Chapter 3
Preference over Worlds: Static Logic

3.1 Introduction

Preferences arise from comparisons between alternatives, say, outcomes, actions, or situations. Such a comparison is typically associated with some ordering, indicating that one alternative is "better" than another. For instance, when playing chess or other games, choosing a move π_1 instead of π_2 is determined largely by a consideration concerning the outcomes that π_1 or π_2 leads to. In general, individual preferences can be used to predict behavior by rational agents, as studied in game theory and decision theory. Preference logics in the literature study the abstract properties of different comparative structures [101].

Preference statements can be weaker or stronger in what they say about alternatives being compared – and also, they may be more "objective" or more "epistemic". A statement like "I prefer sunsets to sunrises" can be cast merely in terms of "what is better for me", or as a more complex propositional attitude involving my beliefs about the relevant events. In this chapter, we take a somewhat objective approach, where a binary primitive preference relation in possible worlds models supports a unary modality "true in some world which is at least as good as the current one" (similar models have been studied in [55] and [94]). We will call this relation "betterness" to distinguish it from richer notions of preference. Then we will use a standard modal language to express preference and its related notions. We will show that this language is very expressive, being able to express various kinds of preference that agents may have between propositions, i.e., types of events.[1]

This chapter is structured as follows. In Section 3.2 we will introduce the language and semantics for modal betterness logic. A complete axiomatization will be stated. In Section 3.3 we will study expressive power, with a special interest in defining preference over propositions ("generic preference") in terms of *lifting* the primitive betterness relation from possible worlds to an ordering over propositions, viewed as sets of possible worlds. Various possible lifts will be considered. Finally, as an illustration, Section 3.4 will focus on one type of lift, namely the $\forall\exists$-version,

[1] In a different setting, [44] showed how such a language, extended with hybrid modalities, defines conditionals, Nash equilibrium, and Backward Induction solutions to games.

F. Liu, *Reasoning about Preference Dynamics*, Synthese Library 354,
DOI 10.1007/978-94-007-1344-4_3,

and explore its logical properties in more detail. Our conclusion summarizes the platform laid in this chapter for the rest of this book.

3.2 Modal Betterness Logic

As we commented in the above, semantical betterness relations are treated as modalities in the language. We now introduce the language of modal betterness logic:

Definition 3.1 (modal betterness language) Take any set of propositional variables Φ, with a variable p ranging over Φ. The modal betterness language $\mathcal{L}_{\mathcal{B}}$ is given by the following inductive syntax rule:

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

The intended reading of $\langle\leq\rangle\varphi$ is "φ is true in some world that is at least as good as the current world", while $\langle<\rangle\varphi$ says that "φ is true in some world that is strictly better than the current world."[2] We will discuss how these two modalities help lift betterness relations to preferences over propositions in Section 3.3, in particular, in the case where the betterness relations lack connectedness. In general, these notions are agent-relative, but in what follows we will mostly suppress this aspect, since it is orthogonal to our main points. In addition, the auxiliary existential modality $E\varphi$ says that "there is some world where φ is true". Combinations of these modalities can capture a wide variety of binary preference statements comparing propositions, again, we will show this soon.

As usual, we will write $[\leq]\varphi$ for the universal modality $\neg\langle\leq\rangle\neg\varphi$, and we will write $[<]\varphi$ and $U\varphi$ for the duals of $\langle<\rangle\varphi$ and E, respectively. Either $[\leq]$ or $\langle\leq\rangle$ can be introduced as a primitive, we will take the technical convenience into account and use both formats interchangeably in this context.

How is this formal language connected to "preference" as it occurs in natural discourse? One may be inclined to read $\langle\leq\rangle\varphi$ as "some agent prefers φ". But as with other logical systems, there is a gap between the formalism and common usage. E.g., just saying that the agent sees some better world where φ holds seems too weak, while the universal modality $[\leq]\varphi$ "in all better worlds" seems much too strong. Cf. [102] for a thorough discussion of senses of preference, and ways in which formal languages do or do not match up.

Here we just point out the following feature. Our approach emphasizes comparisons of worlds, rather than propositions, whereas common notions of preference often play between propositions, or semantically, sets of worlds. Even so, the preferences between propositions can be *defined* in the present language, as a lift from the

[2] We use two independent modalities here for weak and strict betterness. This may seem strange, since strict order was definable in terms of weak order. But the point is that this definition cannot be reproduced in a natural way inside our modal language, which therefore brings out reasoning with both modalities on a par.

betterness relations. We will study those lifts soon in Section 3.3. For the moment, we just take this expressive power of our modal language for granted. The virtue of working with simple base modalities, as we do, is that these “decompose” many more complex preference statements in a perspicuous manner, while allowing for a simple dynamic approach later on.

Definition 3.2 (modal betterness model) A modal betterness model is a tuple $\mathfrak{M} = (S, \leq, V)$ where S is a set of possible worlds, $\leq$ is a reflexive and transitive relation (the ‘betterness’ pre-order) over these worlds, and V is a valuation assigning truth values to proposition letters at worlds.[3]

We read $s \leq t$ as “t is at least as good as s”, or “t is weakly better than s”. If $s \leq t$ but not $t \leq s$, then t is *strictly better* than s, written as $s < t$. If $s \leq t$ and $t \leq s$, then s and t are *indifferent*.

Note that we do not require that our betterness relations be *connected* in the sense of the Lewis sphere models for conditional logic. In general, we want to allow for genuinely incomparable worlds where an agent has no preference either way, not because she is indifferent, but because she has no means of comparing the worlds at all. It is a very natural situation we may often encounter in real life. This is just as in the semantics for the minimal conditional logic. Of course, in special settings, such as the standard utility-based preference orderings of outcomes in a game, connectedness may be quite appropriate.

Definition 3.3 (truth conditions) Given a modal betterness model $\mathfrak{M} = (S, \leq, V)$, and a world $s \in S$, we define $\mathfrak{M}, s \models \varphi$ (formula φ is true in $\mathfrak{M}$ at s) in the usual manner by induction on the construction of the formula φ:

$\mathfrak{M}, s \models \top$	iff	always.
$\mathfrak{M}, s \models p$	iff	$s \in V(p)$.
$\mathfrak{M}, s \models \neg\varphi$	iff	not $\mathfrak{M}, s \models \varphi$.
$\mathfrak{M}, s \models \varphi \wedge \psi$	iff	$\mathfrak{M}, s \models \varphi$ and $\mathfrak{M}, s \models \psi$.
$\mathfrak{M}, s \models \langle\leq\rangle\varphi$	iff	for some t with $s \leq t$, $\mathfrak{M}, t \models \varphi$.
$\mathfrak{M}, s \models \langle<\rangle\varphi$	iff	for some t with $s < t$, $\mathfrak{M}, t \models \varphi$.
$\mathfrak{M}, s \models E\varphi$	iff	for some world t in S, $\mathfrak{M}, t \models \varphi$.

Definition 3.4 (modal equivalence) Two models $\mathfrak{M}, s$ and $\mathfrak{M}', t$ are modally equivalent, written as $\mathfrak{M}, s \leftrightsquigarrow \mathfrak{M}', t$ if they satisfy the same formulas from $\mathcal{L}_{\mathcal{B}}$.

[3] In this chapter, we use pre-orders since we want the generality of possibly non-total preferences. Total orders, the norm in areas like game theory, provide an interesting specialization for the results in this chapter. We will study total ordered preference in Chapter 7.

Definition 3.5 (bisimulation) Two models $\mathfrak{M}, s$ and $\mathfrak{M}', t$ are bisimilar (written $\mathfrak{M}, s \mathrel{\mathring{=}} \mathfrak{M}', t$) if there is a relation $Z \subseteq S \times S'$ such that:

(1) If sZt then for all $p \in \Phi$, $s \in V(p)$ iff $t \in V(p)$.
(2) If sZt and $s \leq s'(s < s')$ then there is a $t' \in S'$ such that $t \leq t'(t < t')$ and $s'Zt'$. (the forth condition).
(3) If sZt and $t \leq t'(t < t')$ then there is a $s' \in S$ such that $s \leq s'(s < s')$ and $s'Zt'$. (the back condition).
(4) For all $s \in S$, there is a $t \in W'$ such that sZt.
(5) For all $t \in S'$, there is a $s \in W$ such that sZt.

This is what is called a total bisimulation, as it includes conditions 4 and 5. It is easy to show that any two bisimular models are modally equivalent with regard to our language, in other words, we say that the language is bisimulation-invariance. Bisimular models are often used to show undefinability of certain operators in some language. We will come back to this issue in Section 3.3.

As for the resulting logics, [39] give a complete axiomatization for the logic of weak and strict betterness modalities. Essentially, the system consists of **S4**-axioms for operator $[\leq]$, **K** for $[<]$, and **S5**-axioms for universal modality U, plus some axioms for interaction between operators. We restate it here, omitting the proof.

Theorem 3.6 *The modal betterness logic is completely axiomatized by the following set of principles:*

(1) *propositional tautologies*
(2) $[\leq](\varphi \to \psi) \to ([\leq]\varphi \to [\leq]\psi)$
(3) $[<](\varphi \to \psi) \to ([<]\varphi \to [<]\psi)$
(4) $[\leq]\varphi \to \varphi$
(5) $[\leq]\varphi \to [\leq][\leq]\varphi$
(6) $U(\varphi \to \psi) \to (U\varphi \to U\psi)$
(7) $U\varphi \to \varphi$
(8) $U\varphi \to UU\varphi$
(9) $\neg U\varphi \to U\neg U\varphi$
(10) $[\leq]\varphi \to [<]\varphi$
(11) $[<]\varphi \to [\leq][<]\varphi$
(12) $[<]\varphi \to [<][\leq]\varphi$
(13) $[\leq]([\leq]\varphi \vee \psi) \wedge [<]\psi \to \varphi \vee [\leq]\psi$
(14) $U\varphi \to [\leq]\varphi$
(15) $E\varphi \leftrightarrow \neg U\neg\varphi$
(16) $\langle\leq\rangle\varphi \leftrightarrow \neg[\leq]\neg\varphi$
(17) $\langle<\rangle\varphi \leftrightarrow \neg[<]\neg\varphi$

The correspondence between the above axioms and the frame properties can be studied just as one does in standard modal logics [24]. Additional axioms in our language impose further frame conditions on models. Nevertheless, we will work with the minimal system described above, leaving such extras aside.

3.3 Expressive Power

Our modal base language seems so simple and standard that it may be hard to see what it can define by way of more complex notions relevant to preference. In this section, we will show that it can express more than the reader might have thought. We will give two examples: "generic preferences" and "conditional preference".

Preference between specific worlds, introduced above, is just as in decision theory and game theory. But preference can be used to compare different sorts of things. In game theory, we do need to compare kinds of situation. Starting with von Wright, logicians have studied "generic preferences" between kinds of object, or kinds of situation. Such scenarios, too, occur in many other fields in the literature, with various interpretations of the basic relation $y \leq x$. It is interpreted as "x is as least as normal (or typical) as y" in [55] on conditional and default reasoning, as "x at least as preferred or desirable as y" in [66], as "x is no more remote from actuality than y" in [127] on counterfactuals, and as "x is as likely as y" in [94] on qualitative reasoning with probability. In all these settings, it makes sense to extend the given order on worlds to an order of propositions φ, ψ. For instance, in real life, students may have preferences concerning courses, but they need to also form an order over kinds of courses, say theoretical versus practical, i.e., over sets of individual courses. Likewise, we may have preferences regarding individual commodities, but we often need a preference over sets of them. And similar aggregation scenarios are abundant in social choice theory, for which an extensive survey is [21].

In what follows, we will show that preference over propositions is definable as a binary operator $P(\varphi, \psi)$ in our modal logic, whose language is rich enough to explicitly define "lifts" of betterness on worlds to a binary ordering on sets of worlds. Studies of such lifts abound (cf. [42, 44] and [39]), and our purpose in this section is merely to streamline some results.

3.3.1 Generic Preference: Quantifier Lifts

One obvious way of lifting world orders $x \leq y$ to proposition or set orders $X \trianglelefteq Y$ uses definitional schemas that can be classified by the quantifiers which they involve. As has been observed by many authors (cf. [39]), there are four obvious two-quantifier combinations for lifting:

(1) $X \trianglelefteq^{\forall\forall} Y \Leftrightarrow \forall x \in X\ \forall y \in Y\colon x \leq y$;

(2) $X \trianglelefteq^{\forall\exists} Y \Leftrightarrow \forall x \in X\ \exists y \in Y\colon x \leq y$;

(3) $X \trianglelefteq^{\exists\forall} Y \Leftrightarrow \exists x \in X\ \forall y \in Y\colon x \leq y$;

(4) $X \trianglelefteq^{\exists\exists} Y \Leftrightarrow \exists x \in X\ \exists y \in Y\colon x \leq y$.

Taking the strict version of the betterness relation gives four more combinations:

(5) $X \lhd^{\forall\forall} Y \Leftrightarrow \forall x \in X\, \forall y \in Y: x < y$;

(6) $X \lhd^{\forall\exists} Y \Leftrightarrow \forall x \in X\, \exists y \in Y: x < y$;

(7) $X \lhd^{\exists\forall} Y \Leftrightarrow \exists x \in X\, \forall y \in Y: x < y$;

(8) $X \lhd^{\exists\exists} Y \Leftrightarrow \exists x \in X\, \exists y \in Y: x < y$.

As usual, we can define $X \lhd Y$ as $X \unlhd Y$ and $\neg Y \unlhd X$. One can argue for any of these as a notion of generic preference. Reference [44] claims that $\lhd^{\forall\forall}$ is the notion of "preference" intended by von Wright in his seminal work on preference logic [197] and provides an axiomatization. But the tradition is much older, and (modal) logics for preference relations over sets of possible worlds have been considered by [55, 127] and [94], and other authors. In particular, [94] studied the above ∀∃-combination.

We are not in the position to claim that one lift is more plausible than another, but our main concern here is the logical properties of lifts, and the expressive power of our modal betterness language.

3.3.2 Expressing Generic Preferences in $\mathcal{L}_{\mathcal{B}}$

So far, our discussion of preference over propositions has been semantic-oriented. A natural question is the following: can we express the above generic preferences in the language $\mathcal{L}_{\mathcal{B}}$? The answer is positive. We start with the following four:

Definition 3.7 (**generic preference: ∀∃ and ∃∃**) The ∀∃-preference and ∃∃-preference can be defined in the language $\mathcal{L}_{\mathcal{B}}$ as follows:

(1) $\varphi \unlhd^{\forall\exists} \psi \;:=\; U(\varphi \to \langle\leq\rangle\psi)$

(2) $\varphi \unlhd^{\exists\exists} \psi \;:=\; E(\varphi \wedge \langle\leq\rangle\psi)$

(3) $\varphi \lhd^{\forall\exists} \psi \;:=\; U(\varphi \to \langle<\rangle\psi)$

(4) $\varphi \lhd^{\exists\exists} \psi \;:=\; E(\varphi \wedge \langle<\rangle\psi)$

We can read $\varphi \unlhd^{\forall\exists} \psi$ as "for each φ-world, there exists a ψ-world which is as good as that φ-world", and read $\varphi \lhd^{\forall\exists} \psi$ as "for each φ-world, there exists a better ψ-world". Once again, this "majorization" is one very natural way of comparing sets of possible worlds – and it has counterparts in many other areas which use derived orders on powerset domains. In particular, [94] took this definition (with an interpretation of "relative likelihood" between propositions) and gave a complete logic for the case in which the basic order on S is a pre-order. It is also well-known that Lewis gave a complete logic for preference relations over propositions in his study of counterfactuals in [127], where the given order on S is quasi-linear.

Now let us consider the remaining combinations. It turns out to be more complex and interesting, having to do with the models under discussion in the following sense. If the betterness relations satisfy the property of *connectedness*, we can make use of the Definition 3.7 and define those cases as follows:

Definition 3.8 (**generic preference: ∀∀ and ∃∀**) The ∀∀-preference and also the ∃∀-preference can be defined in the language $\mathcal{L}_{\mathcal{B}}$ on connected models:

(1) $\varphi \trianglelefteq^{\forall\forall} \psi \ := \ U(\psi \rightarrow [<]\neg\varphi)$

(2) $\varphi \trianglerighteq^{\exists\forall} \psi \ := \ E(\varphi \wedge [<]\neg\psi)$

(3) $\varphi \lhd^{\forall\forall} \psi \ := \ U(\psi \rightarrow [\leq]\neg\varphi)$

(4) $\varphi \rhd^{\exists\forall} \psi \ := \ E(\varphi \wedge [\leq]\neg\psi)$

However, if we drop connectedness, basic modal definability fails:

Fact 3.9 *The ∀∀ and ∃∀-meaning of propositional preference cannot be defined in the modal betterness language $\mathcal{L}_B$ on non-connected models.*

Proof This can be shown by a standard argument, providing two bisimilar models one of which satisfies the lift, and one of which does not (cf. [44]). □

Thus, some lifts would amount to genuine first-order extensions of the modal base language, in the spirit of hybrid logics. While we will not pursue such extensions in our book, many of the results that we develop will survive such generalizations.[4]

3.3.3 Conditional Preference

Here is one more example that shows the expressive power of our basic language. A widespread *maximality operator* $[Best(\psi)]\varphi$ in the literature on conditional or deontic logic says that the "best' ψ-worlds in some relevant order satisfy some proposition φ". The counterpart of this notion may be called *conditional preference*, and it can be expressed as follows in our language:

(1) $$P^{\psi}\varphi := U(\psi \rightarrow \langle\leq\rangle(\psi \wedge [\leq](\psi \rightarrow \varphi))).$$

This says that φ is preferred on condition of ψ, if and only if, for all worlds that satisfy ψ, there is a better ψ-state such that all ψ-states above it are φ. A similar modal definition for conditionals was first proposed by [53] and [55]. Here, we

---

4 One example are the added "intersection modalities" of Chapter 5.

restated it in our language. We omit the simple verification that this formula really has the intended meaning, at least on finite models.

Interestingly, our base language even offers another definition, whose syntax is even closer to the maximality clause for the antecedent. For instance, using the strict betterness modality, [83, Ch. 3] defines conditional preference as

$$P^{\psi}\varphi \;:=\; U((\psi \wedge \neg\langle<\rangle\psi) \rightarrow \varphi).$$

A few comments are in order here. First, for these definitions to capture the intended meaning of the maximality operator, finiteness, or more generally, converse well-foundedness of the ordering should be assumed, to make sure that maximal worlds exist satisfying given formulae. In the absence of converse well-foundedness, as was observed in [55], Formula (1) expresses something more than maximality.

Next, what we showed here does not just apply to betterness and preference. Following the same idea, given an order relation of "relative plausibility", we can define *conditional beliefs*, which will be important in many of our later chapters. Formula (1) above then amounts to the standard syntactic "relativization" of absolute belief (as truth in all most plausible worlds) to just the worlds satisfying the antecedent. Moreover, we will often relativize this still further, restricting everything to just the set of worlds that are *epistemically accessible* from the current one. Thus, in Chapter 5, we will use a knowledge modality K instead of the universal modality U when lifting "epistemically entangled" betterness relations, while dealing analogously with plausibility relations.

For now, we just emphasize what all this has shown. Our basic modal betterness logic can encode quite a few complex properties of preference – as well as, reinterpreting the basic world order, of related notions such as belief.

3.4 Preservation and Characterization of ∀∃-Preference

In this section, mainly an excursion, we will raise a few more semantic issues, to further understanding the lifting phenomenon per se.

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

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

But this does not constrain our lifts very much, since all four quantifier combinations satisfy it. We will not explore further constraints here. Instead, we concentrate on one particular lift, namely, ∀∃-preference, and try to understand better how the lift generally works. One question that comes to mind immediately is this: Can the properties of an underlying preference on worlds be preserved when it is lifted to the level of propositions? In particular, consider reflexivity and transitivity that we assumed for preference in Section 3.2. Can we show $\trianglelefteq^{\forall\exists}(\varphi, \psi)$ has these two properties? In fact, the answer is positive. We can even prove something stronger:

Fact 3.10 *Reflexivity and transitivity of the relation $\leq$ are preserved in the lifted relation $\trianglelefteq^{\forall\exists}$, but also vice versa.*

Proof Reflexivity. To show that $\trianglelefteq^{\forall\exists}(X, X)$, by Definition 3.7, we need that $\forall x \in X \exists y \in X : x \leq y$. Since we have $x \leq x$, take y to be x, and we get the result.

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

Transitivity. Assume that $\trianglelefteq^{\forall\exists}(X, Y)$ and $\trianglelefteq^{\forall\exists}(Y, Z)$. We show that $\trianglelefteq^{\forall\exists}(X, Z)$. By Definition 3.7, this means we have $\forall x \in X \exists y \in U : x \leq y$ and $\forall y \in Y \exists z \in Z : y \leq z$. Then by transitivity of the base relation, we have that $\forall x \in X \exists z (x \leq z$, and this is precisely $\trianglelefteq^{\forall\exists}(X, Z)$.

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

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

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

Next, staying at the level of propositions, suppose we have a preference relation that is a ∀∃-lift from a base relation over possible worlds. What are necessary and sufficient conditions for being such a relation? The following theorem provides a complete characterization:

Theorem 3.11 (characterization) *A binary relation $\trianglelefteq$ over propositions satisfies the following four properties iff it is a ∀∃-lifting of some preference relation over the underlying possible worlds.*

1. $Y \trianglelefteq X \Rightarrow Y \bigcap Z \trianglelefteq X$ (*left downward monotonicity*)
2. $Y \trianglelefteq X \Rightarrow Y \trianglelefteq X \bigcup Z$ (*right upward monotonicity*)
3. $\forall i \in I, Y_i \trianglelefteq X \Rightarrow \bigcup_i Y_i \trianglelefteq X.$ (*left union property*)
4. $\{y\} \trianglelefteq \bigcup_i X_i \Rightarrow \{y\} \trianglelefteq X_i$ *for some* $i \in I.$ (*right distributivity*)

Proof (⇐) Assume that $\trianglelefteq$ is a ∀∃-lifting. We show that $\trianglelefteq$ has the four properties.

(1) Assume $Y \trianglelefteq X$, i.e., $\forall y \in Y \exists x \in X : y \leq x$. Since $Y \bigcap Z \subseteq Y$, we also have $\forall y \in Y \bigcap Z \exists x \in X : y \leq x$, and hence $Y \bigcap Z \trianglelefteq X$.
(2) Assume $\forall y \in Y \exists x \in X : y \leq x$. Since $X \subseteq X \bigcup Z$, we have $\forall y \in Y \exists x \in X \bigcup Z : y \leq x$: that is, $Y \trianglelefteq X \bigcup Z$.
(3) Assume that for all $i \in I$, $\forall y \in Y_i \exists x \in X : y \leq x$. Let $y \in \bigcup_i Y_i$, then for some j: $y \in Y_j$. By the assumption, we have $\forall y \in Y_j \exists x \in X : y \leq x$, so $\exists x \in X : y \leq x$. This shows that $\bigcup_i Y_i \trianglelefteq X$.

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

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

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

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

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

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

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

We have spent so much time with this one generic preference because it occurs quite widely in the literature, providing a typical use of our modal preference logic. But of course, many other lifts could be analyzed in a similar manner.

3.5 Conclusion

We have introduced a static modal betterness language in this chapter, which admits of a complete axiomatization for its valid forms of reasoning with preference. We then showed that the betterness relation over possible worlds can be lifted to generic preferences over propositions using various quantifier combinations. We discussed different lifts, giving us a better impression of the expressive power and limitations of our modal language. Finally, we characterized one lift completely, the widespread and well-behaved notion of "∀∃-preference".

This modal betterness language is just a launching platform in this book. As we will see in the next chapter, it provides a natural model for the dynamics of preference change, both at the basic level of modification in the betterness relations, and derived from that, at the level of defined "lifted" generic preferences between propositions.