Dynamic Logic of Preference Upgrade

Johan van Benthem & Fenrong Liu

November 3, 2005

1 Introduction: Changing Preferences

The notion of preference occurs across many areas, including the philosophy of action, decision theory, optimality theory, and game theory. In these settings, individual preferences between worlds or actions can be used to predict behavior by rational agents. In a more abstract sense, the notion of preference also occurs in conditional logic, non-monotonic logic and belief revision theory, whose semantics involve an ordering of the possible worlds in terms of relative similarity or plausibility, or other preference-like relations.

Preference logics There are various preference logics in the literature which can describe this sort of comparative structure by means of various devices (Hanson 1990). This diversity reflects the rich texture of the subject. Thus, agents' preferences can run between worlds or between actions, preference statements can be weaker or stronger in what they say about those worlds or actions being compared – and also, they may be more 'austere' or more 'lush'. A statement like "I prefer sunsets to sunrises" can be cast merely in terms of 'what is better according to me', or as some more complex propositional attitude involving my beliefs about the relevant events. In this paper, we take an austere modal approach, where a binary preference relation support a unary modality "true in some world which is at least as good" (Boutilier 1994, Halpern 1997). Van Benthem, van Otterloo & Roy 2005 show how such a language, when extended with a few operators from hybrid languages, can define and analyze quite a few notions such as different conditionals, Nash equilibrium, and backward induction solutions to games. This language can express various notions of preference that agents may have between propositions, i.e., types of event. Moreover, eventually, we add explicit epistemic operators to the language, allowing us to express agents' attitudes toward what is good or better for them.

Preference dynamics Our main concern in this paper, however, is one of *dynamics*. Preferences are not static, but they can change through commands of moral authorities, suggestions from friends, or just changes in our own evaluation of the world and our possible actions. This process of change can have various triggers. For instance, intuitively, a command

"See to it that φ!"

makes worlds where φ holds preferred over those where it does not - at least, if we accept the preference induced by the issuer of the command. But also a process of planning, with just our own goals in mind, may gradually introduce preferences over actions as ways toward reaching the goal, as we learn more about the actual world. These and other dynamic aspects of preference have been noted by many authors, including van Benthem, van Eijck & Frolova 1993, Žarnić 2003, Tan & van der Torre 1998, and Yamada 2005.

Moreover, related ideas play in the dynamic semantics for conditional logics (Spohn 1988, Veltman 1996). In its static Lewis-style semantics, a conditional $\varphi \Rightarrow \psi$ says roughly that

ψ is true in all most-preferred φ-worlds $\quad (\natural)$

But one plausible way of taking a conditional is, not as a true/false description of a current preference, but rather as an instruction for *adjusting* that preference so as to make $\natural$ the case. Even more simply, consider a so-called default assertion like

"*Normally* φ"

As Veltman 1996 points out, this does not eliminate $\neg\varphi$-worlds from our current model, in the usual dynamic sense of information update on our current range of worlds. Accommodating this assertion rather makes the $\neg\varphi$-worlds doxastically less preferred than φ-worlds.

Trigger 1: suggestions Our general view is that there are many triggers for preference change, and dynamic preference logics should provide a format for studying these in an appropriate generality. To find such formats, however, in this paper, we start from a simple test scenario that may be called a 'suggestion'. Consider someone who has no preference between taking a trip (p) or staying at home ($\neg p$). Now a person comes along and says

"Let's take a trip!"

The way we think of this, 'taking' that suggestion means that any preference we might have had for staying at home is removed from the current model. Here is an intuitive picture of the sort we have in mind for this:

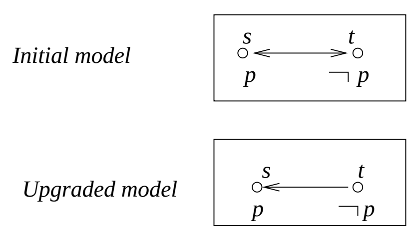

Thus, in our scenario, a suggestion removes already existing preference links: but it does not add new ones. This mechanism will be studied in greater detail later on, and it is our point of entry into more general kinds of preference upgrade. Even so, by way of contrast, here is one illustration of an alternative, which does not remove links, but rather adds them.

Trigger 2: commands In the above picture, the agent now prefers the trip, so this has become her priority, or in another, more deontic reading, even her duty. But in general, suggestions are weaker than commands. Taking the suggestion does not mean, however, that the person will now prefer all p-worlds to the $\neg p$-ones. It all depends on the setting, i.e., the preference structure already in place. If the agent was indifferent between p and $\neg p$ with arrows both ways, then indeed, the suggestion will induce a preference. But the agent may also be unable to compare the two situations, as in the following model with two entirely unrelated worlds:

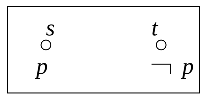

In this case, a suggestion in the preceding relation-decreasing sense does not make the worlds comparable. With real commands "Take that trip!", however, we want to make sure that the agent now prefers p. In that case, we need to *add* preference links to the picture, making the world with $\neg p$ less preferred. Our eventual technical proposals to follow can also deal with upgrades that add links, and this alternative will be mentioned at various places.

Incidentally, the difference between link elimination and link addition may not be all that great from a technical perspective. Adding links to a relation R may also be viewed as eliminating links from the *complement* of R.

Dynamic logics of upgrade Construed either way, eliminative or additive, preference change is highly reminiscent of existing systems for information *update* in dynamic-epistemic logic (Gerbrandy 1999, Baltag, Moss & Solecki 1998, van Benthem 2002, van Ditmarsch, van der Hoek & Kooi 2006). In the latter paradigm, incoming assertions or observations change the domain of the current model and/or its accessibility relations. In our scenario, current preference relations are changed by incoming suggestions or commands. Thus, we will speak henceforth of preference *upgrade* as a counterpart to the better-known term update. The main point of this paper is to show that preference upgrade is a viable phenomenon, which is just as susceptible to systematic modification as information, temporal perspective, or other parameters of 'logical dynamics'. (Cf. van Benthem, van Eijck & Frolova 1993, van Benthem 1996, or in the setting of conditional logic, Spohn 1988, Veltman 1996). We will show how this dynamics can be implemented by the very same methodology that has been developed for information update in dynamic-epistemic logic.

This paper is structured as follows. First we present a new epistemic preference logic, which combines preference conceived modally with standard epistemic operators (Section 2). Its semantics is based on preferences between worlds. This allows us to talk about knowing or not knowing one's preferences, or regretting that the best scenario is not going to happen. Next, in Section 3, we provide formal definitions for preference upgrade, with an emphasis on the above 'suggestions' increasing our preference for one proposition over its negation. Interestingly, this also suggests alternative formulations for information update. Section 4 defines a dynamic version of the static epistemic preference language, where a mechanism of information update lives together with one of preference upgrade. There is a completeness theorem in terms of the usual style of reduction axioms recursively analyzing postconditions of actions. This is our first 'existence proof' for a compositional dynamics of upgrade, in tandem with update of information. Next, in Section 5, we consider more general upgrade scenarios: first with general schemes of link elimination, and finally, with the full strength of 'product update' for information using 'event (action) models'. This requires enriching the event models of dynamic-epistemic logic with agents' preferences between actions or events. Again, complete logics turn out to exist, so the analogy between update and upgrade extends all the way. Even so, in Section 6, we show that some intuitions of upgrade do not fit the product update pattern, as they depend essentially on numerical utilities, representing 'strength of preference'. We then present an alternative dynamic approach, viz. a complete dynamic-epistemic logic which also performs 'utility upgrade', on the analogy of current belief revision systems that manipulate 'graded models' in the sense of Spohn (Aucher 2003, Liu 2004). Section 7 then outlines some applications of our dynamic upgrade logics, in both formats considered here, to default reasoning, deontic logic, and logics of commands. In Section 8, we briefly compare preference upgrade and utility upgrade, though their precise relationship is left open in this paper: they may just be two intuitive styles of thinking that do not reduce to each other. Section 9 is a brief survey of related work, and Section 10 contains our conclusions about what we have right now and what we want to do next.

This paper is a proposal for a certain style of thinking about preference upgrade, and an existence proof for a logical methodology in doing so. We do not claim to have addressed all intuitive senses of the notion of preference, or all logical issues arising in the areas where it plays a role. A more extensive discussion of plausible upgrade mechanisms with various triggers, various senses of preference (short-term and long-term), and further applications, may be found in Liu 2005.

2 Epistemic Preference Logic

2.1 Language and semantics

The main language used in this paper has two components. There is a preference modality as in van Benthem, van Otterloo & Roy 2005, and in addition, we have the standard knowledge operators from epistemic logic.

Definition 2.1 *(Language) Take a set of propositional variables P and a set of agents I, with p ranging over P and i over I. The* ***epistemic preference language*** *is given by:*

$$\varphi ::= \bot \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_i\varphi \mid [pref]_i\varphi \mid U\varphi.$$

Intuitively, $K_i\varphi$ stands for 'agent i knows that φ', while $[pref]_i\varphi$ says that the agent considers all worlds with φ as least as good as the current one [1].

How is this language connected to preference? $\langle pref\rangle_i$ may be read, with some poetic licence, as 'agent i prefers φ'. But this is misleading, as 'preference' has many senses beyond this. A more neutral reading is just as a statement about what is better according to the agent: 'φ holds in some world which i regards as at least as good'. More elaborate senses of preference can then be dealt with by further definitions. Our three modal assertions of what may be called knowledge, 'better-ness' and 'world access', can express quite a few notions of preferences between propositions (van Benthem, van Otterloo & Roy 2005). For instance,

$$U(\psi \to \langle pref\rangle_i\varphi)$$

expresses one strong sense of 'agent i prefers φ to ψ', viz. each ψ-world has at least one epistemic alternative which is at least as good according to the agent. Of course, many more notions can be defined in the same vein. We take this expressive power of our language for granted here, but we continue with our set of base modalities, as these 'decompose' more complex preference statements in a perspicuous manner, while allowing for a simple dynamic approach later on.

Definition 2.2 *(Models) An* ***epistemic preference model*** *is a tuple $\mathcal{M}$=(S, $\{\sim_i \mid i \in I\}$, $\{\preceq_i \mid i \in I\}$, V), with S a set of possible worlds, $\sim_i$ the usual equivalence relation of epistemic accessibility for agent i, and V a valuation for proposition letters. Moreover, $\preceq_i$ is a reflexive and transitive relation over the worlds.*

Here, we read $s \preceq_i t$ as saying that 't is at least as good for agent i as s', or 't is weakly preferred to s'. If $s \preceq_i t$ but not $t \preceq_i s$, then t is *strictly preferred* to s, written as $s \prec_i t$. If $s \preceq_i t$ and $t \preceq_i s$, then agent i is *indifferent* between s and t. Models can also have a distinguished actual world, but we rarely use this feature here.

Note that we do not require that our preference relations be *connected*, say, in the sense of the Lewis sphere models for conditional logic. In general, we want to allow for genuinely incomparable worlds where an agent has no preference either way, not because she is indifferent, but because she has no means of

[1] For technical convenience, we often shift to the corresponding existential modalities $\langle K\rangle_i$, $\langle pref\rangle_i$, and $E\varphi$. These seem more difficult to read in terms of intuitive linguistic expressions. But they help in finding and checking valid principles, and in semantic arguments generally.

comparing the worlds at all. This is just as in the semantics for the minimal conditional logic. Of course, in special settings, such as the utility-based preference orderings of outcomes in a game, connectedness may be quite appropriate.

Definition 2.3 *(Semantics) Given an epistemic preference model* M=$(S, \{\sim_i | i \in I\}, \{\preceq_i | i \in I\}, V)$*, and a world* $s \in S$*, we define* $\mathcal{M}, s \models \varphi$ *(formula* φ ***is true in*** $\mathcal{M}$ ***at*** s*) by induction on* φ*:*

1. $\mathcal{M}, s \models p$ *iff* $s \in V(p)$
2. $\mathcal{M}, s \models \neg\varphi$ *iff not* $\mathcal{M}, s \models \varphi$
3. $\mathcal{M}, s \models \varphi \wedge \psi$ *iff* $\mathcal{M}, s \models \varphi$ *and* $\mathcal{M}, s \models \psi$
4. $\mathcal{M}, s \models K_i\varphi$ *iff for all* $t : s \sim_i t$ *implies* $\mathcal{M}, t \models \varphi$
5. $\mathcal{M}, s \models [pref]_i\varphi$ *iff for all* $t : s \preceq_i t$ *implies* $\mathcal{M}, t \models \varphi$
6. $\mathcal{M}, s \models U\varphi$ *iff for all* t*:* $\mathcal{M}, t \models \varphi$*.*

Expressive power As we noted, van Benthem, van Otterloo & Roy 2005 have shown that the pure modal preference part of this language, with the help of the universal modality, can express a variety of natural notions of preference between propositions. Moreover, following Boutilier 1994, they show that this language can faithfully embed non-iterated conditionals $\varphi \Rightarrow \psi$ using the above preference operator $\langle pref \rangle_i$, as follows:

$$U(\varphi \rightarrow \langle pref \rangle_i(\varphi \wedge [pref]_i(\varphi \rightarrow \psi)).$$

But with our additional epistemic operators, we can express more, viz. the interplay of preference and knowledge. Here are two examples. The first represents an intuition of self-reflection that one may or may not have as part of the notion of 'preference', the second describes an unfortunate but ubiquitous phenomenon:

1. $\langle pref \rangle_i\varphi \rightarrow K_i\langle pref \rangle_i\varphi$: Positive Introspection
2. $\langle pref \rangle_i\varphi \wedge K_i\neg\varphi$: Regret.

We will return to epistemic-preference principles later on. Incidentally, limitations to expressive power can be shown using the obvious notion of *bisimulation* for the epistemic preference language.

2.2 Proof System and Completeness

Our epistemic preference logic can be axiomatized completely in a standard modal style, given our choice of epistemic preference models (cf. Blackburn, de Rijke & Venema 2001).

Theorem 2.4 *(Completeness) Epistemic preference logic is completely axiomatizable w.r.t epistemic-preference-models.*

Proof. The proof of completeness is standard, and an axiom system can easily be extracted from the literature. On top of the principles of the minimal modal logic for all separate operators K_i and $[pref]_i$, it has the axioms of $S4$ for the $[pref]_i$, and $S5$ for the K_i and the universal modality U. Finally, also as usual, there are two axioms relating the operators:

$$\text{(a) } U\varphi \to K_i\varphi, \quad \text{(b) } U\varphi \to [pref]_i\varphi.$$

■

Additional axioms in our language will impose further frame conditions on top of our general models. Here are two examples, which can be proved by standard modal frame-correspondence techniques:

Fact 2.5

- *A preference frame* $\mathcal{F} = (S, \{\sim_i \mid i \in I\}, \{\preceq_i \mid i \in I\})$ *satisfies connectedness* $\forall x \forall y : x \preceq_i y \vee y \preceq_i x$, *iff the following formula is true in the frame:* $(\varphi \wedge E\psi) \to \langle pref\rangle_i\psi \vee E(\psi \wedge \langle pref\rangle_i\varphi)$.
- *An epistemic preference frame* $\mathcal{F}$ *makes the Introspection Axiom* $\langle pref\rangle_i\varphi \to K_i\langle pref\rangle_i\varphi$ *true iff it satisfies the following condition:*

$$\forall s \forall t \forall u : (s \preceq_i t \wedge s \sim_i u \to u \preceq_i t)$$

Nevertheless, we will work with the minimal system described above in this paper, leaving such extras to asides.

3 Modelling Preference Upgrade

3.1 Brief review of epistemic information update

The basic paradigm for epistemic update is public announcement. Suppose that an agent does not know if p is the case, but learns this fact through an announcement $!p$. Then we get the following sort of model change, where the dotted line indicates the agent's uncertainty in the initial situation:

Initial model

Updated model

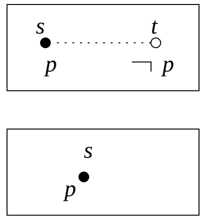

The announcement eliminates the $\neg p$-world, and afterwards, the agent knows that p. There is an extensive literature on dynamic epistemic logics for public announcements and more sophisticated epistemic events, that can modify information in different ways for different agents. We refer to Baltag, Moss & Solecki 1999, van Benthem, van Eijck & Kooi 2005; and the references in Section 5.4 on 'product update'.

These logics all work essentially on the same design principle. First, a class of models is chosen representing the relevant information structures, together with some appropriate static language for describing these. Usually, these are models for some version of standard epistemic logic. Next, an update mechanism is proposed which transforms given models under some chosen set of epistemic actions. For public announcement, this works by simple elimination of worlds, yielding a definable submodel:

> *A public announcement* $!\varphi$ *of a true proposition* φ *turns the current model* $(\mathcal{M}, s)$ *with actual world* s *into the model* $(\mathcal{M}_{!\varphi}, s)$ *whose worlds are just the set* $\{w \in S \mid \mathcal{M}, w \models \varphi\}$.

For more complex actions, a construction is used which forms *products* $\mathcal{M} \times \mathcal{E}$ of the current epistemic model $\mathcal{M}$ with some 'event model' $\mathcal{E}$ containing all relevant events or actions.

Next, the static language gets a dynamic extension where the informative events themselves are displayed and manipulated. For public announcement, a typical assertion of this sort is

> $[!\varphi]K_i\psi$: *after a successful public announcement of* φ, *agent* i *knows that* ψ.

Here the semantic clause for the dynamic modality is simply as follows:

$$\mathcal{M}, s \models [!\varphi]\psi \text{ iff (if } \mathcal{M}, s \models \varphi\text{, then } \mathcal{M}_{!\varphi},\ s \models \psi)$$

Usually, the effects of events can then be described completely in a recursive manner, leading to a compositional analysis of communication and other cognitive processes. As a crucial illustration, here is the key *reduction axiom* in current logics of public announcement for a true assertion resulting in an epistemic possibility for agent i:

$$\langle !\varphi\rangle\langle K\rangle_i\psi \leftrightarrow \varphi \wedge \langle K\rangle_i\langle !\varphi\rangle\psi$$

As discussed in the literature, semantically, this reflects a sort of perfect recall for updating agents. Computationally, axioms like this help drive a reduction algorithm for dynamic epistemic statements to static epistemic statements, allowing us to borrow known decision procedures for the base language.

3.2 Upgrade as relation change

With the paradigm of public announcement in mind, we now define the mechanism of preference change described informally in the above. Our static models are of course the epistemic preference structures of Section 2:

$$M = (S, \{\sim_i \mid i \in I\}, \{\preceq_i \mid i \in I\}, V)$$

Our triggers are events of publicly suggesting φ, written as follows:

$$\sharp\varphi$$

These lead to the following model change, removing preferences for $\neg\varphi$ over φ:

Definition 3.1 *Given any epistemic preference model* $(\mathcal{M}, s)$*, the* ***upgraded model*** $(\mathcal{M}_{\sharp\varphi}, s)$ *is defined as follows.*

(a) $(\mathcal{M}_{\sharp\varphi}, s)$ *has the same domain, valuation, epistemic relations, and actual world as* $(\mathcal{M}, s)$*, but*

(b) the new preference relations are now
$\preceq_i^* = \preceq_i - \{(s,t) \mid \mathcal{M}, s \models \varphi$ *and* $\mathcal{M}, t \models \neg\varphi\}$.

We suppress agent subscripts henceforth whenever convenient. It will be clear that this definition fits our motivating examples, and others are easily found.

This upgrade for events of suggestion replaces a preference relation by a definable subrelation. This may also be written as follows in the standard notation of dynamic logic:

$$R := R - (?\varphi; R; ?\neg\varphi)$$

We will consider more general relation-changing operations in Section 5. For instance, if one wanted to add links, rather than just subtract them, the format would still work. E.g., the relation-extending stipulation

$$R := R \cup (?\neg\varphi; \top; ?\varphi)$$

where $\top$ is the universal relation, would make every φ-world preferable to every $\neg\varphi$-world. With our upgrade defined, we are in a position to define a dynamic language for preference upgrade. But before doing so in Section 4, we consider some features of the mechanism just defined.

Preservation properties of upgrade Perhaps the most pressing issue in our system is whether a proposed model changing operation stays inside the class of intended static models. For the update associated with public announcements $!\varphi$, this was so - and the reason is the general logical fact that submodels preserve *universally defined* relational properties like reflexivity, transitivity, and symmetry. For our notion of upgrade, the properties to be preserved are reflexivity and transitivity of preference relations (epistemic relations remain unchanged). This time, however, no general result comes to the rescue, since we only have the following counterpart to the preservation result for submodels:

Fact 3.2 *The first-order properties preserved under taking subrelations are precisely those definable using negated atoms, $\wedge$, $\vee$, $\exists$, $\forall$.*

But neither reflexivity nor transitivity is of this particular syntactic form. Nevertheless, using some special properties of our proposal, we can prove

Fact 3.3 *The operation $\mathcal{M}_{\sharp\varphi}$ preserves reflexivity and transitivity.*

Proof. Reflexivity is preserved since we never delete loops (s, s). As for transitivity, suppose that $s \preceq^* t \preceq^* u$, while not $s \preceq^* u$. By the definition of $\sharp\varphi$, we must then have $\mathcal{M}, s \models \varphi$ and $\mathcal{M}, u \models \neg\varphi$. Consider the intermediate point t. Case 1: $\mathcal{M}, t \models \varphi$. Then the link (t, u) should have been removed from $\preceq$. Case 2: $\mathcal{M}, t \models \neg\varphi$. In this case, the link (s, t) should have been removed. Either way, we have a contradiction. ■

On the other hand, our upgrades $\sharp\varphi$ can lead to loss of connectedness of the preference order. Our earlier example already showed this. If an agent prefers a φ-world to an $\neg\varphi$-world, and that is all, then the upgrade $\sharp\varphi$ will remove that only link between the two, giving a model in which the two options have become incomparable. Likewise, our upgrades can lead to a loss of positive introspection:

Example 3.4 Consider the following two-world model, where we only draw non-reflexive arrows for the preference relation $\preceq$. The agent does not know whether p holds, but she does know that p may be better for her, as both worlds see a preferred p-world in her preference order:

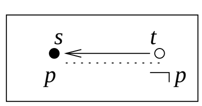

Now, an upgrade command $\sharp\neg p$ suggesting that $\sharp\neg p$ is not so bad after all, would change this model to the following one:

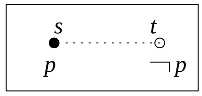

Let us suppose that the p-world is the actual one. Here, $\langle pref\rangle p$ holds, because of the reflexivity of $\preceq$. But in the epistemically indistinguishable $\sharp p$-world, $\langle pref\rangle p$ is false. Thus, even though p may be better for her, the agent does not know this any more. A real scenario where this happens–admittedly, a bit contrived – might be one where we do not know if we are sleeping or awake. Let us say that, originally, in both cases, we prefer being asleep. Now an upgrade happens, suggesting that real waking life might not be so bad after all. In that case, there being a better waking world – as seen from the current one – contains more information than before (as 'better' has become a more demanding notion), and indeed, knowing it would tells us whether we are awake or asleep!

In some settings, preference introspection seems plausible, and a desirable property of models to be preserved. We can then change the above notion of upgrade to deal with this, e.g., by making sure that similar links are removed at epistemically indistinguishable worlds, or study which special sorts of upgrade in our language have the property of always preserving preference introspection. The latter would then be the 'reasonable' or 'sensible' series of suggestions. As it happens, the utility upgrade mechanism of Section 6 is of the latter kind.

Update by link cutting Update and upgrade do not lead wholly separate lives in our setting. For instance, if we want to model the earlier phenomenon of 'regret' about worlds that are no longer viable options, epistemic updates for $!\varphi$ should not remove the $\neg\varphi$-worlds, since we might still want to refer to them, and perhaps even mourn their absence. One way of doing this is by redefining the update for public announcement as a relation-changing operation, too, of 'link cutting'. This time, instead of the above $!\varphi$, we write the relevant update action as follows:

$$\varphi!$$

and we write the updated model as $\mathcal{M}_{\varphi!}$ in order to distinguish it from that we have by eliminating worlds. More precisely, we would have to change notations in our formal languages to reflect the two varieties of exclamation maker now – but we trust the reader can disambiguate in context. The corresponding semantic operation for $\varphi!$ on models is this:

Definition 3.5 *The* ***modified public update model*** $\mathcal{M}_{\varphi!}$ *is the original model* $\mathcal{M}$ *with its worlds and valuation unchanged, but with accessibility relations* $\sim_i$ *replaced by the following version where no crossing is possible between the* φ*- and* $\neg\varphi$*-zones of the model:*

$$(?\varphi; \sim_i; ?\varphi) \cup (?\neg\varphi; \sim_i; ?\neg\varphi)$$

Fact 3.6 *The pure epistemic logic of public announcement is the same with* $!\varphi$ *and with* $\varphi!$.

Proof. The change from world elimination to link elimination makes no difference for dynamic epistemic formulas, since the $\neg\varphi$-zone will become inaccessible from the φ-zone, and hence irrelevant to semantic evaluation of all dynamic epistemic formulas starting from worlds satisfying φ. ■

Nevertheless, the second update stipulation has some advantages. It was first proposed, to our knowledge, in Snyder 2004 (cf. van Benthem & Liu 2004) for modelling the behavior of *memory-free* agents, whose epistemic accessibility relations are quite different from those for the idealized update agents of standard dynamic epistemic logic. Moreover, in the present setting, in stating regrets, we need the consistency of a formula like

$$K_i p \wedge \langle pref \rangle_i \neg p$$

Yes, I know that p, but it would be better if it weren't... Modified update allows us to have this consistently.

But link cutting has some curious features, too. E.g., link cutting in the current model is the same for announcements $\varphi!$ and $(\neg\varphi)!$: both remove links between φ-worlds and $\neg\varphi$-ones. The only difference is that the former can only take place at a current world which satisfies φ, and the latter in one satisfying $\neg\varphi$. This analogy is reflected in valid principles of the logic, but we do not pursue this purely update-related phenomenon here.

Discussion: Update and upgrade Distinguishing the two versions of information update also leads to a subtle distinction in a combined update-upgrade logic. If processing $!\varphi$ eliminates all worlds we know to be non-actual, our preference statements adjust automatically to what we know about the facts. This is the behavior of realists, who never cry over spilt milk. For those realists i, the following combined announcement/preference principle will be valid, at least for atomic statements p which do not change their truth values by being announced

$$[!p][pref]_i p$$

But this principle will not be valid for more poetic souls, who still deplore the way things have turned out to be. For them, update amounts to link-cutting $\varphi!$, they stick to their preferences between all worlds, and the new fact learnt may even introduce regrets:

$$\langle pref \rangle_i \neg p \wedge \langle K \rangle_i \neg p \rightarrow [p!](\langle pref \rangle_i \neg p \wedge K_i p)$$

Remark 3.7 (Diversity of Agents) This is one more case of legitimate differences between updating agents. Liu 2004 presents new update logics that can model agents differing in memory capacities, and policies for belief revision.

4 Dynamic Epistemic Upgrade Logic

4.1 Language and semantics

Now we can introduce an enriched dynamic language for update and upgrade. Its static part is the earlier language of Section 2, but its action vocabulary now contains both public announcements $\varphi!$ and suggestions $\sharp\varphi$. Adding the original world-eliminating announcements $!\varphi$ is a routine matter, so we highlight the behaviour of the less standard variant.

Definition 4.1 *(Language) Take a set of propositional variables P and a set of of agents I, with p ranging over P and i over I. The **dynamic epistemic preference language** is given by:*

$$\varphi ::= \bot \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_i\varphi \mid [pref]_i\varphi \mid U\varphi \mid [\pi]\varphi$$
$$\pi ::= \varphi! \mid \sharp\varphi.$$

We could also add the usual program operations of composition, choice, and iteration from propositional dynamic logic to the action vocabulary - but in this paper, we will have no special use for these. The new language can be interpreted on epistemic preference models in the following way, where we choose the 'regret' variant of update for the novelty:

Definition 4.2 *(Semantics) Given an epistemic preference model $\mathcal{M}$, the **truth definition** for formulas is as before, but now with two additional key clauses for the action modalities:*

$$(\mathcal{M}, s) \models [\varphi!]\psi \textit{ iff } \mathcal{M}_{\varphi!}, s \models \psi$$
$$(\mathcal{M}, s) \models [\sharp\varphi]\psi \textit{ iff } \mathcal{M}\sharp\varphi, s \models \psi$$

4.2 Preference upgrade logic

On epistemic preference models, all valid principles of the static language are still available, as described in Section 2. Moreover, the usual axioms for public announcement hold, be it with one twist. As we saw, the usual updates $!\varphi$ eliminate all $\neg\varphi$-worlds, but updates $\varphi!$ leave all worlds in the model, cutting links instead. This makes no difference with pure epistemic dynamic axioms, but it does with axioms for global existential modalities referring to the whole domain of the model. The usual reduction axiom reads like this:

$$\langle !\varphi\rangle E\psi \leftrightarrow \varphi \wedge E\langle !\varphi\rangle\psi$$

But the axiom below is different, as $E\varphi$ can still refer to worlds after the update which used to be $\neg\varphi$. Further comments on valid axioms will be found below. We now focus on what is new here: upgrade, and its interplay with modified update. It is easy to see the soundness of the following principles, stated with existential modalities for convenience:

Theorem 4.3 *(Soundness) The following formulas are valid:*

1. $\langle\varphi!\rangle p \leftrightarrow (\varphi \wedge p)$
2. $\langle\varphi!\rangle\neg\psi \leftrightarrow (\varphi \wedge \neg\langle\varphi!\rangle\psi)$
3. $\langle\varphi!\rangle(\psi \wedge \chi) \leftrightarrow ((\langle\varphi!e\psi \wedge \langle\varphi!\rangle\chi)$
4. $\langle\varphi!\rangle\langle K\rangle_i\psi \leftrightarrow (\varphi \wedge \langle K\rangle_i\langle\varphi!\rangle\psi)$
5. $\langle\varphi!\rangle\langle pref\rangle_i\psi \leftrightarrow (\varphi \wedge \langle pref\rangle_i\langle\varphi!\rangle\psi)$
6. $\langle\varphi!\rangle E\psi \leftrightarrow \varphi \wedge E(\langle\varphi!\rangle\psi \vee \langle\neg\varphi!\rangle\psi)$

7. $\langle\sharp\varphi\rangle p \leftrightarrow p$
8. $\langle\sharp\varphi\rangle\neg\psi \leftrightarrow \neg\langle\sharp\varphi\rangle\psi$
9. $\langle\sharp\varphi\rangle(\psi \wedge \chi) \leftrightarrow ((\langle\sharp rphi\rangle\psi \wedge \langle\sharp\varphi\rangle\chi)$
10. $\langle\sharp\varphi\rangle\langle K\rangle_i\psi \leftrightarrow \langle K\rangle_i\langle\sharp\varphi\rangle\psi$
11. $\langle\sharp\varphi\rangle\langle pref\rangle_i\psi \leftrightarrow (\neg\varphi \wedge \langle pref\rangle_i\langle\sharp\varphi\rangle\psi) \vee ((\langle pref\rangle_i(\varphi \wedge \langle\sharp\varphi\rangle_i\psi))$
12. $\langle\sharp\varphi\rangle E\psi \leftrightarrow E\langle\sharp\varphi\rangle\psi$

Proof. The first four formulas are the well-known valid reduction axioms for public announcement. The fifth formula, about commutation of $\langle\varphi!\rangle$ and $\langle pref\rangle_i$, expresses the fact that epistemic update does not change any preference relations. The special case of $E\varphi$ has been commented on above.

Next comes a similar set of reduction principles for upgrade. Axiom 7 is like Axiom 1, but simpler - as there is no precondition for $\sharp\varphi$: this operation can always be performed. Given that, we just state that atomic facts do not change under upgrade. The next two axioms express that upgrade is a function. Then comes a commutation principle for preference and knowledge which reflects the fact that upgrade does not change any epistemic relations.

Axiom 11 is crucial, as it encodes precisely by which definition we have changed the preference relation. It says essentially this. After an upgrade for φ, a preference link will lead from the current world to a φ-world iff this same link existed before. This means that it has not been removed, ruling out the case where it led from an actual world which verified φ to some other one verifying $\neg\varphi$. The remaining three cases for that link to have persisted are described succinctly in the two disjuncts on the right-hand side. Of course, as the upgrade may have changed truth values of formulas, we must be careful, and say that, before the upgrade, the link went to a world satisfying $\langle\sharp\varphi\rangle$ rather than φ.

The last axiom in the list is simply a commutation principle for preference and existential modalities. ■

This dynamic epistemic upgrade logic (henceforth, *DEUL*) can explain general effects of changes in information and preference. In particular, as noted earlier, it can deal with combined scenarios like introducing 'regret'. Say, a sequence of instructions

$$\sharp p; \neg p! \quad \text{for atomic } p$$

will first make p attractive, and afterwards, unobtainable. The logic records this as the (derivable) validity of regret principles like that at the end of Section 3:

$$\langle pref \rangle_i p \rightarrow [\sharp p][\neg p!](\langle pref \rangle_i p \wedge K_i \neg p)$$

More generally, *DEUL* can analyze the basic propositional scenarios of obeying successive commands or reasoning toward achieving practical goals proposed in Zarnic 2003, and Yamada 2005.

Theorem 4.4 *DEUL is completely axiomatized by the above reduction axioms.*

Proof. The reduction axioms, whose soundness we have already seen, are clearly sufficient for eventually turning every formula of our language into a static one without announcement or suggestion modalities. Then we can use the completeness theorem for our static language. ■

The same reduction method also shows that *DEUL* is decidable.

We have reached the first major conclusion of this paper:

> **preference upgrade has a complete compositional logic-**
> **just like, and even jointly with, knowledge update.**

4.3 New issues of interest: coherence

Despite the technical analogies between information update and preference upgrade, there are also intuitive differences. One typical illustration is the intuitive notion of 'coherence'. In pure public announcement logics, the only relevant aspects of coherence for a sequence of assertions seem to be these:

> (a) Do not make *inconsistent* assertions, false at the actual world; and perhaps, not to waste anyone's time: (b) Do not make assertions which are *common knowledge*, and which do not change the model.

But in combination with upgrade, we can make other distinctions. E.g., the effect of a sequence with two conflicting suggestions

$$\sharp\varphi; \sharp\neg\varphi$$

is not inconsistency, but it still has some strange aspects. Generally speaking, such a sequence will make the ordering non-connected, as it removes arrows either way between φ-worlds and $\neg\varphi$-worlds. It is an interesting issue which sequences of upgrades do have the coherence property that they leave connected preference relations connected.

Of course, in reality, one often resolves conflicts in suggestions received by means of some authority ranking among the issuers of those suggestions. This is somewhat like the reality of information update. We often get contradictory information from different sources, and we need some notion of *reliability* differentiating between these to get to any sensible total update. Both issues, however, go beyond the ambitions of this paper, as they involve the gap between actual informational events and their translation into the idealized model changes offered by dynamic epistemic logics, whether or not enriched with upgrade mechanisms.

5 Relation Change and Product Upgrade

5.1 Reduction axioms reflect definable operations

To a logician, standard epistemic update $!\varphi$ essentially relativizes a model $\mathcal{M}$ to a definable submodel $\mathcal{M}_{!\varphi}$. The relation between evaluation at both sites is expressed in the following standard result:

Fact 5.1 *(Relativization Lemma) Assertions φ hold in the relatived model iff their syntactically relativized versions were true in the old model:*

$$\mathcal{M}_{!\varphi} \models \psi \textit{ iff } \mathcal{M} \models (\psi)^{\varphi}$$

In this light, the reduction axioms for public announcement merely express the inductive facts about the modal assertion $\langle !\varphi\rangle$ referring to the left-hand side, relating these on the right to relativization instructions creating $(\psi)^{\varphi}$.

This same idea applies to preference upgrade $\sharp\varphi$. This time, the relevant semantic operation on models is *redefinition of base relations.* The same is true, evidently, for the new link-cutting update operation $\varphi!$. Van Benthem 2002 notes how relativization and redefinition make up the standard notion of *relative interpretation* between theories in logic when objects are kept fixed - while product update relates to more complex reductions forming new objects as tuples of old objects. In this light, the reduction axioms for *DEUL* again reflect a simple inductive definition, this time for what may be called *syntactic re-interpretation* of formulas. This operation leaves all logical operators unchanged, but it changes occurrences of the redefined relation symbol by its definition. There is one slight difference though. Relation symbols for preference only occur implicitly in our modal language, through the modalities. This is why the key reduction axiom in the above reflects a format of the following abstract recursive form:

$$\langle R := def(R)\rangle\langle R\rangle\varphi \leftrightarrow \langle def(R)\rangle\langle R := def(R)\rangle\varphi.$$

5.2 Dynamic logic of relation changers

Further relation-changing operations can be defined, and make sense in our dynamic logics. We already mentioned the case of

$$R := R \cup (?\neg\varphi; \top; ?\varphi)$$

Here again, reduction axioms would be immediate, because of the following straightforward validities from propositional dynamic logic:

$$\langle R \cup (?\neg\varphi; \top; ?\varphi)\rangle\psi \leftrightarrow \langle R\rangle\psi \vee \langle ?\neg\varphi; \top; ?\varphi\rangle\psi$$
$$\leftrightarrow \langle R\rangle\psi \vee (\neg\varphi \wedge E(\varphi \wedge \psi))$$

This example suggests a much more general observation, which we state somewhat informally here:

Fact 5.2 *Every relation-changing operation that is definable in PDL without iteration has a complete set of reduction axioms in dynamic epistemic logic.*

Proof. It is easy to see that every definition for a new relation $R^\sharp$ in this format is equivalent to a finite union of finite compositions of

(a) atomic relations R_i, (b) test relations $?\varphi$ for formulas of the base language.

But then, the standard axioms for union, composition, and tests in *PDL* rewrite all existential modal statements $\langle R^\sharp\rangle\varphi$ to obvious compounds in terms of just basic modalities $\langle R_i\rangle\varphi$. ■

This *PDL*-style analysis can even derive reduction axioms automatically:

Example 5.3 In this perspective, our upgrade operation $\sharp\varphi$ is really the following relation-changer:

$$R := (?\neg\varphi; R) \cup (R; ?\varphi)$$

Thus, the key reduction axiom can be derived as follows:

$$\langle\sharp\varphi\rangle\langle R\rangle\psi$$
$$\leftrightarrow \langle(?\neg\varphi; R) \cup (R; ?\varphi)\rangle\langle\sharp\varphi\rangle\psi$$
$$\leftrightarrow \langle ?\neg\varphi; R\rangle\langle\sharp\varphi\rangle\psi \vee \langle R; ?\varphi\rangle\langle\sharp\varphi\rangle\psi$$
$$\leftrightarrow (\neg\varphi \wedge \langle R\rangle\langle\sharp\varphi\rangle\psi) \vee \langle R\rangle(\varphi \wedge \langle\sharp\varphi\rangle\psi)$$

The latter is just the version that we found 'by hand' in the above.

But we can do still better than this, and achieve the same generality as dynamic epistemic logics for information update.

5.3 Product update

The usual generalization of eliminative public announcement is product update (Gerbrandy 1999, Baltag-Moss-Solecki 1999, van Ditmarsch, van der Hoek & Kooi 2006). We briefly recall the basics of this procedure. A general update step, corresponding to some possibly complex informative event, has two components:

(a) an *epistemic model* $\mathcal{M}$ of all relevant possible worlds with agents' uncertainty relations indicated,

(b) an *event model* $\mathcal{E}$ of all relevant events, again with agents' uncertainty relations between them.

Events carry information because they take place with restrictions, encoded by

(c) *preconditions* PRE_a for events a,

The preconditions are usually supposed to be common knowledge among agents. In the simplest case, they are formulated in the pure epistemic language describing facts and agents' (mutual) information about them. Now, we are ready to define the update mechanism:

Definition 5.4 *(Product update model) The* ***product update model*** $\mathcal{M} \times \mathcal{E}$ *is defined as follows:*

- *The domain is* $\{(s, a) \mid s \text{ a world in } \mathcal{M}, a \text{ an event in } E, (\mathcal{M}, s) \models PRE_a\}$.
- *The new uncertainties satisfy* $(s, a) \sim_i (t, b)$ *iff both* $s \sim_i t$ *and* $a \sim_i b$.
- *A world* (s, a) *satisfies a propositional atom* p *iff* s *already did in* $\mathcal{M}$.

Remark 5.5 If we want a version that leaves all old worlds in place, just as with the above new announcement operator $\mathcal{M}_{\varphi!}$, we would need to cut relational links again so as to 'isolate' the pairs (s, a) where $(\mathcal{M}, s)$ fails to satisfy the precondition for action a.

Definition 5.6 *(Semantics) The language has new dynamic modalities* $\langle \mathcal{E}, a \rangle$ *referring to complex epistemic actions, and these are interpreted as follows:*

$$\mathcal{M}, s \models \langle \mathcal{E}, a \rangle \varphi \text{ iff } \mathcal{M} \times \mathcal{E}, (s, a) \models \varphi$$

This is the most powerful epistemic update calculus to date. As with the special case of public announcement, it yields a complete and decidable logic via a set of reduction axioms for all possible forms of postcondition (cf. Baltag, Moss & Solecki 1998, van Benthem, van Eijck & Kooi 2005).

5.4 Product upgrade

This setting can be extended quite simply to preference upgrade over epistemic preference models. First, we enrich epistemic event models with preference relations, indicating which events agents prefer over which others. These preferences may come from actual pay-offs or other benefits, but they may also be abstract relative plausibilities again, as in the semantics of conditional logic.

Definition 5.7 *(Product upgrade with preference) The output for product upgrade on epistemic preference models are again the above epistemic models $\mathcal{M} \times \mathcal{E}$. But this time, we keep all world/action pairs (s, a) represented, as these are the non-realized options that we can still have regrets about. Then it remains to set the new preferences, and here, we can just follow the above direct product rule for relations:*

$$(s,t) \preceq_i (u,v) \textit{ iff } s \preceq_i u \textit{ and } t \preceq_i v$$

This product upgrade covers at least the earlier upgrade instruction $\sharp p$ for suggestions. To see this, consider the following event model:

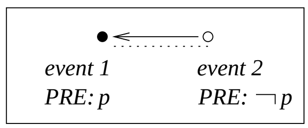

Here the two events cannot be distinguished epistemically by the agent. Moreover, recall that, in addition to the arrow drawn, our preference relations always have all reflexive loops.

Fact 5.8 $\mathcal{M}_{\sharp\varphi} \cong \mathcal{M} \times \mathcal{E}^{\sharp\varphi}$.

Proof. One can think of the special event model $\mathcal{E}^{\sharp\varphi}$ as having two events "*seeing that* φ", "*seeing that not-*φ". But all that matters is how $\mathcal{E}$ works. From a purely epistemic viewpoint, the accessible part of $\mathcal{M} \times \mathcal{E}^{\sharp\varphi}$ merely *copies* the old model $\mathcal{M}$, as only one event can take place at each world. Moreover, the old epistemic accessibilities just get copied with the product rule, since epistemic accessibility holds between all pairs of events. As for the new preference structure, consider, e.g., any pair (s,t) in $\mathcal{M}$ where $\neg\varphi$ holds at s. Then the product model $\mathcal{M} \times \mathcal{E}^{\sharp\varphi}$ contains a unique relevant corresponding pair

$$((s,\ \textit{event 2}), (t,\ \textit{event 1})).$$

The product update rule for preferences gives a preference here from left to right. Indeed, the only case where this copying from $\mathcal{M}$ fails is when the old preference and the event preference do not match up. But this only happens in those cases where $\sharp\varphi$ would reject an existing link, namely, when $s \preceq t$, while $\mathcal{M}, s \models \varphi$ and $\mathcal{M}, t \models \neg\varphi$. ■

Thus, as with public announcement and epistemic product update, one simple event model suffices to mimick our base mechanism for update or upgrade.

Much more generally, every upgrade rule which takes a current preference relation to a PDL-definable subrelation can be dealt with in the same style as above, by putting in enough events and preconditions. There are of course much more complex event models still, with many more worlds and complex preference relations for agents. These represent more refined scenarios for joint update and upgrade.

Given the technical similarity of our product upgrade rule for preference to that for epistemic accessibility, the following is easy to see:

Theorem 5.9 *The dynamic logic of product update plus upgrade can be axiomatized completely by means of dynamic-epistemic-style reduction axioms.*

We do not spell out here what these axioms look like, but it is a routine exercise.

Our second main conclusion in this paper is this:

> **preference upgrade can be combined naturally with the richest knowledge update mechanisms known so far**.

Virtues of the combination This is not just a matter of technical 'lifting'. We think that the above setting has clear independent interest. For instance, in philosophy, there is a well-known distinction between *preferences between states-of-affairs*, associated with 'consequentialist ethics', and *preferences between actions* in 'voluntarist ethics'. Our product update system models both kinds, and is able to study their interplay. Moreover, there is also a computational angle. It has been proposed to create a 'dynamic deontic' version of PDL itself, starting from preferences between worlds, but moving on to preferences between actions (Meijer 1988, van der Meijden 1990). Pucella & Weissmann 2004 follow up on the latter, and propose relation change as a way of 'changing policies'. Rohde 2005 provides a general background for this in so-called 'sabotage modal logic', where arbitrary links can be cut from models.

Thus, we see our product upgrade system also as one principled 'preferentialized' version of propositional dynamic logic.

6 Dynamic Logic of Utility Change

Intuitive problems Even product upgrade has its limitations, due to its conjunctive condition $(s, t) \preceq_i (u, v)$ iff $s \preceq_i u$ and $t \preceq_i v$ for new relations between pairs. Thus, when model size does not increase, it can only take relations between worlds to subrelations. In particular, then, product update cannot define the earlier relation-increasing instruction 'always prefer φ'... which added all pairs in $(?\neg\varphi; \top; ?\varphi)$ to the existing relation R.

But more serious, in our view, is the following problem with product upgrade. The only new preferences allowed between world-event pairs are those which existed before in both components. But this does not seem right in all settings.

Example 6.1 (Upgrade and Preference Strength) Imagine that someone prefers world s hugely over world t, while having just a very slight preference for event v over event u. It seems reasonable to say that, after update/upgrade, this person would still prefer (s, u) over (t, v). But this calculus of strength of preferences is beyond the qualitative framework so far.

This problem is reminiscent of one from dynamic epistemic logic. Think again of abstract preference relations, such as relative plausibility. In that case, our models can interpret notions of belief, say, as truth in all most plausible worlds. But a well-known difficulty in dynamic epistemic logic is this. There is no evident qualitative account of changing relative plausibility between worlds in models for *belief revision*! Perhaps the best solution so far is the quantitative one of Aucher 2003. One changes the static base models to a *graded* variant (introduced in Spohn 1988), where worlds get numerical plausibility grades. Belief update can then be defined on these numerical values for worlds, provided we also work with graded event models. Once this has been done, a dynamic doxastic logic results in the earlier style, whose complete logic is driven by reduction axioms as usual, now also for graded belief.

6.1 Static utility models and their language

Exactly the same shift in perspective makes sense here, changing qualitative preference relations to some formal grading ones. And the interpretation for that lies in hand. After all, numerical *utilities* occur just as widely as preferences, across many fields, such as game theory and decision theory. And as with grading, it is easy to define static epistemic utility models:

Definition 6.2 ***Static epistemic utility models*** *are tuples*

$$\mathcal{M} = (S, \{\sim_i |i \in I\}, \{u_i | i \in I\}, V)$$

where the u_i are now functions assigning each world a utility value ranging from 0 to Max, either in the natural numbers or some other value structure.

In terms of the earlier relational set-up, agents' preference relations are then totally connected orders, assuming that we set

$$s \preceq_i t \text{ iff } u_i(s) \preceq u_i(t)$$

We will take small Greek letters to denote the utility values. Following Aucher 2003, we could then introduce a language of graded preference modalities, indicating their strength for agents. But a simple design is more workable and perspicuous (Liu 2004). For each agent I and each value $\alpha \in \{0, 1, \ldots, Max\}$, we add a propositional constant q_i^α to our language:

Definition 6.3 *(Language) The* ***epistemic utility language*** *is defined inductively as*

$$\varphi := \bot \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_i\varphi \mid q_i^\alpha$$

Definition 6.4 *(Semantics) In a static epistemic uitility model $\mathcal{M} = (S, \{\sim_i |i \in I\}, \{u_i|i \in I\}, V)$, the key new semantic clause is that for the* ***utility constants****:*

$$\mathcal{M}, s \models q_i^\alpha \textit{ iff } u_i(s) = \alpha$$

This static language can still express preferences with formulas like

$$q_i^\alpha \wedge E(q_i^\beta \wedge \varphi)$$

With $\alpha < \beta$, this says that agent i sees a world with higher utility where φ holds. The complete epistemic-utility logic of such models is easily determined (cf. van der Hoek 1992 for a pioneering study), using the techniques of Liu 2004.

6.2 Dynamic logic of utility upgrade

Of more interest in this paper is what happens to epistemic utility models when triggers come in for *utility change*. This time, we only quickly sketch the product update mechanism, as it is quite straightforward and appealing in this setting.

As usual, we have some current epistemic utility model $\mathcal{M}$, and some event model $\mathcal{E}$ containing all relevant events. But now, to get going, we must also assume that $\mathcal{E}$ itself has agents' utilities attached to its events. These utilities may reflect real enjoyment values attached by agents to events, or some more abstract form of evaluation. And then, we need an upgrade rule. Here is about the simplest and most obvious version that one can think of:

Definition 6.5 *(Bare addition rule) The new utilities in product models are defined by the following rule:*

$$\textit{utility of } (s, a) = \textit{utility for } s \textit{ in } \mathcal{M} + \textit{utility for } a \textit{ in } \mathcal{E}$$

Here is a simple illustration of how this stipulation works out, giving another take on our original idea of upgrading in favor of some proposition p.

Example 6.6 (Scoring for p) Consider the following one-agent event model $\mathcal{E}$ with two epistemically indistinguishable events, that we might read as 'register p' (a), 'register $\neg p$' (b):

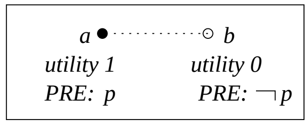

With any model $\mathcal{M}$, this just copies all existing worlds s to either a world (s, *event 1*) (if $\mathcal{M}, s \models p$), or to (s, *event 2*) (if $\mathcal{M}, s \models \neg p$), while all epistemic relations stay the same. The above Addition Rule only creates one difference: each p-world gets an additional unit of utility. This is essentially the view that obeying suggestions 'increases merit'.

In this example, one simple two-event model encodes one particular upgrade procedure for a suggestion. This would also work for stronger commands, now changing the utility in the event model to give them a higher scoring power. Still more complex event models will then encode more sophisticated ways of changing utilities as suggestions or commands come in. As for the dynamic logic of such upgrade mechanisms, again we are well within standard approaches:

Theorem 6.7 *The dynamic epistemic logic of utility upgrade (DLUU) is completely axiomatizable.*

Proof. This can be shown just as in Liu 2004 for the case of belief revision with numerical world values. The only new reduction axiom, over and above the usual ones for the epistemic operators, is that for the utility constants, and it just encodes the Addition Rule:

$$\langle \mathcal{E}, a \rangle q_i^{\alpha} \leftrightarrow q_i^{\alpha - u_i(a)}$$

■

On the analogy with belief revision, we might expect utility upgrade to have several flavors, or 'policies', depending on the relative weights attached to prior utility values and to the last-observed events. But this diversity may be somewhat spurious, as the Addition Rule seems pretty inevitable. In that case, all further diversity would have to be encoded in the models themselves.

7 Illustrations: Defaults and Obligations

So far, we have found two upgrade mechanisms, each representing a plausible view of incoming triggers that change preferences. We now illustrate these frameworks in two concrete settings. Our aim here is not some full-fledged application to existing systems. We rather want to show how the earlier logical issues in this paper correspond to real questions of independent interest.

7.1 Default reasoning

Consider practical reasoning with default rules of the form "*if* φ*, then* ψ":

> "If I take the train right now, I will be home to-night".

These are defeasible conditionals, which recommend concluding ψ from φ, but without excluding the possibility of $\varphi \wedge \neg\psi$-worlds, be it that the latter are now considered exceptional circumstances. Intuitively, the latter are not 'ruled out' from our current model, but only 'downgraded' when a default rule is adopted. Veltman 1996 is an influential dynamic treatment, making a default an instruction for changing the current preference order between worlds. The simplest case has just one assertion φ which is being 'recommended' - in Veltman's terms, there is an instruction "*Normally,* φ". From our perspective, there are two ways of taking this. First we go back to our scenario in Section 7:

Definition 7.1 *(Processing defaults by utility update)* ***Utility update for defaults*** *raises the utility of* φ*-worlds by some fixed number (+1, or higher) indicated in the relevant event model with two actions 'register* φ*', 'register* $\neg\varphi$*'.*

Our intuition here is as follows. Each default assertion φ represents a way in which a world can be 'good'. The number of successes in meeting all requirements on the table (after all, there may be many defaults that we work with at any stage) is the utility of the world. Another way of taking this same scenario is as one of satisfying *constraints from some given list*, computing least-violating situations as in linguistic Optimality Theory. Our dynamic logic for utility upgrade then describes this process in all its epistemic utility effects.

But one can also go another way, using a scenario of *relation change for defaults*, as in our earlier Section 3. Suppose that we want to give an incoming default rule "*Normally,* φ" 'priority', in that after its processing, all best worlds are indeed φ-worlds. This is not guaranteed by the preceding utility update rule, which will only make the best worlds φ in case there are φ-worlds among those with highest utility so far. Here is a more drastic procedure, which will validate the preceding intuition:

Definition 7.2 *(Processing defaults by relation change) We make all* φ*-worlds better than all* $\neg\varphi$*-worlds, and within the* φ*- and* $\neg\varphi$*-areas, we leave the old preferences in place. Formally, this his one of our earlier PDL-style relation-changing instructions: the old preference relation* R *becomes*

$$(?\varphi; R; ?\varphi) \cup (?\neg\varphi; R; ?\neg\varphi) \cup (?\neg\varphi; \top; ?\varphi).$$

Interestingly, this is the union of the earlier link cutting version of public announcements $\varphi!$ plus the upgrade operation with relation extension considered in Section 4.

Fact 7.3 *Relational default processing can be axiomatized completely.*

Proof. By the method of Section 5.2, the key reduction axiom follows automatically from the given PDL-form, yielding

$$\langle upgr(\varphi)\rangle\langle pref\rangle\psi \leftrightarrow (\varphi\wedge\langle pref\rangle(\varphi\wedge\langle upgr(\varphi)\rangle\psi)\vee(\neg\varphi\wedge\langle pref\rangle(\neg\varphi\wedge \langle upgr(\varphi)\rangle\psi)) \vee (\neg\varphi \wedge E(\varphi \wedge \langle upgr(\varphi)\rangle\psi))$$

■

Thus, we have at least two plausible versions of default logic in our upgrade setting. Moreover, their validities are axiomatizable in a systematic style via reduction axioms, rather than more ad-hoc default logics found in the literature.

Things need not stop here, as our two approaches may be generalized. E.g., the relation-changing version puts heavy emphasis on the last suggestion made, giving it the force of a command. This seems too strong in many cases, as it gears everything toward the last thing heard. A more reasonable scenario is this. We are given a sequence of instructions inducing preference changes, but they need not all be equally urgent. We need to find out our total commitments eventually. But the way we integrate these instructions may be partly left up to the *policy* that we choose, partly also to another parameter of the scenario: viz. the relative force or *authority* of the issuers of the instructions. One particular setting where this happens is again Optimality Theory. Ranked constraints determine the order of authority, but within that, one counts numbers of violations. Cf. A. Prince & P. Smolensky 1993 for a good exposition, and de Jongh & Liu 2005 for a logical exploration.

7.2 Deontic logic and commands

Similar considerations apply to deontic logic (Åqvist 1987). Originally, this was the study of assertions of obligation

$O\varphi$: 'it ought to be the case that φ',

as well as statements of conditional obligation $O(\varphi|\psi)$, say, emanating from some moral authority. The cumulative weight of all true O-statements represents all the obligations an agent has at the current stage.

In the standard semantics of deontic logic, $O\varphi$ is treated as a universal modality over some deontic accessibility relation. But the underlying intuition is that those φ ought to be case which are true in *all best possible worlds*, as seen from the current one. Again, this naturally suggests a preference order among worlds.

And then, once more, we can think of this setting dynamically, using either of our two upgrade scenarios.

Initially, there are no preferences between worlds. Then some moral authority comes in and starts 'moralizing': i.e., introducing evaluative distinctions between worlds. If this process works consistently, we end with a new ordering of our worlds from which our current obligations may be computed, as those assertions which are true in all best worlds. Whether a sequence of commands makes sense in this way may depend on more than consistency: thus, the issue of 'coherence' noted briefly in Section 3 comes back with greater force now.

Looking backward, or forward in upgrade But deontic logic also raises other issues that we encountered before. In particular, one strong moral intuition seems to be that, after a command has been given (say, 'Thou shalt not kill'), the relevant proposition becomes true in all best possible worlds. In commands, there is a future-oriented aspect:

> 'See to it that φ' should result in a new situation where $O\varphi$ is true.

But as we have seen in Section 4, not every upgrade $\sharp\varphi$ has the effect that φ becomes true in the new most preferred worlds. Indeed, there is a general difficulty with specifications of the form 'See to it that φ'. Dynamic epistemic logic is mainly about events with their preconditions. Thus, the information one gets from an event is *past-oriented*, describing what was the case at the time the event happened. But, even a simple epistemic event can change the truth value of assertions at worlds - witness public announcements turning ignorance into knowledge. In fact, such changes in truth value are what dynamic logics of update and upgrade are about.

But it is not so easy to just define some action or event as leading to the truth of some proposition. This works for simple factual effects of actions like opening a door (van Benthem, van Eijck & Kooi 2005), but it is not clear what this should even mean with more complex stipulations. The event itself may get in the way. For instance, there are no obvious actions 'seeing to it that' arbitrary mixtures of knowledge and ignorance in groups arise afterwards, and the same seems true of complex deontic assertions, involving the achievement by fiat of different obligations and permissions for different agents. Whether deontic reasoning needs some sort of future-oriented update and upgrade above what we have offered seems an interesting question. For temporal logics of such $STIT$ operators, cf. Belnap, Perlof & Xu 2001.

8 Comparing Preference and Utility Upgrade

The co-existence of two upgrade mechanisms, raises some obvious questions. We confine ourselves to a few observations. First, there is an obvious semantic sense in which the two approaches are related. Clearly, any epistemic utility model induces a totally ordered epistemic preference model, by setting

$s \preceq_i t$ iff $u_i(s) \preceq u_i(t)$

And likewise, any totally ordered preference model generates a matching utility model by numbering equivalence classes through the associated propositions:

$$\kappa(\varphi) = min\{u_i(s) : s \in [\![\varphi]\!]\}.$$

But this still leaves some questions unanswered. In particular, are the two dynamic languages equally powerful?

Utility and preference dynamics First, it seems clear that utility upgrade is not definable in terms of relational preference upgrade. The *DEUL* language of Section 4 simply lacks the expressive power for the *DLUU* language of Section 6. But this poverty may still be circumvented by mimicking numerical values through equivalence classes at higher or lower positions in a total preference order. At present, though, it is not clear to us exactly how this would work.

Also conversely, it seems impossible to define our apparently simpler relation-changing upgrades $\sharp\varphi$ in terms of utility upgrade via event models. The reason is this. Any event model can only upgrade φ-worlds, of perhaps different sorts, by some amount X. But then,

> in a model where some φ-world has utility 0 and some $\neg\varphi$-world utility X, even an X-upgrade for φ-worlds will not change the preference order as required by $\sharp\varphi$.

Again, there may be more sophisticated comparisons here, but our two mechanisms look somewhat different in spirit.

Further issues: bisimulation There are other interesting issues left unexplored in this paper. For instance, relation-changing *DLUU* and utility-changing *DEUL* might differ in their natural bisimulations between models.

Example 8.1 From a qualitative preference perspective, it would make sense to identify the following two models, where we identify worlds by their utilities:

After all, the pure preference pattern is the same in both. But the utilities really make a difference, even without explicit utility language. Consider the default event model $\mathcal{E}$ which upgraded all φ-worlds (t 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 removes it.

This difference shows up in formulas of our dynamic language. Clearly, some sort of bisimulation makes sense for preferential and for utility models, as we want the right abstract pattern, not irrelevant details of numerical assignments. We leave such issues, and their effects on our choice of dynamic upgrade languages, for further investigation.

9 Related Work

The ideas in this paper have a long history, and there are many proposals in the literature having to do with 'dynamification' of preferences, defaults, and obligations. We just mention a few related approaches here, though we do not make any detailed comparisons. Meyer 1988 was probably the first to look at deontic logic from a dynamic point of view, with the result that deontic logics are reduced to suitable versions of dynamic logics. This connection has become a high-light in computer science since, witness the regular *DEON* conference series. In a line that goes back to Spohn 1988, Veltman 1996 presents an update semantics for default rules, locating their meaning in the way in which they modify expectation patterns. This is part of the general program of 'update semantics' for conditionals and other key expressions in natural language. Tan & van der Torre 1998 use ideas from update semantics to formalize deontic reasoning about obligations, but with motivation from computer science. In their view, the meaning of a normative sentence resides in the changes it brings about in the 'ideality relations' of agents to whom the norm applies. Van der Meyden 1996 takes the deontic logic/dynamic logic interface a step further, distinguishing two notions of permission, one of which, 'free choice permission' requires a new 'dynamic logic of permission', where preferences can hold between actions. Completeness theorems with respect to this enriched semantics are given for several systems. Pucella & Weissman 2004 provide a dynamified version of the dynamic logic of permission, in order to deal with building up of agents' policies by adding or deleting transitions. Demri 2005 reduces an extension of van der Meyden's logic to propositional dynamic logic, yielding an EXPTIME decision procedure, and showing how dynamic logic can deal with agents' policies. Following van Benthem's 'sabotage games', Rohde 2005 studies general modal logics with operators that describe effects of deleting arbitrary transitions - without a fixed upgrade definition as in our analysis. Model checking for such logics becomes Pspace-complete, and satisfiable is undecidable. Parikh & Pacuit 2005 observe that an agent's obligations are often dependent on what she knows, and introduce a close relative of our epistemic preference language, but over temporal tree models. Our own approach goes back to van Benthem, van Eijck Frolova 1993, which discusses general formats for upgrading preference relations. Zarnic 2003 uses similar ideas, combined with a simple update logic to formalize natural language imperatives of the form *FIAT* φ, which can be used in describing the search for solutions of given planning problems. Finally, Yamada 2005 takes the update paradigm to logics of commands and obligations, modelling changes brought about by various acts of commanding. It combines a multi-agent variant of the language of monadic deontic logic with a dynamic language for updates and commands. This is closest to what we do, although we think that this paper provides a much more general treatment of possible upgrade instructions, while our later utility update variants are new altogether.

10 Conclusion

Preference upgrade seems a natural and crucial element in logical dynamics. We have shown that it can be modelled as relation change in a standard qualitative dynamic format, up to the expressive level of the best available system, that of product update. Nevertheless, this first approach leaves something to be desired in intuitive scenarios where strengths of preferences seem essential. Therefore, we have also given a second calculus, this time for event models with utility update. It works on the same pattern as the first. As of yet, we have not been able to determine to which extent our two approaches are reducible to each other. But this does not matter: upgrade works fine either way, and the two formats may just reflect interestingly different intuitions.

Of course, we do feel our approach leaves things to be desired. Here are some further directions which appeal to us.

Deeper interactions between knowledge and preference changes Our language separates modalities for knowledge and for 'better' according to agents. Nevertheless, one might make a case that natural notions of 'preference' involve an irreducible mixture of information and world comparison. This may require further languages beyond what we have given here. Already as it stands, our epistemic preference language cannot define the notion of a 'best' world. And more epistemically, it cannot define the notion of a 'better accessible world', which would require an intersection of the epistemic and preferential relations that we have kept separate.

Preference changes in groups In this paper, no genuine interactions are studied between preferences of different agents. But dynamic epistemic logics definitely include social notions of knowledge, and hence, similar notions might make sense for preference. Part of these issues already come up in *social choice theory*: how should individual preferences be integrated into a group ranking? But the issues become even more intriguing when we take the possibly diverse knowledge of group member into account about others' preferences.

Short-term versus long-term preference Next, there is the matter of temporal perspective. In realistic situations, preferences come in different kinds. Long-term preferences define our general modus operandi, and they are not changed by the 'local triggers' studied in our paper. The latter only describe the short-term preferences that we may have in the course of a conversation, problem solving session, or a game. Long-term preferences are crucial to, e.g., evolutionary game theory, where they underlie social norms – or optimality theory, where they encode the stable preferences of language users concerning syntax and semantics. Indeed, the optimality theory idea of 'ranked constraints' may put these various roles into one framework: cf. de Jongh & Liu 2005.

Games We would also like to apply our perspectives and results to the analysis of games. Van Benthem, van Otterloo & Roy 2005 show how adding a static modal preference logic to a propositional dynamic logic provides a perspicuous definition for the backward induction solution to finite extensive games. But we feel that, just like epistemic update, preference change is also a natural local feature of extensive games. In this way, we can think of preferences for outcomes, not as given once and for all, but rather as explained dynamically by what happens in the course of a game: cf Liu 2005.

References

[1] L. Åqvist: *Introduction to Deontic Logic and the Theory of Normative Systems*, Naples: Bibliopolis, 1987.

[2] G. Aucher: *A Combined System of Update Logic and Belief Revision*, Master of Logic Thesis, ILLC University of Amsterdam, 2003.

[3] A. Baltag, L. Moss & S. Solecki: The Logic of Public Announcements, Common Knowledge and Private Suspicions, *Proceedings TARK* 1998, pp. 43-56, Morgan Kaufmann, Los Altos, 1998.

[4] N. Belnap, M. Perloff & M. Xu: *Facing the Future*, Oxford University Press, Oxford, 2001.

[5] J. van Benthem, J. van Eyck & A. Frolova: Changing Preferences, Report CS-93-10, Centre for Mathematics Computer Science, Amsterdam, 1993.

[6] J. van Benthem: *Exploring Logical Dynamics*, CSLI Publications, Stanford, 1996.

[7] J. van Benthem: Games in Dynamic-Epistemic Logic, *Bulletin of Economic Research* 53:4, pp. 219-248 (Proceedings LOFT-4, Torino), 2001.

[8] J. van Benthem: 'One is a Lonely Number': on the Logic of Communication, Logic Colloquium 2002, Amsterdam & Stanford.

[9] J. van Benthem: An Essay on Sabotage and Obstruction, ILLC Report, 2004. To appear in D. Hutter, ed., *Festschrift for Joerg Siekmann.*

[10] J. van Benthem, J. van Eijck & B. Kooi: Logics for Communication and Change, ILLC Amsterdam, CWI Amsterdam, Philosophical Institute, University of Groningen, 2005.

[11] J. van Benthem & F. Liu: Diversity of Logical Agents in Games, Report PP-2004-13, ILLC Amsterdam. Also in *Philosophia Scientiae* 8:2, pp. 163-178, 2004.

[12] J. van Benthem, S. van Otterloo & O. Roy: Preference Logic, Conditionals and Solution Concepts in Games, 2005. To appear.

[13] P. Blackburn, M. de Rijke & Y. Venema: *Modal Logic*, Cambridge University Press, 2001.

[14] C. Boutilier: Conditional Logics of Normality: a Modal Approach. *Artificial Intelligence* 68: pp. 87-154, 1994.

[15] S. Demri: A Reduction From DLP to PDL, an unpublished draft, 2004.

[16] H. van Ditmarrsch: Dynamic Belief Revision, Lecture Slides at University of Liverpool, 2004.

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

[18] J. Gerbrandy: *Bisimulation on Planet Kripke*, Ph.D dissertation, ILLC Amsterdam, 1999.

[19] Joseph Y. Halpern: Defining Relative Likelihood in Partially-ordered Preferential Structure. *Journal of Artificial Intelligence Research* 7: pp. 1-24, 1997.

[20] S. Hanson: Preference-based Deontic Logic. *Journal of Philosophical Logic*, vol. 19, pp. 75-93, 1990.

[21] W. van der Hoek: *Modalities for Reasoning about Knowledge and Quantitaties*, Ph.D dissertation, ILLC Amsterdam, 1992.

[22] Z. Huang: *Logics for Agents with Bounded Rationality*, ILLC Dissertation Series 1994-10, University of Amsterdam, 1994.

[23] D. de Jongh & F. Liu: Constraint Ranking, Belief and Preferences (working paper), 2005.

[24] F. Liu: *Dynamic Variations: Update and Revision for Diverse Agents*, Master of Logic Thesis, ILLC University of Amsterdam, 2004.

[25] F. Liu: Preference Changes in Games (working paper), 2005.

[26] R. van der Meyden: The Dynamic Logic of Permission, *Journal of Logic and Computation*, Vol. 6, No. 3, pp.465-479, 1996.

[27] 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: pp. 109-136, 1988.

[28] M. Osborne & A. Rubinstein: *A Course in Game Theory*, The MIT Press, Cambridge (Mass.) 1994.

[29] R. Parikh, E. Pacuit & E. Cogan: The Logic of Knowledge Based Obligation, 2005.

[30] J. A. Plaza: Logics of Public Announcements, Proceedings 4th International Symposium on Methodologies for Intelligent Systems, 1989.

[31] A. Prince & P. Smolensky: *Optimality Theory: Constraint Interaction in Generative Grammar.* Malden, Ma: Blackwell. Techinical Report, Linguistic Department, Rutgers University, New Brunswick, NJ and Deparment of Computer Science, University of Colorado at Boulder, 1993.

[32] R. Pucella V. Weissmann: Reasoning about Dynamic Policies, *Proceedings FoSSaCS-7*, Springer Lecture Notes in Computer Science 2987, pp. 453-467, 2004.

[33] Ph. Rohde: *On Games and Logics over Dynamically Changing Structures*, Ph.D. Thesis, Department of Informatics, Technische Hochschule Aachen (RWTH), 2005.

[34] W. Spohn: Ordinal Conditional Functions: A Dynamic Theory of Epistemic States, in W. L. Harper et al., eds., *Causation in Decision, Belief Change and Statistics II*, Kluwer, Dordrecht, pp. 105-134, 1988.

[35] F. Veltman: Defaults in Update Semantics, *Journal of Philosophical Logic* 25, pp. 221-261, 1996.

[36] T. Yamada: Commands and Changing Obligations, in preparation, 2005.

[37] B. Žarnić: Imperative Change and Obligation to Do. In: K. Segerberg and R. Sliwinski (eds.) *Logic, Law, Morality: Thirteen Essays in Practical Philosophy in Honour of Lennart Aqvist.* pp.79-95. Uppsala philosophical studies 51 Uppsala: Department of Philosophy, Uppsala University 2003.