
Deontic Logic and Changing Preferences

JOHAN VAN BENTHEM AND FENRONG LIU

ABSTRACT. This chapter generalizes the dynamic epistemic analysis of information flow to deontic logic and normative reasoning, an area where agents' evaluation of worlds or outcomes is essential. The relevant dynamics is taken from current logics of preference structure and preference change. These logics are presented in two versions, one with a primitive 'betterness' relation between worlds, and one with criteria for preference given by explicit 'priority graphs'. The entanglement of preference and informational attitudes like knowledge and belief is then studied in the same style. With this framework in place, we look at deontic issues again. In particular, we can model changing obligations and norms. Finally, we discuss some further issues that arise in our setting, including the role of groups, longer-term procedures, and the need for more finely-grained syntactic representations of information, evaluation, and deontic deliberation.

1 Logical dynamics in general

Current dynamic logics can deal with many aspects of information-driven agency, including acts of inference, observation, communication, and other forms of interaction. But agency equally well involves evaluation, 'colouring' our view of the world, and driving our desires, decisions, and actions in it. A purely informational agent may be rational in the sense of clever reasoning, but a *reasonable* agent is one whose actions are in harmony with what she wants. The two dimensions are intimately related: we are driven by what we believe what we want, we seek information to further goals that we desire.

Now agents must constantly cope with changes in information, either because they learn more, or because the world has changed. But equally well, they constantly undergo changes in evaluation, sometimes by intrinsic changes of heart, but most often through events with normative impact. In line with the preceding, the two forms of dynamics are often entangled: for instance, learning more can change my evaluation of a situation.

In this chapter, we will discuss how current dynamic epistemic logics can also work for preference structure and preference change, two basic aspects of evaluation - and we show how the resulting framework then throws new light on several varieties of deontic reasoning. Throughout, we will keep

a narrow focus on dynamic epistemic methods: but we will point at some other existing approaches toward the end of the paper.

2 Dynamic epistemic logic of knowledge and belief

Before analyzing evaluation dynamics, we first develop the basic methodology of this chapter for the purely informational case.

2.1 Epistemic logic and semantic information

For a start, we briefly review the basics of epistemic logic, referring to the standard epistemic and modal literature for details (cf. [Fagin *et al.*, 1995] and [Blackburn *et al.*, 2001]).

DEFINITION 1 *Let a set of propositional variables Φ be given, and a set of agents A . The epistemic language is defined by the syntax rule*

$$\varphi := \top \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_a\varphi \quad \text{where } p \in \Phi, a \in A.$$

We will focus on single agents in this chapter, leaving epistemically important group notions like ‘common knowledge’ to our discussion at the end. Semantic models for this language encode agents’ ‘information ranges’ in the form of equivalence classes of binary uncertainty relations:

DEFINITION 2 *An epistemic model is a tuple $\mathfrak{M} = (W, \{\sim_a\}_{a \in A}, V)$, where W is a set of epistemically possible states, \sim_a is an equivalence relation on W , and V is a valuation function from Φ to subsets of W .*

DEFINITION 3 *For epistemic models $\mathfrak{M} = (W, \{\sim_a \mid a \in A\}, V)$ with worlds $s \in S$, we define $\mathfrak{M}, s \models \varphi$ (epistemic formula φ is true in \mathfrak{M} at s) by induction on φ :*

1. $\mathfrak{M}, s \models \top$ iff always.
2. $\mathfrak{M}, s \models p$ iff $s \in V(p)$.
3. $\mathfrak{M}, s \models \neg\varphi$ iff not $\mathfrak{M}, s \models \varphi$.
4. $\mathfrak{M}, s \models \varphi \wedge \psi$ iff $\mathfrak{M}, s \models \varphi$ and $\mathfrak{M}, s \models \psi$.
5. $\mathfrak{M}, s \models K_a\varphi$ iff for all $t : \text{if } s \sim_a t$, then $\mathfrak{M}, t \models \varphi$.

In this chapter, we interpret the knowledge operator over an equivalence relation, making the resulting epistemic logic modal **S5**. Just for concreteness, we state this basic system here:

THEOREM 4 *Basic epistemic logic is axiomatized completely by the following axioms and inference rules:*

1. *Tautologies of propositional logic*
2. $K_a(\varphi \rightarrow \psi) \rightarrow (K_a\varphi \rightarrow K_a\psi)$
3. $K_a\varphi \rightarrow \varphi$
4. $K_a\varphi \rightarrow K_aK_a\varphi$
5. $\neg K_a\varphi \rightarrow K_a\neg K_a\varphi$
6. *From $\varphi, \varphi \rightarrow \psi$, infer ψ*
7. *If $\vdash \varphi$, then also $\vdash K_a\varphi$*

Axiom 2 expresses closure of knowledge under known consequences. This form of logical omniscience can be questioned, but we will stick with it for most of this chapter, while briefly discussing its possible deontic import toward the end. Axioms 4 and 5 express what has been called positive and negative introspection. The latter principles are optional, and plausible alternatives exist, employing models with more general accessibility relations. This choice point is orthogonal to the main concerns of this chapter, however, so we ignore it henceforth.

Static epistemic logic describes what agents know on the basis of their current semantic information. But information flows, and a richer story must also include dynamics of actions that produce and modify information. We now turn to the simplest case of this: reliable public announcements, or public observations, that shrink the current information range.

2.2 Dynamic logic of public announcement

Public announcement logic (*PAL*) is a combination of epistemic logic and one dynamic event, namely, *announcement* of new ‘hard information’ expressed in some proposition φ . The corresponding ‘update action’ $!\varphi$ transforms a current epistemic model \mathfrak{M}, s into its submodel $\mathfrak{M}|\varphi, s$ where all worlds that did not satisfy φ have been eliminated. This is the basic scenario of information gain in science and common sense, by shrinking one’s current epistemic range of uncertainty.

To describe this phenomenon, the *language* of *PAL* uses both formulas for propositions and action expressions for announcements:

$$\begin{aligned} \varphi &:= \top \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_a\varphi \mid [A]\varphi \\ A &:= !\varphi \end{aligned}$$

The new dynamic formula $[\varphi]\psi$ says that “after announcing the true proposition φ , formula ψ holds”:

$$\mathfrak{M}, s \models [\varphi]\psi \text{ iff if } \mathfrak{M}, s \models \varphi, \text{ then } \mathfrak{M}|_{\varphi}, s \models \psi.$$

This language can make typical assertions about knowledge change such as $[\varphi]K_a\psi$, which states what agent a will know after having received the hard information that φ . In particular, the knowledge change before and after an update can be captured by so-called *reduction axioms*. Here is the complete logical system for information flow under public announcement ([Gerbrandy, 1999], [Plaza, 1989]):

THEOREM 5 *PAL is axiomatized completely by the usual laws of epistemic logic plus the following reduction axioms:*

1. $[\varphi]q \leftrightarrow (\varphi \rightarrow q)$ for atomic facts q
2. $[\varphi]\neg\psi \leftrightarrow (\varphi \rightarrow \neg[\varphi]\psi)$
3. $[\varphi](\psi \wedge \chi) \leftrightarrow ([\varphi]\psi \wedge [\varphi]\chi)$
4. $[\varphi]K_a\psi \leftrightarrow (\varphi \rightarrow K_a[\varphi]\psi)$

These elegant axioms analyze reasoning about effects of getting hard information, through observation, communication, or other reliable means. In particular, the knowledge reduction axiom reduces knowledge after new information to conditional knowledge the agent had before, but in a subtle recursive manner. This is necessary since the process of information update can change truth values of epistemic assertions: e.g., I did not know if p , but after the event $!p$, I do.

There are two major noteworthy features to this approach. First, the analysis is *compositional*, breaking down the ‘postconditions’ behind the dynamic modalities $[\varphi]$ recursively. Next, the ‘reduction axioms’ take every formula of our dynamic-epistemic language eventually to an equivalent formula inside the static pure epistemic language. In terms of models, this means that the current static model already ‘pre-encodes’ all information about what might happen when agents communicate what they know. In terms of the logic, the reduction procedure means that *PAL* is *decidable*, since its static epistemic base logic is decidable.¹

¹In recent versions of *PAL*, possible sequences of announcements may be constrained by informational ‘protocols’. Then no straightforward reduction is possible, and the logic also encodes a genuine new kind of *procedural information* ([Benthem *et al.*, 2009c]). We return to protocols for deontic purposes at the end of this chapter.

In what follows, *PAL* will serve as a pilot example for many other complex cases, for example, changes in beliefs, preferences, and obligations. In each case, the ‘triggering events’ can be different: for instance, obligations can change through actions of commanding by a normative authority. The methodology of *PAL* remains in force, though in each case, with new twists.

2.3 From knowledge to belief and soft information

Knowledge rests on hard information, but most of the information that we have and act on is soft, giving rise to beliefs, that are not always true, and that can be revised when shown inadequate. One powerful model for soft information and belief reflects the intuition that we believe those things that hold in the *most plausible* worlds in our epistemic range. I believe that this train will take me home on time, even though I do not know that it will not suddenly fly away from the tracks. But the worlds where it stays on track are more plausible than those where it flies off, and among the latter, those where it arrives on time are more plausible than those where it does not. This intuition leads to the following semantics:

DEFINITION 6 *An epistemic-doxastic model is a tuple $\mathfrak{M} = (W, \{\sim_a\}_{a \in A}, \{\leq_a\}_{a \in A}, V)$ with $(W, \{\sim_a\}_{a \in A}, V)$ an epistemic model as before, while the \leq_a are binary comparative plausibility relations for agents between worlds.*

Intuitively, $\leq_{a,s} xy$ says that, in world s , agent a considers world x at least as plausible as y . Our semantics assumes that these orderings are the same for epistemically indistinguishable worlds: agents have fixed plausibility judgments among their epistemically accessible worlds.² We assume that this relation is a pre-order, leaving room for genuinely incomparable worlds – but much of what we say would also hold for the special case of *connected* pre-orders where any two worlds are comparable.³

As with epistemic models, our logical analysis works largely independently from such specific design decisions about the ordering, important though they may be in specific applications.

One can interpret many logical languages in this comparative order structure. In what follows, we work with modal formalisms for the usual reasons of perspicuous formulation and low complexity (cf. [Blackburn *et al.*, 2007]).

First of all, there is plain belief as truth in all most plausible worlds:

²More general versions use ternary world-dependent plausibility relations $\leq_{a,s}$, as in the semantics of conditional logic: cf. [Lewis, 1973], [Spohn, 1988], or [Shoham, 1988].

³Connected orders are equivalent to the ‘sphere models’ of conditional logic or belief revision theory, but in these areas, too, a generalization to pre-orders has been made: [Burgess, 1984], [Veltman, 1985].

$\mathfrak{M}, s \models B_a \varphi$ iff $\mathfrak{M}, t \models \varphi$ for all worlds $t \sim_a s$ that are maximal in the order $\leq_a xy$ in the \sim_a -equivalence class of s .

But the more general notion in our models is that of a *conditional belief*:

$\mathfrak{M}, s \models B_a^\psi \varphi$ iff $\mathfrak{M}, t \models \varphi$ for all worlds $t \sim_a s$ that are maximal for $\leq_a xy$ in the set $\{u \mid s \sim_a u \text{ and } \mathfrak{M}, u \models \psi\}$.⁴

Conditional beliefs generalize absolute beliefs, which are definable as $B_a^\top \varphi$. They *pre-encode* absolute beliefs that we will have *if* we learn certain things. Indeed, the above semantics for $B_a^\psi \varphi$ is formally similar to that for conditionals $\psi \Rightarrow \varphi$. This allows us to use known results:

THEOREM 7 *The logic of $B_a^\psi \varphi$ is axiomatized by standard propositional logic plus the following laws of conditional logic:*

- (a) $\varphi \Rightarrow \varphi$
- (b) $\varphi \Rightarrow \psi$ implies $\varphi \Rightarrow \varphi \vee \chi$
- (c) $\varphi \Rightarrow \psi, \varphi \Rightarrow \chi$ imply $\varphi \Rightarrow \psi \wedge \chi$
- (d) $\varphi \Rightarrow \psi, \chi \Rightarrow \psi$ imply $(\varphi \vee \chi) \Rightarrow \psi$
- (e) $\varphi \Rightarrow \psi, \varphi \Rightarrow \chi$ imply $(\varphi \wedge \psi) \Rightarrow \chi$.

Richer modal logics for belief also exist, such as **KD45**. The structural content of their additional axioms can be determined through frame correspondence techniques (see [Blackburn *et al.*, 2007], [Benthem, 2010b]).

Recently, another operator has emerged as important, in between knowledge and belief, viz. ‘safe belief’: intuitively, belief that we have which cannot be falsified by true information. Over epistemic plausibility models \mathfrak{M} , its force is intermediate:

DEFINITION 8 *The modality of safe belief $B_a^+ \varphi$ is interpreted as follows:*

$\mathfrak{M}, s \models B_a^+ \varphi$ iff for all worlds $t \sim_a s$: if $s \leq_a t$, then $\mathfrak{M}, t \models \varphi$.

Thus, φ is true in all accessible worlds that are at least as plausible as the current one. This includes the most plausible worlds, but it need not include all epistemically accessible worlds, since the latter may have some less plausible than the current one. The logic for safe belief is just **S4**, since it is in fact the simplest modality over the plausibility order.

2.4 Dynamic logics of belief change

Now we want to deal with explicit events that change beliefs.

⁴These intuitive ‘maximality’ formulations must be modified somewhat in models allowing infinite descent in the plausibility ordering. Cf. Footnote 8 below.

Hard information The first obvious kind of triggering event are the earlier public announcements of hard information. The complete logic of belief change under hard information can be developed in analogy with the earlier dynamic epistemic logic *PAL*, by means of world elimination. Its key reduction axiom for new beliefs uses conditional beliefs:

FACT 9 *The following formula is valid in our semantics:*

$$[!\varphi]B_a\psi \leftrightarrow (\varphi \rightarrow B_a[!\varphi]^\varphi\psi)$$

To keep the complete dynamic language in harmony, we also need a reduction axiom for conditional beliefs:

THEOREM 10 *The logic of conditional belief under public announcements is axiomatized completely by*

- (a) *any complete static logic for the model class chosen,*
- (b) *the PAL reduction axioms for atomic facts and Boolean operations,*
- (c) *the following reduction axiom for conditional beliefs:*

$$[!\varphi]B_a^X\psi \leftrightarrow (\varphi \rightarrow B_a[!\varphi]^\varphi \wedge [!\varphi]^X\psi)$$

This analysis also extends to safe belief, with this reduction axiom:

FACT 11 *The following PAL-style axiom holds for safe belief:*

$$[!\varphi]B_a^+\psi \leftrightarrow (\varphi \rightarrow B_a^+(\varphi \rightarrow [!\varphi]\psi)).$$

Using this equivalence, one can show that safe belief in factual propositions (not containing epistemic or doxastic operators) remains safe belief after updates with hard factual information.

Soft information But belief change also involves more interesting events. Soft information ‘upgrade’ does not eliminate worlds, but rather *changes the plausibility order*, ‘promoting’ or ‘demoting’ worlds according to their properties. Here is one way in which this may happen: ⁵

DEFINITION 12 (*Radical, or lexicographic upgrade*). *A radical upgrade $\uparrow\varphi$ changes the current order \leq between worlds in \mathfrak{M}, s to that of a new model $\mathfrak{M}|\varphi, s$ where all φ -worlds in \mathfrak{M}, s become better than all $\neg\varphi$ -worlds, while, within those two zones, the old plausibility ordering remains.*

⁵In this section, we drop epistemic accessibility, and focus on plausibility order only.

Note how no worlds are eliminated here.⁶

We take a matching upgrade modality in our dynamic language:

$$\mathfrak{M}, s \models [\uparrow\varphi]\psi \text{ iff } \mathfrak{M}|\varphi, s \models \psi.$$

THEOREM 13 ([Benthem, 2007]) *The dynamic logic of radical upgrade is axiomatized completely by*

- (a) *a complete axiom system for conditional belief on the static models,*
- (b) *the following reduction axioms:*

$$\begin{aligned} &[\uparrow\varphi]q \leftrightarrow q, \quad \text{for all atomic proposition letters } q \\ &[\uparrow\varphi]\neg\psi \leftrightarrow \neg[\uparrow\varphi]\psi \\ &[\uparrow\varphi](\psi \wedge \chi) \leftrightarrow ([\uparrow\varphi]\psi \wedge [\uparrow\varphi]\chi) \\ &[\uparrow\varphi]B^x\psi \leftrightarrow (E(\varphi \wedge [\uparrow\varphi]\chi) \wedge B([\uparrow\varphi]^{\varphi \wedge [\uparrow\varphi]\chi}\psi)) \vee (\neg E(\varphi \wedge [\uparrow\varphi]\chi) \wedge \\ &\quad B([\uparrow\varphi]^{\uparrow\varphi\chi}\psi)) \end{aligned}$$

Here the operator ‘ E ’ stands for the existential epistemic modality, and hence we need to add a simple reduction axiom for knowledge that we forego here. Likewise, it is easy to extend this analysis again to safe belief.

There are many further policies for changing plausibility order whose dynamic logic can be axiomatized in a similar manner. For instance, ‘conservative upgrade’ $\uparrow\varphi$ only puts the *most plausible* φ -worlds on top in the new model, leaving the rest in their old positions. For general results, see [Benthem, 2007], [Baltag and Smets, 2008] and [Benthem, 2010a]. [Rott, 2006] is an excellent source for variety of policies in belief revision theory.

2.5 General dynamic methodology and its applications

Extensions of the above approach are in circulation for further aspects of information-driven agency. One major new idea is that of *DEL*, dynamic epistemic logic of general observation where different agents can have different access to the current informational event, as in card games, or communication with security restrictions. This approach generalizes *PAL* as well as the above dynamic doxastic logics, using ‘event models’ and ‘product update’ (cf. [Baltag *et al.*, 1998], [Ditmarsch *et al.*, 2007], [Benthem, 2010a]). This approach has been applied to many further aspects of agency: changes in intentions ([Roy, 2008], [Icard *et al.*, 2010], trust [Holliday, 2009]), games ([Otterloo, 2005]), and the dynamics of inferences ([Velazquez-Quesada, 2009]) and questions ([Benthem and Minica, 2009]).

⁶One can also make *PAL*-update non-eliminative, cutting epistemic links rather than removing worlds. Cf. [Liu, 2004], [Snyder, 2004] for details. This is sometimes preferable even for epistemic purposes, when we still want to refer to former alternatives.

3 Deontic logic and preference logic

We now turn to our main interest, the realm of evaluation. In what follows, we focus on the dynamic methodology for preference change proposed in [Liu, 2008], [Girard, 2008] – taking its results to deontic logic later.

First we briefly discuss the static base logic, our usual point of departure. To unclutter notation, here and henceforth, we will suppress agent indices for modal operators and their relations.

3.1 Basics of deontic logic

Deontic logic is the logical study of normative concepts such as obligation, prohibition, permission and commitment. This area was initiated by von Wright in [Wright, 1951] who introduced the logic of absolute obligation. As a reaction to paradoxes with this notion, conditional obligation was then proposed in [Wright, 1956], [Wright, 1964] and [Fraassen, 1972]. Good reviews systematizing the area can be found in [Åqvist, 1987], [Åqvist, 1994].

One often thinks of deontic logic as the study of some accessibility relation from the actual world to the set of ‘ideal worlds’, but the more sophisticated view ([Hansson, 1969], [Fraassen, 1973] and [Jackson, 1985]) has models with a binary comparison relation that we may call ‘betterness’.⁷ Such more general comparisons make sense, for instance, when reasoning about ‘the lesser of two evils’, or about possibilities for ‘moral improvement’ of some given situation, rather than a straight jump to the ideal.

DEFINITION 14 *A deontic model is a tuple $\mathfrak{M} = (W, R, V)$ with W a non-empty set of possible worlds, R a binary relation on W of deontic betterness, and V a valuation for proposition letters.*

One can impose further constraints on the relation, with different deontic logics resulting – but as before, this is not our main concern in this paper.

These models support a wide variety of languages for describing semantic structure and inference patterns, from basic modal languages to richer systems like hybrid logics, first-order logic, or fixed-point logics. In what follows, as with information dynamics, we highlight modal formalisms.

The binary relation R now interprets $O\varphi$ (absolute obligation) as φ *being true in all best worlds*, much like belief with respect to plausibility. Likewise, we interpret conditional obligation $O^\psi\varphi$ like conditional belief: φ holds *in*

⁷Hansson argued that von Wright-type deontic logic can be naturally interpreted in terms of a preference relation ‘is at least as ideal as’ among possible worlds.

the best ψ -worlds.⁸ ⁹ For further information on deontic logic, we refer to [Åqvist, 1994] and various chapters in this Handbook.

From deontic logic to preference logic

The deontic relation of betterness shows intuitive analogies with the notion of *preference*. One can think of deontic evaluation as reflecting the preferences of a moral authority or law-giver. And also formally, the two viewpoints are close: ‘betterness for agent a ’ is close to ‘agent a ’s preference’. Since the evaluation dynamics that we are after has been studied mainly for preference logic (but see [Torre, 1997] and [Torre and Tan, 1999] for exceptions), we now make this turn. Indeed, this preference perspective was advocated forcefully in [Hansson, 1969].

Preference logic as a separate subject, has its own history. For new ideas and results in the area, we refer to [Hansson, 2001a] and [Grune-Yanoff and Hansson, 2009]. In the bulk of this chapter, we will now survey some recent developments in the study of preference statics and dynamics, returning to deontic logic toward the end. Of course, in preference logic, we tend to focus on describing our own preferences, rather than those of others, but in a multi-agent setting, such as that found in games or indeed many deontic scenarios, the distinction loses its edge.

4 Basic static preference logic

In the coming sections, we will discuss some basics of modal preference logic, starting with its statics, and continuing with dynamics of preference change. Our treatment follows ideas from [Boutilier, 1994] and [Halpern, 1997], and for the dynamics, [Benthem *et al.*, 2006] and [Benthem and Liu, 2007].

4.1 Modal preference logic

The first kind of models that we consider are like in decision theory or game theory: there is a set of alternatives (worlds, outcomes, objects) ordered by a primitive relation that we dub ‘betterness’ to distinguish it from richer notions of preference. While each agent has her own betterness order, in what follows, this agent-dependence is often not important, so we suppress indices wherever we can.

DEFINITION 15 (Modal betterness model.) *A modal betterness model is a tuple $\mathcal{M} = (W, \preceq, V)$ with W a set of possible worlds, \preceq a reflexive and*

⁸Returning to an earlier point, this may also be generalized as follows, avoiding trivialization in infinite models (cf. [Girard, 2008]): $\mathfrak{M}, s \models O^\psi \varphi$ iff $\forall t \sim s \exists u (t \preceq u$ and $\mathfrak{M}, u \models \psi$ and $\forall v \sim s ((u \preceq v$ and $\mathfrak{M}, v \models \psi) \Rightarrow \mathfrak{M}, v \models \varphi)$.

⁹There are also more abstract versions of this semantics, using binary deontic betterness relations R^ψ , and setting $\mathfrak{M}, s \models O^\psi \varphi$ iff for all t in W with $sR^\psi t$, $\mathfrak{M}, t \models \varphi$.

transitive relation (the ‘betterness’ pre-order) over these worlds, and V a valuation assigning truth values to proposition letters at worlds.¹⁰

The order relation in these models also induces a strict variant:

If $s \preceq t$ but not $t \preceq s$, then t is *strictly better* than s ,
written as $s \prec t$.

Here is a simple formalism that can say a lot about these structures:

DEFINITION 16 (Basic modal betterness language.) *Take any set of propositional variables Φ , with a variable p ranging over Φ . The modal betterness language is given by the following inductive syntax rule:*

$$\varphi := \top \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid \langle \leq \rangle \varphi \mid \langle < \rangle \varphi \mid E\varphi.$$

The intended reading of $\langle \leq \rangle \varphi$ is “ φ is true in a world that is at least as good as the current world”, while $\langle < \rangle \varphi$ says that “ φ is true in a world that is strictly better than the current world.” In general, these notions are agent-relative, but as we said, we will mostly suppress this aspect here, since it is orthogonal to our main points. In addition, the auxiliary existential modality $E\varphi$ says that “there is a world where φ is true”. Combinations of these modalities can capture a wide variety of binary preference statements comparing propositions, witness the cited literature.

As usual, we write $[\leq]\varphi$ for the universal modality $\neg\langle \leq \rangle\neg\varphi$, and we use $[<]$ and U for the duals of $\langle < \rangle \varphi$ and E , respectively.

DEFINITION 17 (Truth definition.) *Truth conditions for the atomic propositions and Boolean combinations are standard. Modalities work like this:*

- $\mathcal{M}, s \models \langle \leq \rangle \varphi$ iff for some t with $s \preceq t$, $\mathcal{M}, t \models \varphi$.
- $\mathcal{M}, s \models \langle < \rangle \varphi$ iff for some t with $s \prec t$, $\mathcal{M}, t \models \varphi$.
- $\mathcal{M}, s \models E\varphi$ iff for some world t in W , $\mathcal{M}, t \models \varphi$.

The defined modalities become the obvious universal versions.

THEOREM 18 *The modal betterness logic is completely axiomatized by the following principles, where $i \in \{\leq, U\}$:*

1. *propositional tautologies*

¹⁰As we said before, we use pre-orders since we want the generality of possibly non-total preferences. Still, total orders, the norm in areas like game theory, provide an interesting specialization for the results in this chapter – but we will only mention it in passing.

2. $[i](\varphi_1 \rightarrow \varphi_2) \rightarrow ([i]\varphi_1 \rightarrow [i]\varphi_2)$
3. $[i]\varphi \rightarrow \varphi$
4. $[i]\varphi \rightarrow [i][i]\varphi$
5. $\neg U\varphi \rightarrow U\neg U\varphi$
6. $U\varphi \rightarrow [\preceq]\varphi$
7. $\langle i \rangle \varphi \leftrightarrow \neg [i]\neg \varphi$
8. $[<](\varphi_1 \rightarrow \varphi_2) \rightarrow ([<]\varphi_1 \rightarrow [<]\varphi_2)$
9. $\langle < \rangle \varphi \rightarrow \langle \preceq \rangle \varphi$
10. $\langle \preceq \rangle \langle < \rangle \varphi \rightarrow \langle < \rangle \varphi$

4.2 Some special logical features of preference

Next we discuss some special logical features of preference that go beyond standard modal logic of pre-orders, and that will eventually turn out of interest to deontics as well.

Lifting to generic preferences

While betterness relates specific objects or worlds, preference is often used generically for comparing different sorts of things. In particular, starting with Von Wright ([Wright, 1963]), logicians have studied preferences $P(\varphi, \psi)$ between propositions (properties of worlds, or objects).

There is not one such notion, but many, that can be defined by a ‘lift’ of the betterness order among worlds to sets of worlds. For instance, compare your next moves in a game, identified with the set of outcomes that they lead to. Which move is ‘better’ depends on the criterion chosen: maybe we want the one with the highest possible outcome, or the one with the highest minimally guaranteed outcome, etcetera. Such options are reflected in various quantifier combinations for the lifting.

In particular, von Wright had a $\forall\forall$ -type preference between sets P, Q :

$$\forall x \in P \forall y \in Q: x \preceq y.$$

Our own running example will be the a $\forall\exists$ -type

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

Intuitively, this says that for any P -world, there is a Q -world which is at least as good as that ψ -world.¹¹

¹¹In the earlier game setting, this stipulation would say that the most preferred moves have the highest possible outcomes.

$\forall\exists$ preference can be defined in the above modal preference language, using a universal modality over all worlds:

$$P^{\forall\exists}(\varphi, \psi) := U(\psi \rightarrow \langle \leq \rangle \varphi).$$

This satisfies the usual properties for preference, reflexivity and transitivity: for instance, $P\varphi\psi$ and $P\psi\chi$ imply $P\varphi\chi$.

Other quantifier combinations and modal definitions lead to other notions of generic preference. This proliferation may be a problem ('doing what is best' then depends on one's stipulation as to 'best'), but there is no consensus on an optimal choice in the literature. A logical approach should just make the options clear, and analyze them all.

Ceteris paribus features

Next, preference seldom comes in pure form: the comparison between alternatives is often entangled with other considerations. Again, games are an example. Usually, players do not compare moves was the sets of all their possible outcomes, but rather, they compare the *most plausible* outcomes of their moves. This is the 'normality sense' of the pervasive *ceteris paribus* character of preference: we do not compare all the φ and ψ -worlds, but only the 'normal ones' in some relevant sense. This theme will return in our discussion of doxastic entanglement in Section 8.

But there are also other natural senses of *ceteris paribus*. Von Wright ([Wright, 1963]) already noticed the 'equality sense' in which there is a hidden assumption of independence: we only make comparisons between worlds that give the same truth values to some specified set of atomic propositions (or even complex formulas). This logic is axiomatized and analyzed in detail in cf. [Bentham *et al.*, 2009b]. Important as it is, we will not pursue this theme of propositional (in-)dependence in this chapter.

Richer preference languages

As we already mentioned, modal languages are just one step on a ladder of formalisms. Especially in the case of preference, it has been claimed that richer languages are needed to faithfully render basic notions and reasoning. We refer to the literature for details: cf. [Jongh and Liu, 2009] on first-order preference languages, [Grandi and Endriss, 2009] on first-order languages of social choice, [Bentham *et al.*, 2006] on hybrid modal preference languages for defining backward induction solutions in games, and the new hybrid modal preference language of [Seligman, 2010] of 'desire' and 'freedom'.

5 World based dynamics of preference change

Now let us look at how given preferences can change. Intuitively, there are many acts and events that can have such effects. Maybe the purest

form is a radical *command* by some moral authority to do something: this makes the worlds where we act better than those where we do not (cf. [Yamada, 2006]), at least, if we ‘take’ the order as legitimate, and change our evaluation accordingly, overriding any preferences that we ourselves might have had. Technically, this dynamics will change a current betterness relation in a model. This can be studied entirely along the lines already sketched for information dynamics. There are also earlier treatments, of which we mention [Torre and Tan, 1999], based on [Veltman, 1996].

5.1 Betterness change

[Benthem and Liu, 2007] is a first systematic study of betterness change using methods from dynamic-epistemic logic. The running example in their approach is a ‘public suggestion’ that a proposition φ be the case. This relatively modest ordering change leaves the set of worlds the same, but it removes any preferences the agent might have for $\neg\varphi$ over φ among these. The main point is that ‘events with evaluative import’ can act as triggers that change some current betterness relation:

$\# \varphi$

These lead to the following model change:

DEFINITION 19 *Given any modal preference model (\mathfrak{M}, s) , the upgraded model $(\mathfrak{M}_{\# \varphi}, s)$ has the same domain, valuation, and actual world as (\mathfrak{M}, s) , but the new preference relations are now*

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

Given a preference model \mathfrak{M} , a new action modality is interpreted as:

$$(\mathcal{M}, s) \models [\# \varphi] \psi \quad \text{iff} \quad \mathcal{M}_{\# \varphi}, s \models \psi.$$

THEOREM 20 ([Benthem and Liu, 2007]) *The dynamic preference logic of suggestion is completely axiomatized by the following principles:*

1. $\langle \# \varphi \rangle p \leftrightarrow p$.
2. $\langle \# \varphi \rangle \neg \psi \leftrightarrow \neg \langle \# \varphi \rangle \psi$.
3. $\langle \# \varphi \rangle (\psi \wedge \chi) \leftrightarrow (\langle \# \varphi \rangle \psi \wedge \langle \# \varphi \rangle \chi)$.
4. $\langle \# \varphi \rangle \langle \leq \rangle \psi \leftrightarrow (\neg \varphi \wedge \langle \leq \rangle \langle \# \varphi \rangle \psi) \vee (\langle \leq \rangle (\varphi \wedge \langle \# \varphi \rangle \psi))$.
5. $\langle \# \varphi \rangle E \psi \leftrightarrow E \langle \# \varphi \rangle \psi$.

Similar results hold for dynamic logics with many other kinds of strong or weak command, some like the radical or conservative plausibility changes found with soft information and belief revision (cf. [Liu, 2008]).

This is an analysis of betterness change and direct modal statements about it. But it also applies to the earlier lifted generic preferences. As an illustration, consider the $\forall\exists$ -lift defined earlier:

FACT 21 *The following equivalence holds for generic $\forall\exists$ preference:*

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

We omit the simple calculation. Similar results may be obtained for other set liftings such as Von Wright's $\forall\forall$ -version.

5.2 Digression: general betterness change

There is a much more general theory of relation change that works for a wide class of triggering events that change evaluation. One way of achieving greater generality uses program notation from propositional dynamic logic (*PDL*). For instance, suggesting that φ ($\sharp\varphi$) is defined by the program:

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

where R is the given input relation, while the operations $? \varphi$ test whether the relevant proposition φ or related ones hold. In particular, the disjunct $(? \varphi; R; ? \varphi)$ means that we keep all old betterness links that run from φ -worlds to φ -worlds. This definition is equivalent in *PDL* to

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

Again this says, but now more compactly, that we keep all old R -links, except for those that ran from φ -worlds to $\neg\varphi$ -worlds.

Likewise, basic betterness changers from belief revision can be defined in this format. For instance, the earlier *radical upgrade*, incorporating the wish of some over-riding authority, is defined by

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

Here the constant \top denotes the universal relation that holds between any two worlds. This reflects the meaning of this transformation: all φ -worlds become better than all $\neg\varphi$ -worlds, whether or not they were better before, and within these two zones, the old ordering remains.

The earlier conservative upgrade can be dealt with similarly. Again, this makes sense as a command changing betterness order: the new best worlds will satisfy φ , but we now leave the agent more of her original betterness

order. Differences with radical commands will show up in judgments of ‘conditional betterness’. Such judgments occur in the literature on conditional obligation: see [Hansson, 1969].

Given any *PDL* definition of this sort, one can automatically write reduction axioms for the complete dynamic logic of its induced model change. This is proved in [Benthem and Liu, 2007], while an alternative based on product update and ‘event models’ occurs in [Baltag and Smets, 2008]. As an example, here is the computation for our ‘suggestions’:

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

This theoretical variety matches the wealth of evaluative actions that we encounter in daily life.

5.3 Further scenarios: changes in numerical values

While the main thrust of this chapter are qualitative approaches, there are also numerical approaches to preferences, employing utilities ([Rescher, 1966], [Trapp, 1985]) or more abstract ‘grades’ for worlds ([Spohn, 1988]). We briefly indicate how dynamic ideas can work in this setting, too.

The static base language extends our earlier modal one with graded modalities indicating the strength of preference (cf. [Aucher, 2003] on product update for numerical plausibility models). Our brief sketch follows the stream-lined version of [Liu, 2004]).

DEFINITION 22 (Language) *The preference evaluation language \mathcal{L} is given by the inductive syntax rule*

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

where $m \in \mathbb{Z}$ is an integer indicating a ‘grade’ of preference.

The propositional constants q_a^m say, intuitively, that ‘agent a assigns the current world a value of at most m ’.

DEFINITION 23 (Evaluation models) *An evaluation model is a modal preference model as before, expanded with an evaluation function v_a assigning to each world an element from $\{-\infty\} \cup \mathbb{Z} \cup \{\infty\}$ ¹²*

¹²In [Aucher, 2003], the range is natural numbers up to a maximal element *Max*.

Evaluation functions induce a total ordering, our general case of pre-orders would require more abstract ‘grades’ for worlds.

The semantics for our language is as usual, with one new clause:

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

Now we define dynamic structures of events that can change evaluation, extending the standard format of *DEL* product update ([Baltag *et al.*, 1998], [Ditmarsch *et al.*, 2007]):

DEFINITION 24 (Evaluation event model) *An evaluation event model is a tuple $\mathcal{E} = (E, \sim_a, v_a, PRE)$ where E is a non-empty set of events, \sim_a a binary epistemic accessibility relation on E , v_a an evaluation function assigning each action an element from \mathbb{Z} , and PRE is a ‘precondition function’ from E to the set of all epistemic propositions in our language.*

The evaluation functions v_a indicate degrees of preference for agents over events. This is a major extension of standard views: we do not just evaluate static states of affairs, but also actions or events.

DEFINITION 25 (Evaluation product update) *Let an evaluation model $\mathfrak{M} = (S, \sim_a, v_a, V)$ and an evaluation event model $\mathfrak{E} = (E, \sim_a, v_a, PRE)$ be given. The evaluation product update model is the structure $\mathfrak{M} \otimes \mathfrak{E} = (S \otimes E, \sim'_a, v'_a, V')$ with*

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

The resulting dynamic logic of numerical evaluation can be axiomatized like the systems we discussed earlier. More complex numerical evaluation or utility dynamics can also be dealt with: cf. [Liu, 2004], [Liu, 2009b].

6 Proposition-based preference logic

Primitive betterness relations reflect what are sometimes called ‘intrinsic preferences’. But very often, our preferences have underlying structure, and we compare according to criteria: our preferences are then reason-based, or ‘extrinsic’. In this section we develop the latter view, that has motivations in linguistics (Optimality Theory, cf. [Prince and Smolensky, 1993]), belief

revision (entrenchment, cf. [Rott, 2003]), and deontic logic (cf. [Fraassen, 1973], [Goble, 2000] and [Jackson, 1985]), as we shall see in Section 9.

Our analysis starts with linear orders of relevant properties that determine our evaluation of objects or worlds.

6.1 Priority based preference

[Jongh and Liu, 2009] start from a given primitive ordering among propositions (‘priorities’), and then deriving a preference among objects.

DEFINITION 26 *A priority sequence is a finite ordered sequence of formulas (‘priorities’) written as follows:*

$$C_1 \gg C_2 \cdots \gg C_n \quad (n \in \mathbb{N}),$$

where the C_m comes from a language describing objects (or worlds), with exactly one free variable x in each C_m .

DEFINITION 27 *Given a priority sequence of length n , and two objects x and y , $\text{Pref}(x, y)$ is defined in the usual lexicographic manner: at the first property C_i in the given sequence where x, y show a different truth value, $C_i(x)$ holds, but $C_i(y)$ fails.*

[Jongh and Liu, 2009] analyze the logic of this framework. [Benthem *et al.*, 2010] provide applications to deontic logic.

Still, this is only one of many ways of deriving a preference from a priority sequence. A good overview can be found in [Coste-Marquis *et al.*, 2004].

6.2 Pre-orders

In general, comparison order may be non-connected, and then the preceding idea needs a significant generalization. This was done, in a setting of ‘generalized social choice’ and belief merge, in [Andréka *et al.*, 2002], which we adapt slightly here (cf. [Liu, 2010] for a more extensive study). The following definitions contain a ‘free parameter’ for a language L that can be interpreted in the earlier modal betterness models \mathfrak{M} . For simplicity only, we will take this to be a simple propositional language of properties.

DEFINITION 28 (Priority graph.) *A priority graph $\mathcal{G} = \langle P, < \rangle$ is a strictly partially ordered set of propositions in the relevant language of properties L .*

Here is how one derives a betterness relation from a priority graph:

DEFINITION 29 (Betterness from a priority graph.) *Let $\mathcal{G} = \langle P, < \rangle$ be a priority graph, and \mathfrak{M} a model in which the language L defines properties of objects. The induced betterness relation $\preceq_{\mathcal{G}}$ is defined as follows:*

$$y \preceq_{\mathcal{G}} x := \forall P \in \mathcal{G} ((Py \rightarrow Px) \vee \exists P' < P (P'x \wedge \neg P'y)).$$

Here, in principle, $y \preceq_{\mathcal{G}} x$ requires that x has every property in the graph that y has. But there is a possibility of ‘compensation’: if y has P while x does not, this is admissible, provided there is some property P' with higher priority in the graph where x does better: x has P' while y lacks it.¹³

One can think of priority graphs of propositions as hierarchically ordered information sources, but equally well as complex hierarchies of laws, or of norm givers with relative authority.

6.3 Static logic and representation theorem

Here is a crucial technical property of this framework:

THEOREM 30 ([Liu, 2010]) *Let $\mathfrak{M} = (W, \preceq, V)$ be any modal model, without constraints on its relation. The following two statements are equivalent:*

- (a) *The relation $y \preceq x$ is a reflexive and transitive order,*
- (b) *There is a priority graph $\mathcal{G} = (P, <)$ such that, for all worlds $x, y \in W$, $y \preceq x$ iff $y \preceq_{\mathcal{G}} x$.*

This representation theorem says that the general logic of derived extrinsic betterness orderings is still just that of pre-orders. But it also tells us that any pre-order can be ‘rationalized’ as an extrinsic reason-based one without disturbing the model as it is.

6.4 Priority dynamics and graph algebra

Now, we have a new locus for preference change: the structure of the underlying propositions. For linear priority sequences, these are the obvious operations $[+C]$ of adding a new proposition C to the right, $[C^+]$ of adding C to the left, and various functions $[-]$ dropping first, last or intermediate elements of a priority sequence. [Jongh and Liu, 2009] give complete dynamic logics for these. Here is one typical valid principle:

$$[+C]Pref(x, y) \leftrightarrow Pref(x, y) \vee (Eq(x, y) \wedge C(x) \wedge \neg C(y))$$

This set of natural operations is even richer with priority graphs. However, in this setting an elegant mathematical alternative arises, in terms of two basic operations that combine arbitrary graphs:

- $\mathcal{G}_1; \mathcal{G}_2$, adding a graph to another in top position
- $\mathcal{G}_1 \parallel \mathcal{G}_2$, adding two graphs in parallel.

¹³It is easy to see that this stipulation subsumes the earlier priority sequences: linear priority graphs lead to lexicographic order.

One can think of this as ‘sequential’ versus ‘parallel’ composition, and the special case where one of the graphs consists of just one proposition models simple update actions. This calculus has been axiomatized completely in [Andréka *et al.*, 2002] by algebraic means, while [Girard, 2008] has a modal-style axiomatization. We display its principles, since they express the essential recursion underlying priority graph dynamics (here, the proposition letter n is a ‘nominal’ denoting one single world):

$$\begin{aligned} \langle \mathcal{G}_1 \parallel \mathcal{G}_2 \rangle^{\leq n} &\leftrightarrow \langle \mathcal{G}_1 \rangle^{\leq n} \wedge \langle \mathcal{G}_2 \rangle^{\leq n}. \\ \langle \mathcal{G}_1 \parallel \mathcal{G}_2 \rangle^{< n} &\leftrightarrow (\langle \mathcal{G}_1 \rangle^{< n} \wedge \langle \mathcal{G}_2 \rangle^{\leq n}) \vee (\langle \mathcal{G}_1 \rangle^{\leq n} \wedge \langle \mathcal{G}_2 \rangle^{< n}). \\ \langle \mathcal{G}_1; \mathcal{G}_2 \rangle^{\leq n} &\leftrightarrow (\langle \mathcal{G}_1 \rangle^{\leq n} \wedge \langle \mathcal{G}_2 \rangle^{\leq n}) \vee \langle \mathcal{G}_1 \rangle^{< n}. \\ \langle \mathcal{G}_1; \mathcal{G}_2 \rangle^{< n} &\leftrightarrow (\langle \mathcal{G}_1 \rangle^{\leq n} \wedge \langle \mathcal{G}_2 \rangle^{< n}) \vee \langle \mathcal{G}_1 \rangle^{< n}. \end{aligned}$$

These axioms reduce complex priority relations to simple ones, after which the whole language reduces to the modal logic of weak and strict atomic betterness orders. In particular, this modal graph logic is decidable.

7 A two-level view of preference

Now we have two ways of looking at preference: one through intrinsic betterness order on modal models, the other through priority structure inducing extrinsic betterness orders. One might see this as calling for a reduction from one level to another, but instead, combining the two perspectives seems attractive as providing a richer modeling tool for agency.

7.1 Harmony between world order and propositional order

In many cases, the two levels are in close harmony, allowing for easy switches from one to the other (cf. [Liu, 2008]):

DEFINITION 31 *Let $\alpha: (\mathcal{G}, A) \rightarrow \mathcal{G}'$, with $\mathcal{G}, \mathcal{G}'$ priority graphs, and A a new proposition. Let σ be a map from (\preceq, A) to \preceq' , where \preceq and \preceq' are betterness relations over worlds. We say that α induces σ , if always:*

$$\sigma(\preceq_{\mathcal{G}}, A) = \preceq_{\alpha(\mathcal{G}, A)}$$

Here are two results for our earlier key betterness transformers:

FACT 32 *Taking a suggestion A is the map induced by the priority graph update $\mathcal{G} \parallel A$. More precisely, the following diagram commutes:*

$$\begin{array}{ccc}
\langle \mathcal{G}, < \rangle & \xrightarrow{\parallel^A} & \langle (\mathcal{G} \parallel A), < \rangle \\
\downarrow & & \downarrow \\
\langle W, \preceq \rangle & \xrightarrow{\#^A} & \langle W, \#A(\preceq) \rangle
\end{array}$$

Next, consider a priority graph $(\mathcal{G}, <)$, with a new proposition A added on top. The dynamics at the two levels is now correlated as follows:

FACT 33 *Prefixing a new proposition A to a priority graph $(\mathcal{G}, <)$ induces the radical upgrade operation \uparrow^A on possible worlds models. More precisely, the following diagram commutes:*

$$\begin{array}{ccc}
\langle \mathcal{G}, < \rangle & \xrightarrow{A; \mathcal{G}} & \langle (A; \mathcal{G}), < \rangle \\
\downarrow & & \downarrow \\
\langle W, \preceq \rangle & \xrightarrow{\uparrow^A} & \langle W, \uparrow^A(\preceq) \rangle
\end{array}$$

Thus the two kinds of preference dynamics can dovetail very well.

7.2 Co-existence of two dynamics and language change

Even so, there are good reasons for working with both levels.

One reason is that not all natural operations on graphs have corresponding betterness transformers at all. An example from [Liu, 2010] is *deletion* of the topmost elements from graphs. This syntactic operation is not invariant for replacing graphs by others inducing the same betterness order, and hence it is a genuine extension of preference change.

But also conversely, there is no general match. Not all *PDL*-definable operations from Section 5.2 are graph-definable. One striking reason is that not all betterness transformers preserve the base properties of reflexivity and transitivity. Consider $?A; R$, that is: ‘keep the old relation only when A is true’. This change does not preserve reflexivity of R , as $\neg A$ -worlds have no relations any more. $?A; R$ amounts to a refusal to henceforth make betterness comparisons at worlds that lack property A . Though somewhat idiosyncratic, this seems a bona fide mind change for an agent.

All this argues for a more general policy of modeling both intrinsic and extrinsic preference for agents, with reasons for the latter encoded in priority graphs that are part of the model. Here one might think that intrinsic

betterness relations merely reflect an agent’s raw feelings or prejudices. But the contrast is relative, not absolute.¹⁴

In particular, intrinsic betterness can become extrinsic through a dynamics that has been largely outside the scope of dynamic-epistemic logic so far, viz. *language change*. One mechanism here is the proof of the earlier representation theorem (Theorem 30). It partitions the given betterness pre-order into clusters, and if these are given names as new ‘properties’, then it is easy to see that the resulting strict order of clusters is a priority graph inducing the given order. This may be thought of as rationalization, but alternatively, someone else might observe an agent’s preferences, and postulating reasons for them. This is closer to the notion of ‘revealed preference’ from the economics literature: cf. [Houser and Kurzban, 2002].

In general, reasons for given preferences may have to come from some other, richer language than the one that we started with: we are witnessing a dynamic act of language creation.

8 Entanglement of evaluation and information

The major agency systems of information and evaluation do not live in isolation: they interact all the time. A rational agent can process information well in the sense of proof or observation, but is also ‘reasonable’ in a broader sense. This *entanglement* of knowledge, belief, and preference shows in many specific settings. We will look at a few cases, and their impact on the dynamics of preference change. For a general discussion, we refer to the chapter by Rohit Parikh in this Handbook.¹⁵

8.1 Generic preference involving knowledge

Recall that in Section 4.2, we defined generic preference as follows:

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

This refers to possibilities in the whole model, including even those that an agent might know to be excluded. [Benthem and Liu, 2007] defend this scenario in terms of ‘regret’, but there is also a reasonable intuition that preference only runs among situations that are epistemically possible.

¹⁴A nice illustration is deontic logic. If I obey the command of a higher moral authority, I may acquire an extrinsic preference, whose reason is obeying a superior. But for that higher agent, that same preference may be intrinsic: “The king’s whim is my law”.

¹⁵Though we will mainly discuss how information dynamics encroaches upon preference and deontic notions, the opposite influence has also been studied. Information flow depends on *trust* and *authority*: which are clearly deontic notions. Moreover, following Wittgenstein, Brandom ([Brandom, 1994] has argued that language use can only be fully understood in terms of commitments that carry rights and obligations.

In a first step, this suggests the entangled notion that, for any ψ -world that is *epistemically accessible* to agent a in the model, there exists a world which is as good as that world, where φ is true. This can be written with an epistemic modality:

$$Pref^{\forall\exists}(\varphi, \psi) := K_a(\psi \rightarrow \langle \leq \rangle \varphi). \quad (K_{bett})$$

But this is not yet what we are after, since we want the ‘better world’ to be epistemically accessible itself. [Liu, 2009a] shows how this cannot be defined in a simple combined language of knowledge and betterness, and that instead, a language is needed with an *intersection modality* for epistemic accessibility and betterness. The latter can be axiomatized, and it also supports a dynamic logic of preference change as before.

Issues of entanglement become even more appealing with generic preference and belief, where the two relational styles of modeling were very similar. We discuss this in a bit more detail.

8.2 Generic preferences involving belief

Again, we might start with a mere combination formula

$$Pref^{\forall\exists}(\varphi, \psi) := B_a(\psi \rightarrow \langle \leq \rangle \varphi). \quad (B_{bett})$$

This says that, among the most plausible worlds for the agent, for any ψ -world, there exists a world which is as good φ is true.¹⁶

Again, this seems not quite right, since we want the better worlds relevant to preference to stay inside the most plausible part of the model.¹⁷ To express this, we again need a stronger merge of the two relations by intersection. The key clause for a corresponding new modality then reads:¹⁸

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

Again, the static and dynamic logic of this system yield to the general dynamic-epistemic machinery explained in earlier sections.

8.3 Other ways of combining preference and normality

Other entangled versions of plausibility and betterness abound in the literature. For instance, [Boutilier, 1994] has models $\mathcal{M} = (W, \leq_P, \leq_N, V)$ with W a set of possible worlds, V a valuation function and \leq_P, \leq_N two transitive connected relations $x \leq_P y$ (‘ y is as good as x ’) and $x \leq_N y$ (‘ y is as normal as x ’). He then defines an operator of *conditional ideal goal* (IG):

¹⁶Incidentally, one might also think here of using a *conditional belief* $B^\psi \langle \leq \rangle \varphi$, but this seems an intuitively implausible form of entanglement.

¹⁷Intuitively, this means that we are ‘informational realists’ in our desires.

¹⁸One might read this new notion as a sort of ‘wishful safe belief’.

$$\mathcal{M} \models IG^\psi \varphi \text{ iff } \text{Max}(\leq_P, \text{Max}(\leq_N, \text{Mod}(\psi))) \subseteq \text{Mod}(\varphi)$$

This says that the best of the most normal ψ worlds satisfy φ . Such entangled notions are still expressible in the modal systems of this chapter.

$$\text{FACT 34 } IG^\psi \varphi := (\psi \wedge \neg \langle B^< \rangle \psi) \wedge \neg \langle < \rangle (\psi \wedge \neg \langle B^< \rangle \psi) \rightarrow \varphi^{19}$$

Following up on this tradition in agency studies, [Lang *et al.*, 2003] define the following entangled notion of preference:

DEFINITION 35 $\mathcal{M} \models \text{Pref}^*(\varphi, \psi)$ iff for all $w' \in \text{Max}(\leq_N, \text{Mod}(\psi))$ there exists $w \in \text{Max}(\leq_N, \text{Mod}(\varphi))$ such that $w' <_P w$.

This reflects the earlier-mentioned *ceteris paribus* sense of preference, where one compares only the ‘normal’ worlds of the relevant kinds. In particular, this also makes sense in epistemic game theory, where ‘rationality’ means comparing moves by their most plausible consequences (according to the player’s beliefs) and then choosing the best.

The preceding notions are similar to our earlier one with an intersection modality, but not quite. They only compare the two most plausible parts for each proposition. We give no deeper technical analysis here, but do note that we are still within the bounds of this paper.

FACT 36 *Pref** is definable in our doxastic preference language.

8.4 Degrees of entanglement and probability

Preference, knowledge, and belief can be entangled in many ways. These range from simple combination of modalities to defining new merged modalities as in the above. And even more deeply entangled notions exist, such as *expected value* in probability theory, whose definition mixes beliefs and evaluation even more strongly. We believe that such notions can be dealt with, too, by combining the numerical dynamics of our Section 5.3 with the *DEL*-style probability dynamics of [Benthem *et al.*, 2009a] – but at present, there is no such system.

8.5 Preference change, update and belief revision

Even just juxtaposed, our treatment of the statics and dynamics of belief and preference shows many similarities. These will be so obvious to the reader that we need not spell them out here. It is an interesting test, then, if the earlier dynamic methods transfer well to entangled notions of preference. After all, intuitively, entangled preferences can change both because

¹⁹Here, $B^<$ is a modality of *strict belief* that we do not define here.

of two kinds of trigger: evaluative acts like suggestions or commands, and informative acts changing our beliefs.

Indeed, the earlier dynamic methods apply. As an illustration, we quote one result from [Liu, 2008]:

THEOREM 37 *The dynamic logic of intersective preference H is axiomatizable, with the following key reduction axioms:*

1. $\langle \#A \rangle \langle H \rangle \varphi \leftrightarrow (A \wedge \langle H \rangle (A \wedge \langle \#A \rangle \varphi)) \vee (\neg A \wedge \langle H \rangle \langle \#A \rangle \varphi)$.
2. $\langle \uparrow A \rangle \langle H \rangle \varphi \leftrightarrow (A \wedge \langle H \rangle (A \wedge \langle \uparrow A \rangle \varphi)) \vee (\neg A \wedge \langle H \rangle (\neg A \wedge \langle \uparrow A \rangle \varphi)) \vee (\neg A \wedge \langle \text{bett} \rangle (A \wedge \langle \uparrow A \rangle \varphi))$.
3. $\langle A! \rangle \langle H \rangle \varphi \leftrightarrow A \wedge \langle H \rangle \langle A! \rangle \varphi$.

This may not be all, though, since there may also be *entangled triggering events* that do not easily reduce to purely informational or purely evaluative actions.²⁰ Also, there may be switches between the two pure viewpoints, where the same scenario may be modeled either in terms of preference change, or as information change (see the Example “Buying a House” in [Liu, 2009b] and “Visit by the Queen” in [Lang and Torre, 2008]). We leave these matters to other chapters in this Handbook.

9 Deontic reasoning, changing norms and obligations

In this section, we return to deontic logic, and take a look at some major issues concerning obligations and norms from the standpoint of our dynamic systems for preference change.²¹

Maybe the most obvious task here would be charting the large variety of deontic actions in daily life that affect normative betterness orderings, from commands to promises and permissions. We will not do such a survey here, but the examples presented in what follows should convince the reader that a dynamic action perspective on deontic issues is natural, and that much can be done with the tools presented here.

9.1 Unary and dyadic obligation

Our static logics were motivated by the deontic semantics of ‘ideality’, in particular in its binary versions. In fact, deontic logic came first with this idea. To see this, consider the pioneering [Hansson, 1969], where dyadic

²⁰For a formal analogy, see the events with ‘parallel composition’ of information and issue dynamics in [Bentham and Minica, 2009].

²¹Our treatment largely follows the DEON paper [Bentham *et al.*, 2010]

obligations of the type ‘it is obligatory that φ under condition ψ ’ are interpreted by an ‘ideality relation’ and the notion of maximality:

$$(1) \quad \mathcal{M}, s \models \mathbf{O}(\varphi \mid \psi) \iff \text{Max}(\|\psi\|_{\mathcal{M}}) \subseteq \|\varphi\|_{\mathcal{M}}$$

where models are like the ones we have used. This account builds on the observation—stemming from ethical theory—that the deontic notions of obligation, permission and prohibition can be naturally made sense of in terms of an *ideality ordering* \preceq on possible worlds:

“[...] to assert that a certain line of conduct is [...] absolutely right or obligatory, is obviously to assert that more good or less evil will exist in the world, if it is adopted, than if anything else be done instead.” [Moore, 1903]²²

Depending on the properties of \preceq , different logics are obtained: [Hansson, 1969] starts with a \preceq which is only reflexive, moving then to total pre-orders. This is of course the same idea that has also emerged in conditional logic, belief revision, and the linguistic semantics of generic expressions.

One criticism of the above account has been that it made conditional obligation lack the property of antecedent strengthening (see [Tan and Torre, 1996]). This, however, makes perfect sense in our view, as it reflects precisely the non-monotonicity inherent in the dynamics of information change, where the ‘most ideal worlds’ can change.²³

In this light, our chapter has taken up an old idea in the semantics of deontic logic, and then added some recent themes: criterion-based priority structure, dynamics of evaluative acts and events, and extended logical languages making these explicit. This seems a natural continuation of deontic logic, while also linking it up with developments in other fields.

9.2 Dynamic perspectives on deontic paradoxes

The ‘dynamic turn’ has thrown fresh light on many issues in epistemology and other areas of philosophy. Deontic logic is no exception, as we shall see.

The scenario below is a classic of deontic logic and illustrates in a straightforward way the problem of ‘contrary-to-duty’ obligations (*CTDs*).

EXAMPLE 38 (Gentle murder; [Forrester, 1984]) “*Here is the problem: Let us suppose a legal system which forbids all kinds of murder, but which considers murdering violently to be a worse crime than murdering gently. [...] The system then captures*

²²Cited in [Fraassen, 1973, p. 6].

²³Variations of this modeling gave rise to various preference-based semantics of deontic logic: see [Torre, 1997] for an overview of this development.

its views about murder by means of a number of rules, including these two:

1. *It is obligatory under the law that Smith not murder Jones.*
2. *It is obligatory that, if Smith murders Jones, Smith murders Jones gently.” [Forrester, 1984, p. 194]*

We show how the priority view of Section 6.1, even just the format of linear priority sequences, can represent this scenario in a natural way. Recall that a linear priority sequence P_1, \dots, P_n essentially combines bipartitions $\{\mathcal{I}(p_i), \neg\mathcal{I}(p_i)\}$ of the domain of discourse S . The more we move towards the right-hand side of the sequence, the more atoms p_i are falsified. In a deontic reading this simply means that, the more we move towards the right-hand side of the sequence, the more violations hold.

Concretely, in the Gentle Murder scenario, the result is two classes of ideality: a class l_1 in which Smith does not murder Jones, i.e., $l_1 := \neg m$; and another l_2 in which either Smith does not murder Jones or he murders him gently, i.e., $l_2 := \neg m \vee (m \wedge g)$. In other words, we have a priority sequence \mathcal{B} with $l_2 \prec l_1$. Such a sequence orders the worlds via its induced relation $\preceq_{\mathcal{B}}^{IM}$ in three clusters. The most ideal states are those satisfying l_1 , worse but not worst states satisfy $V_1 := \neg l_1$ but at the same time l_2 , and, finally, the worst states satisfy $V_2 := \neg l_2$.

EXAMPLE 39 (Gentle murder (continued)) *Consider the priority sequence for Gentle murder given just now: $\mathcal{B} = (l_1, l_2)$. We can naturally restrict this sequence to an occurrence of the first violation by intersecting all formulas with V_1 . As a result, the first proposition is a contradiction that can never distinguish worlds. Thus, the best among the still available worlds are those with $Max^+(\mathcal{B}^{V_1}) = l_2 \wedge V_1$. A next interesting restriction in the Gentle Murder context is \mathcal{B}^{V_2} , which describes what the original priority sequence prescribes under the assumption that also the CTD obligation “kill gently” has been violated. In this case we end up in a set of all equally bad states.*

This brief sketch may suffice to show our approach provides a simple perspective on the robustness of norms and laws viewed as CTD structures: they can still function when transgressions have taken place.²⁴

²⁴Representing CTD structures as finite chains of properties is not a new idea: it already occurs informally in [Fraassen, 1973]. The first formal account is found in [Governatori and Rotolo, 2005], where an elegant Gentzen calculus is developed for handling formulae of the type $\varphi_1 \otimes \dots \otimes \varphi_n$ with \otimes a connective representing a sort of ‘sub-ideality’ relation. Unlike this proof-theoretic approach, our approach is geared towards semantics. It is an interesting open problem if the proof calculus of [Governatori and Rotolo, 2005] can be embedded in the modal logics of this chapter.

Other major puzzles in the deontic literature, such as the Chisholm Paradox, can be given an similar representation (cf. [Benthem *et al.*, 2010]).

9.3 Typology of change in a two-level approach

We have shown how two-level betterness/priority structure provides a natural medium for modeling deontic notions. Likewise, it yields a rich account of deontic changes. In Section 7, we developed a theory of both informational and evaluative changes, either directly on possible world order, or on priority structure underlying such orders. This also makes sense here.

As an illustration, we add a dynamic twist to the above running example:

EXAMPLE 40 (Gentle murder dynamified) *Let us start with the priority sequence consisting of $\mathcal{B} = (\neg m)$. This generated a total pre-order where all $\neg m$ states are above all m states: “It is obligatory under the law that Smith not murder Jones”. Suppose this is the given deontic state-of-affairs. Now, how can we refine this order so as to introduce the sub-ideal obligation to kill gently: “it is obligatory that, if Smith murders Jones, Smith murders Jones gently”? Or, in other words, how can we model the introduction into a legal code of a contrary-to-duty obligation?*

Intuitively, this change can happen in one of two ways:

1. *We refine the given betterness ordering ‘on the go’ by requesting a further bipartition of the violation states, putting the $m \wedge g$ -states above the $m \wedge \neg g$ -states. This can be seen as the successful execution of a command of the sort “if you murder, then murder gently!”.*
2. *We introduce a new law ‘from scratch’, where $m \rightarrow g$ is now explicitly formulated as a class of possibly sub-ideal states. This can be seen as the enactment of a new priority sequence $(\neg m, m \rightarrow g)$, one that we have encountered before.²⁵*

The example illustrates how a *CTD* sequence can be dynamically created either by uttering a sequence of commands stating what ought to be the case in a sub-ideal situation, or by enacting a whole new priority sequence.

Now Theorem 32 from Section 7 applies: the two kinds of change amount to the same thing! In other words, in this scenario, the same deontic change can be obtained both by ‘refining’ the order dictated by a given law, and by enacting a new ‘law’.

Of course, this is just a beginning. Our discussion of two-level dynamics in Section 7.2 also suggests that some changes in laws, such as *abrogation* (a counterpart to the earlier ‘deletion’) have no obvious counterpart at the pure worlds level. This brings us to a more general topic:

²⁵To see this, notice that $m \rightarrow g$ is equivalent to $\neg m \vee (m \wedge g)$.

9.4 Norm change

Dynamic aspects of norms—the so-called *norm change* problem—have recently gained attention from researchers in deontic logic, legal theory and multi-agent systems. Before concluding this section on deontic dynamics, we briefly put his chapter in perspective with some recent literature.

Existing approaches to norm change fall into two main groups. In syntactic approaches—inspired by legal practice—norm change is an operation performed directly on the explicit provisions in the ‘code’ of the normative system [Governatori and Rotolo, 2008a], [Governatori and Rotolo, 2008b], [Boella *et al.*, 2009].²⁶ In semantic approaches, norm change follows the update paradigm of this chapter (cf. also [Aucher *et al.*, 2009]). Our betterness dynamics belongs to this latter group. Thus, it can be naturally related to context dynamics in norm systems (cf. [Aucher *et al.*, 2009]).²⁷

Also, as observed already in Section 8 on entanglement (cf. [Lang *et al.*, 2003] for a deontic source), the dynamic logic connection allows for a unified treatment of *two* kinds of change that mix harmoniously in deontic reasoning: ‘information change’ given a fixed normative order, and ‘evaluation change’ modifying such an order. Their interplay reflects the entanglement of obligation, knowledge and belief studied in [Liu, 2008], and with more sophisticated scenarios, in [Pacuit *et al.*, 2006]. The latter includes issues that we have ignored here, such as the subtle difference between ‘knowing one’s duty’ versus ‘having duties to know’.

10 Further directions

In this final section, we briefly indicate a number of further directions where deontic logic meets with current trends in logical dynamics of agency.

10.1 Language, speech acts, and fine-structure of agency

Logical dynamics as developed in this chapter is triggered by informative and evaluative events. Very often, these events are *speech acts* of telling, asking, commanding, and so on. Indeed, natural language has a huge and sophisticated repertoire of speech acts with a deontic flavour (commanding, promising, allowing, and so on that invite further study in our framework, taking earlier studies in meta-ethics and Speech Act Theory ([Searle and Vanderbeken, 1985]) to the next level. Such studies will also need a more fine-grained account of the *agency* in dynamic triggers, that has been ignored throughout this chapter. For instance, things are said by someone to someone, and their uptake depends on relations of authority or trust. Like-

²⁶For more on these perspectives, we refer to the relevant chapters in this Handbook.

²⁷The bridge here is our earlier analysis: obligations defined via ideality and maximality are special kinds of classifications of an Andersonian-Kangerian type.

wise, promises, commands, or permissions are given by someone to someone, and their normative effect depends in subtle ways on who does and is what. [Yamada, 2010] is a first study of this fine-structure of normative action in terms of dynamic-epistemic logic.

10.2 Multi-agency and groups

Perhaps the most conspicuous turn in studies of information dynamics has been its emphasis on social scenarios ([Benthem, 2010a]. The paradigm for language use is communication between different agents²⁸, the paradigm for logic is argumentation between different parties, social behaviour is kept in place by mutual expectations, and so on. ‘Many-mind problems’ are the key phenomenon, just as in physics, ‘many-body problems’ are the heart of understanding Nature.

This trend also shows in epistemic and doxastic logics focusing on multi-agent interaction. Part of this can be represented by mere iteration of single-agent modalities, as in a ’s knowing that b does, or does not, know some fact. The languages and logics of this chapter are well-suited for that as they stand. But the next stage is the introduction of *groups as agents*, and corresponding notions such as ‘common knowledge’ or ‘distributed knowledge’ in groups, and likewise for beliefs (cf. [Fagin *et al.*, 1995], [Meyer and Hoek, 1995], etc.) This same social turn seems highly relevant to deontic logic. Indeed, from the start, deontic notions and morality seems all about *others*: my duties are usually toward other people, my norms come from outside sources: my boss, a lawgiver, etcetera.²⁹

In principle, our logics can deal with this like in the informational case, though one may have to think about even basic issues.³⁰ Just think of what it means to have a group-based ‘common obligation’: is this more like common belief, or more like some demand for *joint action* of the group? Other issues will arise with entanglement of informational and evaluative acts for groups: cf. [Hartog, 1985], [Kooi and Tamminga, 2006] and [Konkka, 2000] on morality as held together by ‘mutual expectations’. This links up with dynamic-epistemic logics of *trust* ([Holliday, 2009]).

As for social deontic dynamics, new types of event come up now. In the informational realm, a major topic of study is *belief merge*: how do groups form, and what social preferences and moral duties arise for plural subjects? Here dynamic epistemic logic and deontic logic meet with Social Choice

²⁸Think of asking some one a question, or telling someone something

²⁹The first community to realize this essential social aspect may have been the computer scientists, specially those involved in Multi-Agent Systems: cf. [Meyer, 1988], [Wooldridge, 2000], and [Rao and Georgeff, 1991]

³⁰It is easy to interpret informational iterations like $K_a K_b p$, but what, for instance, is the meaning of an ‘iterated obligation’ $O_a O_b p$?

Theory, where exciting new interfaces are building right now (cf. the chapter on social choice in this Handbook). Indeed, the priority graphs of Section 6 are a promising model for both belief merge and deontic merges into group obligations when group members have different authorities concerning various issues.

Eventually, all this may evolve into a mix of social choice theory, legal reasoning, and logics of institutions and organizations ([Grossi, 2007]).

10.3 Games and dependent behaviour

Multi-agency is tied together not just by social knowledge or beliefs, but also by dependent action. Therefore, logics of agency have close connections with game theory ([Shoham and Leyton-Brown, 2008], [Bentham, 2010a]), the study of strategic behaviour and its equilibria. In deontic practice, dependent action is crucial, too (think of sanctions or rewards), and games are again a congenial paradigm. Many specific topics reinforce this analogy. We already saw how belief-entangled set lifting is crucial to player's choices and their 'rationality' in the sense of game theory, making preference logics a natural tool in the analysis of games ([Roy, 2008], [Dégremont, 2010]). Conversely, ideas from game theory are making their way into deontic logic, witness the use of game solution procedures as moral deliberation procedures in [Loohius, 2009]. One might even argue that dependent behaviour is the source of morality, and in that sense, games would be the really natural next stage after the single-episode driven dynamic logics of this chapter.

10.4 Temporal perspective and long-term deontics

Games are one example of a long-term activity. The general logical setting here are temporal logics ([Fagin *et al.*, 1995], [Parikh and Ramanujam, 2003]) where new phenomena come to the fore having to do with long-term process structure. In our view, deontics and morality is not just about single episodes, but it only makes sense in a setting of action and interaction over time. Indeed, early work in deontic logic used temporal logics (cf. the pioneering dissertation [Eck, 1982]) where events happen in infinite histories, and obligations come and go. Likewise, in the multi-agent community, logics have been proposed for preferences over time, and planning behaviour leading to most preferred histories (cf. [Meyden, 1996], [Sergot, 2004]). For the case of information dynamics, connections between dynamic logics and temporal logics have been well-studied ([Bentham *et al.*, 2009c]), with an interesting emerging role for *protocols* as a new object of study, i.e., available procedures for reaching goals. Plans and protocols seem to have a clear normative dimension as well, and it might be worth-while to incorporate them into the evaluation dynamics of this chapter.

10.5 Syntactic perspectives: fine-structure of deontics

Finally, most dynamic logics for agency so far, whether about information dynamics or evaluation dynamics, are semantic in nature. The states changed by the process are semantic models. Now in epistemic logic, there is a continuing debate about the right representation of *information* here. Semantic information, though common to many areas, including decision theory and game theory, is coarse-grained, identifying logically equivalent propositions, and thereby suppressing the very act of logical analysis as an information-producing process. Zooming in on the latter, agents engage in many activities, such as inference, memory retrieval, introspection, or other forms of ‘awareness management’ that require a more fine-grained notion of information, closer to syntax. Several dynamic logics of this kind have been proposed in recent years (cf. [Benthem and Martínez, 2008], [Jago, 2006], [Velazquez-Quesada, 2009]), though no consensus has emerged.

Again, these issues return with equal force with deontic practice. Indeed, this chapter showed signs of this in its use of priority sequences and priority graphs. One thinks of these as syntactic objects than get manipulated syntactically with insertions, deletions, permutations, and the like. But also, deontic logic has its own very practical counterpart to the epistemological problem of ‘omniscience’. My duties to you cannot be based on my seeing every consequence of my commitments at once. I owe you ‘careful deliberation’, not omniscience. Citizens are supposed to know the law, but they need not be professional lawyers in seeing every deductive consequence. Hence, here too, there is a need for more fine-grained dynamic representations, closer to — though not necessarily identical with— actual syntax.

11 Short guide to the literature

The ideas in this chapter have a long history, and there are many proposals in the literature having to do with combination and ‘dynamification’ of preferences, defaults, and obligations. In addition to those included in this chapter, we would like to mention some other related approaches here.

Computation and agency [Meyer, 1988] was probably the first to look at deontic logic from a dynamic point of view, with the result that deontic logics are reduced to suitable versions of dynamic logics. This connection has become a high-light in computer science since, witness the regular *DEON* conference series. In the same computational tradition, [Meyden, 1996] takes the deontic logic/dynamic logic interface a step further, distinguishing two notions of permission, one of which, ‘free choice permission’ requires a new dynamic logic of permission, where preferences can hold between actions. Completeness theorems with respect to this enriched semantics are

given for several systems. [Pucella and Weissmann, 2004] provide a dynamified version of the logic of permission that builds action policies for agents by adding or deleting transitions. [Demri, 2005] reduces an extension of van der Meyden’s logic to *PDL*, yielding an EXPTIME decision procedure, and showing how *PDL* can deal with agents’ policies.

Semantics of natural language In another line, that goes back to [Spohn, 1988], [Veltman, 1996] presents an update semantics for default rules, locating their meaning in the way in which they modify expectation patterns. This is part of the general program of ‘update semantics’ for conditionals and other key expressions in natural language. [Torre and Tan, 1999] use ideas from update semantics to formalize deontic reasoning about obligations. In their view, the meaning of a normative sentence resides in the changes it brings about in the ‘ideality relations’ of agents to whom a norm applies. [Mastop, 2005] extends the general update semantic analysis of imperatives to include third person and past tense imperatives, while also applying it to the notion of free choice permission. [Parent, 2003] outlines a preference-based account of communication, which brings the dynamics of changing obligations for language users to the fore.

Philosophical logic and beyond Looking in another direction, with belief change as a starting point, [Hansson, 1995] identifies four types of changes in preference, namely revision, contraction, addition and subtraction, and shows that they satisfy plausible postulates for rational changes in preferences. The collection [Grune-Yanoff and Hansson, 2009] brings together the latest approaches on preference change from philosophy, economics and psychology. As we already mentioned, Hansson wrote a series of seminal papers combining ideas from preference logic and deontic logic, see e.g. [Hansson, 1990b], [Hansson, 1990a] and [Hansson, 2001b].

Philosophy and computation We just mention a few items from a rich literature. [Wellman and Doyle, 1991] gives a preference-based semantics for goals in decision theory. This provides a criterion for verifying the design of goal-based planning strategies, and a new framework for knowledge-level analysis of planning systems. [Horty, 1993] considers commonsense normative reasoning, arguing that the techniques of non-monotonic logic provide a better framework than the usual modal treatments. The paper has applications to conflicting obligations and conditional obligations.

Dynamic epistemic approaches [Zarnic, 2003] uses a simple dynamic update logic to formalize natural language imperatives of the form *FIAT* φ , which can be used in describing the search for solutions of planning problems. [Yamada, 2006] takes the update paradigm to logics of commands and obligations, modeling changes brought about by various acts of commanding. It combines a multi-agent variant of the language of monadic deontic

logic with a dynamic language for updates and commands. One step further, [Yamada, 2008] distinguishes the illocutionary acts of commanding from the perlocutionary acts that affect preferences of addressees, proposes a new dynamic logic which combines preference upgrade and deontic update, and discusses some deontic dilemmas in this setting.

Entanglement and deontics over time [Pacuit *et al.*, 2006] observe that an agent’s obligations are often dependent on what she knows, and introduce an epistemic preference language over temporal tree models. They study distinctions such as knowing one’s duty versus having a duty to know. [Lang *et al.*, 2003] propose a logic of desires whose semantics contains two ordering relations of preference and normality, respectively, and then interpret “in context A , I desire B ” as “the best among the most normal $A \wedge B$ worlds are preferred to the most normal $A \wedge \neg B$ worlds”, thus providing a new entanglement of preference and normality.

12 Conclusion

We have shown how current dynamic logics of agency can deal with deontic change on the model of preference change. In doing so, we have identified a number of perspectives and further directions where deontic and informational concerns seem a very natural fit.

BIBLIOGRAPHY

- [Andréka *et al.*, 2002] H. Andréka, M. Ryan, and P-Y. Schobbens. Operators and laws for combining preferential relations. *Journal of Logic and Computation*, 12:12–53, 2002.
- [Åqvist, 1987] L. Åqvist. *Introduction to Deontic Logic and the Theory of Normative Systems*. Naples: Bibliopolis, 1987.
- [Åqvist, 1994] L. Åqvist. Deontic logic. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, volume 2. Dordrecht: Kluwer, 1994.
- [Aucher *et al.*, 2009] G. Aucher, D. Grossi, A. Herzig, and E. Lorini. Dynamic context logic. In X. He, J. Horty, and E. Pacuit, editors, *Proceedings of LORI 2009*, volume 5834 of *FOLLI-LNAI*, pages 15–26. Springer, 2009.
- [Aucher, 2003] G. Aucher. A combined system for update logic and belief revision. Master’s thesis, MoL-2003-03. ILLC, University of Amsterdam, 2003.
- [Baltag and Smets, 2008] A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In M. Wooldridge G. Bonanno, W. van der Hoek, editor, *Logic and the Foundations of Game and Decision Theory*, volume 3 of *Texts in Logic and Games*. Amsterdam University Press, 2008.
- [Baltag *et al.*, 1998] A. Baltag, L.S. Moss, and S. Solecki. The logic of common knowledge, public announcements, and private suspicions. In I. Gilboa, editor, *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 98)*, pages 43–56, 1998.
- [Benthem and Liu, 2007] J. van Benthem and F. Liu. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logic*, 17:157–182, 2007.
- [Benthem and Martínez, 2008] J. van Benthem and M. Martínez. The stories of logic and information. In J. van Benthem and P. Adriaans, editors, *Handbook of Philosophy of Information*. Elsevier, 2008.

- [Benthem and Minica, 2009] J. van Benthem and S. Minica. Toward a dynamic logic of questions. In X. He, J.Horty, and E. Pacuit, editors, *Proceedings of LORI 2009*, volume 5834 of *FOLLI-LNAI*, pages 27–41, 2009.
- [Benthem *et al.*, 2006] J. van Benthem, S. van Otterloo, and O. Roy. Preference logic, conditionals and solution concepts in games. In H. Lagerlund, S. Lindström, and R. Sliwinski, editors, *Modality Matters: Twenty-Five Essays in Honour of Krister Segerberg*, pages 61–77. Uppsala Philosophical Studies 53, 2006.
- [Benthem *et al.*, 2009a] J. van Benthem, J. Gerbrandy, and B. Kooi. Dynamic update with probabilities. *Studia Logica*, 93(1):67–96, 2009.
- [Benthem *et al.*, 2009b] J. van Benthem, P. Girard, and O. Roy. Everything else being equal: A modal logic approach for ceteris paribus preferences. *Journal of Philosophical Logic*, 38(1):83–125, 2009.
- [Benthem *et al.*, 2009c] van J. Benthem, J. Gerbandy, T.Hoshi, and E.Pacuit. Merging frameworks for interaction. *Journal of philosophical logic*, 38(5):491–521, 2009.
- [Benthem *et al.*, 2010] J.van Benthem, D. Grossi, and F. Liu. Deontics = betterness + priority. In *DEON’2010 Proceedings*, 2010.
- [Benthem, 2007] J. van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logic*, 17:129–156, 2007.
- [Benthem, 2010a] J. van Benthem. *Logic Dynamics of Information and Interaction*. Cambridge University Press. To appear, 2010.
- [Benthem, 2010b] J.van Benthem. *Modal Logic for Open Minds*. CSLI, Stanford, 2010.
- [Blackburn *et al.*, 2001] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.
- [Blackburn *et al.*, 2007] P. Blackburn, J. van Benthem, and F. Wolter. *Handbook of Modal Logic*. Elsevier, 2007.
- [Boella *et al.*, 2009] G. Boella, G. Pigozzi, and L. van der Torre. Normative framework for normative system change. In Decker P., Sichman J., Sierra C., and Castelfranchi C., editors, *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, 2009.
- [Boutilier, 1994] C. Boutilier. Conditional logics of normality: A modal approach. *Artificial Intelligence*, 68:87–154, 1994.
- [Brandom, 1994] R. Brandom. *Making it Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press, 1994.
- [Burgess, 1984] J. Burgess. *Basic Tense Logic*, volume 2 of *Handbook of Philosophical Logic*, chapter 2, pages 89–133. Reidel, 1984.
- [Coste-Marquis *et al.*, 2004] S. Coste-Marquis, J. Lang, P. Liberatore, and P. Marquis. Expressive power and succinctness of propositional languages for preference representation. In *Proceedings of KR 2004*. AAAI Press, 2004.
- [Dégremont, 2010] C. Dégremont. *The Temporal Mind: Observations on the Logics of Belief Change in Interactive System*. PhD thesis, ILLC, Univesity of Amsterdam, 2010.
- [Demri, 2005] S. Demri. A reduction from DLP to PDL. *Journal of Logic and Computation*, 15:767–785, 2005.
- [Ditmarsch *et al.*, 2007] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Berlin: Springer, 2007.
- [Eck, 1982] J. van Eck. *A System of Temporally Relative Modal and Deontic Predicate Logic and its Philosophical Applications*. PhD thesis, 1982.
- [Fagin *et al.*, 1995] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about Knowledge*. The MIT Press, 1995.
- [Forrester, 1984] J. Forrester. Gentle murder, or the adverbial samaritan. *Journal of Philosophy*, 81:193–197, 1984.
- [Fraassen, 1972] B. van Fraassen. The logic of conditional obligation. *Journal of Philosophical Logic*, 1:417–438, 1972.
- [Fraassen, 1973] B. van Fraassen. Values and the heart’s command. *The Journal of Philosophy*, 70(1):5–19, 1973.

- [Gerbrandy, 1999] J. Gerbrandy. *Bisimulation on Planet Kripke*. PhD thesis, ILLC, University of Amsterdam, 1999.
- [Girard, 2008] P. Girard. *Modal Logics for Belief and Preference Change*. PhD thesis, Stanford University, 2008.
- [Goble, 2000] L. Goble. Multiplex semantics for deontic logic. *Nordic Journal of Philosophical Logic*, 5(2):113–134, 2000.
- [Governatori and Rotolo, 2005] G. Governatori and A. Rotolo. Logic of violations: A gntzen system for reasoning with contrary-to-duty obligations. *Australasian Journal of Logic*, 3, 2005.
- [Governatori and Rotolo, 2008a] G. Governatori and A. Rotolo. Changing legal systems: Abrogation and annulment. part 1: Revision and defeasible theories. In Ron van der Meyden and Leendert van der Torre, editors, *Proceedings of the 9th International Conference on Deontic Logic in Computer Science (DEON 2008), Luxembourg, Luxembourg, July 15-18, 2008*, volume 5076 of *LNAI*, pages 3–18. Springer, 2008.
- [Governatori and Rotolo, 2008b] G. Governatori and A. Rotolo. Changing legal systems: Abrogation and annulment. part 2: Temporalised defeasible logic. In G. Boella, G. Pigozzi, M. P. Singh, and H. Verhagen, editors, *Proceedings of the 3rd International Workshop on Normative Multiagent System (NormAS 2008), Luxembourg, Luxembourg, July 14-15, 2008*, pages 112–127, 2008.
- [Grandi and Endriss, 2009] U. Grandi and U. Endriss. First-order logic formalisation of arrows theorem. In X. He, J. Horty, and E. Pacuit, editors, *Proceedings of LORI 2009*, volume 5834 of *FOLLI-LNAI*, pages 251–262, 2009.
- [Grossi, 2007] D. Grossi. *Designing Invisible Handcuffs. Formal Investigations in Institutions and Organizations for Multi-Agent Systems*. PhD thesis, Utrecht University, 2007. SIKS Dissertation Series 2007-16.
- [Grune-Yanoff and Hansson, 2009] T. Grune-Yanoff and S.O. Hansson, editors. *Preference Change: Approaches from Philosophy, Economics and Psychology*. Theory and Decision Library. Springer, 2009.
- [Halpern, 1997] J.Y. Halpern. Defining relative likelihood in partially-ordered preferential structure. *Journal of Artificial Intelligence Research*, 7:1–24, 1997.
- [Hansson, 1969] B. Hansson. An analysis of some deontic logics. *Nous*, 3:373–398, 1969.
- [Hansson, 1990a] S.O. Hansson. Defining ‘good’ and ‘bad’ in terms of ‘better’. *Notre Dame of Journal of Formal Logic*, 31:136–149, 1990.
- [Hansson, 1990b] S.O. Hansson. Preference-based deontic logic. *Journal of Philosophical Logic*, 19:75–93, 1990.
- [Hansson, 1995] S.O. Hansson. Changes in preference. *Theory and Decision*, 38:1–28, 1995.
- [Hansson, 2001a] S.O. Hansson. Preference logic. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, volume 4, chapter 4, pages 319–393. Dordrecht: Kluwer, 2001.
- [Hansson, 2001b] S.O. Hansson. *The Structure of Values and Norms*. Cambridge University Press, 2001.
- [Hartog, 1985] G. den Hartog. *Wederkerige Verwachtingen [Mutual Expectations]*. PhD thesis, University of Amsterdam, 1985.
- [Holliday, 2009] W. Holliday. Trust and the dynamics of testimony. In D. Grossi, L. Kurzen, and F. Velzquez-Quesada, editors, *Logic and Interactive Rationality. Seminars Yearbook 2009*, pages 147–178. Institute for Logic, Language and Computation, University of Amsterdam, 2009.
- [Horty, 1993] J. Horty. Deontic logic as founded on nonmonotonic logic. *Annals of Mathematics and Artificial Intelligence*, 9:69–91, 1993.
- [Houser and Kurzban, 2002] D. Houser and R. Kurzban. Revealed preference, belief, and game theory. *Economics and Philosophy*, 16:99–115, 2002.
- [Icard et al., 2010] T. Icard, E. Pacuit, and Y. Shoham. Joint revision of belief and intentions. In *Proceedings of KR’2010*, 2010.

- [Jackson, 1985] F. Jackson. On the semantics and logic of obligation. *Mind*, XCIV:177–195, 1985.
- [Jago, 2006] M. Jago. *Logics for Resource-Bounded Agents*. PhD thesis, University of Nottingham, 2006.
- [Jongh and Liu, 2009] D. de Jongh and F. Liu. Preference, priorities and belief. In T. Grune-Yanoff and S.O. Hansson, editors, *Preference Change: Approaches from Philosophy, Economics and Psychology*, Theory and Decision Library. Springer, 2009.
- [Konkka, 2000] J. Konkka. Funk games: Approaching collective rationality. E-Theis, Helsinki, 2000.
- [Kooi and Tamminga, 2006] B. Kooi and A. Tamminga. Conflicting obligations in multi-agent deontic logic. In John-Jules Ch. Meyer and Lou Goble, editors, *Deontic Logic and Artificial Normative Systems: 8th International Workshop on Deontic Logic in Computer Science*, volume 4048 of *LNCS*, pages 175–186. Springer, 2006.
- [Lang and Torre, 2008] J. Lang and L. van der Torre. From belief change to preference change. In *Proceedings of ECAI-2008*, 2008.
- [Lang et al., 2003] J. Lang, L. van der Torre, and E. Weydert. Hidden uncertainty in the logical representation of desires. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI'03)*, 2003.
- [Lewis, 1973] D. Lewis. *Counterfactuals*. Oxford: Blackwell, 1973.
- [Liu, 2004] F. Liu. Dynamic variations: Update and revision for diverse agents. Master's thesis, MoL-2004-05. ILLC, University of Amsterdam, 2004.
- [Liu, 2008] F. Liu. *Changing for the Better: Preference Dynamics and Agent Diversity*. PhD thesis, ILLC, University of Amsterdam, 2008.
- [Liu, 2009a] F. Liu. Logics for interaction between preference and belief. Manuscript, Department of Philosophy, Tsinghua University, 2009.
- [Liu, 2009b] F. Liu. Preference change: A quantitative approach. *Studies in Logic*, 2(3):12–27, 2009.
- [Liu, 2010] F. Liu. A two-level perspective on preference. Manuscript, Department of Philosophy, Tsinghua University, 2010.
- [Loohius, 2009] L. O. Loohius. Obligations in a responsible world. In X. He, J. Horty, and E. Pacuit, editors, *Proceedings of LORI 2009*, number 5834 in FoLLI-LNAI, pages 251–262. Springer, 2009.
- [Mastop, 2005] R. Mastop. *What Can You Do? Imperative Mood in Semantic Theory*. PhD thesis, ILLC, University of Amsterdam, 2005.
- [Meyden, 1996] R. van der Meyden. The dynamic logic of permission. *Journal of Logic and Computation*, 6:465–479, 1996.
- [Meyer and Hoek, 1995] J.-J.Ch. Meyer and W. van der Hoek. *Epistemic Logic for Computer Science and Artificial Intelligence*. Number 41 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1995.
- [Meyer, 1988] J.-J.Ch. Meyer. A different approach to deontic logic: Deontic logic viewed as a variant of dynamic logic. *Notre Dame Journal of Formal Logic*, 29:109–136, 1988.
- [Moore, 1903] G. E. Moore. *Principia Ethica*. Cambridge University Press, 1903.
- [Otterloo, 2005] S. van Otterloo. *A Strategic Analysis of Multi-agent Protocols*. PhD thesis, Liverpool University, UK, 2005.
- [Pacuit et al., 2006] E. Pacuit, R. Parikh, and E. Cogan. The logic of knowledge based on obligation. *Knowledge, Rationality and Action*, 149:311–341, 2006.
- [Parent, 2003] X. Parent. Remedial interchange, contrary-to-duty obligation and commutation. *Journal of Applied Non-Classical Logics*, 13(3/4):345–375, 2003.
- [Parikh and Ramanujam, 2003] R. Parikh and R. Ramanujam. A knowledge-based semantics of messages. *Journal of Logic, Language and Information*, 12:453–467, 2003.
- [Plaza, 1989] J.A. Plaza. Logics of public announcements. In *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, 1989.
- [Prince and Smolensky, 1993] A. Prince and P. Smolensky. *Optimality Theory: Constraint Interaction in Generative Grammar*. Malden, Ma: Blackwell, 1993.

- [Pucella and Weissmann, 2004] R. Pucella and V. Weissmann. Reasoning about dynamic policies. In *Proceedings FoSSaCS-7*, Lecture Notes in Computer Science 2987, pages 453–467, 2004.
- [Rao and Georgeff, 1991] A.S. Rao and M.P. Georgeff. Modeling rational agents within a BDI-architecture. In J. Allen, R. Fikes, and E. Sandewall, editors, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, pages 473–484. Morgan Kaufmann publishers Inc.: San Mateo, CA, USA, 1991.
- [Rescher, 1966] N. Rescher. Notes on preference, utility, and cost. *Synthese*, 16:332–343, 1966.
- [Rott, 2003] H. Rott. Basic entrenchment. *Studia Logica*, 73:257–280, 2003.
- [Rott, 2006] H. Rott. Shifting priorities: Simple representations for 27 iterated theory change operators. In H. Langerlund, S. Lindström, and R. Sliwinski, editors, *Modality Matters: Twenty-Five Essays in Honour of Krister Segerberg*, pages 359–384. Uppsala Philosophical Studies 53, 2006.
- [Roy, 2008] O. Roy. *Thinking before Acting: Intentions, Logic and Rational Choice*. PhD thesis, ILLC, University of Amsterdam, 2008.
- [Searle and Vanderbeken, 1985] J. Searle and D. Vanderbeken. *Foundations of Illocutionary Logic*. Cambridge University Press, Cambridge, 1985.
- [Sergot, 2004] M. Sergot. (C+)⁺⁺: An action language for modelling norms and institutions. Technical Report 8, Department of Computing, Imperial College, London, 2004.
- [Shoham and Leyton-Brown, 2008] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game Theoretic and Logical Foundations*. Cambridge University Press, 2008.
- [Shoham, 1988] Y. Shoham. *Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence*. The MIT Press: Cambridge, MA, 1988.
- [Snyder, 2004] J. Snyder. Product update for agents with bounded memory. Manuscript, Department of Philosophy, Stanford University, 2004.
- [Spohn, 1988] W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W.L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change and Statistics II*, pages 105–134. Kluwer, Dordrecht, 1988.
- [Tan and Torre, 1996] Y.-H. Tan and L. van der Torre. How to combine ordering and minimizing in a deontic logic based on preferences. In M. Brown and J. Carmo, editors, *Deontic Logic, Agency and Normative Systems, DEON '96: Third International Workshop on Deontic Logic in Computer Science*, pages 216–232, 1996.
- [Torre and Tan, 1999] L. van der Torre and Y. Tan. An update semantics for deontic reasoning. In P. McNamara and H. Prakken, editors, *Norms, Logics and Information Systems*, pages 73–90. IOS Press, 1999.
- [Torre, 1997] L. van der Torre. *Reasoning about Obligations: Defeasibility in Preference-based Deontic Logic*. PhD thesis, Rotterdam, 1997.
- [Trapp, 1985] R.W. Trapp. Utility theory and preference logic. *Erkenntnis*, 22:301–339, 1985.
- [Velazquez-Quesada, 2009] F.R. Velazquez-Quesada. Inference and update. *Synthese*, 169:283–300, 2009.
- [Veltman, 1985] Frank Veltman. *Logics for Conditionals*. PhD thesis, University of Amsterdam, 1985.
- [Veltman, 1996] F. Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25:221–261, 1996.
- [Wellman and Doyle, 1991] M. Wellman and J. Doyle. Preferential semantics for goals. In *Proceedings of the National Conference on Artificial Intelligence*, pages 698–703, 1991.
- [Wooldridge, 2000] M. Wooldridge. *Reasoning about Rational Agents*. The MIT Press, 2000.
- [Wright, 1951] G.H. von Wright. Deontic logic. *Mind*, 60:1–15, 1951.

- [Wright, 1956] G.H. von Wright. A note on deontic logic and derived obligation. *Mind*, 65:507–509, 1956.
- [Wright, 1963] G.H. von Wright. *The Logic of Preference*. Edinburgh, 1963.
- [Wright, 1964] G.H. von Wright. A new system of deontic logic. In *Danish Yearbook of Philosophy*, pages 173–182. 1964.
- [Yamada, 2006] T. Yamada. Acts of commands and changing obligations. In K. Inoue, K. Satoh, and F. Toni, editors, *Proceedings of the 7th Workshop on Computational Logic in Multi-Agent Systems (CLIMA VII)*, 2006. Revised version appeared in LNAI 4371, pages 1-19, Springer-Verlag, 2007.
- [Yamada, 2008] T. Yamada. Logical dynamics of some speech acts that affect obligations and preferences. *Synthese*, 165(2):295–315, 2008.
- [Yamada, 2010] T. Yamada. Scorekeeping and dynamic logics of speech acts. Manuscript, Hokkaido University, 2010.
- [Zarnic, 2003] B. Zarnic. Imperative change and obligation to do. In K. Segerberg and R. Sliwinski, editors, *Logic, Law, Morality: Thirteen Essays in Practical Philosophy in Honour of Lennart Aqvist*, pages 79–95. Uppsala philosophical studies 51. Uppsala: Department of Philosophy, Uppsala University, 2003.

Johan van Benthem

Institute for Logic, Language and Computation (ILLC),
University of Amsterdam,
P.O. Box 94242,
1090 GE AMSTERDAM,
The Netherlands.
Email: johan.vanbenthem@uva.nl

Fenrong Liu

Department of Philosophy,
Tsinghua University,
Beijing, 100084,
China
Email: fenrong@tsinghua.edu.cn