the introduction again, though the evaluation event model there is quite simple and it contains only one event each time.

Example 3.6 Assume that in the initial model S_0, agent a has the same evaluations towards s and t where d_1 would be chosen at s and d_2 at t. She gives 0 to both of them, pictured below:

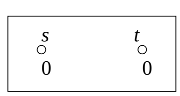

S0

[3] To distinguish between preference change and knowledge change, in this paper we use the word 'upgrade' for the former, and 'update' for the latter.

Afterward, the newspaper brings in a new information "the government is planning to build a park near d_1" (denoted by p), it positively effects the value of s in the model S_0, but has no effect on t. The initial model S_0 is upgraded to S_1:

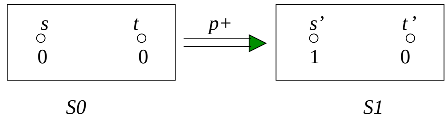

In the model S_1, clearly, a would prefer d_1 over d_2 since the value for s' is greater than that for t'. The story goes on, the new information "the criminal rate is going up in the neighborhood of d_1"(denoted by q) causes values to decrease. The model changes in the following way:

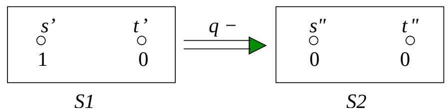

With the evaluation changes, preference changes accordingly, agent a has no preference over d_1 and d_2.

This example shows us how incoming information changes our values of the states. Although the event can be very complex, such a process goes on continuously, and eventually we prefer things with a higher score. However, several issues remain to be discussed: First of all, the *sources* of information. As discussed extensively in various contexts, not all incoming information is equally reliable. In order to propose a realistic evaluation upgrade rule, the *reliability* of information must be taken into account. Also, another key issue concerns the relative different *forces* of information. In multi-agent system, the same information may have different force for different agents. For instance, the agent a may take a piece of information seriously, while the agent b does not do so. These two aspects are parameterized in the following new upgrade rule.

Definition 3.7 (Parameterized rule) *Let $\mu(e)$ be a reliability function, and $\lambda(e)$ a relative force function. The domains of these two functions are the set of events, and the ranges of these functions are $\mathbb{N}$.*[4] *Given the value for the previous state s and event e, the new value for state (s, e) is defined by the following:*

$$v_a(s, e) = v_a(s) + v_a(e) \cdot \mu(e) \cdot \lambda(e).$$

Back to the first step of Example 3.6, suppose agent a only half trusts what the newspaper said, namely $\mu(e) = 5$. And the relative force of the park building information is 4, i.e. $\lambda(e) = 4$, which shows she thinks it is rather important. Then the value of s' in the model S_1 would be calculated as

$$v_a(s, e) = 0 + 1 \cdot 5 \cdot 4 = 20$$

With the Parameterized rule, we can better understand how information is being processed. But things need not stop here, one could propose other types of evaluation upgrade rules to interpret more complex situations. For example, the agent may give more weight to the previous state (behave conservatively), which seems to call for a parameter associated with the value for s in the above rule, as was proposed for belief revision of diverse agents in [Liu04] and [Liu06]. Or in some situations, one needs to consider the dependence between information that comes later and that comes earlier. We will leave these issues for further investigation.

3.3 Dynamic epistemic evaluation logic

We are now ready to define a logic for dynamical evaluation upgrade mechanisms. But in this section we confine ourselves to the Addition rule only.

[4] In practice, one can choose a natural number between 0 and 10 to denote the reliability or the relative force.

Definition 3.8 (Dynamic epistemic evaluation language) *Let a finite set of proposition variables Φ, a finite set of agents G, and a finite set of events E be given. The dynamic epistemic language is defined by the rule*

$$\varphi := \top \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid q_a^m \mid K_a\varphi \mid [e]\varphi$$

where $p \in \Phi$, $a \in G$, $e \in E$, and $m \in \mathbb{Z}$.

Again, we will not include the usual action operations like composition, choice, etc. But we have formulas of the form $[e]q_a^m$, for which we will find reduction axioms as follows.

Theorem 3.9 (Soundness) *Dynamic epistemic evaluation logic $(DEEL)$ consists of the following formulas, and it is sound w.r.t. evaluation product upgrade models:*

1. $[e]p \leftrightarrow p$
2. $[e]\neg\varphi \leftrightarrow \neg[e]\varphi$
3. $[e](\varphi \wedge \psi) \leftrightarrow [e]\varphi \wedge [e]\psi$
4. $[e]K_a\varphi \leftrightarrow PRE(e) \rightarrow \bigwedge_{f\in E}\{K_a[f]\varphi : e \sim_a f\}$
5. $[e]q_a^m \leftrightarrow q_a^{m-v_a(e)}$.

Proof. To prove the validity of the above axioms, we consider two models: $(\mathcal{M}, s)$ and $(\mathcal{M}\otimes\mathcal{E}, s)$ before and after the upgrade. Axiom 1 says that the upgrade will not change the objective valuation of atomic propositions. And axioms 2 and 3 are just Boolean operations, easy to see.

For axiom 4, the formula $[e]K_a\varphi$ says that, in $\mathcal{M}\otimes\mathcal{E}$, all worlds $\sim_a$-accessible from s satisfy φ. The corresponding worlds in $\mathcal{M}$ are those worlds which are $\sim_a$-accessible from s and which satisfy $PRE(e)$. Moreover, given that truth values of formulas may change in an update step, the correct description of these worlds in $\mathcal{M}$ is not that that they satisfy φ (which they do in $\mathcal{M}\otimes\mathcal{E}$), but rather $[e]\varphi$: they become φ after the update. Finally, $[e]$ is a partial operation, as $PRE(e)$ has to be true in order to execute e. Putting this together, $[e]K_a\varphi$ says the same as $PRE(e) \rightarrow K_a(PRE(e) \rightarrow [e]\varphi)$. We can simplify this to $PRE(e) \rightarrow K_a[e]\varphi$. Finally, incorporating the uncertainty agents may have concerning events into our consideration, we get axiom 4.

Likewise, the formula $[e]q_a^m$ says that, in $\mathcal{M}\otimes\mathcal{E}$, the agent a assign the value m to the world s where she stands. According to the Addition rule, the value of s in $(\mathcal{M}\otimes\mathcal{E}, s)$ is the sum of the value for s in $\mathcal{M}$ and that for e in $\mathcal{E}$. Thus the right value for the world s in $\mathcal{M}$ is $m - v_a(e)$. This is what axiom 5 says. ■

Theorem 3.10 (Completeness) *The logic $DEEL$ is completely axiomatized by the above reduction axioms.*

Proof. We have seen the soundness of the above reduction axioms. Note that they are all equivalences, so they are clearly sufficient for eventually turning every formula from the dynamic language into a static one. Then we can use the completeness theorem for our static evaluation language in section 2. ■

One final issue remains to be discussed: do other upgrade rules define a complete logic, and in particular, the Parameterized rule? There is no general results here. But the Parameterized rule does suggest the following reduction axiom. Although it seems a bit clumsy, its validity can be proved in a similar way to axiom 5:

$$[e]q_a^m \leftrightarrow PRE(e) \rightarrow q_a^{m-v_a(e)\cdot\mu(e)\cdot\lambda(e)}$$

However, once we introduce a weight for the previous state, this job becomes harder. If the upgrade rule is functionally expressible, we can still get a complete logic, though clearly substraction will no longer work.

4 Illustration: commands and obligations

So far, we have found a mechanism which represents a plausible view of incoming information that changes preferences. We now illustrate this framework in a different setting, namely deontic logic. Our aim is to show how the logical issues discussed in this paper correspond to real questions of independent interest.

Originally, deontic logic (Åqvist 1987) was the study of assertions of obligation like 'it ought to be the case that φ' (denoted as $O\varphi$) emanating from some moral authority. The standard truth condition for the expression $O\varphi$ is

$$\mathcal{M}, s \models O\varphi \text{ iff for all } t \in S \text{ such that } s \sim t \text{ and } \mathcal{M}, t \models \varphi.$$

The underlying intuition is that φ ought to be case which are true in *all best possible worlds*, as seen from the current one. This naturally suggests an ordering among worlds, and we will see that this allows for a quantitative interpretation.

Likewise, we can think of the deontic setting dynamically: obligations may be changed due to incoming information, or they can be treated as programs or actions themselves. So far, much research in these dynamic aspects has been carried out by [Mey88], [TT99], [Mey96], [Zar03] and so on. The most recent work is [Yam06] (accepted by CLIMA VII, 2006) which takes the dynamic epistemic logic paradigm to obligation changes brought about by acts of commanding in the multi-agent context. Here is the reduction axiom proposed in [Yam06]:

$$[!_a\varphi]O_a\psi \leftrightarrow O_a(\varphi \rightarrow [!_a\varphi]\psi)$$

where the intended interpretation of $O_a\varphi$ is 'it is obligatory for the agent a ($\in G$) that φ', and $[!_a\varphi]$ is intended to represent the action of commanding an agent a to see to it that φ.

It is no surprise that Yamada's system can be translated into the qualitative relation-changing version of preference upgrade proposed in [BL06]. This result hinges on the fact that deontic semantics suggest an ordering among possible worlds. Naturally, the mechanism of evaluation upgrade applies to obligation change as well, but with a more refined view. We can now indicate the 'weight' of a command in terms of the numerical points, as pictured in the following event model:

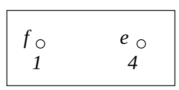

where command e has more strength than f does.

In particular, the current approach also is an improvement in the sense that it brings out insights to the issue of *conflicting commands*, which has been discussed in many papers. Let us first look at a variation of the example in [Yam06]:

Example 4.1 *Suppose you are reading an article in the office you share with your two bosses and a few other colleagues. It is a hot summer noon, the temperature is above 30 degree Celsius. You can open the window, turn on the air conditioner, or concentrate on your reading and ignore the heat. Then your boss A commands you to open the window, your boss B commands you not to do that. What effects do their commands have on the current situation? Which command would you obey?*

A theorem of the form $[!_a(\varphi \wedge \neg\varphi)]O_a\psi$ (Dead End) in [Yam06] handles this problem. It says that contradictory commands lead to an obligational dead end. But this implicitly rules one important aspect, i.e. *the hierarchy of authorities*, out of our scope. Your two bosses may well stand at different authority levels, you may refuse to open the window if your boss B is in a higher position than A. This shows that in a deontic setting, managing conflict is much more than managing consistency. To model the possible contradictory commands carried by different authorities, our current system provides at least one new way of doing this by the following rephrased upgrade rule.

Definition 4.2 (Deontic parameterized rule) *Let $\eta(e)$ be an authority function, and $\lambda(e)$ a relative force function. The domains of these two functions are the set of events, and the ranges are $\mathbb{N}$. Given the value for the previous state s and event e, the new value for state (s, e) is defined by the following:*

$$v_a(s,e) = v_a(s) + v_a(e) \cdot \eta(e) \cdot \lambda(e).$$

Since we are still in the multi-agent context, the relative force applies here very well. Again the agent a may take the boss's commands seriously, whereas agent b may not.

Note that by introducing hierarchy of authorities into the above upgrade rule, we actually deal with the problem *within* the logic. One promising way to handle this issue is to think of the hierarchy as sort of *outside* meta constraints ordering. The idea is from Optimality theory (cf. [PS93]) in which constraints are strictly ordered according to their importance. For a logical investigation concerning constraints and preference change, we refer to [JL06].

One final remark: we have discussed how evaluation upgrade can deal with deontic reasoning in a dynamic style, adding some new twists, such as evaluation of actions of commanding, and resolving conflicts between commands from different agents. This style of analysis is quite general, and it can also be applied to *default reasoning*. Here agents receive incoming information which does not necessarily eliminate worlds, but changes their evaluations of those worlds: more precisely, the *plausibilities* which they assign to these worlds. A typical example is the instruction 'Normally, φ' in [Vel96], which changes the preference ordering between worlds so as to give the φ worlds a higher position. For this same purpose, from the perspective of evaluation upgrade, we can take an event model $\mathcal{E}$ including two events "see φ", "see $\neg\varphi$" with different values (say +1, 0) to model a default 'Normally φ'. Executing the upgrade with $\mathcal{E}$ leads to a new model where the φ-worlds have all gained one point, upgrading their position in the agent's expectation pattern encoded in the plausibilities. In this way, the dynamic evaluation language becomes a sort of default language, where

The expression $[\text{“}see\ \varphi\text{”}]\psi$ plays the role of a default conditional 'if φ then ψ'.

A complete evaluation default logic (EDL) can be deduced directly from our general logic $DEEL$. This new insight leads to the following question, namely, how to compare the overall $DEEL$ to default logic in [Vel96]? My conjecture would be that $DEEL$ seems to be much richer, because by varying the event values in $\mathcal{E}$, one can describe the behavior of a whole *family* of different 'default conditionals'. It all depends on which strengths the agent wishes to assign to the antecedents of those default conditionals.

5 Further logical issues

To get a good understanding of the expressiveness of the evaluation language presented in section 2 we look at some issues concerning bisimulation, a fundamental notion in modal logics. First we formulate the standard bisimulation definition for evaluation models below. The conditions for the epistemic relations $\sim_a$ are omitted, as they are routine.

Definition 5.1 (Evaluation bisimulation) *Let $\mathcal{M} = (S, v_a, V)$ and $\mathcal{M}' = (S', v'_a, V')$ be two evaluation models. A non-empty binary relation $Z \subseteq S \times S'$ is called an evaluation bisimulation between $\mathcal{M}$ and $\mathcal{M}'$ if the following conditions are satisfied:*

(i) *If sZs' then s and s' satisfy the same propositional variables.*

(ii) *If sZs' and $v_a(s) \leq v_a(t)$ (or $s \preccurlyeq_a t$), then there exists t' in $\mathcal{M}'$ such that tZt' and $v'_a(s') \leq v'_a(t')$ (or $s' \preccurlyeq_a t'$) (the forth condition).*

(iii) *If sZs' and $v'_a(s') \leq v'_a(t')$ (or $s' \preccurlyeq_a t'$), then there exists t in $\mathcal{M}$ such that tZt' and $v_a(s) \leq v_a(t)$ (or $s \preccurlyeq_a t$) (the back condition).*

Example 5.2 From the view point of the above evaluation bisimulation, it would make sense to identify the following two models, where we identify worlds by their evaluations:

After all, the pure preference pattern is the same in both. But the evaluations make a difference in the evaluation language. Consider the event model $\mathcal{E}$ which upgrades all φ-worlds (s in the pictures) with 1 each time it is applied. Applying $\mathcal{E}$ once to the model on the left keeps the preference intact, but on the right, it voids it. All this seems to suggest that we need a new bisimulation definition for evaluation models to express the intensity of preferences. Here is one proposal.

Definition 5.3 (Distance) *The distance between two possible states s and t in an evaluation model is defined as $\mathcal{D}_a(s,t) = \mid v_a(s) - v_a(t) \mid$.*

In Example 5.2 the distance between s and t is 2 in the model on the left, but it is 1 on the right.

Definition 5.4 (Distance bisimulation) *Let $\mathcal{M} = (S, v_a, V)$ and $\mathcal{M}' = (S', v'_a, V')$ be two evaluation models. A non-empty binary relation $Z \subseteq S \times S'$ is called distance bisimulation between $\mathcal{M}$ and $\mathcal{M}'$ if the following conditions are satisfied:*

(i) *If sZs' then s and s' satisfy the same propositional variables.*

(ii) *If sZs', $s \leq t (t \leq s)$ and $\mathcal{D}_a(s,t) = k$, then there exists t' in $\mathcal{M}'$ such that tZt', $s' \leq t'(t' \leq s')$ and $\mathcal{D}_a(s',t') = k$ (the forth condition).*

(iii) *If sZs', $s' \leq t'(t' \leq s')$ and $\mathcal{D}_a(s',t') = k$, then there exists t in $\mathcal{M}$ such that tZt', $s \leq t(t \leq s)$ and $\mathcal{D}_a(s,t) = k$ (the back condition).*

As usual, we say two evaluation models are *bisimilar* when there is some evaluation bisimilation linking two states in the two models. Intuitively, if the *same efforts* (same distance) are made to get from one state to another in each model, then the two models are bisimilar.

This means that with the notion of comparative distance, we can say sentences like 'd_1 is preferable over d_2 more than d_1 is preferable over d_3', which simply means $\mathcal{D}(s_1, s_2) > \mathcal{D}(s_1, s_3)$ in the model, where d_1, d_2 and d_3 are chosen in s_1, s_2 and s_3, respectively. This is what most languages of qualitative preference are not able to do. Following this line may be related somehow to the modal languages for 'geometry' studied in [BGKV06].

6 Conclusions

We have presented here a quantitative semantic of preference in terms of evaluation functions. A new language with propositional constants was proposed and it turned out to be both concise and expressive. Moreover, such a quantitative perspective suggests a different way to deal with preference changes when processing new information. We followed the standard mechanism of product update, and proposed a new Addition rule and a new Parameterized rule to characterize the subtleties of value changes. A complete dynamic epistemic evaluation logic was presented for the evaluation upgrade. We then shifted to the deontic setting and showed that the current mechanism applies there as well, in particular, it provides a way to solve the issue of contradictory obligations. Finally, we ended up with a new technical result concerning bisimulation for evaluation models. As an immediate follow-up we would like to pursue how these abstract results can be used to analyze further problems in decision theory and game theory.

Acknowledgement

Special thanks go to J. van Benthem, U. Endriss, D. de Jongh, E. Pacuit, F. Roelofsen, B. Semmes, T. Yamada, J. Zvesper for their helpful comments.

References

[Auc03] G. Aucher. A combination system for update logic and belief revision. Master's thesis, ILLC, University of Amsterdam, 2003.

[Ben96] J. van Benthem. *Exploring Logical Dynamics*. CSLI Publication, Stanford, 1996.

[BGKV06] P. Balbiani, V. Goranko, R. Kellerman, and D. Vakarelov. Logics for geometric structures. In M. Aiello, I. Pratt-Hartmann and J. van Benthem, Handbook of Spatial Logics, see http://dit.unitn.it/ aiellom/hsl/index.html, 2006.

[BL06] J. van Benthem and F. Liu. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logic*, 2006. To appear.

[BMS98] A. Baltag, L.S. Moss, and S. Solecki. The logic of common knowledge, public announcements, and private suspicions. In I. Gilboa, editor, *Proceedings of the 7th conference on theoretical aspects of rationality and knowledge (TARK 98)*, pages 43–56, 1998.

[DHK06] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer, Berlin, 2006. To appear.

[Ger99] J. Gerbrandy. *Bisimulation on Planet Kripke*. PhD thesis, ILLC, Amsterdam, 1999.

[JL06] D. de Jongh and F. Liu. Optimality, belief and preference. Tech Report PP-2006-38, ILLC, University of Amsterdam, 2006.

[Liu04] F. Liu. Dynamic variations: Update and revision for diverse agents. Master's thesis, ILLC, University of Amsterdam, 2004.

[Liu06] F. Liu. Diversity of agents. Tech Report PP-2006-37, ILLC, University of Amsterdam, 2006.

[Mey88] J-J.Ch. Meyer. A different approach to deontic logic: Deontic logic viewed as a variant of dynamic logic. *Notre Dame Journal of Formal Logic*, 29:109–136, 1988.

[Mey96] R. van der Meyden. The dynamic logic of permission. *Journal of Logic and Computation*, 6:3:465–479, 1996.

[Pla89] J. A. Plaza. Logics of public announcements. In *Proceedings 4th International Symposium on Methodologies for Intelligent Systems*, 1989.

[PS93] A. Prince and P. Smolensky. *Optimality Theory: Constraint Interaction in Generative Grammar*. Malden, Ma: Blackwell, 1993.

[Spo88] W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W. L. Harper et al., editor, *Causation in Decision, Belief Change and Statistics II*, pages 105–134. Kluwer, Dordrecht, 1988.

[TT99] L. van der Torre and Y. Tan. An update semantics for deontic reasoning. In P. McNamara and H. Prakken, editors, *Norms, Logics and Information Systems*, pages 73–90. IOS Press, 1999.

[Vel96] F. Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25:221–261, 1996.

[Yam06] T. Yamada. Commands and changing obligations. Accepted by the Seventh International Workshop on Computational Logic in Multi-Agent Systems (CLIMA VII). Department of Philosophy, University of Hokkaido, 2006.

[Zar03] B. Zarnic. Imperative change and obligation to do. In K. Segerberg and R. Sliwinski, editors, *Logic, Law, Morality: Thirteen Essays in Practical Philosophy in Honour of Lennart Aqvist*, pages 79–95. Uppsala philosophical studies 51. Uppsala: Department of Philosophy, Uppsala University, 2003.