Changing for the better : preference dynamics and agent diversity

Liu, F.

Publication date
2008

Chapter 4

Comparisons and Combinations

In the preceding two chapters, we have presented two different approaches to preference structure and preference change. These proposals were based on different intuitions, both plausible and attractive. Even so, the question naturally arises how the two perspectives are related. The aim of the present chapter is to draw a comparison, connect the modal logic based view of Chapter 2 with the priority-based view of Chapter 3, and try to integrate them. To see why this makes sense, let us start by briefly summarizing some key ideas.

The approach taken in Chapter 2 had the following main points:

- The *basic structures* were models $(W, \sim, \preceq, V)$ with a set of worlds (or objects) with a reflexive and transitive binary 'betterness' relation $\preceq$ ('at least as good as'), while we also assumed a standard epistemic accessibility relation for agents.

- The *language* used was an epistemic language extended with a universal modality U (or its existential dual E) plus a standard unary modality $[bett]$ using betterness as its accessibility relation.[1]

- *Preference* was treated as a relation over propositions, with the latter viewed as sets of possible worlds. Precisely, it is a *lifting* of the betterness relation to such sets. There are various ways of lifting, determined by quantifier combinations. One typical example uses $\forall\exists$, and it was defined as:

  $$Pref^{\forall\exists}(\varphi, \psi) ::= U(\psi \to \langle bett \rangle \varphi).$$

  An alternative would be to take the epistemic modality K instead of U here, making preference a partly betterness-based, partly epistemic notion, subject to introspection in the usual way.

[1] To distinguish from notations we will use later on for preference, we write the operator here as $[bett]$ instead of $[pref]$ in Chapter 2. We will mostly omit the subscript for agents.

- *T*he language also had a *dynamic* aspect, and information update affects knowledge just as in standard DEL. Likewise, changes in preference were dealt with, by first defining changes in the basic betterness relations at the possible world level. These were handled by the standard DEL methodology. A typical example was the new reduction axiom for changes in betterness modalities after the action of 'suggesting that A':

$$\langle \sharp\varphi \rangle \langle bett \rangle \psi \leftrightarrow (\neg\varphi \wedge \langle bett \rangle \langle \sharp\varphi \rangle \psi) \vee (\langle bett \rangle (\varphi \wedge \langle \sharp\varphi \rangle \psi)).$$

  Given this reduction axiom plus that for the universal modality U, we were then able to derive a reduction axiom for the propositional preference operator $Pref^{\forall\exists}$. In case we use the epistemic modality to define preference between propositions, this will actually make *two* types of event, and two corresponding reduction axioms relevant. In addition to explicit betterness transformers such as suggestions, also, a pure information update affects K operators, and hence also the preferences involving them.

Thus, Chapter 2 provides an account of betterness ordering of objects, and its dynamics, intertwined with agents' knowledge and information update. Preference comes out as a *defined* concept.

Prima facie, Chapter 3 took a quite different approach, starting from a given priority order among propositions ('priorities'), and then deriving a preference order among objects. Again, we review the main contributions:

- A *priority base* is given first, consisting of strictly ordered properties:

$$P_1(x) \gg P_2(x) \gg \cdots \gg P_n(x)$$

  Next, a *preference order* $\preceq$ over objects is *derived* from this priority base, depending on whether the objects have the properties in the priority sequence or not. There are many ways for such a derivation, but we have taken one inspired by Optimality Theory (OT) saying that the earlier priorities in the given sequence count strictly heavier than the later ones. The preference derived this way is a quasi-linear order, not just reflexive and transitive, but also 'connected'.

- Here, too, preference was intertwined with information and agents' propositional attitudes, but this time, focused on their *beliefs*. In the case of incomplete information about which properties in the priority sequence objects possess, preference was defined in terms of which properties agents believe the objects to possess.

- On the syntax side, to speak about preference over objects, a fragment of a *first-order language* has been used, with a binary relation $Pref$. A doxastic belief operator was then added to explicitly describe object preferences based on beliefs.

- In this setting, too, we analyzed *changes in preference*. There are two possible sources for this, as before. Either preference change is caused by a change in the priority sequence, leading to a new way of ordering objects, or it is caused by a change in beliefs. We have given reduction axioms for both these scenarios, now for the languages appropriate here.

Clearly, despite the difference in starting point, the agendas of Chapters 2 and 3 are very similar. The purpose of this Chapter is to make this explicit, and see what questions arise when we make the analogies more precise.

Our discussion will be mostly semantics-oriented, though we will get to a comparison with syntax later in this chapter. Next, we will make the following simplification, or rather abstraction. At the surface, Chapter 2 speaks about ordering over possible worlds, and propositions as set of possible worlds, while Chapter 3 is about ordering 'individual objects' using their properties. In what follows, we will take all 'objects' to be 'worlds' in modal models - but this is just a vivid manner of speaking, and nothing would be lost if the reader were just to think of 'points' and 'properties' instead of worlds and propositions.[2] Finally, in order to be neutral on the different perspectives of Chapters 2 and 3, we start with two orderings at different levels. One is the betterness relation over possible worlds, written as $(W, \preceq)$, the other a preference or priority relation over propositions, viewed as sets of possible worlds, denoted by $(\mathscr{P}, <)$.

The greater part of this chapter will be devoted to exploring the deeper connection between these two orderings. The main questions we are going to pursue are the following:

- How to *derive* a preference order of 'betterness' over possible worlds from an ordered priority sequence?

- In the opposite direction, how to *lift* a betterness relation on worlds to an ordering over propositions?

The following diagram shows these two complementary directions:

$$\begin{array}{c} (\mathscr{P}, <) \\ \text{lift} \uparrow \quad \downarrow \text{derive} \\ (W, \preceq) \end{array}$$

[2]There are interesting intuitive differences, however, between *object preference* and *world preference*, which will be discussed briefly at the end of this chapter. See also the remark after Theorem 3.7.12.

Besides these connections between the two levels, what will be of interest to us is how to relate dynamical changes at the two levels to each other.

The chapter is organized as follows. In Section 4.1 we propose a simple structure called *structured model* containing both a preference over possible worlds, and an ordering over propositions. In Section 4.2, we first study ways of deriving object preferences from a priority base, including some representation theorems. We will relate our method to other approaches in the literature, in particular, the preference merge in [ARS02]. In Section 4.3, we look at the opposite direction: how to lift an object preference relation to an ordering over propositions. A characterization theorem will be proved for the natural lifting of type $\forall\exists$. In Section 4.4 we study how concrete order-changing operations at the two levels correspond to each other. Section 4.5 is to connect the PDL-definable preference change to an alternative approach from the recent literature, product update on belief revision, which uses 'event model' as in dynamic-epistemic logic. Then in Section 4.6, we move from semantic structures to formal languages, and provide a comparison of the various logical languages used in the previous two chapters. Finally, in Section 4.7 we compare the different ways of preference interacting with belief in the previous two chapters, and we end this chapter with a proposal of putting all our systems together in one 'doxastic preferential predicate logic' of object and world preference.

4.1 Structured models

Preference over propositions $(\mathscr{P}, <)$ and betterness over possible worlds $(W, \preceq)$ can be brought together as follows:

4.1.1. Definition. A *structured model* $\mathcal{M}$ is a tuple $(W, \preceq, V, (\mathscr{P}, <))$, where W is a set of possible worlds, $\preceq$ a preference relation over W, V a valuation function for proposition letters, and $(\mathscr{P}, <)$ an ordered set of propositions, the 'important properties' or priorities.[3]

Structured models extend standard modal models, which may be viewed as the special case where $\mathscr{P}$ equals the powerset of W.

Here are some further notational stipulations. As in Chapter 2, $y \preceq x$ means that the world x is 'at least as good as' the world y or 'preferable over' y, while

[3]Compared with ordinary modal models $(W, \preceq, V)$, structured models have a new component, viz. a set of distinguished propositions $\mathscr{P}$. Agenda-based modal models introduced in [Gir08] have a similar structure, as an agenda is a set of distinguished propositions, too. It would then be very natural to look at modal languages over worlds where the valuation map only assigns propositions from $\mathscr{P}$ as values to atomic proposition letters. Moreover, we can let $\mathscr{P}$ determine the world preference relations. It would be interesting to find out what happens to standard modal logics on such restricted models, as these will now encode information about the structure of $\mathscr{P}$. This issue will not be pursued in this chapter.

$y \prec x$ means that x is strictly preferable over y: i.e. $y \preceq x$ but not $x \preceq y$. To emphasize preference relations induced by a priority sequence $(\mathscr{P}, <)$, we will write $y \preceq_{\mathscr{P}} x$. In general, the set $\mathscr{P}$ will be a partial order - but it is useful to also have a simpler case as a warm-up example. Suppose that $\mathscr{P}$ is a flat set, without any ordering. We then write the structured model simply as $(W, \preceq, V, \mathscr{P})$.

Some explanations on notation: We use capital letters $S, T, X, Y, P_i, \ldots$ for arbitrary propositions in the set $\mathscr{P}$, small letters $x, y, z, \ldots$ for arbitrary possible worlds, while Px means that $x \in P$. As for the other level, we write $Y \trianglelefteq X$ (or $Y \triangleleft X$) when proposition X is at least as good as (or strictly preferable to) Y.

Structured models simply combine the approaches in the previous two chapters. In particular, the syntactic priorities of Chapter 3 are now moved directly into the models, and sit together with an order over worlds. We will see how these two layers are connected in the next sections.

4.2 Moving between levels: From propositional priorities to object preferences

We first look at the derivation of preference ordering from a primitive priority sequence. This is a common scenario in many research areas. For instance, given a goal base as a finite set of propositions with an associated rank function, [CMLLM04] extends this priority on goals to a preference relation on alternatives. Likewise, in the theory of belief revision, an epistemic 'entrenchment relation' orders beliefs, those with the lowest entrenchment being the ones that are most readily given up ([GM88], [Rot03]). But our main motivating example in Chapter 3 was linguistic *Optimality Theory* ([PS93]). Here a set of alternative structures is generated by the grammatical or phonological theory, while an order over that set is determined by given strictly ordered constraints. Language users then employ the optimal alternative that satisfies the relevant constraints best.

What interests us in this chapter is the formal mechanism itself, i.e. how to get a preference order from a priority sequence. There are many proposals to this effect in the literature. In what follows, we will discuss a few, including the one adopted in Chapter 3, to place things in a more general perspective, and facilitate comparison with Chapter 2.

Chapter 3 considered only finite linearly ordered priority sequences, and object order was derived via an *OT*-style lexicographic stipulation. In terms of structured models $(W, \preceq, V, (\mathscr{P}, <))$, $(\mathscr{P}, <)$ is a finite linear order. Now, we first spell out the *OT*-definition:[4]

[4] The formulation here is slightly, but inessentially different from the inductive version we had in Chapter 3.

$$y \preceq^{OT}_{\mathscr{P}} x ::= \forall P \in \mathscr{P}(Px \leftrightarrow Py) \vee \exists P' \in \mathscr{P}(\forall P < P'(Px \leftrightarrow Py) \wedge (P'x \wedge \neg P'y)).$$

To recall how this works, we repeat an earlier illustration from Chapter 3:

4.2.1. Example. Alice is going to buy a house. In doing so, she has several things to consider: the cost, the quality, and the neighborhood. She has the following priority sequence:

$$C(x) \gg Q(x) \gg N(x),$$

where $C(x)$, $Q(x)$ and $N(x)$ stand for 'x has low cost', 'x is of good quality' and 'x has a nice neighborhood', respectively. Consider two houses d_1 and d_2 with the following properties: $C(d_1), C(d_2), \neg Q(d_1), \neg Q(d_2), N(d_1)$ and $\neg N(d_2)$. According to the OT-definition, Alice prefers d_1 over d_2 strictly, i.e. $d_2 \prec d_1$.

This OT-definition is by no means new. It was also investigated in other literature, e.g. [BCD+93] on priority-based handling of inconsistent sets of classical formulas, and [Leh95] on getting new conclusions from a set of default propositions. It has been called *leximin ordering* in this and other literature.

But besides the OT-definition, there are other ways of deriving object preferences from a priority base. To see this, first consider the above-mentioned simplest case where $\mathscr{P}$ is *flat*. The following definition gives us a very natural order (we call it the '$*$-definition'):

$$y \preceq^{*}_{\mathscr{P}} x ::= \forall P \in \mathscr{P}(Py \to Px).$$

This is found, e.g. in the theory of default reasoning of Veltman ([Vel96]), or in the topological order theory of Chu Spaces ([Ben00]). Incidentally, the same order would arise on our OT-definition of object preference if we take the flat set to have the trivial universal ordering relation. Next, as noted earlier, given a non-strict order $\preceq$, one can define its strict version $\prec$ in the following:

$$y \prec x ::= y \preceq x \wedge \neg(x \preceq y).$$

So the strict version of $y \preceq^{*}_{\mathscr{P}} x$ can be written as:

$$y \prec^{*}_{\mathscr{P}} x ::= \forall P \in \mathscr{P}(Py \to Px) \wedge \exists P' \in \mathscr{P}(P'x \wedge \neg P'y).$$

In the following we will only present non-strict versions of preference.

Returning to general structured models with the OT-definition, several questions arise naturally. Which orders over possible worlds are produced? Can we always find some priority sequence that produces such a given object order? The following representation result gives a precise answer. We had a similar result in Chapter 3, but this time, we will provide a new proof, while dropping the finiteness assumption.

4.2.2. THEOREM. *For any standard model $\mathcal{M} = (W, \preceq, V)$, the following two statements are equivalent:*
(a) there is a structured model $\mathcal{M}' = (W, \preceq, V, (\mathscr{P}, <))$ s.t.

$$y \preceq x \quad \textit{iff} \quad y \preceq^{OT}_{\mathscr{P}} x \quad \textit{for all} \quad x, y \in W.$$

(b) $y \preceq x$ is a quasi-linear order.

Proof. (a)⇒(b). Chapter 3 showed that the OT-definition always generates a quasi-linear order $\preceq^{OT}$.
Now for the converse direction (b)⇒(a). First, define a 'cluster' as a maximal subset X of W such that $\forall y, z \in X$: $y \preceq z$. Clusters exist by Zorn's Lemma, and different clusters are disjoint by their maximality. Each point x belongs to a cluster, which we will call C_x. First, we define a natural ordering of clusters reflecting that of the worlds:

$$C' \trianglelefteq C \quad \text{if} \quad \exists y \in C', \exists x \in C : y \preceq x.$$

We first prove the following connection with the given underlying object order:

4.2.3. LEMMA.

$$y \preceq x \quad \textit{iff} \quad C_y \trianglelefteq C_x.$$

Proof. (⇒). By definition, $x \in C_x$ and $y \in C_y$, so $C_y \trianglelefteq C_x$.
(⇐). If $C_y \trianglelefteq C_x$, then by definition $\exists u \in C_x$, $v \in C_y$ with $v \preceq u$. So we have $u \preceq x (x \in C_x)$ and $y \preceq v (y \in C_y)$ – and hence by transitivity, we get $y \preceq x$. □

Importantly, this order on the clusters is not just quasi-linear: it is a *strict linear order*, as required in our definition of a priority base. Accordingly, we define the set $\mathscr{P}^\bullet$ as the set of all clusters. Moreover, we let the order of greater priority run in the upward direction of the cluster order. (This choice of direction is just a convention - but one has to pay attention to it in the following arguments.)

We are now ready to prove our main statement, for all worlds y, x:

$$y \preceq x \quad \text{iff} \quad y \preceq^{OT}_{\mathscr{P}^\bullet} x.$$

(⇒). Assume that $y \preceq x$. We have to show that

$$\forall P \in \mathscr{P}^\bullet (Px \leftrightarrow Py) \vee \exists P' \in \mathscr{P}^\bullet (\forall P < P' (Px \leftrightarrow Py) \wedge (P'x \wedge \neg P'y)).$$

To see this, note that by Lemma 4.2.3, $C_y \trianglelefteq C_x$. Then we distinguish two cases. If $C_y = C_x$, then x, y share this 'property' and no other, and hence the left disjunct holds. If $C_y \neq C_x$, then $C_y \triangleleft C_x$ (by linearity), and then $x \in C_x$ and $y \notin C_x$, and therefore, the right disjunct holds.

($\Leftarrow$). Now assume that $y \preceq^{OT}_{\mathscr{P}^\bullet} x$. There are again two cases. First, let x, y share the same 'properties' in $\mathscr{P}^\bullet$, i.e. $Px \leftrightarrow Py$. Then in particular, $x, y \in C_x$, and hence $y \preceq x$. Next let there be some $P' \in \mathscr{P}^\bullet$ with $P'x \wedge \neg P'y$, while x, y share the same properties $P < P'$. Since $P'x$, we must have $P' = C_x$, while $C_x \neq C_y$. Since we have $\forall P < P'(Px \leftrightarrow Py)$, we conclude that $\neg(C_x \lhd C_y)$. Therefore $C_y \unlhd C_x$, and hence by Lemma 4.2.3, $y \preceq x$. □

Interestingly, the relevant propositions constructed in this argument are all *mutually disjoint*. This may not be the most obvious scenario from the Optimality Theory perspective, but as we shall see later, it is a technically convenient setting, which involves no loss of generality.

Alternative definitions of world orderings

Another appealing definition of preference over alternatives from a priority sequence is called *best-out ordering*. It was used in [BCD+93] as an alternative way of priority-based handling of inconsistent set of classical formulas:

$$y \preceq^{best}_{\mathscr{P}} x ::= \forall P \in \mathscr{P}(Px \wedge Py) \vee \exists P'(\forall P < P'(Px \wedge Py) \wedge (P'x \wedge \neg P'y)).$$

The best-out-format is similar to the OT-definition, except that instead of having the equivalence $(Px \leftrightarrow Py)$ it only requires a conjunction $(Px \wedge Py)$. Intuitively, this means that only the positive cases matter when deriving a preference order from a priority base. Looking at Example 4.2.1 again, we get $d_2 \preceq d_1$ and $d_1 \preceq d_2$, d_1 and d_2 are equally preferable for Alice, because after observing that $\neg Q(d_1)$ and $\neg Q(d_2)$, she won't consider N at all. While these alternatives are interesting, we now move to a slightly more general comparative perspective.

Preference merge and partial order

In what follows, we will consider another formal approach for deriving preference from a priority base, proposed in [ARS02]. The ideas here work differently, but as we shall see, it is a natural generalization of our $\preceq^{OT}_{\mathscr{P}}$ in two ways:

- one merges arbitrary given relations on objects, not just those given by propositions or properties.
- the definition for the merged preference works with partial orders on sets of propositions, which includes our linear priority orders as a special case.

First, we briefly review the basic notions in [ARS02]. A preference relation is any reflexive transitive relation ('pre-order'). Suppose there is a family of such preference relations $(R_x)_{x \in V}$, all on the same set W, V is a set of variables. A question arising in many settings, from social choice theory to 'belief merge' in belief revision theory, is how to combine these relations into a single relation on the same set W. In particular, the given preferences can originate from different criteria that we wish to combine according to their importance - as in 'many-dimensional decision problems'. But the initial orders do not come as a mere set. We need further structure to arrive at a plausible notion of merge. The core notion here is a *priority graph*, defined in [ARS02] as follows:

4.2.4. Definition. A *priority graph* is a tuple $(N, <, v)$ where N is a set of nodes, $<$ is a strict partial order on N (the 'priority relation') and v is a function from N to a set of variables standing for the given binary relations to be merged. The set N may be infinite, though we will mainly look at finite cases.

A priority graph may be viewed as an ordering of variables for 'input relations'. Some variables may be represented several times in the ordering, simply by repeating their occurrences in the priority graph. Now, any priority graph denotes the following operator on the given preference relations:

4.2.5. Definition. The V-ary operator $\circ$ denoted by the priority graph $(N, <, v)$ is given by

$$m \circ ((R_x)x \in V)n \iff \forall i \in N.(mR_{v(i)}n \vee \exists j \in N.(j < i \wedge mR^{<}_{v(j)}n))$$

where $V = v[N]$, the set of variables that occur in the graph, and $R^{<}_{v(j)}$ is the strict version of the partial order $R_{v(j)}$.

The operator $\circ$ takes a set of preference relations $(R_x)_{x \in V}$ and returns a single one. The concrete intuition behind Definition 4.2.5 is this:

> $m \preceq_G n$ if for all separate relations $R_{v(i)}$: either $mR_{v(i)}n$ or if n fails this 'test' with respect to m, it 'compensates' for this failure by doing better than m on some more important test, i.e. some relation $R^{<}_{v(j)}$ holds with $j < i$, i.e. with higher priority in the graph.

To see how all this works, we look at the following example from [ARS02]:

4.2.6. Example. The priority graph $g_1 = (N, <, v)$ has $N = \{1, 2, 3\}$ with $1 < 2$ and $1 < 3$ and $v(1) = y$, $v(2) = x$ and $v(3) = y$. See Figure 4.1.

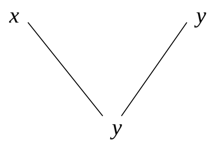

Figure 4.1: Priority graph

Here lower down means higher priority. This graph denotes a binary operator since there are only two distinct variables. It takes two preference relations, say

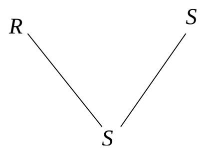

Figure 4.2: Representations of preference relations

R and S, and returns one which represents their combination with the given priority. Thus, if σ_1 is the operator denoted by the graph, then $\sigma_1(R, S)$ is the following prioritized combination of R and S, see Figure 4.2.

Working out what this means by definition 4.2.5, we obtain the relation $\sigma_1(R, S) = (R \cap S) \cup S^{<}$.

To understand further how the *ARS* system works, note that the above relation $(R \cap S) \cup S^{<}$ is exactly the same as that induced by the simpler linear priority graph shown in Figure 4.3.

Figure 4.3: Linear graph

This is no coincidence. There is a rather interesting algebraic structure behind all this. In particular, [ARS02] proved that every merge operation defined by a priority graph can be derived as an algebraic composition of the following two basic relations:

(i) $(R \cap S) \cup S^{<}$,

(ii) $R \cap S$.

Interestingly, the priority graph inducing the second operation is not linear, but a simple disjoint union, and hence a partial order:

Figure 4.4: Partial order

It is easy to see that by definition 4.2.5, this works out to the intersection of R and S. Indeed, this suggests a much more general observation showing how operations on priority graphs affect the merged outcomes. It can be proved by simple inspection of the above definition:

4.2.7. FACT. For any two priority graphs $\mathscr{P}$ and $\mathscr{P}'$, the following equivalence holds

$$y \preceq_{\mathscr{P} \uplus \mathscr{P}'} x \quad \text{iff} \quad (y \preceq_{\mathscr{P}} x) \text{ and } (y \preceq_{\mathscr{P}'} x),$$

where $\uplus$ denotes the *disjoint union* of the two graphs.[5]

Comparing *ARS* to *OT*

The *ARS* way of thinking matches well with the *OT*-style definition for $\preceq^{OT}$. The latter worked with an ordered set of propositions rather than relations R_i. But it can easily be recast in the latter manner. We merely associate each proposition A with an ordering relation $\preceq (A)$ derived as follows:

$$y \preceq (A)x \quad \text{iff} \quad (Ay \rightarrow Ax) \vee (\neg Ay \wedge Ax).$$

This is precisely the sort of world ordering encountered in belief revision when a signal comes in that A is the case (cf. [Rot07], [Ben07a]). Indeed, it is completely interchangeable whether we talk about propositions A or relations $\preceq(A)$: both contain the same information. But the relational format is more general, since not all given object orders need to be generated from propositions in this simple manner. Also, intuitively, the priority order of propositions in the *OT*-format corresponds to the order of relations in the priority graph (where we will disregard the issue of repeated occurrences of variables, which would correspond to repeated occurrences of the same property in a priority sequence).

We can now write the *ARS*-definition as it applies to our proposition orders:

$$y \preceq^{ARS}_{\mathscr{P}} x ::= \forall P \in \mathscr{P}((Py \rightarrow Px) \vee \exists P' < P(P'x \wedge \neg P'y)).$$

For a contrast, recall the definition of $\preceq^{OT}_{\mathscr{P}}$:

$$y \preceq^{OT}_{\mathscr{P}} x ::= \forall P \in \mathscr{P}(Px \leftrightarrow Py) \vee \exists P' \in \mathscr{P}(\forall P < P'(Px \leftrightarrow Py) \wedge (P'x \wedge \neg P'y)).$$

Syntactically, this looks quite different, with an inversion in quantifier scope. But actually, the following equivalence result holds:

4.2.8. THEOREM. *For any finite linearly ordered set of propositions $\mathscr{P}$,*

$$y \preceq^{OT}_{\mathscr{P}} x \quad \textit{iff} \quad y \preceq^{ARS}_{\mathscr{P}^*} x \quad \textit{for all worlds} \quad x, y,$$

where $\mathscr{P}^*$ is the priority graph derived from $\mathscr{P}$ by replacing each proposition A by its relation $\preceq (A)$ and keeping the old order from $\mathscr{P}$.

[5] In fact, this is one of the axioms of the graph calculus studied in [Gir08], it is formulated as $\langle G_1 \uplus G_2 \rangle s \leftrightarrow \langle G_1 \rangle s \cap \langle G_2 \rangle s$ there (s is a nominal).

Proof. ($\Rightarrow$). Let $y \preceq^{OT}_{\mathscr{P}} x$. Suppose $y \preceq^{ARS}_{\mathscr{P}^*} x$ does not hold. Then we have

$$\exists P \in \mathscr{P}^*((Py \wedge \neg Px) \wedge \forall P' < P(P'x \rightarrow P'y)).$$

Let P^* be such a P. Notice that $\neg(P^*x \leftrightarrow P^*y)$. By $y \preceq^{OT}_{\mathscr{P}} x$, let $P^\bullet$ be the smallest $P \in \mathscr{P}^*$, s.t. $P^\bullet x \wedge \neg P^\bullet y$. Then $P^\bullet$ cannot come before P^*, since we have $\forall P' < P^*(P'x \rightarrow P'y)$. The only possible case is then $P^\bullet \geq P^*$. Here $P^\bullet = P^*$ leads to a contradiction, as $P^\bullet x$ but $\neg P^*x$. But if $P^\bullet$ comes after P^*, then by the OT-definition, $P^*x \leftrightarrow P^*y$, and again we get a contradiction.

($\Leftarrow$). Let $y \preceq^{ARS}_{\mathscr{P}^*} x$. Suppose it is not the case that $y \preceq^{OT}_{\mathscr{P}} x$, then we have

$$\exists P \in \mathscr{P}\neg(Px \leftrightarrow Py) \wedge \forall P'(\exists P < P'\neg(Px \leftrightarrow Py) \vee (P'x \rightarrow P'y))$$

From $\exists P \in \mathscr{P}\neg(Px \leftrightarrow Py)$, without loss of generality, take $P^* = P$, the smallest $P \in \mathscr{P}$ where $\neg(P^*x \leftrightarrow P^*y)$. Applying the conjunct $\forall P'(\exists P < P'\neg(Px \leftrightarrow Py) \vee (P'x \rightarrow P'y))$ to P^*, which was chosen to be smallest with the non-equivalence, we get that $\neg(P^*y \rightarrow P^*x)$, i.e. $P^*y \wedge \neg P^*x$. Applying the ARS-definition to P^*, we have that $(P^*y \rightarrow P^*x) \vee \exists P' < P^*(P'x \wedge \neg P'y)$. Here the second disjunct does not hold because of P^*'s minimality for equivalence failure. But the first disjunct does not hold either: $(P^*y \rightarrow P^*x)$ cannot occur since $P^*y \wedge \neg P^*x$. So we conclude that $\neg(y \preceq^{ARS} x)$, a contradiction, and we are done. □

The ARS-definition also applies to situations in which the order of $\mathscr{P}$ is partial, and/or infinite, so we can think of it as a natural generalization of our earlier OT-definition.

More on partial orders and pre-orders

There are many technical results in [ARS02], including an algebraic axiomatization of merge operations and a characterization of priority graph-based merge with respect to some conditions from social choice theory. But our reason for considering this system is simply this: the base preference relations over possible worlds considered in Chapter 2 are reflexive, transitive pre-orders, rather than quasi-linear ones. Moreover, in the context of dynamic relation transformations, reflexivity and transitivity were preserved (for a proof, see Chapter 2), but not in general connectedness. We will return to this issue later, but for the moment, we state a representation result comparable to Theorem 4.2.2, which works for object pre-orders and partial graph order, thereby generalizing our earlier case of quasi-linear object order and strictly linear constraint order.

4.2.9. Theorem. *Let $\mathcal{M} = (W, \preceq, V)$ be an ordinary object model. Then the following two statements are equivalent:*

(a) there is a structured model $\mathcal{M}' = (W, \preceq, V, (\mathscr{P}, <))$*, with* $(\mathscr{P}, <)$ *a priority graph over given propositional relations* $\preceq (A)$ *such that*

$$y \preceq x \quad \textit{iff} \quad y \preceq^{ARS}_{\mathscr{P}} x \quad \textit{for all} \quad x, y \in W.$$

(b) $y \preceq x$ *is a reflexive transitive pre-order.*

Proof. (a)⇒ (b). It is easy to see from the earlier definition that all relations $\preceq^{ARS}_{\mathscr{P}}$ generated by priority graphs decorated with pre-orders must be reflexive and transitive.
Conversely, consider the direction (b)⇒ (a). Similarly to the proof of Theorem 4.2.2, we can define the set of all clusters C_x. We define an order over clusters in the same way:

$$C \trianglelefteq C' \quad \text{if} \quad \exists y \in C, \exists x \in C' : y \preceq x.$$

Lemma 4.2.3 holds for pre-orders as well, since its proof did not rely on the quasi-linearity of the object order. Indeed, there is again an 'improvement': the cluster ordering becomes a *partial order*: i.e. two clusters which mutually precede each other must be the same. (The latter property is not true for pre-orders in general, and not even for quasi-linear orders.)

As we have seen in the preceding, each 'cluster proposition' P is associated with an ordering relation $\preceq (P)$, so the partial order of clusters gives us a partially ordered priority graph. Again we need to show that the relation induced by this matches up with the given one, that is:

$$y \preceq x \quad \text{iff} \quad y \preceq^{ARS}_{\mathscr{P}} x.$$

(⇒). Assume that $y \preceq x$. We need to show that $y \preceq^{ARS}_{\mathscr{P}} x$, i.e., $\forall P \in \mathscr{P}((Py \to Px) \vee \exists P' < P \wedge P'x \wedge \neg P'y))$. So, consider any cluster proposition P, or its associated relation. If $Py \to Px$, we are done. So suppose $Py \wedge \neg Px$. Then we must have $P = C_y$, since as before, our cluster propositions form a disjoint partition. Moreover, we have $C_y \trianglelefteq C_x$ by our Lemma applied to $y \preceq x$, and since $C_y \neq C_x$ (they are disjoint since x is not in C_y), we get $C_y \triangleleft C_x$. It is easy to see that C_x is the 'compensating' P' for x that we need to verify the second disjunct in the ARS-definition.
(⇐). Assume that $y \preceq^{ARS}_{\mathscr{P}} x$. Consider the predicate $P = C_y$, for which clearly Py holds. There are two cases. First assume that Px. Since $P(= C_y)$ is a cluster, we have that $y \preceq x$. Next, assume that not Px. By the 'compensation clause' of $y \preceq^{ARS} x$, $\exists P' < P : P'x \wedge \neg P'y$. Clearly, this P' can only be C_x, and hence we have $C_y \triangleleft C_x$, $C_y \trianglelefteq C_x$, and by Lemma 4.2.3, we get $y \preceq x$. □

We will continue our discussion of the ARS-format in later sections. For now we just make one simple but useful observation:

4.2.10. FACT. Let $\mathscr{P}; A$ be a set of ordered priorities with a priority A at the end, then we have

$$y \preceq^{ARS}_{\mathscr{P};A} x \Rightarrow y \preceq^{ARS}_{\mathscr{P}} x.$$

Proof. If there is a predicate P' in $\mathscr{P}; A$ such that $P'y \wedge \neg P'x$, then by the ARS-definition, there must be a 'higher compensating predicate in $\mathscr{P}; A$, but given that A comes last, this compensation can only happen inside $\mathscr{P}$. □

4.3 Going from world preference to propositional priorities

Now let us take the opposite perspective to that of the preceding Section 4.2. This time, a primitive order over worlds is given, and we would like to lift it to an order over propositions, so that we can compare sets of possible worlds.

This scenario, too, occurs in many places in the literature, with various interpretations of the basic relation $y \preceq x$. It is interpreted as 'x is as least as normal (or typical) as y' in [Bou94] on conditional and default reasoning, as 'x at least as preferred or desirable as y' in [DSW91], as 'x is no more remote from actuality than y' in [Lew73] on counterfactuals, and as 'x is as likely as y' in [Hal97] on qualitative reasoning with probability. In all these settings, it makes sense to extend the given order on worlds to an order of propositions P, Q. For instance, in real life, students may have preferences concerning courses, but they need to also form an order over kinds of courses, say theoretical versus practical, i.e. over sets of individual courses. Likewise, we may have preferences regarding individual commodities, but we often need a preference over sets of them. And similar aggregation scenarios are abundant in social choice theory, for which an extensive survey is [BRP01].

Quantifier lifts

One obvious way of lifting world orders $x \preceq y$ to proposition or set orders $P \preceq Q$ uses definitional schemas that can be classified by the quantifiers which they involve. As has been observed by many authors (cf. [BRG07]), there are four obvious two-quantifier combinations:

$$\forall x \in P \forall y \in Q : x \preceq y; \quad \forall x \in P \exists y \in Q : x \preceq y;$$
$$\exists x \in P \forall y \in Q : x \preceq y; \quad \exists x \in P \exists y \in Q : x \preceq y.$$

One can argue for any of these. [BOR06] claims that $\forall\forall$ is the notion of 'preference' intended by von Wright in his seminal work on preference logic ([Wri63]) and provides an axiomatization.[6] But the tradition is much older, and (modal)

[6][BOR06] need to assume quasi-linearity of the world preference relation to define the lifted relation within their modal language.

logics for preference relations over sets of possible worlds have been considered by [Lew73], [Bou94] and [Hal97], and other authors. In particular, [Hal97] studied the above combination of $\forall\exists$, defined more precisely as follows:

4.3.1. DEFINITION. Let $(W, \preceq, V, (\mathscr{P}, <))$ be any structured model. For $X, Y \in \mathscr{P}$, we define $Y \trianglelefteq^{\forall\exists} X$ if for all $y \in Y$, there exists some $x \in X$ with $y \preceq x$.

As usual, we define the strict variant $Y \lhd X$ as '$Y \trianglelefteq^{\forall\exists} X$ and not $X \trianglelefteq^{\forall\exists} Y$'. Similarly, we can define all the other quantifier combinations.

As we have seen in Chapter 2, some of these combinations can be expressed directly in a modal language for the models $\mathcal{M} = (W, \preceq, V)$ by combining a betterness modality with a universal modality.[7] As for the $\forall\exists$-preference, it can be defined in such a modal language as follows:

$$Pref^{\forall\exists}(\varphi, \psi) ::= U(\psi \to \langle bett \rangle \varphi).$$

This says that, for any ψ-world in the model, there exists a better φ-world. Once again, this 'majorization' is one very natural way of comparing sets of possible worlds - and it has counterparts in many other areas which use derived orders on powerset domains. In particular, [Hal97] took this definition (with an interpretation of 'relative likelihood' between propositions) and gave a complete logic for the case in which the basic order on W is a pre-order. It is also well-known that Lewis gave a complete logic for preference relations over propositions in his study of counterfactuals in [Lew73], where the given order on W is quasi-linear. In what follows, we side-step these completeness results,[8] but raise a few more semantic issues, closer to understanding the lifting phenomenon per se.

More on $\forall\exists$-preference

A natural question to ask is: Which lift is 'the right one'? This is hard to say, and the literature has never converged on any unique proposal. There are some obvious necessary conditions, of course, such as the following form of 'conservatism':

`Extension rule`: For all $x, y \in X$, $\{y\} \trianglelefteq \{x\}$ iff $y \preceq x$.

But this does not constrain our lifts very much, since all four quantifier combinations satisfy it. We will not explore further constraints here. Instead, we concentrate on one particular lift, and try to understand better how it works. One question that comes to mind immediately is this: Can the properties of

[7] [BOR06] shows that the standard modal language plus universal modality is not sufficiently expressive to define the intended meaning of $\forall\forall$ or $\exists\forall$. The language has to be extended further by a strict preference operator $[bett^s]$.

[8] See however our discussion in Section 3.7.

an underlying preference on worlds be preserved when it is lifted to the level of propositions? In particular, consider reflexivity and transitivity that we assumed for preference in Chapter 2. Can we show $\trianglelefteq^{\forall\exists}(\varphi, \psi)$ has these two properties? We can even prove something stronger:

4.3.2. Fact. Reflexivity and transitivity of the relation $\preceq$ are preserved in the lifted relation $\trianglelefteq^{\forall\exists}$, but also vice versa.

Proof. *Reflexivity.* To show that $\trianglelefteq^{\forall\exists}(X, X)$, by Definition 4.3.1, we need that $\forall x \in X \exists y \in X : x \preceq y$. Since we have $x \preceq x$, take y to be x, and we get the result.

In the other direction, we take $X = \{x\}$. Then apply $\trianglelefteq^{\forall\exists}(X, X)$ to it to get $\forall x \in X \exists x \in X : x \preceq x$. Since x is the only element of X, we get $x \preceq x$.

Transitivity. Assume that $\trianglelefteq^{\forall\exists}(X, Y)$ and $\trianglelefteq^{\forall\exists}(Y, Z)$. We show that $\trianglelefteq^{\forall\exists}(X, Z)$. By Definition 4.3.1, this means we have $\forall x \in X \exists y \in U : x \preceq y$ and $\forall y \in Y \exists z \in Z : y \preceq z$. Then by transitivity of the base relation, we have that $\forall x \in X \exists z (x \preceq z$, and this is precisely $\trianglelefteq^{\forall\exists}(X, Z)$.

In the other direction, let $x \preceq y$ and $y \preceq z$. Take $X = \{x\}, Y = \{y\}$ and $Z = \{z\}$. Applying $\trianglelefteq^{\forall\exists}$, we see that $X \trianglelefteq Y$ and $Y \trianglelefteq Z$, and hence by transitivity for sets, $X \trianglelefteq Z$. Unpacking this, we see that we must have $x \preceq z$. □

Likewise, we can prove that if $\trianglelefteq^{\forall\exists}$ is quasi-linear, then so is $\preceq$. But the converse direction does not hold.

Besides the three properties mentioned, many others make sense. In fact, the preceding simple argument suggests a more extensive correspondence between relational properties for individual orderings and their set liftings, which we do not pursue here.

Next, staying at the level of propositions, consider an analogue to the representation theorems of the preceding section. Suppose we have a preference relation that is a $\forall\exists$-lift from a base relation over possible worlds. What are necessary and sufficient conditions for being such a relation? The following theorem provides a complete characterization.

4.3.3. Theorem. *A binary relation $\trianglelefteq$ over propositions satisfies the following four properties iff it is a $\forall\exists$-lifting of some preference relation over the underlying possible worlds.*

(1) $Y \trianglelefteq X \Rightarrow Y \bigcap Z \trianglelefteq X$ (*left downward monotonicity*)

(2) $Y \trianglelefteq X \Rightarrow Y \trianglelefteq X \bigcup Z$ (*right upward monotonicity*)

(3) $\forall i \in I, Y_i \trianglelefteq X \Rightarrow \bigcup_i Y_i \trianglelefteq X$. (*left union property*)

(4) $\{y\} \trianglelefteq \bigcup_i X_i \Rightarrow \{y\} \trianglelefteq X_i$ *for some* $i \in I$. (*right distributivity*)

Proof. ($\Leftarrow$). Assume that $\trianglelefteq$ is a $\forall\exists$-lifting. We show that $\trianglelefteq$ has the four properties.

(1). Assume $Y \trianglelefteq X$, i.e., $\forall y \in Y \exists x \in X : y \preceq x$. Since $Y \bigcap Z \subseteq Y$, we also have $\forall y \in Y \bigcap Z \exists x \in X : y \preceq x$, and hence $Y \bigcap Z \trianglelefteq X$.

(2). Assume $\forall y \in Y \exists x \in X : y \preceq x$. Since $X \subseteq X \bigcup Z$, we have $\forall y \in Y \exists x \in X \bigcup Z : y \preceq x$: that is, $Y \trianglelefteq X \bigcup Z$.

(3). Assume that for all $i \in I$, $\forall y \in Y_i \exists x \in X : y \preceq x$. Let $y \in \bigcup_i Y_i$, then for some j: $y \in Y_j$. By the assumption, we have $\forall y \in Y_j \exists x \in X : y \preceq x$, so $\exists x \in X : y \preceq x$. This shows that $\bigcup_i Y_i \trianglelefteq X$.

(4). Assume that $\forall y \in \{y\} \exists x \in \bigcup_i X_i : y \preceq x$. Then there exists some X_i with $\exists x \in X_i : y \preceq x$, that is: $\{y\} \trianglelefteq X_i$ for some $i \in I$.

($\Rightarrow$). Going in the opposite direction, we first define an object ordering

$$y \preceq x \quad \text{iff} \quad \{y\} \trianglelefteq \{x\}. \quad (\dagger)$$

Next, given any primitive relation $Y \trianglelefteq X$ with the above four properties, we show that we always have

$$Y \trianglelefteq X \quad \text{iff} \quad Y \trianglelefteq^{\forall\exists} X.$$

where the latter relation is the lift of the just-defined object ordering.
($\Rightarrow$). Assume that $Y \trianglelefteq X$. For any $y \in Y$, $\{y\} \subseteq Y$ by reflexivity. Then, by Property (1) we get $\{y\} \trianglelefteq X$. But then also $\{y\} \trianglelefteq \bigcup_{x \in X} \{x\}$, as $X = \bigcup_{x \in X} \{x\}$. By Property (4), there exists some $x \in X$ with $\{y\} \trianglelefteq \{x\}$, and hence, by Definition ($\dagger$), $y \preceq x$. This shows that, for any $y \in Y$, there exists some $x \in X$ s.t. $y \preceq x$, which is to say that $Y \trianglelefteq^{\forall\exists} X$.

($\Leftarrow$). Assume that $\forall y \in Y \exists x \in X : y \preceq x$. By definition ($\dagger$), $y \preceq x$ is equivalent to $\{y\} \trianglelefteq \{x\}$. Since $\{x\} \subseteq X$, by Property (2) we get that $\{y\} \trianglelefteq X$. Thus, for any $y \in Y$, $\{y\} \trianglelefteq X$. By Property (3) then, $\bigcup_{y \in Y} \{y\} \trianglelefteq X$, and this is just $Y \trianglelefteq X$.[9] □

Finally, we ask how orderings of propositions produced by set lifting relate to the priority orderings which were central in Chapter 3 and Section 4.2. Intuitively, there need not be any strong connection here, since priority ordering is about relative importance, rather than preference. Nevertheless, in some special cases, we can say more. We have some results of this sort on the the $\trianglelefteq^{\forall\exists}$-lifting, but instead, we cite an observation from Chapter 3 which is relevant here. Priority order and lifted object order can coincide when we work with special sets of worlds. Here is how:

[9] [Hal97] gave a complete logic for $Pref^{\forall\exists}$, whose axiomatization looks different from our characterization. But one should be able to show they are essentially equivalent.

4.3.4. Definition. A set X is *upward closed* if

$$\forall x, y \in X(y \in X \wedge y \preceq x \to x \in X).$$

4.3.5. Fact. Consider only sets X that are upward closed. We define

$$y \preceq x \quad \text{iff} \quad \forall X(x \in X \leftrightarrow y \in X) \vee \exists X(x \in X \wedge y \notin X).$$

Then the $\trianglelefteq^{\forall\exists}$-lifting of this object ordering becomes equivalent to set inclusion, and the latter is equivalent to the priority sequence:

$$X \subseteq Y \Leftrightarrow X \gg Y.$$

Much more general questions arise here about connections when transformations are repeated:

- Given a priority order $(\mathscr{P}, <)$ and its induced world order $\preceq^{ARS}_{\mathscr{P}}$(or $\preceq^{OT}_{\mathscr{P}}$), when can $<$ on $\mathscr{P}$ be retrieved as a quantifier lift of $\preceq^{ARS}_{\mathscr{P}}$?
- Given a world order $(W, \preceq)$, and some lift, say $\trianglelefteq^{\forall\exists}$, used as a priority order on the powerset $\mathcal{P}(W)$, when can the relation $\preceq$ on W be retrieved as the derived order of $\trianglelefteq^{\forall\exists}$?

Answering these questions would show us further connections between the two levels of worlds and propositions, when both relative importance and preference are involved in lifting and deriving. We will not pursue these matters here - but there is certainly more harmony than what we have uncovered so far.

4.4 Dynamics at two levels

Dynamics has been one of the core issues investigated in the previous two chapters. We have modeled changes in the betterness relation over possible worlds in Chapter 2. And we also considered possible changes in priority sequences in Chapter 3. As we pointed out after introducing the structured models in Section 4.1, we want to relate the changes at the two levels in a systematic manner.

Relation transformers at the world level

In Chapter 2 betterness relations over possible worlds are the locus of dynamics. New information or other triggers come in which rearrange this order. We recapitulate a few concrete operations from Chapter 2.

The simplest operation was $Cut(A)$. It cuts only the accessibility links between worlds in A and $\neg A$, but keeps all possible worlds around.[10] The following

[10] This is different from eliminative update for public announcement, where worlds may disappear. With preference change, belief revision, or even information update for memory-bounded agents, it is reasonable to keep all possible worlds, as shown in [BL07], [Ben07a], and [Liu07].

definition of $Cut(A)$ is stated using some self-explanatory notation from propositional dynamic logic:

4.4.1. Definition. For any relation R and proposition A, the new relation $Cut(A)(R)$ is defined as:

$$Cut(A)(R) ::= (?A; R; ?A) \cup (?\neg A; R; ?\neg A).$$[11]

Next, 'suggesting A' (written as $\sharp A$) was the main action considered in Chapter 2, changing (preference) relations in the following manner:

4.4.2. Definition. For any relation R and proposition A, the new relation $\sharp A(R)$ is defined as:

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

Thus, in the updated models no $\neg A$-worlds are preferable to A-worlds. We can also define the suggestion relation by

$$Cut(A)(R) \cup (?\neg A; R; ?A).$$

In fact, $Cut(A)$ is a basic operation, and it will return below. One slightly more complex relevant operation is 'upgrade with A', written as $\Uparrow A$:

4.4.3. Definition. The new relation $\Uparrow A(R)$ is defined as:

$$\Uparrow A(R) ::= Cut(A)(R) \cup (?\neg A; \top; ?A).$$

After the new information A has been incorporated, the upgrade places all A-worlds on top of all $\neg A$-worlds, keeping all other comparisons the same. This time, besides $Cut(A)$ as before, new links may be added by the disjunct $(?\neg A; \top; ?A)$ with the *universal relation* $\top$. Alternatively, going back to Section 4.2, $\Uparrow A(R)$ can also be defined as

$$Cut(A)(R) \cup \preceq (A)(R).$$

Here, it is important to note that the above definitions make the ordering over possible worlds hold between whole 'zones' of the model given by propositions. Worlds which satisfy exactly the same propositions behave the same. Thus, we are ordering 'kinds of worlds' (through a partition of the domain of worlds), rather than worlds per se. Keeping this way of thinking in mind helps understand many technicalities in what follows.

[11]Note that the link-cutting operation is found as *agenda expansion* in [BRG07] and *PDL test action* in [HLP00].

We have already used a simple program fragment of the *PDL* language to define new relations over possible worlds after some definable operation has taken place. Let us look at this more generally. The basic elements that build up the new relations are:

$$?\varphi \mid ; \mid \cup \mid R \mid \top.$$

These are the standard *PDL* operations. $?\varphi$ is a test, while ; and $\cup$ denote sequential composition and choice, respectively. R is the given input relation, treated as an atom, and the constant $\top$ denotes the universal relation that holds everywhere. The following fact provides a useful 'normal form':

4.4.4. Fact. Every *PDL* operation in the above style has a definition as a union of finite 'trace expressions' of the form '$?A_1; \{R, \top\}; ?A_2; \{R, \top\}; ...$', where $\{R, \top\}$ means either R or $\top$.

Proof. To turn arbitrary program expressions into normal form, we apply the following equivalences that drive unions outwards:

$$\begin{aligned} (S \cup T); U &= S; U \cup T; U, \\ S; (T \cup U) &= S; T \cup S; U, \end{aligned}$$

where S, T, U, and V are of any form of $?\varphi$, or R or $\top$ from the *PDL* language. This may still leave us with sequences of tests, instead of single ones. But the former can be contracted using the valid identity

$$?A_1; ?A_2 \quad = \quad ?(A_1 \wedge A_2). \qquad \square$$

In fact, our definitions of the operations $Cut(A)$, $\Uparrow A$, and $\sharp A$ were already in normal form. But in principle, many relations can be defined in a *PDL* format, covering a large space of possible relation transformers. While some of these make intuitive sense, others are just mathematical curiosities. We will return to this issue later on.

Propositional level transformers of priority orders

At the level of linearly ordered finite sets of priorities $(\mathscr{P}, <)$, some natural operations have been considered in Chapter 3. These were:

- $[^+A]$ adds A to the right of the sequence,
- $[A^+]$ adds A to the left,
- $[-]$ drops the last element of the sequence,
- $[i \leftrightarrow i+1]$ interchanges the i-th and i+1-th elements.

We have shown that any one of the first two operations, plus the last two are sufficient to make any changes to a finite ordered priority sequence. Some operations are considered in [ARS02] as well, now of course on priority graphs, which generalize finite sequences. Examples are taking the disjoint union of two relations, deleting a node from a priority graph, and putting a new link j below i if this does not change the down-set of i. To keep things simple, in what follows, we consider only graph analogues of the preceding 'postfixing' and 'prefixing' operators $[^+A]$ and $[A^+]$, which we will write as $\mathscr{P}; A$ and $A; \mathscr{P}$. In the special case of flat sets $\mathscr{P}$, both collapse to one operation $\mathscr{P} + A$.

Relating dynamics at the two levels

Having reviewed some basic operations at the two levels, let us now try to find systematic correspondences between them. First, we state two trivial but general observations that are easily obtained from the definitions and representation theorems in the previous sections. These work globally in that we do not need to identify which specific dynamic changes have taken place. For convenience, we state several results in the *OT*-setting, but our results hold for *ARS*-style pre-orders as well in general. We start by taking object-level transformers to propositional ones.

4.4.5. CLAIM. *Given the OT-definition, if a relation change over possible worlds models respects quasi-linear order, then there exists a corresponding change on the set of ordered propositions.*

Proof. Assume the old relation is R_1 and after a change it becomes R_2. Since R_1 and R_2 are both quasi-linear orders, Theorem 4.2.2, gives corresponding priority sequences $\mathscr{P}_1$ and $\mathscr{P}_2$. This is the propositional change we are after. □

Of more interest is the question whether the induced change between $\mathscr{P}_1$ and $\mathscr{P}_2$ can be *defined* using the given one from R_1 to R_2. We discuss this issue later on. For now, here is the converse general observation, which is even more trivial.

4.4.6. CLAIM. *Given the OT-definition, if a relation change over propositions respects the linear order, then there is a corresponding change in preference over possible worlds.*

This follows at once from the earlier definitions.

Uniform definable connections

Next, we consider more uniform connections between transformations at the two levels. We start with the following notion about relation-transforming functions.

4.4.7. Definition. Let F: $(\mathscr{P}, A) \to \mathscr{P}'$, where $\mathscr{P}$ and $\mathscr{P}'$ are set of propositions, and A is a new proposition. Let σ: $(\preceq, A) \to \preceq'$, where $\preceq$ and $\preceq'$ are relations over possible worlds, and A is a new proposition. We say that *the map F induces the map σ*, given a definition of deriving object preferences from propositions, if, for any set of propositions $\mathscr{P}$ and new proposition A, we have

$$\sigma(\preceq_{\mathscr{P}}, A) = \preceq_{F(\mathscr{P},A)} .$$

We start our discussion of such connections with a simplest case, viz. adding a new proposition to a flat priority set. Recall the $(*)$-definition for inducing object order in Section 4.2.

4.4.8. Fact. Given the $(*)$-definition, taking a suggestion A given some relation over possible worlds is induced by the following operation at the propositional level: adding a new proposition A to a flat $\mathscr{P}$. More precisely, the following diagram commutes:

$$\begin{array}{ccc} \langle W, \mathscr{P}\rangle & \xrightarrow{+A} & \langle W, \mathscr{P} \cup A\rangle \\ {\scriptstyle *}\downarrow & & \downarrow{\scriptstyle *} \\ \langle W, \preceq\rangle & \xrightarrow{\sharp A} & \langle W, \sharp A(\preceq)\rangle \end{array}$$

Proof. We need to prove the following equivalence:

$$y \preceq^*_{\mathscr{P}+A} x \quad \text{iff} \quad y \sharp A(\preceq^*_{\mathscr{P}})x.$$

$(\Leftarrow)$. We know that after $\sharp A$, the relation between y and x can be expressed as:

$$\sharp A(\preceq^*_{\mathscr{P}}) ::= (?A; \preceq^*_{\mathscr{P}}; ?A) \cup (?\neg A; \preceq^*_{\mathscr{P}}; ?\neg A) \cup (?\neg A; \preceq^*_{\mathscr{P}}; ?A)$$

In terms of a relation between arbitrary worlds x and y, the above three cases give the implication $Ay \to Ax$. By $y \preceq^*_{\mathscr{P}} x$, we also have that $\forall P \in \mathscr{P}$: $Py \to Px$. Hence $\forall P \in \mathscr{P} + A$: $Py \to Px$: i.e., $y \preceq^*_{\mathscr{P}+A} x$.
$(\Rightarrow)$. Assume that $y \preceq^*_{\mathscr{P}+A} x$, i.e., $\forall P \in \mathscr{P} + A$: $y \in P \to x \in P$. In particular, it cannot the case that $y \in A \wedge x \notin A$. Thus, out of all pairs in the given relation R, those satisfying $(?A; \preceq^*_{\mathscr{P}}; ?\neg A)$ can no longer occur. This is precisely how we defined the relation $y \sharp A(\preceq^*_{\mathscr{P}})x$. □

Simple as it is, this argument shows that natural operations at both levels can be tightly correlated.

Next we consider the case of an *ordered* set $(\mathscr{P}, <)$, where a new proposition A is added in front. The dynamics at the two levels is correlated as follows:

4.4.9. Fact. Given the OT-definition, upgrade $\Uparrow A$ over possible worlds is induced by the following operation on propositional priority orders: prefixing a new A to an ordered propositional set $(\mathscr{P}, <)$. More precisely, the following diagram commutes:

$$\begin{array}{ccc} \langle W, (\mathscr{P}, <)\rangle & \xrightarrow{A;\mathscr{P}} & \langle W, (A;\mathscr{P}, <)\rangle \\ {\scriptstyle OT}\downarrow & & \downarrow{\scriptstyle OT} \\ \langle W, \preceq\rangle & \xrightarrow{\Uparrow A} & \langle W, \Uparrow A(\preceq)\rangle \end{array}$$

Proof. Again, we have to prove a simple equivalence:

$$y \preceq^{OT}_{A;\mathscr{P}} x \quad \text{iff} \quad y \Uparrow A(\preceq^{OT}_{\mathscr{P}})x.$$

($\Leftarrow$). We know that after the operation $\Uparrow A$, the relation between y and x can be expressed as:

$$\Uparrow A(\preceq^{OT}_{\mathscr{P}}) ::= (?A; \preceq^{OT}_{\mathscr{P}}; ?A) \cup (?\neg A; \preceq^{OT}_{\mathscr{P}}; ?\neg A) \cup (?\neg A; \top; ?A).$$

Call these Cases (a), (b) and (c), respectively. We show that $y \preceq^{OT}_{A;\mathscr{P}} x$, i.e.,

$$\forall P \in A; \mathscr{P}(Px \leftrightarrow Py) \vee \exists P' \in A; \mathscr{P}(\forall P < P'(Px \leftrightarrow Py) \wedge (P'x \wedge \neg P'y)).$$

In Case (a) and (b), the new predicate A in top position does not distinguish the worlds x, y , and hence their order is determined by just that in $\mathscr{P}$. In Case (c), since $(Ax \wedge \neg Ay)$, for any pair of y and x, A is the compensating predicate P' in $A; \mathscr{P}$ that we need for the OT-definition. Thus in all cases, we have $y \preceq^{OT}_{A;\mathscr{P}} x$.

($\Rightarrow$). Assume that $y \preceq^{OT}_{A;\mathscr{P}} x$. Consider the following two cases. (i) For all $P \in A; \mathscr{P}$ $Px \leftrightarrow Py$. In particular then, $Ax \leftrightarrow Ay$, and we get Cases (a) and (b). (ii) There exists some $P' \in A; \mathscr{P}$ such that for all $P < P'(Px \leftrightarrow Py)$ while $(P'x \wedge \neg P'y)$. Then $P' = A$ or $P' \in \mathscr{P}$. If $P' = A$, $Ax \wedge \neg Ay$, and we get Case (c). If $P' \in \mathscr{P}$, then, by the prefixing, $A < P'$, by assumption we have $Ax \leftrightarrow Ay$, and again we get Case (a) and (b). □

We have now proved two results in a similar format, linking operations at the possible world level to operations at the priority level. Actually, one can think of such connections in two ways:

(i) Given any priority-level transformer, we define a matching world-level relation transformer.

(ii) Given any world-level relation transformer, we define a matching priority-level transformer.

As an instance of direction (i), let us consider the natural operation $\mathscr{P}; A$ of postfixing a proposition to an ordered propositional set. It turns out that we do not have a very simple corresponding operation at the possible world level. We need some relational algebra beyond the earlier PDL-format, witness the following observation.

4.4.10. FACT. $y \preceq^{ARS}_{\mathscr{P};A} x$ iff $y \prec^{ARS}_{\mathscr{P}} x \vee (y \preceq^{ARS}_{\mathscr{P}} x \wedge (Ay \to Ax))$.

Proof.($\Rightarrow$). Assume that $y \preceq^{ARS}_{\mathscr{P};A} x$. By Fact 4.2.10, this implies that $y \preceq^{ARS}_{\mathscr{P}} x$. Now consider the following two cases:
(i) $Ay \to Ax$. Then the right disjunct $(y \preceq^{ARS}_{\mathscr{P}} x \wedge (Ay \to Ax))$ holds.
(ii) $Ay \wedge \neg Ax$. Recall that $y \prec^{ARS}_{\mathscr{P}} x$ was defined as

$$y \preceq^{ARS}_{\mathscr{P}} x \wedge \neg(x \preceq^{ARS}_{\mathscr{P}} y).$$

We have to prove this conjunction, of which we have the left conjunct already. Suppose that $x \preceq^{ARS}_{\mathscr{P}} y$, we will derive a contradiction. Since $Ay \wedge \neg Ax$, according to the ARS-definition applied to $y \preceq^{ARS}_{\mathscr{P};A} x$, we have $\exists P' \in \mathscr{P}; A(P' < A \wedge P'x \wedge \neg P'y)$. Note that $P' \in \mathscr{P}$, and hence we get $\exists P'' \in \mathscr{P}(P'' < P' \wedge P''x \wedge \neg P''y)$. Repeating these two steps, we get an infinite downward sequence (or a finite cycle) of 'compensations' in $\mathscr{P}$, and this contradicts the well-foundedness property of the priority graph.
($\Leftarrow$). Again consider two cases:
(i) $y \prec^{ARS}_{\mathscr{P}} x$. Then $\exists P \in \mathscr{P}$: $Px \wedge \neg Py$. By the definition of $\mathscr{P}; A$, any such P satisfies $P < A$. So this P is the compensation in $\mathscr{P}; A$. Hence $y \prec^{ARS}_{\mathscr{P};A} x$.
(ii) $y \preceq^{ARS}_{\mathscr{P}} x \wedge (Ay \to Ax)$. We have to show $y \preceq^{ARS}_{\mathscr{P};A} x$ for the extended priority graph $\mathscr{P}; A$. Now if $Py, \neg Px$ holds for any $P \in \mathscr{P}; A$, then either $P \in \mathscr{P}$ or $P = A$. If $P \in \mathscr{P}$, then it has a compensation P' in the set $\mathscr{P}$ such that $P'x$ and $\neg P'y$. P' is also in $\mathscr{P}; A$. Hence $y \preceq^{ARS}_{\mathscr{P};A} x$. If $P = A$, then we would have $Ay, \neg Ax$: but this contradicts $(Ay \to Ax)$. □

Next we illustrate direction (ii) from given object-relation-transformers to priority operations. Consider our 'suggestion' operation $\sharp A$ at the level of possible worlds. First, observe what this does for linearly ordered priorities:

4.4.11. EXAMPLE. Let $\mathscr{P} = \{P\}$, which is trivially linearly ordered. This makes any P-world more preferable than any $\neg P$-world. After the suggestion A comes in, $\neg A$-worlds can no longer be preferable to A-worlds. This results in what is shown in Figure 4.5.

In particular, looking at worlds s and t, $\neg P \wedge A$ is true in s, and $P \wedge \neg A$ is true in t. But s and t are not comparable in this new model, while in the old model, t was preferable to s (since t was a P-world and s was not). Thus, we lose the connectedness.

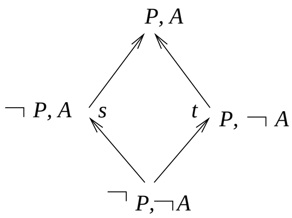

Figure 4.5: Loss of connectedness

Thus, given some relation change at the level of possible worlds, it is by no means the case that there must be a simple corresponding priority operation. As we have just seen:

4.4.12. FACT. Given a linear propositional set $(\mathscr{P}, <)$. $\sharp A$ is not induced by any F that preserves linearity.

Proof. Linear propositional sets induce quasi-linear orders, and we have just seen how suggestions can lose the quasi-linearity. □

However, if we take a partially ordered constraint set, i.e. if we move to priority graphs once more, then we get a positive result.

4.4.13. FACT. The operation $\sharp A$ is induced by the following operation F on priority graphs $\mathscr{P}$:

$$F(A, \mathscr{P}) = (A; \mathscr{P}) \uplus (\mathscr{P}; A).$$

Proof. Note that we take a disjoint union here where the same A occurs twice, but at different positions. We show the following equivalence:

$$y \preceq^{ARS}_{A;\mathscr{P} \uplus \mathscr{P};A} x \quad \text{iff} \quad y \sharp A(\preceq^{ARS}_{\mathscr{P}})x.$$

By Fact 4.2.7, we need to show:

$$y \preceq^{ARS}_{A;\mathscr{P}} x \textit{ and } y \preceq^{ARS}_{\mathscr{P};A} x \quad \text{iff} \quad y \sharp A(\preceq^{ARS}_{\mathscr{P}})x.$$

On the left-hand side, by Facts 4.4.9 and 4.4.10, we get

$$y \Uparrow A(\preceq_{\mathscr{P}})x \textit{ and } (y \prec_{\mathscr{P}} x \vee (y \preceq_{\mathscr{P}} x \wedge (Ay \rightarrow Ax)).$$

Figure 4.6 depicts what happens with the first conjunct of this formula. And Figure 4.7 depicts what happens with the second conjunct.

Clearly, the intersection of these two relations gets us precisely what we had in Example 4.4.11. The partial order on the priority graph allows for intersection, and hence incomparable situations. □

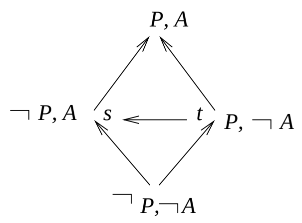

Figure 4.6: A-worlds preferable to $\neg A$-worlds

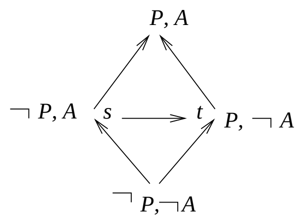

Figure 4.7: P-worlds preferable to $\neg P$-worlds

General formats of definition

These examples raise some general questions. In particular, can every PDL-definable world-level relation transformer which takes some old relation R and new proposition A as input be generated by some simple operation on priority graphs? This is not easy to answer in general, and we will merely provide some discussion. To match our examples so far, we first restrict the PDL-format in the following two aspects:

(i) We only consider atomic tests $?A$, $?\neg A$.

(ii) We only consider normal forms with single occurrences of the relation R. Forms such as $?A; R; R; ?\neg A$ will be disregarded.[12]

Given the above restrictions, we can enumerate all possible Cases:

4.4.14. Fact. There are at most 2^8 basic PDL-transformers over possible worlds.

Proof. Disjuncts in the normal form look like this:

$$\{?A, ?\neg A\}; \{R, \top\}; \{?A, ?\neg A\}$$

where at each position, there are 2 possible options, giving 8 basic cases. To define the new relation, any of these may or may not occur, giving us the exponent. □

For instance, $\Uparrow A(R)$ was a union of the three basic cases

[12] Technically, this restriction makes our operation completely distributive in its R-argument.

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

As observed before, some of these relation changes can be induced by a change in priorities, some cannot. Moreover, not all preserve the base properties of reflexivity and transitivity. For a counter-example, take $?A; R$, that is: 'if A is true, keep the old relation'. This does not preserve reflexivity, as $\neg A$-worlds have no relations any more. So this relation-transformer cannot be defined even using a partial priority graph. $\sharp A$ does yield reflexive and transitive orders (see our brief proof in Chapter 2), but not always connected ones (see Example 4.4.11), and we have just seen it is not definable by a map on linear priority graphs - though we defined it with one on partial graphs. The general question is then: *When* can an operation on relations over possible worlds be induced from some operation at the level of priority graphs? We merely state a conjecture:

4.4.15. Conjecture. *All definitions in PDL-format which preserve quasi-linear order on possible worlds are induced by definable operations on partially ordered priority graphs.*

4.5 An alternative format: 'Priority Product Update'

By now, we have looked at several operations that change given binary relations on possible worlds, or on objects in general. As observed at the beginning of this chapter, these relations can stand for many different things, from plausibility ordering in belief revision (cf. [Rot06]) to relative preference - and hence, techniques developed for one interpretation can often be used just as well for another. In particular, as we have noted, [Ben07a] used preference change as studied in our Chapter 2 for modeling various policies for belief revision. To achieve greater generality here, we have employed PDL-style definitions in the preceding Section 4.4 for relation transformers. In this section, we briefly consider an alternative approach from the recent literature on belief revision, which uses 'event models' as in dynamic-epistemic logic. Our discussion in Chapter 2 has introduced what event models $\mathcal{E}$ are, and the 'product update' $\mathcal{M} \times \mathcal{E}$ over given epistemic models $\mathcal{M}$ that is associated with them.

For a start, consider the following simple example, taken from [BS08]. The transformer $\Uparrow A$ may be naturally associated with an event model with two public announcements, or better in this setting: two 'signals' $!A$ and $!\neg A$, as shown in Figure 4.8, where the signal $!A$ is more plausible, or 'better', than $!\neg A$.

This describes a situation where we are not sure that A holds, but we do think that it is much more plausible than $\neg A$. How do we update with this event model? We need to define the new plausibility order on the pairs ⟨*old world, new*

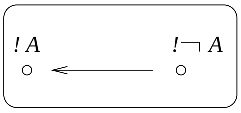

Figure 4.8: $\mathcal{E} = (!\neg A \leq !A)$

event$\rangle$ in $\mathcal{M} \times \mathcal{E}$. The general answer is the following rule of 'priority product update'. One might say that it works 'anti-lexicographically', giving priority to the event order (i.e. the last observation made), and world order only when the event order is indifferent.

4.5.1. Definition. (priority update). In product models $\mathcal{M} \times \mathcal{E}$, the plausibility relation on pairs (s, σ) is as follows:

$$(s, \sigma) \leq (s', \sigma') \quad \text{iff} \quad \sigma < \sigma' \quad \text{or} \quad \sigma \sim \sigma', s \leq s'.$$

[BS08] shows how this stipulation generalizes the examples in [Ben07a], while also dealing with a large variety of multi-agent belief revision scenarios. Moreover, the authors provide a dynamic doxastic language for which they prove completeness. Interestingly, one key element in their analysis is the use of simple modalities $\langle bett \rangle$ for the plausibility order, as well as an existential modality E, both as in our Chapter 2 ([BL07]) – as they in fact point out.[13]

Now the key innovation in this approach is the following shift. Instead of modeling different belief revision policies by different definitions, as we have done, there is just *one single update rule* which works in all circumstances. All further information about how to re-order the worlds more specifically has to be contained in the 'input signal', viz. the event model $\mathcal{E}$ with signals and priority ordering. One benefit of this approach is that one reduction axiom suffices for the basic modalities, instead of the different ones we gave for different policies.[14] This intriguing move shifts the generality from formats of definition for relation transformers to an account of the relevant event models. Moreover, event models have some formal analogies with our earlier priority graphs (for more discussion on this, we refer to [Gir08].).

There are some obvious questions about the relation between this event model format and our earlier PDL-style definitions. While things are not totally clear, and there may be a non-inclusion both ways, we can at least notice a few obvious facts. For the purpose of comparison, we stick with quasi-linear orders.

[13]Priority update also has the flavor of the above priority graph merge, and [Ben07b] gives a generalized formulation which also works for pre-orders.

[14]Of course, the latter still provide more concrete information about specific belief changes. And also, the precise status of the event models in this approach is a bit unclear, since they will often be no longer about real events from the original DEL motivation, but abstract signal combinations designed to encode revision policies.

4.5.2. Fact. Every PDL base definition defines a relation transformer whose action can also be defined as taking products $\mathcal{M} \times \mathcal{E}$ with some suitable event model $\mathcal{E}$ using Priority Product Update.

Proof. Here is a syntactic procedure for turning PDL base definitions into an equivalent event model. Since every such definition can be written in normal form by Fact 4.4.4, it is a finite union of relations between two possible worlds. The procedure goes as follows:

Step 1 Using standard propositional equivalences, rewrite the test conditions in the definition to become disjoint 'state descriptions' from some finite partition of the set of all worlds. We will write $?SD$ when referring to these. The definition then becomes a union of clauses $?SD; R; ?SD'$ and $?SD; \top; ?SD'$.

Step 2 Take the state descriptions as signals in an event model, and put an indifference relation between SD, SD' when the PDL definition has a clause $?SD; R; ?SD'$. Put a directed link from SD to SD' when the PDL definition has a clause $?SD; \top; ?SD'$.

In this way, information about the operation σ is moved into the event model $\mathcal{E}_\sigma$. Now it is easy to see that the following equivalence holds:

For any PDL definable operation σ, $\sigma(R^{\mathcal{M}}) = \leq^{BS}_{\mathcal{M}\times\mathcal{E}_\sigma}$.

The reason is that the indifference clause lets the old model decide, whereas the directed clause imposes new relations as required in the case of a clause $?SD; \top; ?SD'$.[15] □

Conversely, when is a priority update given by some event model $\mathcal{E}$ definable in our PDL format? Again, we just look at a special case, where the set of worlds does not change. The preconditions of the events in $\mathcal{E}$ then form a partition, and also, each event has a unique precondition.[16] In this setting, we also have a converse reduction. Here is the procedure:

Step 1 Let the event preconditions form a partition which is in one-to-one correspondence with the events themselves. This generates a PDL definition of the relation transformer where we test for the preconditions.

Step 2 Put $?SD; R; ?SD'$ when SD is indifferent with SD' in the event model.

[15]Note that this event model just copies the original model $\mathcal{M}$: unlike in general product update, no duplications occur of worlds via events, since all SD are mutually exclusive and together exhaustive.

[16]This is a bit like the scenario for 'protocols' in [BGK06].

Step 3 Put $?SD; \top; ?SD'$ when SD' is preferred to SD in the event model.

It is easy to see, using the analogy with the converse procedure above, that this proves the following

4.5.3. Fact. Priority product update on partition event models induces relation changes which are definable in the basic PDL format.

More can be said here. For instance, if the order on $\mathcal{E}$ is quasi-linear, then only PDL definitions are relevant which preserve quasi-linearity. Indeed, many PDL-definitions produce only pre-orders, and hence priority product update would need to be generalized (cf. [Ben07b]). And on top of that, definitions with iterated occurrences of the input relation R, while perfectly fine from the viewpoint of finding reduction axioms (cf. [BL07]) have no obvious product update counterparts. Conversely, if we were to write PDL-definitions for whole event models, we would need to generalize the [BL07] relation transformers to a setting where the operation does not just change the relation among the existing objects, but also may create new objects and drop old ones.

Our analysis suggests that PDL operations and priority product update have related but still somewhat different intuitions about achieving generality, and their connection is not yet totally clear.

4.6 Comparing logical languages

So far we have compared the proposals from the previous two chapters mainly from a *semantical* point of view. But 'logic' only arises when we also introduce *formal languages* to talk about these semantic structures. This section is meant to draw some comparisons on the formalisms that have been used. We will use some tables to summarize the main features.

The language in Chapter 2 is a modal language over combined epistemic - betterness models, with K as its standard epistemic operator, and a less standard modal operator $[bett]$ for describing 'local preferences'. Following [BOR06], a further hybrid universal modality U was added, to better express, in combination with $[bett]$, various notions of preference between propositions. This is the static part of the complete language. It was used to describe standard modal models $(W, \sim, \preceq, V)$, where, as usual, W is a set of possible worlds, $\sim$ an equivalence relation for knowledge, $\preceq$ the 'at least as good as' relation, and V the atomic valuation function.

In addition, we high-lighted dynamic changes of models in Chapter 2. To do so, two dynamic operators were included in the language: viz. modalities for public announcements $[A!]$ and suggestions $[\sharp A]$. For instance, the formula $[\sharp A]\varphi$ expresses that 'after a suggestion A, φ holds'. Typically, we had a complete set of reduction axioms to speak about the changes before and after a dynamic action. This format for relation change was even extended to PDL-definable changes. In a diagram, here are all the mentioned ingredients:

Static	`modal`	`new operator`	`hybrid`	`preference defined`
	$p \mid \neg \mid \wedge \mid K$	$[bett]$	U	$Pref^{\forall\exists}$,$Pref^{\forall\forall}$, etc.
Dynamic	$A!$	$\sharp A$		

The language in Chapter 3 was not modal, but rather a fragment of a first-order doxastic language. As shown again in the table below, it has two levels. The 'reduced language' consists of propositional formulas, and expressions of preference over constants: e.g. $Pref(d_i, d_j)$ said that 'd_i is preferable over d_j'. In the extended language, we then added first-order quantifiers, predicates, and variables, allowing us to talk about priorities explicitly. Moreover, we introduced an operator for agents' beliefs. Here the intended semantic models are first-order doxastic structures $\langle W, D, R, \{\preceq_w\}_{w \in W}, V\rangle$, where W is a set of worlds, D a set of distinguished constant objects, and R a euclidean and serial accessibility relation on W. For each w, $\preceq_w$ is a quasi-linear order on D, which is the same throughout each euclidean equivalence class, and V is again an atomic valuation function.

Similarly to Chapter 2, Chapter 3 also explored dynamics. Beliefs get changed through change in plausibility structure, just as in [Ben07a], with the revision operator $\Uparrow A$ as an example. Typically, belief revision leads to a preference change, since we defined preference in terms of beliefs. But also, the priority sequence can be changed directly, leading to a second kind of preference change. For this purpose, we introduced modalities for the earlier-mentioned four operations: $[^+A]$ for adding A to the right, $[A^+]$ for adding A to the left, $[-]$ for dropping the last element of a priority sequence, and $[i\leftrightarrow i{+}1]$ for interchanging the i-th and $i{+}1$-th elements. Again, a complete set of reduction axioms has been given for these operators. This time, the relevant Table is:

Static	`reduced language`	`extended language`	`priorities`
	$p \mid \neg \mid \wedge \mid Pref \mid B$	$P(d_i) \mid x \mid \forall$	$P(x) \gg Q(x)$
Dynamic	$\Uparrow A$	$[^+A], [A^+], [-], [i \leftrightarrow i+1]$	

This table may actually suggest a linguistic 'gap'. We have used priority sequences extensively in Chapter 3, but they have never become first-class citizens in the language. If one wanted to develop our ideas more radically, one might use a language for talking about the priorities themselves, their order, principles for reasoning about or with them, or even for changing them.

While this looks like a stark omission, things are actually much brighter. In fact, many of the previous semantic observations are already valid principles of a *calculus of priority sequences*, or even of priority graphs. For instance, the following equivalences were shown valid (Fact 4.2.7, Fact 4.4.9):

1. $\preceq_{\mathscr{P} \uplus \mathscr{P}'} = \preceq_{\mathscr{P}} \cap \preceq_{\mathscr{P}'}$.
2. $\preceq_{A;\mathscr{P}} = \Uparrow A(\preceq_{\mathscr{P}})$.

But is is clear that these are algebraic laws of some kind, provided we introduce the right notation. We see such a calculus as a natural follow-up to Chapter 3. Actually, [Gir08] has started a related investigation on 'agenda change' based on [ARS02]. We refer to Chapter 3 and [Gir08] for further details.

What about comparisons between the logical formalisms employed in Chapters 2 and 3? The difference between modal and first-order is not crucial here, as is well-known from modal correspondence theory (cf. [Ben99], [ABN98] and [BB07]). In Chapter 2, our basic concern is the betterness relation on the possible world level, and the modal language describes its 'local properties' at individual worlds. If one wants to make more global assertions, e.g. about propositional preference, unrestricted quantifiers are needed, and the universal modality U was a half-way station to full first-order logic here. But one can also use the first-order language of Chapter 3 in the end, as well as various fragments of it, modal or non-modal. Of course, things get more complex when we consider dynamic model-changing operators, as we would have to compare dynamic modal and first-order languages. Finally, when a language is added which talks about priorities, i.e. about propositions as objects, then we can view this either as a mild form of second-order logic, or as a *two-sorted first-order language* over our two-level structured models of Section 4.1. And the latter language will again have obvious modal fragments. Thus we conclude that, appearances notwithstanding, the languages used in Chapters 2 and 3 are very close.

Finally, the more interesting question is maybe this. Given the semantic connections between the structures employed in Chapters 2 and 3, and the connections between their languages, can we also find explicit *reductions between the logics* that we have proposed in these two separate investigations? We think one can, but we leave this matter to future investigation. Instead, we conclude this chapter with two further topics. One is the question how the entanglement of preference and belief, which was so central in Chapter 3, should play a role in the modal languages of Chapter 2. The other is the issue how one could merge all ideas from Chapters 2 and 3 into one logical system that might have the power to address preference in much greater generality.

4.7 Preference meets belief

We have seen just now in Section 4.6 that preference is not just a matter of pure 'betterness'. In addition, it involved epistemic operators K of knowledge in

Chapter 2 and doxastic operators B of belief in Chapter 3. Understanding this entanglement of preference with knowledge and belief is of importance, especially when we study how agents make choices under uncertainty. This section compares the perspectives of the previous two chapters in this regard.

Preference is usually defined explicitly in terms of beliefs when we only have incomplete information. Here is a brief review of how this worked with the main notions of Chapter 3:

(i) We distinguished preference over *objects* and preference over *propositions*. First, preference over objects was defined by beliefs on whether objects have certain properties. We then applied the same method to preference over propositions.

(ii) For ease of reading, we repeat some basic definitions here:

Given a priority sequence $\mathscr{P}$ of the following form

$$P_1(x) \gg P_2(x) \gg \cdots \gg P_n(x) \quad (n \in \mathbb{N}),$$

where each of the $P_m(x)$ is a formula from the language, all with one common variable x, we define preference over objects as follows:

$$Pref(d_1, d_2) ::= \exists P' \in \mathscr{P} (\forall P < P' (BPd_1 \leftrightarrow BPd_2) \wedge (BP'd_1 \wedge \neg BP'd_2)). \qquad (Pref\text{-}obje)$$

(iii) Preference over propositions was defined similarly. Given a propositional priority sequence of length n of the following form

$$\varphi_1(x) \gg \varphi_2(x) \gg \cdots \gg \varphi_n(x) \quad (n \in \mathbb{N}),$$

where each $\varphi_m(x)$ is a propositional formula with an additional propositional variable, we define preference over propositions ψ and θ as follows:

$$Pref(\psi, \theta) \quad \text{iff} \quad \text{for some } i \ (B(\varphi_1(\psi) \leftrightarrow B(\varphi_1(\theta)) \wedge \cdots \wedge (B(\varphi_{i-1}(\psi)) \leftrightarrow B(\varphi_{i-1}(\theta))) \wedge (B(\varphi_i(\psi) \wedge \neg B(\varphi_i(\theta))). \qquad (Pref\text{-}prop)$$

(iv) We took the line that preference is a state of mind (i.e. it is subjective, and subject to introspection) and therefore, one prefers one alternative over another if and only if one believes one does. So typically, $Pref(d_1, d_2) \leftrightarrow BPref(d_1, d_2)$ was an axiom of our preference logic, which therefore includes positive introspection.

In contrast, beliefs are not a part of the language considered in Chapter 2, which only has the universal modality U, knowledge operator K and the betterness modality $[bett]$. This vocabulary allows us to express many notions of preference over propositions, depending on different combinations of quantifiers ('liftings'). Since the betterness operator is based on an objective relation in the model, the related preference modality will in general not be subjective. However, we can prefix it with epistemic operators, and achieve attitude-dependence after all. Thus, we discussed the connection between preference and knowledge, asking, e.g. whether epistemized preference validates positive introspection. Also, we have looked at the particular notion of 'regret', interpreted as 'agent a knows that p but she prefers that $\neg p$', which is only possible by combining objective and subjective aspects.

To make the preceding two approaches more comparable, we will first add beliefs to Chapter 2 and develop the system a bit further. After that, we will draw a comparison between Chapter 2 and Chapter 3, and their different notions of preference. First, we briefly review some important technical points from Chapter 2 that will return later.

(i) The language can express preferences over propositions, viewed as sets of possible worlds. We defined the central notion of $Pref^{\forall\exists}$ as follows:

$$Pref^{\forall\exists}(\varphi, \psi) ::= U(\psi \to \langle bett\rangle\varphi). \quad (Ubett)$$

(ii) A new reduction axiom for the operator $\langle bett\rangle$ was proposed to model changes in preference relations, with the running example of a 'suggestion' A ($\sharp A$). Given the reduction axioms below, this also determines the complete logic of changing preferences between propositions.

1. $\langle\sharp A\rangle\langle bett\rangle\varphi \leftrightarrow (\neg A \wedge \langle bett\rangle\langle\sharp A\rangle\varphi) \vee (\langle bett\rangle(A \wedge \langle\sharp A\rangle\varphi))$.
2. $\langle A!\rangle\langle K\rangle\varphi \leftrightarrow (A \wedge \langle K\rangle\langle A!\rangle\varphi)$.
3. $\langle A!\rangle E\varphi \leftrightarrow (A \wedge E(\langle A!\rangle\varphi \vee \langle\neg A!\rangle\varphi))$.

(iii) Since $Pref^{\forall\exists}$ is defined by the operators U and $\langle bett\rangle$, once we have reduction axioms for these two operators separately, we immediately get one for $Pref^{\forall\exists}$. The precise calculation went as follows:

$$\begin{aligned}
&\langle\sharp A\rangle Pref^{\forall\exists}(\varphi, \psi) \leftrightarrow \langle\sharp A\rangle U(\psi \to \langle bett\rangle\varphi)\\
&\leftrightarrow U(\langle\sharp A\rangle(\psi \to \langle bett\rangle\varphi))\\
&\leftrightarrow U(\langle\sharp A\rangle\psi \to \langle\sharp A\rangle\langle bett\rangle\varphi)\\
&\leftrightarrow U(\langle\sharp A\rangle\psi \to (\neg A \wedge \langle bett\rangle\langle\sharp A\rangle\varphi) \vee (\langle bett\rangle(A \wedge \langle\sharp A\rangle\varphi)))\\
&\leftrightarrow U(\langle\sharp A\rangle\psi \wedge \neg A \to \langle bett\rangle\langle\sharp A\rangle\varphi) \wedge U(\langle\sharp A\rangle\psi \wedge A \to \langle bett\rangle(\langle\sharp A\rangle\varphi \wedge A))\\
&\leftrightarrow Pref^{\forall\exists}(\langle\sharp A\rangle\varphi, \langle\sharp A\rangle\psi) \wedge Pref^{\forall\exists}((\langle\sharp A\rangle\varphi \wedge A), (\langle\sharp A\rangle\psi \wedge A)).
\end{aligned}$$

Preference and knowledge

As we saw in the above, the preference in Definitions (*Pref-obje*) and (*Pref-prop*) is based completely on beliefs, and hence Chapter 3 took a subjective stance. But intuitively, the preference defined in (*Ubett*) is more *objective*. In comparing this, for the moment, we ignore the difference between objects and possible worlds, but we will come back to it at the end of this section. Intuitively, (*Ubett*) says the following:

> for any ψ-world in the model, there exists a world which is at least as good as that world, where φ is true.

This can be pictured as follows:

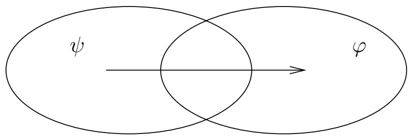

Figure 4.9: Preference defined by U and betterness relations

Essentially this is a comparison between ψ-worlds and φ-worlds in the model, with no subjective attitude involved yet. But even in the setting of Chapter 2, we can create connections between preference, knowledge, and beliefs - as we are going to show now.

Consider the following situation. Instead of picking *any* ψ-world in the model (as the universal modality U does), we now only look at those ψ-worlds that are *epistemically accessible* to agent i. This suggests the following intuition:

> For any ψ-world that is *epistemically accessible* to agent i in the model, there exists a world which is as good as that world, where φ is true.

This can be pictured in the following manner:

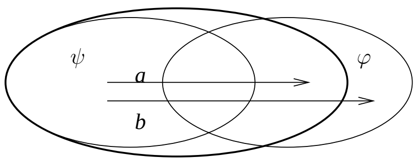

Figure 4.10: Preference defined by K and betterness relations

The part inside the black circle stands for the epistemically accessible worlds. The other two circles in the picture stand for the set of ψ-worlds, and the set of φ-worlds, respectively. So only some of the φ-worlds are epistemically accessible. The betterness relation has two possible cases: either it consists of a-arrows, which means that the better φ-world is itself in the accessible part of the model, or it also consists of b-arrows, which means that the better φ-world need *not* be in the accessible part of the model.

We write the above explanation in the formal language as:

$$Pref^{\forall\exists}(\varphi, \psi) ::= K_i(\psi \to \langle bett \rangle \varphi). \quad (Kbett)$$

Comparing the definitions ($Kbett$) and ($Ubett$), we have simply replaced U with K_i. In fact, looking back at Chapter 2, this is a straightforward step to take, since we had a knowledge operator in the language. The models proposed in Chapter 2 can be used directly for this purpose, and likewise, their complete logics.

Next, regarding dynamics, preference change is now triggered by changes in both epistemic and betterness relations. And we can obtain the right reduction axiom for epistemic preference in the same manner as for the universal modality, by a calculation from the reduction axioms for K_i and $\langle bett \rangle$. It is easy to spell out the details.

Introducing beliefs

In many situations, however, we do not have solid knowledge, but only beliefs. Still we want to compare situations in terms of betterness. In other words, we would like to say things like this:

> for any ψ-world that is most plausible to agent i in the model, there exists a world which is as good as that world, where φ is true.

This requires introducing beliefs – formally at first:

$$Pref^{\forall\exists}(\varphi, \psi) ::= B_i(\psi \to \langle bett \rangle \varphi). \quad (Bbett)$$

Figure 4.11 illustrates what we have in mind now:

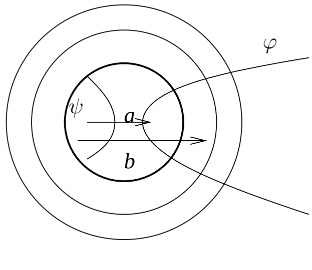

Figure 4.11: Preference defined by B and betterness relations

In the picture, the worlds lie ordered according to their plausibility, as in Lewis' spheres for conditional logic. The part inside the black circle depicts the most plausible worlds. We consider the ψ-worlds in this area, and again distinguish two sorts of preference relations: relations of 'type a' stay inside the most plausible region, relations of 'type b' go outwards to the less plausible, or even implausible region of the model.

To interpret beliefs more formally, we define these models as follows:

4.7.1. Definition. A *belief preference model* is a tuple $\mathcal{M} = (W, \leq, \preceq, V)$, with W a set of possible worlds, $\leq$ a doxastic relation of 'at least as plausible as', and $\preceq$ our earlier relation of 'at least as good as', with V again a valuation for proposition letters.[17]

The truth conditions for the absolute belief operator B and more general conditional beliefs $B^{\psi}\varphi$ are defined as follows:

$\mathcal{M}, s \models B\varphi$ iff $\mathcal{M}, t \models \varphi$ for all worlds t which are minimal for the ordering $\lambda xy. \leq_s xy$.

$\mathcal{M}, s \models B^{\psi}\varphi$ iff $\mathcal{M}, t \models \varphi$ for all worlds t which are minimal for $\lambda xy. \leq_s xy$ in the set $\{u \mid \mathcal{M}, u \models \psi\}$.

Here the truth condition for the unary operator B is essentially the same as in $KD45$-models, with the 'accessible worlds' of the latter system being the most plausible ones. But here we can compare less plausible worlds, too, and this is crucial to understanding conditional belief.

4.7.2. Remark. Given the notion of conditional belief, there is actually an alternative formulation for our formulation of belief-based preference. The above version (*Bbett*) looks at all normal or optimal worlds in the model, and then compares φ-worlds to ψ-worlds there in terms of betterness. The other option would be this: take the preference for ψ over φ itself as a *conditional belief*, using the following formula

$$B^{\psi}\langle bett\rangle\varphi \qquad (Bbett').$$

As is well-known, this is not equivalent to (*Bbett*), and it might be another candidate for belief-based preference. Personally, we think that preference should not involve the conditional scenario of 'having received the information that ψ'. However, both definitions can be treated in the logic we have proposed, and both are amenable to the style of dynamic analysis that we will consider next.

Again, we can now model changes in two ways, through changes in the plausibility relation and through changes in the betterness relation of the model. The

[17] [BS06b] and [BS08] also uses the 'as plausible as' relation to interpret the notion of *safe beliefs* which hold in all worlds that are at least as plausible as the current one. This notion is like our universal betterness modality, but then of course for belief rather than preference.

DEL methodology still applies here, since the two cases are formally very similar. Accordingly, [Ben07a] proposed valid reduction axioms for two sorts of belief change, so-called *radical* revision and *conservative* revision, where the former involves our earlier relation transformer $\Uparrow A$. These technical results can be used here directly. For instance, the reduction axiom for beliefs after radical revision is:

$$[\Uparrow A]B\varphi \leftrightarrow (EA \wedge B([\Uparrow A]\varphi|A)) \vee B[\Uparrow A]\varphi.$$

For a complete system, we also need a reduction axiom for conditional belief. The details are in [Ben07a].

In the same line, once we have reduction axioms for the belief and betterness operators, we can calculate what should be the reduction axiom for defined preference over propositions, e.g., the $Pref^{\forall\exists}$ notion in (*Bbett*). We then obtain a complete logic for dynamical change of belief-based preference.

Definitions (*Kbett*) and (*Bbett*) share a common feature: an arrow to a better φ-world can lead *outside of the accessible or most plausible part* of the model, witness the earlier arrows of type b. The intuition behind this phenomenon is clear, and reasonable in many cases. It may well be that there exists better worlds, which the agent does not view as epistemically possible, or most plausible. But *if* we want to have the two base relations entangled more intimately, we might want to just look at better alternatives inside the relevant epistemic or doxastic zone. Such considerations are found in the study of normative reasoning in [LTW03] where a normality relation and a preference relation live in one model. Likewise, [BRG07] discuss the 'normality sense' of ceteris paribus preference, restricting preference relations to just the normal worlds for the agents. In what follows, we will explore this more intimate interaction of the two base relations a bit more.

Merging relative plausibility and betterness

Now we require that the better worlds relevant to preference stay inside the most plausible part of the model. Intuitively, this means that we are 'informational realists' in our desires. To express this, we need a merge of the two relations, viz. their intersection. Here is how:

4.7.3. Definition. A *merged preference model* is a tuple $\mathcal{M} = (W, \leq, \preceq, \leq\cap\preceq, V)$, with W a set of possible worlds with doxastic and betterness relations, but also $\leq\cap\preceq$ as the intersection of the relations 'at least as plausible as' and 'at least as good as', with V again a valuation for proposition letters.

The original language had separate modal operators B and $[bett]$, but now we extend it with a new modality H. The formula $H\varphi$ is interpreted as 'it is hopeful that φ'. The truth condition for such formulas is as follows:

$\mathcal{M}, s \models H\varphi$ iff for all t with both $s \leq t$ and $s \preceq t$, it holds that $\mathcal{M}, t \models \varphi$.

With this new language over these new models, we can define one more natural notion of preference over propositions, which is actually much closer to Definition (*Pref-prop*) from Chapter 3:

$$Pref^{\forall\exists}(\varphi, \psi) ::= B(\psi \to \langle H \rangle \varphi). \quad (BH)$$

In words, this says that:

> For any most plausible ψ-world in the model, there exists a world which is *as good as* this world, and at the same time, *as plausible as* this world, where φ is true.

Obviously, we can now talk about preferences restricted to the most plausible part of the model. In terms of Figure 4.11, only arrows of 'type a' remain.

Actually, this same move would apply to Definition (*Kbett*) as well. Requiring that the better worlds stay inside the accessible worlds, we would have:

$$Pref^{\forall\exists}(\varphi, \psi) ::= K(\psi \to \langle \sim \cap \preceq \rangle \varphi) \quad (Kbett').$$

This means that we keep only the a-arrows in Figure 4.10.

Definition (BH) gives us a subjective notion of preference, as we consider only the most plausible part of the model. As we said, it is getting closer to Definition (*Pref-prop*). More precisely, Remark 3.7.10 showed that the $\forall\exists$ version of the preference defined by (*Pref-prop*) is equivalent to the following

$$B(\psi \to \langle bett \rangle \varphi).$$

Since then $[bett]\varphi ::= \underline{Pref}(\varphi, \top)$ and $\underline{Pref}$ is defined in terms of beliefs, the betterness relation automatically stays within the plausible part of the models. Hence, $\forall\exists$-(*Pref-prop*) is actually equivalent to Definition (BH).

Now let us quickly look at the expressive power of the modal language with a new operator H. Can the notion of preference in (BH) be defined in the original language with modal operators B and $[bett]$ only? In other words, can iterations of separate doxastic and betterness modalities achieve the same effect as intersection? As we know from general modal logic, this is very unlikely, since intersection modalities are not invariant under bisimulation (cf. [BRV01]). Indeed, the answer is negative:

4.7.4. FACT. $B(\psi \to \langle H \rangle \varphi)$ $(*)$ is not definable in the standard bimodal language with modal operators B and $[bett]$.

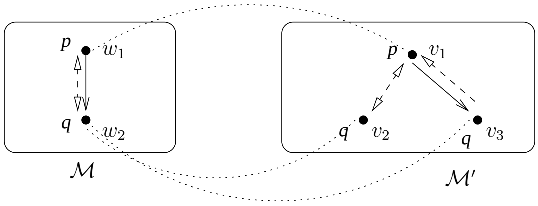

Figure 4.12: Bisimilar models

Proof. Suppose $(*)$ were definable. Then there would be a formula φ in the language without H such that $\varphi \leftrightarrow (*)$ holds in every model. Now consider the two models in Figure 4.12.

The betterness relation $\preceq$ is pictured by solid lines with arrows, and the plausibility relation $\leq$ by dashed lines with arrows. The evaluation of the proposition letters p and q can be read off from the picture. It is easy to see that these two models are bisimilar with respect to both betterness and relative plausibility, with the bisimulation indicated by the dotted lines.

Now, we have $\mathcal{M}, w_1 \models B(p \rightarrow \langle H \rangle q)$, since the p-world w_1 can see a world w_2 which is both better and plausible where q is true. Then we should get $\mathcal{M}, w_1 \models \varphi$, since $\varphi \leftrightarrow (*)$. Because $\mathcal{M}$ and $\mathcal{M}'$ are bisimilar, we would then have $\mathcal{M}', v_1 \models \varphi$. So we should also have $\mathcal{M}', v_1 \models B(p \rightarrow \langle H \rangle q)$. But instead, we have $\mathcal{M}', v_1 \not\models B(p \rightarrow \langle H \rangle q)$, because the p-world v_1 can see v_2 which is plausible but not better, and v_3 which is better but not plausible. So there is no world which is both better and plausible, while satisfying q. This is a contradiction. □

This argument shows that the new language indeed has richer expressive power. While this is good by itself, it does raise the issue whether our earlier methods still work.

In particular, we consider possible dynamic changes to the merged relation. We will only look at our three characteristic actions: radical revision $\Uparrow A$ that changes the plausibility relations, suggestion $\sharp A$ that changes the betterness relations, and the standard public announcement $A!$ that changes the domain of worlds. As it happens, the DEL-method of reduction axioms still applies:

4.7.5. Theorem. *The following equivalences are valid:*

1. $\langle \sharp A \rangle \langle H \rangle \varphi \leftrightarrow (A \wedge \langle H \rangle (A \wedge \langle \sharp A \rangle \varphi)) \vee (\neg A \wedge \langle H \rangle \langle \sharp A \rangle \varphi)$.
2. $\langle \Uparrow A \rangle \langle H \rangle \varphi \leftrightarrow (A \wedge \langle H \rangle (A \wedge \langle \Uparrow A \rangle \varphi)) \vee (\neg A \wedge \langle H \rangle (\neg A \wedge \langle \Uparrow A \rangle \varphi)) \vee (\neg A \wedge \langle bett \rangle (A \wedge \langle \Uparrow A \rangle \varphi))$.
3. $\langle A! \rangle \langle H \rangle \varphi \leftrightarrow A \wedge \langle H \rangle \langle A! \rangle \varphi$.

Proof. We only explain the most interesting Axiom 2 as an illustration. Assume that $\langle \Uparrow A \rangle \langle H \rangle \varphi$. Recall that radical revision $\langle \Uparrow A \rangle$ only changes the plausibility relation, leaving the preference relation intact. The new plausibility relation was written as follows:

$$(?A; R; ?A) \cup (?\neg A; R; ?\neg A) \cup (?\neg A; \top; ?A)$$

Seen from the initial model, we can therefore distinguish three cases, and these are just the three disjuncts on the right-hand side. Note that for the last one we only need to insert the old preference relation $\langle bett \rangle$, since the plausibility relation $(?\neg A; \top; ?A)$ is new. □

As for axiomatizing the complete logic of the new modality H, there are various techniques. For instance, one can introduce 'nominals' from hybrid logic, as is done in axiomatizing modal logics with intersection modalities (cf. [Bla93], [Kat07]). The preceding observation then shows that the complete dynamic logic can be obtained by adding just these three reduction axioms. Thus, the DEL-methodology also works in this extended setting.

For the reader's convenience, we tabulate the many different definitions of preference that we have seen so far, and the implicational relationships between them in the following:

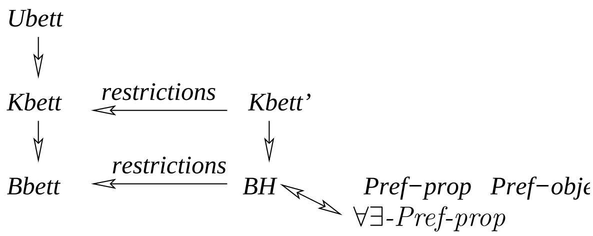

Figure 4.13: Different definitions of preference

We started with the pure betterness Definition ($Ubett$) from Chapter 2. On the same pattern, we 'epistemized' the universal modality, and proposed Definitions ($Kbett$) and ($Bbett$). The relation between these is as follows: ($Ubett$) implies ($Kbett$), ($Kbett$) implies ($Bbett$), but no reverse implication holds. Then, we set restrictions to make the betterness comparisons stay inside the accessible or most plausible region of the model, obtaining new Definitions ($Kbett'$) and (BH). Their relations to the previous notions are as indicated. Finally, as for connections with Chapter 3, defining preference as in Definitions ($Pref$-$prop$) makes them comparable to Definition (BH). This is true technically, but also intuitively, since all these notions are subjective. In particular, we have shown that Definition (BH) is equivalent to a ∀∃ lifted version of Definition ($Pref$-$prop$).

4.8 Combining preference over objects and preference over possible worlds

The preceding section concludes our comparisons between the two levels for defining preference in Chapters 2 and 3. But there is one more viewpoint which we want to mention briefly. Instead of taking these as different approaches that need to be contrasted and compared, one could equally well say that they bring out equally natural aspects of preference which need to be *merged*. And indeed, it is easy to do so. In this final section, we merely show how this may be done consistently, and what further questions would arise.

First, consider the following difference in 'spirit'. Priority sequences naturally fit with preferences between *objects*, while the propositional modal languages used so fit better with preferences between *worlds*. As we have said, technically, this does not make much difference. Theorem 3.7.12 even shows an equivalence between preference over objects and preference over propositional variables. Nevertheless, in real life we often compare objects and situations at the same time, neither exist exclusively. And these express different things. Consider the following example:

4.8.1. EXAMPLE. Alice prefers living in Amsterdam over living in Beijing. In Amsterdam there are two houses d_1 and d_2 for her to choose from, and she prefers d_1 over d_2. In Beijing, she prefers house d_3 over d_4.

A more abstract example of the distinction would be this. In some worlds, we may have many preferences between objects of desire, while in others, we have none at all. In some philosophies and religions, we would prefer the worlds where we have few object preferences (or better: none) to those where we have many. Moreover, things get even more complex when we bring in relative plausibility and beliefs. E.g. Alice may prefer living in Amsterdam, while still thinking it more plausible that she will end up domiciled in Beijing - or vice versa.

To talk about such examples, we need to combine preference over objects and preference over possible worlds in one semantic structure. The proper vehicle for this would join our earlier languages into one *doxastic preferential predicate language* defined as follows:

4.8.2. DEFINITION. Object-denoting terms t are variables $x_1, x_2, \ldots$ and constants $d_1, d_2, \ldots$, $P_1, P_2, \ldots$ are predicates over objects. The *language* is defined in the following syntax format:

$$\varphi ::= Pt_1, \ldots, t_n \mid \neg\varphi \mid \varphi \wedge \psi \mid \underline{Pref}(t_i, t_j) \mid B\varphi \mid [pref]\varphi$$

This is only a small part of a complete doxastic preferential predicate logic, but it is already adequate for many purposes. Of course, one can extend this language with quantifiers $\exists x\varphi$ in a straightforward manner.

Semantic models appropriate to this language may be defined as follows:

4.8.3. Definition. A *preferential doxastic predicate model* is a tuple $\mathcal{M} = (W, \leq, \preceq, \{D_w \mid w \in W\}, \{\preceq_w \mid w \in W\}, V)$, with S a set of possible worlds, $\leq$ and $\preceq$ a plausibility relation and a betterness relation over these. Next, D_w is the domain of objects for each possible world $w \in W$, with $\preceq_w$ a distinguished relation 'at least as good as' over objects in these domains. Finally, V is a valuation or interpretation function for the constants and predicate atoms of the language.

This means that we have preferences between worlds, and inside the possible worlds, we have preference over objects. Moreover, worlds are also ordered according to their plausibility. In this way, preference lives at different levels of the semantic models. The operators of the language can now be interpreted as usual. In particular, if we were to add quantifiers, these would range over the local domains D_w. Notoriously, there are difficult issues in the semantics and proof theory of modal predicate logic (cf. [FM98], [HC96], and [BG07]), but even so, a framework like this seems needed to discuss more subtle points of preference.

For instance, in a language like this we can now state assertions like:

(a) agents believe that objects with property P are always better than objects with property Q,

(b) agents prefer situations where object d is not preferred to object e,

(c) agents prefer situations where they do not know if P to situations where they do know if P.

This would seem to be the expressive power needed to do justice to discussions in the philosophical literature like, [Han90b], [Åqv94] and [Han01a]. Also, this setting gets closer to complex scenarios like the 'deontics of being informed' investigated in [PPC06].

But we can go even further. One might even compare objects across possible worlds, as in Russell's famous sentence "I thought your yaught was longer than it is". Alice might prefer her favorite house in Amsterdam when she lives there to her favorite house in Beijing when she does not live there: 'the grass is always greener on the other side'. And finally, we could also make the priorities of Chapter 3 into explicit elements of the semantics, letting the language speak about whether or not agents believe the priorities which determine their preferences over objects.

Clearly, this richer setting also raises issues of how to do dynamics. For instance, in the Example 4.8.1, both preference over objects and preference over possible worlds may change when new information comes in. We believe that the dynamic logics in our previous chapters can be generalized to deal with this, but there is hardly any work in this direction. For a recent DEL-style approach to modal predicate logic, see [Koo07] on the dynamic semantics of changing assignments as well as worlds.[18]

Our conclusion is this. Merging the approaches in Chapters 2 and 3 is quite feasible, and comparisons at a semantic and modal propositional level reveal many analogies and compatible features. But doing this in full generality would require a modal predicate-logical framework, which seems feasible, but beyond the horizon of this study.

[18]Also relevant are earlier analysis in *PDL* terms by [EC92] and [Ben96].