



UvA-DARE (Digital Academic Repository)

Changing for the better : preference dynamics and agent diversity

Liu, F.

Publication date
2008

[Link to publication](#)

Citation for published version (APA):

Liu, F. (2008). *Changing for the better : preference dynamics and agent diversity*. ILLC.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 2

Dynamic Logic of Preference Upgrade

2.1 Introduction: changing preferences

The notion of preference occurs across many research areas, such as philosophy of action, decision theory, optimality theory, and game theory. Individual preferences between worlds or actions can be used to predict behavior by rational agents. More abstract notions of preference also occur in conditional logic, non-monotonic logic and belief revision theory, whose semantics order worlds by relative similarity or plausibility.

Preference logics Preference logics in the literature describe different comparative structures by means of various devices ([Han90b]). Agents' preferences can run between worlds or between actions, preference statements can be weaker or stronger in what they say about worlds or actions being compared – and also, they may be more 'objective' or more 'epistemic'. A statement like “I prefer sunsets to sunrises” can be cast merely in terms of ‘what is better for me’, or as a more complex propositional attitude involving my beliefs about the relevant events. In this chapter, we take an objective approach, where a binary preference relation supports a unary modality “true in some world which is at least as good as the current one” ([Bou94], [Hal97]). [BOR06] show how such a language, when extended with a few operators from hybrid languages, can define several conditionals, Nash equilibrium, and backward induction solutions to games. The language also expresses various kinds of preference that agents may have between propositions, i.e., types of events. Moreover, we add explicit epistemic operators, allowing us to express agents' attitudes toward what is good or better for them.

Preference dynamics Our main concern in this chapter, however, is one of *dynamics*. Preferences are not static, but they change through commands of moral authorities, suggestions from friends who give good advice, or just changes

in our own evaluation of worlds and actions. Such changes 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 [BEF93], [Han95], [Zar03], [TT99], and [Yam06].

Related ideas all play in the dynamic semantics for conditional logics (for instance, [Spo88], [Vel96]). In its static Lewis-style semantics, a conditional $\varphi \Rightarrow \psi$ says roughly the following

ψ is true in all most-preferred φ -worlds (‡)

But one plausible way of accepting 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 (‡) the case. Even more simply, consider a default assertion like

“*Normally* φ .”

As [Vel96] points out, this does not eliminate $\neg\varphi$ -worlds from our current model, in the usual dynamic sense of information update. Accommodating this assertion rather makes the $\neg\varphi$ -worlds doxastically less preferred than φ -worlds.

Trigger 1: suggestions There are many triggers for preference change in real life, and dynamic preference logics should provide a format for studying these in an appropriate generality. To find such formats, in this chapter, we start from a simple test scenario that may be called a ‘suggestion’. Consider someone who is indifferent between taking a trip (p) and staying at home ($\neg p$). Now his friend comes along and says

“Let’s take a trip!”

‘Taking’ this suggestion means that any preference we might have had for staying at home is removed from the current model. Figure 2.1 shows what we have in mind.

Thus, in our scenario, a suggestion removes already existing preference links: but it does not add new ones. Note that, in addition to arrows drawn, our preference relations always have reflexive loops. This mechanism will be studied in greater detail later on, as an entry into more general kinds of preference upgrade. Even so, by way of contrast, here is one alternative, which does not remove links, but rather adds them.

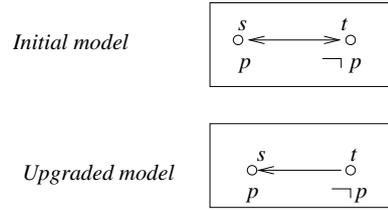


Figure 2.1: Initial model and upgraded model

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



Figure 2.2: A model with two unrelated worlds

A suggestion in the relation-decreasing sense does not make the worlds comparable. With real commands “Take that trip!”, however, we want to make sure the agent now prefers p . Then, we need to *add* preference links to the picture, making the world with $\neg p$ less preferred. Our proposals also deal with upgrades that add links between worlds.

Dynamic logics of upgrade Whether eliminative or additive, preference change is reminiscent of existing systems for information *update* in dynamic-epistemic logic ([Ger99], [BMS98], [Ben07a], [DHK07]). 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 chapter is that preference upgrade is a viable phenomenon, just as susceptible to systematic modification as information, temporal perspective, or other parameters of ‘logical dynamics’ ([BEF93], [Ben96], or in the setting of conditional logic, [Spo88], [Vel96]). 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 chapter is structured as follows. First we present a new joint epistemic preference logic (Section 2.2). 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 2.3, 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 2.4 defines a dynamic version of the static epistemic preference language, where information update lives together with 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. In Section 2.5, we consider more general upgrade scenarios: first with general schemes of link elimination, and then, with the full strength of ‘product update’ for information using ‘event (action) models’. This requires enriching the action models of dynamic-epistemic logic with agents’ preferences between events. Section 2.6 then outlines some applications of our dynamic upgrade logics, to default reasoning, deontic logic, and logics of commands. Section 2.7 is a brief survey of related work, and Section 2.8 contains our conclusions and further directions.

This chapter proposes a certain style of thinking about preference upgrade, and an existence proof for a logical methodology in doing so. We do not address all intuitive senses of preference, or all logical issues arising in the areas where it plays a role. A more extensive discussion of upgrade mechanisms with various triggers, various senses of preference, and further applications, is found in [JL06] and later chapters in this thesis.

2.2 Epistemic preference logic

Language and semantics

The main language used in this chapter has two components: a preference modality as in [BOR06], and the standard knowledge operators from epistemic logic.

2.2.1. DEFINITION. 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 the following rule:

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

Intuitively, $K_i\varphi$ stands for ‘agent i knows that φ ’, while $[pref]_i\varphi$ says that all worlds which the agent considers as least as good as the current one satisfy φ . U is an auxiliary universal modality.¹

¹For 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.

How is this formal language connected to ‘preference’ as it occurs in natural discourse? One may be inclined to read $\langle pref \rangle_i \varphi$ as ‘agent i prefers φ ’. But as with other logical systems, there is a gap between the formalism and common usage. E.g., just saying that the agent sees some better world where φ holds seems too weak, while the universal modality $[pref]_i \varphi$ ‘in all better worlds’ seems much too strong. Cf. [Han01b] for a thorough discussion of senses of preference, and ways in which formal languages do or do not match up. Here we just point out the following facts. First, our formal language can also express intermediate senses of ‘betterness’ for preference, using *combinations* of modalities. E.g., $[pref]_i \langle pref \rangle_i \varphi$ will express, at least on finite connected models, that some *best* world has φ . And one can also express that *all* best worlds satisfy φ : cf. [BOR06]. Moreover, our approach emphasizes comparisons of worlds, i.e., objects, rather than propositions, whereas common notions of preference often play between propositions, or semantically, sets of worlds. Such preferences between propositions can be *defined* on our approach (see again [BOR06]). For instance,

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

expresses one strong sense of ‘agent i prefers φ to ψ ’, viz. each ψ -world s has at least one epistemic alternative which is φ and which is at least as good as s according to the agent. But one can also define the original notion of preference in [Wri63] which says that the agent prefers all φ -worlds to all ψ -worlds (cf. [BOR06]; [BRG07] also deals with Von Wright’s ‘*ceteris paribus*’ clause in the relevant comparisons between worlds). For the moment, we take this expressive power of our simple-looking modal language for granted. The virtue of our simple base modalities is that these ‘decompose’ more complex preference statements in a perspicuous manner, while allowing for a simple dynamic approach later on.

2.2.2. DEFINITION. 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 ,² and V a valuation for proposition letters. Moreover, \preceq_i is a reflexive and transitive relation over the worlds.

We read $s \preceq_i t$ as ‘ 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* in the sense of the Lewis sphere models for conditional logic. In general, we want to

But they help in finding and checking valid principles, and in semantic arguments generally.

²Interpreting the knowledge operator with the equivalence relation is optional in an approach. There are many philosophical discussions about its justification. Various alternatives haven been proposed in terms of model classes. For complete epistemic logics over equivalence relations or other model classes, see the standard references, e.g. [FHMV95] or [BRV01].

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 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 standard utility-based preference orderings of outcomes in a game, connectedness may be quite appropriate.

2.2.3. DEFINITION. Given an epistemic preference model $\mathcal{M} = (S, \{\sim_i \mid i \in I\}, \{\preceq_i \mid i \in I\}, V)$, 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. $\mathcal{M}, s \models \langle K \rangle_i \varphi$ iff for some $t : s \sim_i t$ and $\mathcal{M}, t \models \varphi$.
5. $\mathcal{M}, s \models \langle pref \rangle_i \varphi$ iff for some $t : s \preceq_i t$ and $\mathcal{M}, t \models \varphi$.
6. $\mathcal{M}, s \models E\varphi$ iff for some $t : \mathcal{M}, t \models \varphi$.

Expressive power As we noted, [BOR06] 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, including the original one proposed by Von Wright, as well as other natural options qua quantifier combinations. Moreover, following [Bou94], 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 also express the interplay of preference and knowledge. The following examples represent (a) an intuition of self-reflection of ‘preference’, and (b) an unfortunate but ubiquitous phenomenon:

- $\langle pref \rangle_i \varphi \rightarrow K_i \langle pref \rangle_i \varphi$: Preference Positive Introspection
- $\langle pref \rangle_i \varphi \wedge K_i \neg \varphi$: Regret.

We will return to mixed epistemic-preference principles later on.

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. [BRV01]).

2.2.4. THEOREM. *Epistemic preference logic is completely axiomatizable w.r.t epistemic-preference-models.*

Proof. The proof is entirely by standard techniques. □

Additional axioms in our language impose further frame conditions on models. Here are two examples to show the spirit. They are based on standard modal frame-correspondence techniques:

2.2.5. FACT.

- A preference frame $\mathcal{F} = (S, \{\sim_i \mid i \in I\}, \{\preceq_i \mid i \in I\})$ satisfies *connectedness*, i.e., $\forall x \forall y : x \preceq_i y \vee y \preceq_i x$, iff the following formula is true in the frame:

$$(\varphi \wedge E\psi) \rightarrow \langle pref \rangle_i \psi \vee E(\psi \wedge \langle pref \rangle_i \varphi).$$

- An epistemic preference frame \mathcal{F} makes the *Preference Introspection Axiom* $\langle pref \rangle_i \varphi \rightarrow K_i \langle pref \rangle_i \varphi$ true iff it satisfies the following condition:

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

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

2.3 Modelling preference upgrade

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 pictured in Figure 2.3, where the dotted line in the initial static model indicates the agent's uncertainty in the initial situation.

The announcement eliminates the $\neg p$ -world from the epistemic model, 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. See [BMS98], [Ben07a], and Section 2.5 below.

These logics all work essentially on the same design principle. First, a class of models is chosen representing the relevant information structures, together with

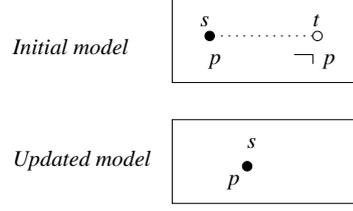


Figure 2.3: Initial model and updated model

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 simply means a possible world elimination, 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 $\{w \in S \mid \mathcal{M}, w \models \varphi\}$. And accessibility relations and valuations are retained on the restricted domain.

More complex actions update to *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 static-dynamic assertion of this sort is

$[\!]\varphi]K_i\psi$: *after a truthful public announcement of φ , the 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.

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 2.2:

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

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 :

2.3.1. DEFINITION. 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

$$\preceq_i^* = \preceq_i - \{(s, t) \mid \mathcal{M}, s \models \varphi \text{ and } \mathcal{M}, t \models \neg\varphi\}.$$
³

We suppress agent subscripts henceforth whenever convenient.

Upgrade for suggestion events replaces a preference relation by a definable subrelation. This may be written as follows in the standard notation of propositional dynamic logic (e.g. [HKT00]):

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

We will consider more general relation-changing operations in Section 2.5. 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 2.4, we consider some features of the mechanism just defined.

³[Har04] analyzes newly defined preference relations in a set-theoretic format.

Preservation properties of upgrade Perhaps the most pressing issue 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, no general result comes to the rescue, since we only have the following counterpart to the preservation result for submodels:

2.3.2. FACT. 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

2.3.3. FACT. 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. \square

On the other hand, our upgrades $\sharp\varphi$ can lead to loss of connectedness of the preference order. Our earlier example already showed this in Section 2.1 (see Figure 2.2). Likewise, our upgrades can lead to a loss of positive introspection, see the following scenario:

2.3.4. EXAMPLE. Consider Figure 2.4 below. There are two worlds ‘asleep’ and ‘awake’. In both models, we do not know if we are sleeping or awake. Initially, we prefer being asleep, and we know our preference. Now an upgrade happens, suggesting that real waking life is not so bad after all. Then we still do not know if we are sleeping or awake, but at the ‘awake’ world we prefer being awake (thought not to be the case at the ‘asleep’ world). Focusing on the ‘asleep’ world in the new model, we still prefer being asleep there. But we no longer know that we prefer it – since we might be in the ‘awake world’. Introspection fails!

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.

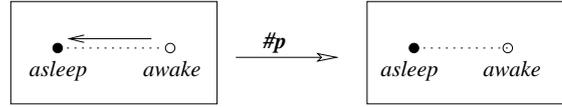


Figure 2.4: Upgrade leading to a loss if positive introspection

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 of ‘link cutting’. This time, instead of the above $!\varphi$, we write the relevant update action as follows, note that the notation ‘!’ is now behind φ :

$\varphi!$

and we write the updated model as $\mathcal{M}_{\varphi!}$ in order to distinguish it from what we obtain by eliminating worlds. We should really change notations to reflect the two kinds of exclamation mark – but we trust the reader can disambiguate in context. The correct semantic operation for $\varphi!$ on models is this:

2.3.5. DEFINITION. 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 a version without any crossing between the φ - and $\neg\varphi$ -zones of \mathcal{M} :

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

2.3.6. FACT. The pure epistemic logic of public announcement is the same with $!\varphi$ and with $\varphi!$.

Nevertheless, the second update stipulation has some advantages. It was first proposed, in [Sny04] (cf. Chapter 5 and 6) 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.

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 is reflected in valid principles of the logic, but we do not pursue this issue 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 is not valid for more nostalgic souls, who still deplore the way things turned out to be. For them, update amounts to the link-cutting operation $\varphi!$, they stick to their preferences between all possible worlds, and the new fact may even introduce regrets:

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

2.4 Dynamic epistemic upgrade logic

Language and semantics

Now we introduce an enriched dynamic language for update and upgrade. Its static part is the earlier language of Section 2.2, but its action vocabulary contains both link-cutting announcements $\varphi!$ and suggestions $\#\varphi$. Adding the original world-eliminating announcements $!\varphi$ is a routine matter, so we highlight the latter less standard variant only.

2.4.1. DEFINITION. Let P be a set of proposition letters and I a set of agents, with p ranging over P , i over I . The *dynamic epistemic preference language* is given by the following rule:

$$\begin{aligned} \varphi &::= \perp \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_i \varphi \mid [pref]_i \varphi \mid U\varphi \mid [\pi]\varphi \\ \pi &::= \varphi! \mid \#\varphi. \end{aligned}$$

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

2.4.2. DEFINITION. Given an epistemic preference model \mathcal{M} , the *truth definition* for formulas is as before, but with two new key clauses for the action modalities:

$$\begin{aligned} (\mathcal{M}, s) \models [\varphi!]\psi &\text{ iff if } \mathcal{M}, s \models \varphi, \text{ then } \mathcal{M}_{\varphi!}, s \models \psi. \\ (\mathcal{M}, s) \models [\#\varphi]\psi &\text{ iff } \mathcal{M}_{\#\varphi}, s \models \psi. \end{aligned}$$

Preference upgrade logic

On epistemic preference models, all valid principles of the static language of Section 2.2 still hold. 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 purely epistemic dynamic axioms, but it does with global existential modalities over the whole domain of the model. The usual reduction axiom for operator E is 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 will be found below. We focus on what is new here: upgrade, and its interplay with modified update. It is not hard to see the soundness of the following principles, stated with existential modalities for convenience:

2.4.3. THEOREM. *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. $\langle \varphi! \rangle (\psi \wedge \chi) \leftrightarrow (\langle \varphi! \rangle \psi \wedge \langle \varphi! \rangle \chi)$.
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. $\langle \# \varphi \rangle p \leftrightarrow p$.
8. $\langle \# \varphi \rangle \neg\psi \leftrightarrow \neg\langle \# \varphi \rangle \psi$.
9. $\langle \# \varphi \rangle (\psi \wedge \chi) \leftrightarrow (\langle \# \varphi \rangle \psi \wedge \langle \# \varphi \rangle \chi)$.
10. $\langle \# \varphi \rangle \langle K \rangle_i \psi \leftrightarrow \langle K \rangle_i \langle \# \varphi \rangle \psi$.
11. $\langle \# \varphi \rangle \langle pref \rangle_i \psi \leftrightarrow (\neg\varphi \wedge \langle pref \rangle_i \langle \# \varphi \rangle \psi) \vee (\langle pref \rangle_i (\varphi \wedge \langle \# \varphi \rangle \psi))$.
12. $\langle \# \varphi \rangle E\psi \leftrightarrow E\langle \# \varphi \rangle \psi$.

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 how we change the preference relation. It says essentially this. After an upgrade for φ , a preference link leads from the current world to a φ -world if and only if this same link existed before. This means that it has not been removed, ruling out the case where it led from an actual world verifying φ to some other one verifying $\neg\varphi$. The three cases where the link does persist are described succinctly in the two disjuncts on the right-hand side. Finally, 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 commutativity principle for preference and existential modalities. \square

This dynamic epistemic upgrade logic (henceforth, *DEUL*) can explain general effects of changes in information and preference. In particular, we can think of our upgrade system as transforming underlying *world-* or *object-*comparison relations, but then, in the matching logic, recording also what changes take place because of this at the level of *propositions*. Thus, given the earlier-noted expressive power of the modal language for notions of preference between propositions, we can derive principles telling us what new propositional preferences obtain after an upgrade action, and relate these to the propositional preferences that we had before. As an illustration, consider the $\forall\exists$ -notion of preference stated earlier:

$$P^{\forall\exists}(\varphi, \psi) \quad \text{iff} \quad U(\psi \rightarrow \langle\text{pref}\rangle_i\varphi).$$

2.4.4. FACT. The following equivalence holds

$$\langle\sharp A\rangle P^{\forall\exists}(\varphi, \psi) \quad \text{iff} \quad P^{\forall\exists}(\langle\sharp A\rangle\varphi, \langle\sharp A\rangle\psi) \wedge P^{\forall\exists}((\langle\sharp A\rangle\varphi \wedge A), (\langle\sharp A\rangle\psi \wedge A)).$$

Proof. This is a simple calculation showing how the dynamic epistemic upgrade logic axiom system works in practice:

$$\begin{aligned} \langle\sharp A\rangle P^{\forall\exists}(\varphi, \psi) &\leftrightarrow \langle\sharp A\rangle U(\psi \rightarrow \langle\text{pref}\rangle_i\varphi) \\ &\leftrightarrow U(\langle\sharp A\rangle(\psi \rightarrow \langle\text{pref}\rangle_i\varphi)) \\ &\leftrightarrow U(\langle\sharp A\rangle\psi \rightarrow \langle\sharp A\rangle\langle\text{pref}\rangle_i\varphi) \\ &\leftrightarrow U(\langle\sharp A\rangle\psi \rightarrow (\neg A \wedge \langle\text{pref}\rangle_i\langle\sharp A\rangle\varphi) \vee (\langle\text{pref}\rangle_i(A \wedge \langle\sharp A\rangle\varphi))) \\ &\leftrightarrow U(\langle\sharp A\rangle\psi \wedge \neg A \rightarrow \langle\text{pref}\rangle_i\langle\sharp A\rangle\varphi) \wedge U(\langle\sharp A\rangle\psi \wedge A \rightarrow \\ &\quad \langle\text{pref}\rangle_i(\langle\sharp A\rangle\varphi \wedge A)) \\ &\leftrightarrow P^{\forall\exists}(\langle\sharp A\rangle\varphi, \langle\sharp A\rangle\psi) \wedge P^{\forall\exists}((\langle\sharp A\rangle\varphi \wedge A), (\langle\sharp A\rangle\psi \wedge A)). \end{aligned}$$

\square

A similar analysis applies Von Wright’s ‘All All’ notion of preference between propositions, relating new preferences in this sense to earlier ones – but we leave this calculation to the reader.⁴

In addition, as noted earlier, our epistemic upgrade logic 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 2.3:

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

Dynamic epistemic upgrade logic can also analyze the basic propositional scenarios of obeying successive commands or reasoning toward achieving practical goals proposed in [Zar03] and [Yam06].

2.4.5. THEOREM. *Dynamic epistemic upgrade logic 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. \square

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

We have reached the first major conclusion of this chapter:

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

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* and false assertions at the actual world; and, do not waste anyone’s time.

⁴One might want to be more radical here, and insist on dynamic preference-changing actions directly *at the level of propositions*, without any dependence on an underlying world-level. This is in line with versions of belief revision theory where one is instructed to come to believe certain propositions. We have some thoughts on this alternative; but it would involve both entrenchment and preference relations on sets of propositions, a more syntactic perspective which raises as many design issues as the world-based semantic framework used in this chapter.

- (b) Do not make assertions which are *common knowledge* in the whole group, 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

$$\sharp p; \sharp \neg p$$

is not inconsistency, but it still has some strange aspects. Generally speaking, such a sequence makes the ordering non-connected, as it removes arrows either way between p -worlds and $\neg p$ -worlds. It is an interesting issue which sequences of upgrades are coherent, in that they preserve the property of connectedness.

In reality, one often resolves conflicts in suggestions 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 go beyond the ambitions of this chapter, as they involve the gap between actual informational events and their translation into the idealized model changes offered by dynamic epistemic logics, whether for update or upgrade.

2.5 Relation change and product upgrade

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 sides is expressed in the following standard result:

2.5.1. FACT. Assertions φ hold in the relativized model iff their syntactically relativized versions were true in the old model:

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

In this light, the reduction axioms for public announcement merely express the inductive facts about the modal assertion $\langle !\varphi \rangle \varphi$ 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 for the new link-cutting update operation $\varphi!$. [Ben07a] 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* 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 sort:

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

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:

$$\begin{aligned} \langle R \cup (? \neg \varphi; \top; ? \varphi) \rangle \psi &\leftrightarrow \langle R \rangle \psi \vee \langle ? \neg \varphi; \top; ? \varphi \rangle \psi \\ &\leftrightarrow \langle R \rangle \psi \vee (\neg \varphi \wedge E(\varphi \wedge \psi)). \end{aligned}$$

The example suggests a much more general observation, which we state informally in the following:

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

Proof. Clearly, 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. The standard *PDL* axioms for union, composition, and tests in *PDL* then rewrite all statements $\langle R^\sharp \rangle \varphi$ to compounds in terms of just basic modalities $\langle R_i \rangle \varphi$. \square

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

2.5.3. EXAMPLE. Our upgrade operation $\sharp \varphi$ is really the relation-changer:

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

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

$$\begin{aligned} \langle \sharp \varphi \rangle \langle R \rangle \psi &\leftrightarrow \langle (? \neg \varphi; R) \cup (R; ? \varphi) \rangle \langle \sharp \varphi \rangle \psi \\ &\leftrightarrow \langle ? \neg \varphi; R \rangle \langle \sharp \varphi \rangle \psi \vee \langle R; ? \varphi \rangle \langle \sharp \varphi \rangle \psi \\ &\leftrightarrow (\neg \varphi \wedge \langle R \rangle \langle \sharp \varphi \rangle \psi) \vee \langle R \rangle (\varphi \wedge \langle \sharp \varphi \rangle \psi). \end{aligned}$$

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 – as will be shown briefly now.

Product update

The usual generalization of eliminative public announcement is product update ([Ger99], [BMS98], [DHK07]). We briefly recall the basics.

2.5.4. DEFINITION. An *event model* is a tuple $\mathcal{E} = (E, \sim_i, PRE)$ such that E is a non-empty set of events, \sim_i 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 a can be performed at world s only if the world s fulfills the precondition $PRE(a)$.

2.5.5. DEFINITION. Given an *epistemic model* \mathcal{M} , an *event model* \mathcal{E} , the *product update model* $\mathcal{M} \times \mathcal{E}$ is defined as follows:

- The *domain* is $\{(s, a) \mid s \text{ a world in } \mathcal{M}, a \text{ an event in } \mathcal{E}, (\mathcal{M}, s) \models PRE(a)\}$.
- 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} .

2.5.6. REMARK. For a version leaving all old worlds in place, as with the above new announcement operator $\varphi!$, we need to cut relational links again (instead of eliminating worlds) so as to ‘isolate’ those pairs (s, a) where (\mathcal{M}, s) fails to satisfy the precondition for action a .

2.5.7. DEFINITION. 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 public announcement, it yields a complete and decidable logic via a set of reduction axioms for all possible forms of postcondition (cf. [BMS98], [BEF93], [BEK06]).

Product upgrade

Next, we enrich epistemic event models with preference relations, indicating which events agents prefer over which others. These preferences may come from pay-offs or other benefits, but they may also be abstract relative plausibilities again, as in models of conditional logic.

2.5.8. DEFINITION. 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:

$(s, t) \preceq_i (u, v)$ iff $s \preceq_i u$ and $t \preceq_i v$.

This product upgrade covers at least the earlier upgrade instruction $\sharp p$ for suggestions. To see this, consider the event model of Figure 2.4:

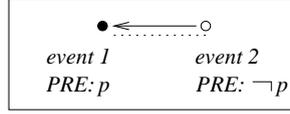


Figure 2.5: Two indistinguishable events

Here the two events cannot be distinguished epistemically by the agent. Recall that the reflexive loops of preference relations are omitted.

2.5.9. FACT. $\mathcal{M}_{\sharp\varphi} \cong \mathcal{M} \times \mathcal{E}^{\sharp\varphi}$, where the event model $\mathcal{E}^{\sharp\varphi}$ has two events “seeing that φ ” (event 1), “seeing that not- φ ” (event 2), with event 2 \preceq event 1.

Proof. From an 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. The old epistemic accessibilities just get copied with the product rule, since accessibility holds between all pairs of events. As for the new preference structure, consider 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 corresponding pair

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

Our product upgrade rule gives a preference here from left to right. 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 $\mathcal{M}, t \models \neg\varphi$. \square

Thus, as with public announcement and epistemic product update, one simple event model suffices to mimic 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. We will return this issue later in Chapter 4.

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

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

We are not going to spell out here what these axioms look like, but it is a routine exercise. Our second main conclusion in this chapter is this:

Preference upgrade can be combined naturally with the richest knowledge update mechanisms known so far.

Virtues of the combination We think the above setting has independent interest. 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’ (cf. [Sch97]). Our product update system models both kinds, and is able to study their interplay. Moreover, there is a computational angle, viz. ‘dynamic deontic’ versions of *PDL* itself, starting from preferences between worlds, but moving on to preferences between actions ([Mey88], [Mey96]). [PW04] follows up on the latter, and propose relation change as a way of ‘changing policies’. [Roh05] 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.

2.6 Illustrations: defaults and obligations

We have presented an upgrade mechanism for incoming triggers that change preferences. We now illustrate this framework in two concrete settings. Our aim is not some full-fledged application to existing systems. We merely show how the logical issues in this chapter correspond to real questions of independent interest.

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 tonight”.

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. [Vel96] 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, one can go this way, using a scenario of *relation change for defaults*, as in our earlier Section 2.3. 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. Here is a more drastic procedure, which will validate the preceding intuition:

2.6.1. DEFINITION. 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.⁵ Formally, this is one of our earlier *PDL*-style relation-changes: 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 the preceding Section 2.4.

2.6.2. FACT. Relational default processing can be axiomatized completely.

Proof. By the method of Section 2.5, the key reduction axiom follows automatically from the given *PDL*-form, yielding

$$\langle \# \varphi \rangle \langle \text{pref} \rangle \psi \leftrightarrow (\varphi \wedge \langle \text{pref} \rangle (\varphi \wedge \langle \# \varphi \rangle \psi) \vee (\neg \varphi \wedge \langle \text{pref} \rangle (\neg \varphi \wedge \langle \# \varphi \rangle \psi)) \vee (\neg \varphi \wedge E(\varphi \wedge \langle \# \varphi \rangle \psi)). \quad \square$$

Thus, we have a plausible version 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. 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. [PS93] for a good exposition, and Chapter 3 for a logical exploration.

From default logic to belief revision Default logic is naturally connected with *belief revision*, since new facts may change earlier conclusions. More generally, an analysis of preference change seems very congenial to analyzing belief revision, with world ordering by relative plausibility (cf. [Gro88], [Rot06]). Indeed, the paper [Ben07a] shows that the techniques for handling relation change developed in this chapter can be used to analyze various belief revision policies, and axiomatize their properties completely.

⁵This is known as the ‘lexicographic’ change in the belief revision community. The idea was first suggested in [Nay94].

Deontic logic and commands

Similar considerations apply to deontic logic ([Åqv87]). 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 sum total 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 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 suggests a preference order among worlds. And then, once more, we can think of this setting dynamically, using our upgrade scenario.

Initially, there are no preferences between worlds. Then some moral authority starts ‘moralizing’: introducing evaluative distinctions between worlds. If this process works well, we get a new ordering of 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, and the issue of ‘coherence’ in Section 2.3 comes back again with greater force now.

Looking backward, or forward in upgrade Deontic logic also raises new issues. One semantic intuition is that, after a command (say, ‘Thou shalt not kill’), the core proposition becomes true in all best possible worlds. Thus, 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 2.4, 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 ’. *DEL* 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.

But it is not so easy to just define an action as achieving the truth of some proposition. This works for simple factual effects of actions like opening a door ([BOR06]), but it is not clear what this should even mean with more complex stipulations. E.g., there is no obvious ‘seeing to it that’ arbitrary mixtures of knowledge and ignorance in groups arise, and the same seems true of complex deontic commands. Whether deontic reasoning needs some sort of future-oriented update and upgrade seems an interesting question. For temporal logics of such *STIT* operators, cf. [BPX01].

2.7 Related work

The ideas in this chapter 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. [Mey88] 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 [Spo88], [Vel96] 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. [TT99] use ideas from update semantics to formalize deontic reasoning about obligations, but with motivations 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. [Mey96] 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. Taking belief change as its starting point, [Han95] identified four types of changes in preference, namely revision, contraction, addition and subtraction, and showed that they satisfy plausible postulates for rational changes in preferences. [PW04] 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. [Dem05] 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’, [Roh05] 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 satisfiability is undecidable. [PPC06] 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. They provide distinctions, like knowing one’s duty versus having a duty to know, whose dynamics invites a merge with our system. Our own approach goes back to [BEF93], which discusses general formats for upgrading preference relations. [Zar03] 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. More generally, [Yam06] takes the update paradigm to logics of commands and obligations, modeling 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. Yamada’s command operator for propositions A can be modeled *exactly* as an upgrade sending R to $R; ?A$ in our system. But this chapter provides a much more general treatment of possible upgrade instructions. Finally, [Rot06] presents a format for relation change which can handle all major current policies for belief revision. [Ben07a] shows how one can axiomatize such policies completely using the methods in Section 2.5 of this chapter.

A full-fledged comparison doing justice to all these approaches is unfortunately beyond the scope of this chapter.

2.8 Conclusion

In this chapter we have shown that preference upgrade is a natural and crucial part of logical dynamics. It can be modeled essentially as relation change in a standard dynamic format, up to the expressive level of the best available system, that of epistemic product update.

Still, our approach leaves things to be desired. In particular, many settings call for more finely-grained distinctions as to *intensity* of preferences, as happens in quantitative versions of social choice theory. [Liu06b] has proposed a mechanism of *utility update*, inspired by [Spo88], [Auc03] and [Liu04], which combines utilities of old worlds and of events to compute utilities of new worlds. With such a system, we can upgrade defaults, duties, or preferences in games in a more controlled local fashion, by adding or subtracting ‘points’. The relationship between our relational upgrade and more quantitative utility update also poses some interesting technical issues, for which we refer to the more extensive exploration in [Liu06b].