Reasoning about general preference relations

Grossi, Davide; van der Hoek, Wiebe; Kuijer, Louwe B.

Published in:
Artificial Intelligence

DOI:
10.1016/j.artint.2022.103793

Publication date:
2022

Reasoning about general preference relations☆

Davide Grossi[1], Wiebe van der Hoek[2], Louwe B. Kuijer[2]

Abstract

Preference relations are at the heart of many fundamental concepts in artificial intelligence, ranging from utility comparisons, to defeat among strategies and relative plausibility among states, just to mention a few. Reasoning about such relations has been the object of extensive research and a wealth of formalisms exist to express and reason about them. One such formalism is conditional logic, which focuses on reasoning about the "best" alternatives according to a given preference relation. A "best" alternative is normally interpreted as an alternative that is either maximal (no other alternative is preferred to it) or optimal (it is at least as preferred as all other alternatives). And the preference relation is normally assumed to satisfy strong requirements (typically transitivity and some kind of well-foundedness assumption).

Here, we generalize this existing literature in two ways. Firstly, in addition to maximality and optimality, we consider two other interpretations of "best", which we call unmatchedness and acceptability. Secondly, we do not inherently require the preference relation to satisfy any constraints. Instead, we allow the relation to satisfy any combination of transitivity, totality and anti-symmetry.

This allows us to model a wide range of situations, including cases where the lack of constraints stems from a modeled agent being irrational (for example, an agent might have preferences that are neither transitive nor total nor anti-symmetric) or from the interaction of perfectly rational agents (for example, a defeat relation among strategies in a game might be anti-symmetric but not total or transitive).

For each interpretation of "best" (maximal, optimal, unmatched or acceptable) and each combination of constraints (transitivity, totality and/or anti-symmetry), we study the sets of valid inferences. Specifically, in all but one case we introduce a sound and strongly complete axiomatization, and in the one remaining case we show that no such axiomatization exists.

☆An early version of this work, containing **MOU** and **ACC** but not the various extensions of those axiomatizations, was presented at KR2020 [11].

[1]University of Groningen and University of Amsterdam

[2]University of Liverpool

1. Introduction

1.1. Preference orders

A preference order over alternatives is a fundamental notion in the formal representation of knowledge, inference and decision making. Such relations can be used to represent, among other things, the relative plausibility of states in knowledge representation, the value or utility of outcomes in social choice or the defeat relation between strategies in game theory.

There is a long tradition, starting in the 1960s [36], of reasoning about such preference stuctures using logic. This programme has been carried out with different focuses and tools, in many different disciplines: within artificial intelligence (AI), for the systematization of many different forms of 'common sense' reasoning as non-standard (in particular, non-monotonic) inference relations [19, 4, 5]; within epistemology, for the representation and analysis of processes of information acquisition by rational agents, in the AGM tradition [1], the conditionals tradition [30, 21, 18, 6, 7, 12], and more recently the dynamic-epistemic logic tradition [33, 3, 8, 34]; within deontic logic, for the representation and analysis of conditional obligations [13, 22, 29, 25, 26]. And the above list is surely non-exhaustive. Conceptual and technical similarities among some of these fields have been object of extensive scrutiny, in particular in the early 1990s [32, 24, 16, 15, 23] and more recently by [35], which was guided by a similar systematization aim as the one we are pursuing in this paper. Recently, a series of papers discussing preferences, ways to represent them and how to reason about them, was published as a special issue in *Artificial Intelligence* [9].

Our paper positions itself within the above tradition and, from a technical point of view, focuses on the further development of the toolbox of conditional logics, which so far has been an object of research especially within the areas of deontic logic and epistemic logic. The applicability of conditional logics is not restricted to those domains, however; any system of preferences can be studied through the lens of conditional logic.

1.2. Conditional logics

The focus of conditional logics is on the "best" alternatives. After all, one should generally believe the best (most plausible) states, choose the best (most preferred) outcome and play the best (least defeated) strategies. Furthermore, we are not interested only in the alternatives that are the best in general, but also those that are *conditionally* the best.

The term "best" in the preceding paragraph is somewhat vague, and deliberately so: depending on the situation it may be disambiguated in different ways. Traditionally (see, e.g., [13, 21, 26]), the "best" alternatives are considered to be the *maximal* ones, i.e., those not strictly dis-preferred to any other, or the *optimal* ones, i.e., those that are weakly preferred over all others. Here, we also consider two other options: the "best" alternatives as the *unmatched* ones, i.e., those not weakly dispreferred to any other, or the *acceptable* ones, i.e., those that are a member of at least one minimal retentive set.[3]

For a given disambiguation of "best", one can then define a conditional logic with an operator $B(\varphi \mid \psi)$, which holds if, of the alternatives satisfying ψ, all the best ones satisfy φ. The dual $P(\varphi \mid \psi) := \neg B(\neg\varphi \mid \psi)$ holds if among the best alternatives satisfying ψ, there is at least one that satisfies ψ. The meaning of $B(\varphi \mid \psi)$ and $P(\varphi \mid \psi)$ depends on the meaning of the preference relation $\preceq$. Here are some examples:

- If $\preceq$ is an agent's plausibility ordering, then $B(\varphi \mid \psi)$ represents that agent's *conditional belief*: under the condition that ψ is true, it believes that φ is true (so if the agent were to learn ψ, it would conclude that φ is likely true). The formula $P(\varphi \mid \psi)$ represents *conditional perceived possibility*: under the condition that ψ is true, the agent does not believe $\neg\varphi$. See, e.g., [33, 3].

- If $\preceq$ is the preference relation of an agent or society, then $B(\varphi \mid \psi)$ represents *conditional obligation* for that agent or society: given that ψ is true, φ should be true. The formula $P(\varphi \mid \psi)$ then represents *conditional permission*: φ is not banned by the truth of ψ, and therefore permitted (although it need not be permitted under a stronger condition $\psi \wedge \chi$). See, e.g., [13, 26].

- If $\preceq$ is a defeat relation among strategies, then $B(\varphi \mid \psi)$ and $P(\varphi \mid \psi)$ represent *conditional rationality*: $B(\varphi \mid \psi)$ holds if, in the strategy space restricted to ψ, playing a strategy satisfying φ is necessary to be rational while $P(\varphi \mid \psi)$ holds if, in the ψ-restricted strategy space, playing a strategy satisfying φ can be rational. See, e.g., [14, 17].

These examples are far from exhaustive; $\preceq$ can represent any relation that indicates one alternative being, in a very broad sense, "better than" or "preferred over" another. Every meaning of $\preceq$ has its own associated meaning of $B(\varphi \mid \psi)$ and $P(\varphi \mid \psi)$, although that meaning is often harder to put into words than the three examples given above.

1.3. Constraints and Validity

Depending on the situation, it may be reasonable to assume that the relation $\preceq$ satisfies certain assumptions. For example, if $\preceq$ represents the plausibility ordering of a single rational agent, then it should be transitive. In other situations, such an assumption would not be warranted, however. For example, the plausibility ordering of a single irrational agent, the aggregated plausibility ordering of a group of agents (rational or not) and the defeat relation among strategies are generally not transitive.

[3] A retentive set can be thought of as a set of alternatives that is collectively maximal. See Definitions 2–3 for details.

As discussed above, these kinds of conditional logics were originally studied as deontic logics. As such, the authors of those studies made assumptions that are appropriate in that context, but not, for example, in the context of comparing strategies in a game. The most important of these assumptions is a kind of well-foundedness, which requires every non-empty set of alternatives to have at least one "best" state. The result of such an assumption is that $B(\varphi \mid \psi) \rightarrow P(\varphi \mid \psi)$ becomes valid, representing the principle that obligation implies permission.

In general, which constraints hold simply depends on the setting one has in mind for the application of the logic. In order to be as general as possible, we therefore do not inherently require the relation to satisfy any constraints and study the effects on the logic of adding various constraints. Specifically, we consider any combination of transitivity, totality and anti-symmetry. Our reasons for choosing these properties are twofold. Firstly, for a given situation it is relatively easy to determine whether it is reasonable to assume that the preference relation is transitive, total and/or anti-symmetric. Secondly, these properties appear to have a clear technical connection to the differences between maximality, optimality, unmatchedness and acceptability (see Proposition 8).

Each choice of a disambiguation of "best" (maximal, optimal, unmatched or acceptable) and set of constraints (transitive, total or anti-symmetric) yields a variant of conditional logic. For each variant, we investigate which entailments are valid. Understanding such valid entailments helps us understand the logic in general, and is especially important in cases where the underlying relation $\preceq$ is not known. For example, an agent may be able to express some or all of their conditional beliefs without being explicitly aware of $\preceq$. A characterization of valid entailment will allow us to draw further conclusions about the agent's beliefs, without ever fully knowing $\preceq$. In particular, if the agent's set of beliefs is inconsistent, and therefore entails $\bot$, we know that the agent is irrational and should drop some of their beliefs.

For each variant, we try to characterize the valid entailments using a sound and strongly complete axiomatization. In all cases but one, we do exactly that. In the one remaining case, we instead prove an impossibility result: we show that a strongly complete axiomatization does not exist.

1.4. Paper contribution

In summary, our contribution is as follows. We define conditional semantics based on unmatchedness and acceptability, in addition to the traditional definitions based on maximality and optimality. We consider situations where the relation $\preceq$ is unconstrained, as well as situations where $\preceq$ is assumed to satisfy any combination of transitivity, totality and anti-symmetry.

We investigate the set of valid inferences for each type of semantics and every set of properties of $\preceq$. In all cases except one, we introduce a sound and strongly complete axiomatization. For the final remaining case we show that no such axiomatization exists. All axiomatizations introduced in this paper are new, and build on axiomatizations from [2, 25, 26].

1.5. Paper outline

In Section 2 we formally define preference structures, as well as the notions of maximality, optimality, unmatchedness and acceptability. In that section we also show how the four definitions compare to one another, and we prove a few further results that are useful later on in the paper. In Section 3 we then define the language and semantics for conditional logics. After that, in Section 4, we take an in-depth look at four examples that illustrate conditional belief, obligation and playability.

In Sections 5–7 we define the axiomatizations, and prove them to be sound and strongly complete. Specifically, in Section 5 we define the shared "core" of axioms that will be included in each of the axiomatizations. In Section 6 we discuss the axiomatizations where $\preceq$ is unconstrained. Finally, in Section 7, we discuss the axiomatizations where $\preceq$ is assumed to be transitive, total and/or anti-symmetric.

In total, we consider 8 different axiomatizations. For the sake of brevity we do not include full soundness and completeness proofs for each axiomatization in the main text. Instead, we include full proofs for some of the more interesting variants. For the remaining variants, we include only a proof sketch in the main text, and place the full proof in the appendix.

2. Preliminaries

2.1. Models

Definition 1 (Models)**.** A *model* is a triple $M = (S, \preceq, V)$ where

- S is a set of states,
- $\preceq \subseteq S \times S$ is a relation and
- $V : \mathbf{P} \to 2^S$ assigns to each atom a set of states.

We define $\prec, \succeq$ and $\succ$ from $\preceq$ in the usual way. Furthermore, we write $s_1 \perp s_2$ for $s_1 \not\preceq s_2$ and $s_2 \not\preceq s_1$, and $s_1 \approx s_2$ for $s_1 \preceq s_2$ and $s_2 \preceq s_1$. Finally, we use a superscript $*$ to denote the transitive closure of a relation.

We interpret $s \in V(p)$ as "the basic fact p is true in state s", and $s_1 \preceq s_2$ as "s_2 is (weakly) preferred to s_1". We take a very broad view of what a state is, and what it means for one state to be preferred to another. In the introduction, we mentioned three examples: $s_1 \preceq s_2$ could mean that the state s_2 is at least as plausible as the state s_1, that the state s_2 is at least as desirable as the state s_1, or that the strategy s_2 defeats the strategy s_1. But there are more possibilities.

For example, $s_1 \preceq s_2$ could mean that s_1 and s_2 are participants in a sports tournament and s_2 has defeated s_1. Or $s_1 \preceq s_2$ could mean that s_1 and s_2 are possible laws that a politician could choose from, and that if they choose s_1 then the majority of the electorate will blame them for not choosing s_2 instead. Or maybe $s_1 \preceq s_2$ means that s_1 and s_2 are species and that s_2 is better adapted to a given environment than s_1 is. Any interpretation of states and preferences is fine, as long as $s_1 \preceq s_2$ indicates that s_2 is, in some sense, (weakly) better than s_1.

On a technical level, there are two important things that we should note about Definition 1. Firstly, $\preceq$ is entirely unconstrained. In particular, it is not required to be transitive, anti-symmetric, total or reflexive. Each of these constraints can reasonably be assumed in certain contexts, but not in others. We therefore do not assume that any of these assumptions hold in general. For transitivity, anti-symmetry and totality, we do consider the situation where they are assumed to hold as a special case, and we study the consequences of such an assumption, see Section 7.

Reflexivity is a slightly different case. There are certainly situations where reflexivity can be assumed, such as when $\preceq$ represents a utility comparison between worlds, and situations where it cannot be assumed, such as when $\preceq$ represents victories of one participant over another in a tournament. However, for every interpretation of $\preceq$ that we can think of, whether s is among the best states is independent of whether $s \preceq s$. If a boxer has defeated all tournament participants other than themselves, they have defeated everyone. In the terminology we use in this paper, that means they are "optimal" among the participants. Likewise, if the only world that provides at least as much utility as a given world is that world itself, then that world should be considered "unmatched", despite it being "matched" by itself.

As such, all logics studied in this paper are invariant under reflexivity. We do not assume that $\preceq$ is reflexive, but in those contexts where reflexivity would be appropriate it can freely be assumed.

The second thing that we should note about Definition 1 is that we use only a single relation $\preceq$. This can easily be generalized to a set $\{\preceq_i \mid i \in I\}$ of relations where I is an index set. Indices $i \in I$ can be used to represent different agents or authorities, with $\preceq_i$ being the preference relation according to i. Then $B_i(\varphi \mid \psi)$ means that among the ψ-states, all of the best ones, according to $\preceq_i$, satisfy φ.

Another way to use multiple relations is to let the relation depend on the state. So instead of a single $\preceq$, we would use a set $\{\preceq_s \mid s \in S\}$. Then $s_1 \preceq_s s_2$ means that s_2 is at least as good as s_1, from the perspective of s. Unlike an index i, the state s is not added as a subscript to the B-operator, but instead used as a point of evaluation. So when evaluated in state s, the formula $B(\varphi \mid \psi)$ holds if all of the best ψ-states, according to $\preceq_s$, satisfy φ.

State dependent preferences can be used to model situations where the preference order differs based on the current state. For example, consider a child that is more excited about toys that they don't have; if they have a toy car but no ball then they prefer the ball over the car, while if they have a ball but no toy car then they prefer the car over the ball. Such differences based on the current state can include *meta-preferences*, where an agent has preferences over their own preferences. For example, an agent might likes sugary drinks while, for health reasons, preferring to like them a bit less. The index set I and state-dependent preferences can also be combined, resulting in a set $\{\preceq_{i,s} \mid i \in I, s \in S\}$, which can be used to represent meta-preferences among multiple agents or authorities, like a parent trying to raise a child in such a way that it adopts a given set of ethical rules.

Using multiple preference relations is unproblematic from a technical point of view; we believe that all of the results in this paper still apply for multiple relations, with the only necessary changes being the addition of indices in the relevant places and, with state-dependent preferences, the removal of the **Abs** axiom (see Section 5). However, while such a generalization would not make any of the results conceptually more difficult, keeping track of the different relations would require a considerable amount of additional notation and bookkeeping. As such, we confine ourselves to remarking that a generalization to multiple relations would be straightforward, while using a single relation in all proofs. We do, however, consider one example using state-dependent preferences in Section 4.

2.2. Four definitions of "best"

Now that we have defined models, we can consider the formal definitions of the four ways in which a state can be among the "best" states. First, however, we need one auxiliary definition.

Definition 2 (Retentive set)**.** Let a model $M = (S, \preceq, V)$ and a set $X \subseteq S$ be given. A set $Y \subseteq X$ is *retentive* in X if there are no $y \in Y, x \in X \setminus Y$ such that $y \prec x$.

A set $Y \subseteq X$ is *minimal retentive* in X if it is retentive in X and there is no non-empty $Y' \subset Y$ that is retentive in X.

Where X is understood we often write "minimal retentive" for "minimal retentive in X." In more mathematical terminology, a retentive set is one that is closed under $\prec$.

Definition 3. Let a model $M = (S, \preceq, V)$ and a set $X \subseteq S$ be given.

- An element $s_1 \in X$ is *maximal in* X if for every $s_2 \in X$, $s_1 \not\prec s_2$.
- An element $s_1 \in X$ is *optimal in* X if for every $s_2 \in X \setminus \{s_1\}$, $s_2 \preceq s_1$.
- An element $s_1 \in X$ is *unmatched in* X if for every $s_2 \in X \setminus \{s_1\}$, $s_1 \not\preceq s_2$.
- An element $s_1 \in X$ is *acceptable in* X if there is a minimal retentive set $Y \subseteq X$ such that $s_1 \in Y$.

We denote the set of maximal elements of X in M by $\mathsf{max}^M(X)$, the set of optimal elements of X in M by $\mathsf{opt}^M(X)$, the set of unmatched elements of X in M by $\mathsf{unm}^M(X)$ and the set of acceptable elements of X in M by $\mathsf{acc}^M(X)$. Where this should not cause confusion we will omit reference to M and speak of $\mathsf{max}(X), \mathsf{unm}(X), \mathsf{opt}(X)$ and $\mathsf{acc}(X)$.

Note that, for every $\beta \in \{\mathsf{max}, \mathsf{opt}, \mathsf{unm}, \mathsf{acc}\}$, whether $s \in \beta(X)$ is independent of whether $s \preceq s$, for the reasons discussed above.

We use "best" as a generic term that can mean any of the four notions, i.e., the best states are the maximal, optimal, unmatched or acceptable states, depending on the context. The functions max, opt, unm and acc are closely related to the *choice functions* [28] studied in social choice and economics, except that a choice function C is generally assumed to satisfy $C(X) \neq \emptyset$ when $X \neq \emptyset$, a property that is not satisfied by any of max, opt, unm and acc.

Maximality and optimality are well studied, as discussed in the introduction. Unmatchedness in preference structures has not, to the best of our knowledge, been studied previously, but it is a relatively straightforward variation of optimality. With optimality, the emphasis with $s_1 \preceq s_2$ is on s_2 being (weakly) preferred over s_1, with the best states being those that are weakly preferred over all other states. With unmatchedness, the emphasis with $s_1 \preceq s_2$ is instead on s_1 being (weakly) dis-preferred to s_2, and the best states are those that are not dis-preferred to any other state.

For example, consider the previously described scenario where $s_1 \preceq s_2$ indicates that if a politician were to implement a law s_1, they would be blamed by the majority of the electorate for not implementing s_2 instead. From an electoral perspective, this politician would be well advised to look for unmatched laws, since those are the ones that do not anger the electorate.

Similarly, unmatchedness can be the appropriate kind of "best" in certain sports, such as boxing, where being undefeated is given great importance. If we write $s_1 \preceq s_2$ when s_1 has lost a match to s_2, so $s_1 \approx s_2$ holds when s_1 and s_2 have defeated each other in separate matches, then the unmatched boxers are the ones that remain undefeated.

Acceptability has not previously been studied in the context of conditional logics, but minimal retentive sets have been studied in various other contexts, under various names. The term *minimal retentive set* is from tournament theory see, e.g., [20]. In the related field of voting theory, minimal retentive sets are also known as *Schwartz sets*, after [27]. In game theory, minimal retentive sets are known as *sink equilibria* [10].

The intuition behind retentive sets is that they are collectively maximal. After all, a set Y is retentive in X if there are no $y \in Y$ and $x \in X \setminus Y$ such that $y \prec x$, much like a state y is maximal in X if there is no $x \in X \setminus \{y\}$ such that $y \prec x$. Another way to think of this is that the retentiveness condition guarantees that if x is acceptable and $x \prec y$, then y is also acceptable. So "bestness", when interpreted as acceptability, is closed under $\prec$.

Merely being part of a retentive set is no evidence of being "best", however, because every $x \in X$ is part of at least one set that is retentive in X, namely X itself. If y is part of some *minimal* retentive set Y, however, then y is a necessary component of the collective maximality of Y. In some situations this does provide evidence of y being among the best alternatives.

An alternative characterization of acceptability is that $s \in \mathsf{acc}(X)$ if for every $s' \in X$ such that $s \prec s'$, there are $s_1, \cdots, s_n \in X$ such that $s' \prec s_1 \prec \cdots \prec s_n \prec s$. So if s is strictly beaten by any other state s', then s must transitively strictly beat s'.

Acceptability tends to be the appropriate notion of "best" in situations where

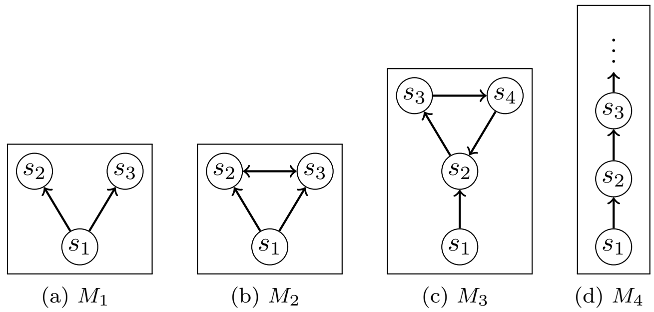

Figure 1: Four example models.

transitivity of $\preceq$ cannot reasonably be assumed. It is no coincidence that acceptability is closely related to notions from social choice, where Condorcet's paradox shows that a collective preference can be non-transitive, and game theory, where cyclical dominance among strategies is common.

The following example illustrates the difference between the four types of "best".

Example 4. Consider the graphs M_1–M_4 drawn in Figure 1, where an arrow from s_1 to s_2 indicates that $s_1 \preceq s_2$. The maximal, optimal, unmatched and acceptable states of M_1, M_2, M_3 and M_4 are as follows.

	M_1	M_2	M_3	M_4
max	$\{s_2, s_3\}$	$\{s_2, s_3\}$	$\emptyset$	$\emptyset$
opt	$\emptyset$	$\{s_2, s_3\}$	$\emptyset$	$\emptyset$
unm	$\{s_2, s_3\}$	$\emptyset$	$\emptyset$	$\emptyset$
acc	$\{s_2, s_3\}$	$\{s_2, s_3\}$	$\{s_2, s_3, s_4\}$	$\emptyset$

Note that this example shows that no $\beta \in \{\mathsf{max}, \mathsf{opt}, \mathsf{unm}, \mathsf{acc}\}$ satisfies the property that $\beta(X) \neq \emptyset$ whenever $X \neq \emptyset$. We argue that this is in line with intuition; in the infinitely ascending chain of M_4, no state s_i is best, since it is always strictly better to choose s_{i+1} over s_i. On finite models, however, such infinitely ascending chains are impossible, so we have the following proposition.

Proposition 5. *Let $M = (S, \preceq, V)$ be given, and let $X \subseteq S$ be a finite, non-empty set. Then* $\mathsf{acc}(X) \neq \emptyset$.

Proof. Let $\preceq_X$ be the restriction of $\preceq$ to X, and let $\prec_X$ be the strict part of $\preceq_X$. Now, let $\prec_X^*$ be the reflexive transitive closure of $\prec_X$.[4] Because it is a transitive relation on a finite set, there is at least one maximal element of X with respect

[4] Note that $\prec_X^*$ is the reflexive transitive closure of the strict part of the relation $\preceq_X$. This is not the same as the strict part of the transitive reflexive closure of $\preceq_X$. In particular, $\prec_X^*$ is generally not anti-symmetric.

to $\prec_X^*$. Let x be such a maximal element, and take $Y = \{y \mid x \prec_X^* y\}$. We will show that Y is minimal retentive in X.

First, to see that Y is retentive, suppose towards a contradiction that there are $y \in Y$ and $z \in X \setminus Y$ such that $y \prec z$. Then also $y \prec_X z$. Furthermore, $y \in Y$, so by the definition of Y we have $x \prec_X^* y$. It then follows that $x \prec_X^* y \prec_X z$ and therefore $x \prec_X^* z$, contradicting $z \notin Y$.

Next, to see that Y is minimal retentive, suppose towards a contradiction that some non-empty $Z \subset Y$ is also retentive in X. Take any $z \in Z$ and $y \in Y \setminus Z$. Because $Z \subset Y$, we have $x \prec_X^* z$. Because x is a maximal element of X under $\prec_X^*$, we must then also have $z \prec_X^* x$, so $z \prec_X^* x \prec_X^* y$, and therefore $z \prec_X^* y$. This implies that $z \prec_X \cdots \prec_X y$. By retentiveness of Z, this implies that $y \in Z$, contradicting $y \in Y \setminus Z$.

We have shown that Y is minimal retentive, and we have $x \in Y$. So $x \in \mathsf{acc}(X)$, which implies that $\mathsf{acc}(X) \neq \emptyset$. □

From Example 4 we can also see that max, opt, unm and acc are all different. They are not, however, unrelated.

Proposition 6. *For every $M = (S, \preceq, V)$ and every $X \subseteq S$,*

1. $\mathsf{opt}(X) \subseteq \mathsf{max}(X)$,
2. $\mathsf{unm}(X) \subseteq \mathsf{max}(X)$ *and*
3. $\mathsf{max}(X) \subseteq \mathsf{acc}(X)$.

Furthermore, apart from $\mathsf{opt}(X) \subseteq \mathsf{acc}(X)$ *and* $\mathsf{unm}(X) \subseteq \mathsf{acc}(X)$ *(which follow from 1.–3. by transtivity of $\subseteq$) no other inclusions hold in general, i.e., there are $X_1, \cdots, X_5$ such that*

- $\mathsf{unm}(X_1) \not\subseteq \mathsf{opt}(X_1)$,
- $\mathsf{opt}(X_2) \not\subseteq \mathsf{unm}(X_2)$,
- $\mathsf{max}(X_3) \not\subseteq \mathsf{opt}(X_3)$,
- $\mathsf{max}(X_4) \not\subseteq \mathsf{unm}(X_4)$ *and*
- $\mathsf{acc}(X_5) \not\subseteq \mathsf{max}(X_5)$ *(and therefore also* $\mathsf{acc}(X_5) \not\subseteq \mathsf{opt}(X_5)$ *and* $\mathsf{acc}(X_5) \not\subseteq \mathsf{unm}(X_5)$*).*

Proof. Suppose $s_1 \in \mathsf{opt}(X)$. Then for every $s_2 \in X$, we have $s_2 \preceq s_1$ and therefore, in particular, $s_1 \not\prec s_2$. By definition, this implies $s_1 \in \mathsf{max}(X)$. It follows that $\mathsf{opt}(X) \subseteq \mathsf{max}(X)$.

Suppose then that $s_1 \in \mathsf{unm}(X)$. Then for every $s_2 \in X$, we have $s_1 \not\preceq s_2$ and therefore, in particular, $s_1 \not\prec s_2$. By definition this implies that $s_1 \in \mathsf{max}(X)$, so $\mathsf{unm}(X) \subseteq \mathsf{max}(X)$.

Finally, suppose that $s_1 \in \mathsf{max}(X)$. Then for every $s_2 \in X$, we have $s_1 \not\prec s_2$. The singleton set $\{s_1\}$ is therefore retentive and, because it is a singleton, also minimal retentive in X. It follows that $s_1 \in \mathsf{acc}(X)$, so $\mathsf{max}(X) \subseteq \mathsf{acc}(X)$.

To see that no other inclusions hold, note that every other inclusion has a counterexample among the models M_1, M_2, M_3 and M_4 from Example 4. □

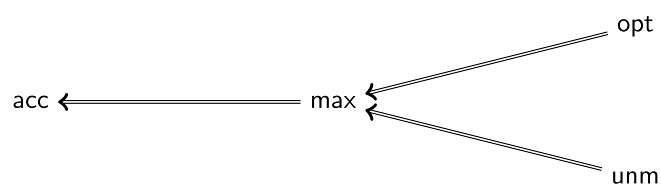

Figure 2: Inclusions among the four "bestness" concepts.

Under certain conditions there are more inclusions that hold, however.

Definition 7. A model $M = (S, \preceq, V)$ is transitive, anti-symmetric or total whenever $\preceq$ is, i.e.,

- M is transitive if for all $s_1, s_2, s_3 \in S$, if $s_1 \preceq s_2$ and $s_2 \preceq s_3$ then $s_1 \preceq s_3$,
- M is anti-symmetric if for all $s_1, s_2 \in S$, if $s_1 \neq s_2$ and $s_1 \preceq s_2$ then $s_2 \not\preceq s_1$ and
- M is total if for all $s_1, s_2 \in S$, if $s_1 \neq s_2$ then either $s_1 \preceq s_2$ or $s_2 \preceq s_1$.

Each of these three conditions collapses one of the distinctions between the bestness concepts.

Proposition 8. *Let $M = (S, \preceq, V)$ be a model and let $X \subseteq S$.*

- *If M is transitive then* $\mathsf{max}(X) = \mathsf{acc}(X)$*,*
- *if M is anti-symmetric then* $\mathsf{max}(X) = \mathsf{unm}(X)$ *and*
- *if M is total then* $\mathsf{max}(X) = \mathsf{opt}(X)$*.*

Before proving Proposition 8, let us consider an auxiliary lemma.

Lemma 9. *Let $M = (S, \preceq, V)$ be transitive, let $X \subseteq S$ and let $Y \subseteq X$ be minimal retentive. Then Y is a singleton.*

Proof. By contradiction. Suppose that $s_1, s_2 \in Y$ such that $s_1 \neq s_2$. Consider the sets $Y_1 = \{t \in Y \mid s_1 \prec t\}$ and $Y_2 = \{t \in Y \mid s_2 \prec t\}$. We will show that Y_1 and Y_2 are also retentive; for reasons of symmetry it suffices to show this for one of the two sets. So take any $t \in Y_1$ and $x \in X \setminus Y_1$. We distinguish two cases:

- if $x \in Y$ then $s_1 \not\prec x$, since otherwise we would have $x \in Y_1$. Because $t \in Y_1$ we have $s_1 \prec t$, so by transitivity of M we obtain $t \not\prec x$,
- if $x \notin Y$ then $t \not\prec x$, since $t \in Y$ and Y is retentive.

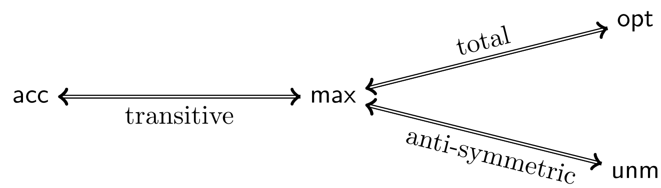

Figure 3: Conditional inclusions among the four "bestness" concepts.

In either case, $t \not\prec x$. This holds for every $t \in Y_1$ and $x \in X \setminus Y_1$, so Y_1 is retentive in X.

We have shown that Y_1 and, by symmetry, Y_2 are retentive. By minimality of Y this implies that $Y_1 = Y_2 = Y$. But that would imply $s_1 \in Y_2$ and $s_2 \in Y_1$ and therefore, by definition, $s_1 \prec s_2$ and $s_2 \prec s_1$. We have arrived at a contradiction and therefore conclude that there are no two distinct elements in Y, so Y is a singleton. □

We now continue to prove Proposition 8.

Proof of Proposition 8. Suppose that M is transitive and that $s \in \mathsf{acc}(X)$. Then there is some minimal retentive set $Y \subseteq X$ such that $s \in Y$. By Lemma 9 it follows that $Y = \{s\}$. Because Y is retentive this implies that for all $x \in X \setminus \{s\}$ we have $s \not\prec x$. So $s \in \mathsf{max}(X)$. It follows that $\mathsf{acc}(X) \subseteq \mathsf{max}(X)$. Proposition 6 showed that $\mathsf{max}(X) \subseteq \mathsf{acc}(X)$, so we have $\mathsf{max}(X) = \mathsf{acc}(X)$.

Suppose that M is anti-symmetric and that $s \in \mathsf{max}(X)$. Then for every $x \in X$, $s \not\prec x$. Due to anti-symmetry this implies that $s \not\preceq x$ for all $x \in X \setminus \{s\}$ and therefore that $s \in \mathsf{unm}(X)$. It follows that $\mathsf{max}(X) \subseteq \mathsf{unm}(X)$. Proposition 6 showed that $\mathsf{unm}(X) \subseteq \mathsf{max}(X)$, so we have $\mathsf{max}(X) = \mathsf{unm}(X)$.

Suppose that M is total and that $s \in \mathsf{max}(X)$. Then for every $x \in X$, $s \not\prec x$. Due to totality this implies that $x \preceq s$ for all $x \in X \setminus \{s\}$ and therefore that $s \in \mathsf{opt}(X)$. It follows that $\mathsf{max}(X) \subseteq \mathsf{unm}(X)$. Proposition 6 showed that $\mathsf{opt}(X) \subseteq \mathsf{max}(X)$, so we have $\mathsf{max}(X) = \mathsf{unm}(X)$. □

In addition to the inclusions shown in Figure 2 we therefore have the conditional inclusions shown in Figure 3.

The main lesson of Propositions 6 and 8 is that for a well-behaved relation $\preceq$ (i.e., one that is transitive, anti-symmetric and total) it does not matter which concept of bestness we use, but that for less well-behaved relations the different bestness notions give different outcomes, making it important to choose the correct notion for a given context.

We should also remark that opt and unm are, in some sense, each other's duals. More precisely, we have the following lemma.

Lemma 10. *Let $M_1 = (S, \preceq_1, V)$ and $M_2 = (S, \preceq_2, V)$ be models with the property that for every $s, t \in S$,*

- $s \prec_1 t$ *if and only if* $s \prec_2 t$,
- $s \approx_1 t$ *if and only if* $s \perp_2 t$ *and*
- $s \perp_1 t$ *if and only if* $s \approx_2 t$.

Then for every $X \subseteq S$, $\mathsf{opt}^{M_1}(X) = \mathsf{unm}^{M_2}(X)$ *and* $\mathsf{unm}^{M_1}(X) = \mathsf{opt}^{M_2}(X)$.

Proof. Because the requirements on M_1 and M_2 are symmetrical, it suffices to show that $\mathsf{opt}^{M_1}(X) = \mathsf{unm}^{M_2}(X)$.

Suppose $s \in \mathsf{opt}^{M_1}(X)$. Then for every $x \in X \setminus \{s\}$ we have $x \prec_1 s$ or $x \approx_1 s$. It follows that $x \prec_2 s$ or $x \perp_2 s$, so $s \in \mathsf{unm}^{M_1}(X)$.

Suppose then that $s \in \mathsf{unm}^{M_2}(X)$. Then for every $x \in X \setminus \{s\}$ we have $x \prec_2 s$ or $x \perp_2 s$. It follows that $x \prec_1 s$ or $x \approx_1 s$, so $s \in \mathsf{opt}^{M_1}(X)$. □

Note that M_1 is total if and only if M_2 is anti-symmetric. So the duality between opt and unm mirrors a similar kind of duality between totality and anti-symmetry.

Before continuing with the language and semantics, let us first consider three more lemmas that will come in useful later. The first lemma states that max and acc depend only on $\prec$, i.e., if $s_1 \perp s_2$ we can change this to $s_1 \approx s_2$ without changing maximality or acceptability, and vice versa.

Lemma 11. *Let* $M_1 = (S, \preceq_1, V)$ *and* $M_2 = (S, \preceq_2, V)$ *be models with the property that for all* $s, t \in S$, $s \prec_1 t$ *if and only if* $s \prec_2 t$. *Then for every* $X \subseteq S$, $\mathsf{max}^{M_1}(X) = \mathsf{max}^{M_2}(X)$ *and* $\mathsf{acc}^{M_1}(X) = \mathsf{acc}^{M_2}(X)$.

Proof. This follows immediately from the fact that s_1 is maximal in X if and only if there is no $s_2 \in X$ such that $s_1 \prec s_2$, and that a set $Y \subseteq X$ is retentive if there are no $y \in Y$ and $x \in X \setminus Y$ such that $y \prec x$. □

The second lemma relates to a condition known as *Sen's* α, after [28]. A choice function β satisfies this condition if for every $Y \subseteq X$ and $y \in Y$, if $y \in \beta(X)$ then $y \in \beta(Y)$. In other words, if y is best in X, then y must also be best in every subset Y of X that contains y. We show that $\mathsf{max}, \mathsf{opt}$ and unm satisfy Sen's α.

Lemma 12. *Let* $M = (S, \preceq, V)$ *be any model and* $\beta \in \{\mathsf{max}, \mathsf{opt}, \mathsf{unm}\}$. *For every* $Y \subseteq X \subseteq S$ *and every* $y \in Y$, *if* $y \in \beta(X)$ *then* $y \in \beta(Y)$.

Proof. We give the proof for $\beta = \mathsf{max}$, the proofs for $\beta = \mathsf{opt}$ and $\beta = \mathsf{unm}$ are analogous. Suppose therefore that $s \in \mathsf{max}(X) \cap Y$. Then for every $t \in X$, $s \not\prec t$. Because $Y \subseteq X$, it follows that for every $t \in Y$, $s \not\prec t$. Together with $s \in Y$ this implies that $s \in \mathsf{max}(Y)$. □

Note that this property does not hold for acc. If Z is minimal retentive in X then $Z \cap Y$ will be retentive in Y, but it need not be minimal retentive. Consider, for example, the model M_3 from Figure 1. In the unrestricted model there is one minimal retentive set, namely $\{s_2, s_3, s_4\}$. So s_2 is acceptable in this model. Yet if we restrict the model to $Y = \{s_1, s_2, s_3\}$ the only minimal retentive set is $\{s_3\}$, so s_2 is not acceptable in this restriction.

Recall, however, that acceptability can be thought of as collective maximality. It should therefore not surprise us that acc does satisfy a collective variant of Sen's α.

Lemma 13. *Let $M = (S, \preceq, V)$ be any model. For every $Y \subseteq X \subseteq S$ and every $Z \subseteq Y$, if Z is miminal retentive in X then Z is minimal retentive in Y.*

Proof. Because Z is retentive in X, it is also retentive in $Y \subseteq X$. Suppose then, towards a contradiction, that Z is not minimally retentive in Y. Then there is a subset $U \subset Z$ that is retentive in Y. Because Z is minimal retentive in X, U cannot be retentive in X, so there are some $u \in U$ and $x \in X \setminus U$ such that $u \prec x$. Furthermore, because $U \subset Z$ we have $u \in Z$, so the retentiveness of Z in X together with $u \prec x$ imply that $x \in Z$.

But we have $Z \subseteq Y$, so this implies that $x \in Y$. So there are $u \in U$ and $x \in Y \setminus U$ such that $u \prec x$, contradicting the retentiveness of U in Y. From this contradiction, we conclude that Z is minimal retentive in Y. $\square$

So suppose that $y \in \mathsf{acc}(X)$. Then $y \in Y \subseteq X$ is not sufficient to conclude that $y \in \mathsf{acc}(Y)$. But if $Z \subseteq X$ is the minimal retentive that witnesses the acceptability of y, then $Z \subseteq Y \subseteq X$ does suffice to conclude that Z is minimal retentive in Y, and therefore that $y \in \mathsf{acc}(Y)$.

3. Language and Semantics

The models developed in the preceding section allow us to represent belief, preference or game-theoretic rationality. Here we introduce a logical language that allows us to reason about properties of such models, and thereby about the situations represented by those models.

In addition to the "bestness" operator $B(\varphi \mid \psi)$ this language also uses the standard Boolean operators, and a universal modality $\Box\varphi$ which holds if φ is true in every state.

Definition 14. Let $\mathbf{P}$ be a countable set of propositional atoms. The language $\mathcal{L}$ is given by the following normal form:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid B(\varphi \mid \varphi) \mid \Box\varphi,$$

where $p \in \mathbf{P}$.

The operators $\wedge, \rightarrow, \leftrightarrow, \top, \bot, \bigwedge$ and $\bigvee$ are defined as abbreviations in the usual way. Furthermore, $P(\varphi \mid \psi)$ abbreviates $\neg B(\neg\varphi \mid \psi)$ and $\Diamond$ abbreviates $\neg\Box\neg$. Finally, if $\Gamma \subseteq \mathcal{L}$ we write $\neg\Gamma$ and $\Box\Gamma$ for $\{\neg\gamma \mid \gamma \in \Gamma\}$ and $\{\Box\gamma \mid \gamma \in \Gamma\}$, respectively.

Note that $B(\varphi \mid \psi)$ and $P(\varphi \mid \psi)$ depend on which states are the best ψ-states. The truth of these formulas therefore depends on how we disambiguate "best". This means that we get four variants of our semantics. Based on the informal description of the meaning of $B(\varphi \mid \psi)$ and $P(\varphi \mid \psi)$ given above, the following definition should be entirely unsurprising.

Definition 15. Let $\beta \in \{\mathsf{max}, \mathsf{unm}, \mathsf{opt}, \mathsf{acc}\}$ and let $M = (S, \preceq, V)$ be a model. The satisfaction relation $\models_\beta$ is given recursively by

$$\begin{array}{lcl} M, s \models_\beta p & \Leftrightarrow & s \in V(p) \\ M, s \models_\beta \neg\varphi & \Leftrightarrow & M, s \not\models_\beta \varphi \\ M, s \models_\beta \varphi \vee \psi & \Leftrightarrow & M, s \models_\beta \varphi \text{ or } M, s \models_\beta \psi \\ M, s \models_\beta B(\varphi \mid \psi) & \Leftrightarrow & \forall s' \in \beta([\![\psi]\!]_\beta) : M, s' \models_\beta \varphi \\ M, s \models_\beta \Box\varphi & \Leftrightarrow & \forall s' \in S : M, s' \models_\beta \varphi \end{array}$$

where $[\![\psi]\!]_\beta := \{t \in S \mid M, t \models_\beta \psi\}$.

If $M, s \models_\beta \varphi$ for every model M and state s, we say that φ is *valid* with respect to β and write $\models_\beta \varphi$. If $\Gamma \subseteq \mathcal{L}$ and for every M, s such that $\forall \gamma \in \Gamma : M, s \models_\beta \gamma$ we also have $M, s \models_\beta \varphi$, then we write $\Gamma \models_\beta \varphi$. Where β is clear from context, we write $[\![\varphi]\!]$ for $[\![\varphi]\!]_\beta$ and $\models$ for $\models_\beta$.

We use superscript tr, to and as to indicate validity when restricted to models that are transitive, total and anti-symmetric, respectively. So for $\tau \subseteq \{tr, to, as\}$, we write $\Gamma \models_\beta^\tau \varphi$ if for every model M satisfying the indicated properties and every state s of that model, if $M, s \models_\beta \gamma$ for all $\gamma \in \Gamma$ then $M, s \models_\beta \varphi$.

Note that if $\beta([\![\psi]\!]_\beta)$ is empty, then $M, s \models B(\varphi \mid \psi)$ for every φ, including $\varphi = \bot$. This phenomenon is quite common in modal logic and first-order logic, and it is not very problematic from a technical point of view. We do, however, have to be somewhat careful with the informal interpretation of $B(\bot \mid \psi)$. While it is undeniably true that if $\beta([\![\psi]\!]_\beta) = \emptyset$ then any formula, including $\bot$, holds in every one of the best ψ-states, it would perhaps not be entirely fair to say that in this case the agent conditionally believes $\bot$, or is conditionally obliged to bring about $\bot$. Instead, it may be better to consider $B(\bot \mid \psi)$ as a warning that specific recommendations about what to do or to believe cannot be drawn from the preference relation. Pragmatic considerations will then need to be used instead. For example, if $\preceq$ represents utility comparisons and $\beta([\![\psi]\!]) = \emptyset$ then one would not try to bring about the best outcome (since there is no best outcome), but instead a good enough outcome. How to make such pragmatic decisions will differ greatly depending on the meaning of $\preceq$ and the context, so we consider it out of scope for the current contribution.

One further thing that we should remark upon is that nesting of the $\Box$ and $B(\cdot \mid \cdot)$ operators can be eliminated, i.e., any formula φ is equivalent to a formula φ_1 that does not contain nested $\Box$ and $B(\cdot \mid \cdot)$. This is because the relation $\preceq$ does not differ per state, so for every choice of φ, ψ and M we have either $M, s \models B(\varphi \mid \psi)$ for all $s \in S$ or $M, s \models \neg B(\varphi \mid \psi)$ for all $s \in S$.

This implies that either $[\![B(\varphi \mid \psi)]\!] = [\![\top]\!]$ or $[\![B(\varphi \mid \psi)]\!] = [\![\bot]\!]$. Because the operators are extensional, we can then substitute either $\top$ or $\bot$ for $B(\varphi \mid \psi)$. For example, $B(p \wedge B(\varphi \mid \psi) \mid \chi)$ is equivalent to

$$(B(\varphi \mid \psi) \to B(p \wedge \top \mid \chi)) \wedge (\neg B(\varphi \mid \psi) \to B(p \wedge \bot \mid \chi)).$$

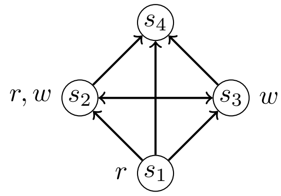

Figure 4: The model M_A for Example 16.

Another way of thinking of this elimination is that a single relation $\preceq$ cannot meaningfully represent meta-preferences, so it is unsurprising that nested $B(\cdot \mid \cdot)$ operators are superfluous. If, however, we were to use state-dependent preferences, as discussed in Section 2, it would be possible to model meta-preferences, so nested $B(\cdot \mid \cdot)$ operators could not be eliminated. An example of such meta-preferences is given as Example 19.

4. Examples

Now that we have defined both models and language, we can consider a few examples. We begin with a simple example of conditional belief.

Example 16. Alice is currently inside, in a position where she cannot directly observe the outside. She is reasoning about whether it is currently raining (r) and whether the street is wet (w), and considers four possible states. These states, and Alice's plausibility order among them, are shown as the model M_A in Figure 4. Note that Alice considers s_2 and s_3 to be equally plausible, so if the street is wet she considers the situation where it is currently raining to be equally plausible as the situation where it is not currently raining.

In any of the four semantics, Alice considers $\neg r \wedge \neg w$ plausible. Furthermore, she considers nothing else plausible, so she also believes $\neg r \wedge \neg w$, so we have $M_A \models P(\neg r \wedge \neg w \mid \top) \wedge B(\neg r \wedge \neg w \mid \top)$. In other words, Alice believes that the ground is not wet and that it is not raining. Her conditional beliefs do depend on the exact semantics that we use, however.

Consider $[\![w]\!] = \{s_2, s_3\}$. Both s_2 and s_3 are acceptable, maximal and optimal in $\{s_2, s_3\}$. As such, for $\beta \in \{\mathsf{max}, \mathsf{opt}, \mathsf{acc}\}$ we have $M_A \models_\beta P(r \mid w) \wedge P(\neg r \mid w)$. So in any of those three semantics, Alice considers both rain (r) and lack of rain ($\neg r$) to be plausible given that the street is wet (w).

Neither s_2 nor s_3 is unmatched in $\{s_2, s_3\}$, however. So under these semantics Alice considers nothing plausible. In particular, $M_A \models_{\mathsf{unm}} \neg P(r \mid w) \wedge \neg P(\neg r \mid w)$.

It seems reasonable to say that both r and $\neg r$ are, in fact, plausible given w. So unm is not the appropriate semantics in this case. More generally, it seems that unm semantics is rarely the correct choice if $\preceq$ represents a plausibility ordering, although it is sometimes the appropriate semantics for different interpretations of $\preceq$. The other three semantics are more suitable when $\preceq$ represents

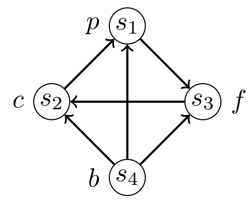

Figure 5: The model M_B for Example 17.

belief, and correspond to different standards of evidence for then something is to be considered plausible. With opt semantics, few things are considered plausible, and therefore many things are believed. In contrast, acc semantics result in many things being considered plausible and few things being believed. Finally, max semantics lies between those extremes. The appropriate choice of semantics therefore depends on the relative risk of considering too many things plausible or too few things plausible. In this example the decision about what to believe does not seem especially weighty, so each of max, opt and acc would be a reasonable choice.

Next, let's consider an obligation example, with a social choice flavour.

Example 17. Among a group of friends a majority prefers pizza (p) over curry (c), curry over fries (f), and fries over pizza. The majority prefers all these alternatives to burgers (b). Bob is supposed to order food for the group, and has a (pretty weak) obligation to order the best food. These preferences are shown as the model M_B in Figure 5.

In max, opt and unm semantics, the cycle between p, c and f means that there is no best option. We then have $B(\bot \mid \top)$, which is best read as "there are no coherent obligations". In particular, with those semantics there is no obligation for Bob to choose p, c or f over b, despite the latter being reviled by everyone.

In acc semantics p, c and f are collectively the best option, so we have $P(p \mid \top)$, $P(c \mid \top)$ and $P(f \mid \top)$. So Bob can order any of those three foods. But he is not allowed to order burgers.

If it turns out that the pizza place is closed, an unequivocal best option appears under any of the four semantics: among c, f and b, a majority prefers c over both other options, so Bob should order curry: $P(c \mid \neg p) \wedge B(c \mid \neg p)$.

Finally, let us look at a strategy example. This example is based on the game *Hearthstone*, as it was in the spring of 2019.

Example 18. *Heartstone*[5] is a digital card game developed by Blizzard Entertainment. Here we look at the pre-match strategy in *Hearthstone*, which

[5] https://playhearthstone.com/en-gb

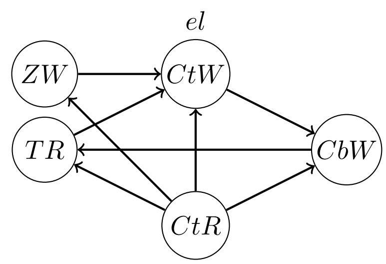

Figure 6: The model M_C for Example 18.

consists of choosing 30 virtual cards to form a deck.[6] Since players select 30 out of hundreds of possible cards, there are, in theory, over 10^{60} different decks. In practice, however, decks can be described by a combination of a playstyle (e.g., "combo", "zoo") and a class (e.g., "warrior", "rogue").

Here we consider a heavily simplified version of the game, and look at five such decks: "**C**on**t**rol **W**arrior" (CtW), "**C**om**b**o **W**arrior" (CbW), "**C**on**t**rol **R**ogue" (CtR), "**T**empo **R**ogue" (TR) and "**Z**oo Warlock" (ZL). In a match between different decks one is usually favoured over the other. For example, CtW usually defeats TR. There are also some cases where no side is favoured, however. For example, ZW is neither favoured nor disfavoured against CbW. See Figure 6 for the model M_C that shows the defeat relations between all five decks. Note that TR, CtW and CbW form a "rock-paper-scissor"-like cycle, where each strategy is defeated by the next one. Also note that CtR loses to each of the other strategies. Playing CtR is therefore a terrible idea, and no experienced player would do so other than as a joke. We include CtR in this example not because it is likely to be played, but because we should keep in mind that bad strategies do exist.

We treat the names of the strategies as atoms that hold only for that strategy. Furthermore, we use an atom el for "Elysiana", one of the cards used in the CtW deck.

No strategy is maximal, optimal or unmatched. The strategies TR, CtW and CbW are acceptable, however. So under **ACC** semantics these three strategies, and only these, are playable: $M_C, \models_{\mathsf{acc}} P(TR \mid \top) \wedge P(CtW \mid \top) \wedge P(CbW \mid \top)$ and $M_C \models_{\mathsf{acc}} B(TR \vee CtW \vee CbW \mid \top)$. The strategy ZW is not playable because it loses to CtW and ties against CbW and CtR. The strategy CtR is not playable because it is awful and loses against all other strategies.

A complicating factor is that some tournaments decided to forbid the use of Elysiana, since playing that card tends to make matches last too long. Consider therefore the restriction of M_C to $[\![\neg el]\!]$, shown in Figure 7. The reason CbW was previously playable was that it beats CtW. In the restricted strategy space

[6] There is also strategy involved in the moment to moment play *during* a match, but here we focus on the pre-match part.

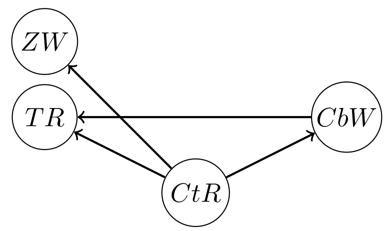

Figure 7: The restriction of M_C to $[\![\neg el]\!]$.

CtW cannot be played, so there is no longer any reason to play CbW, and it is no longer playable. The strategy TR, on the other hand, was playable and remains playable: removal of CtW only benefits TR. The strategy ZW was only not playable because it loses to CtW, so in the restricted strategy space it becomes playable. Finally, CtR was unplayable and remains so, because it loses to every other strategy.

In other words, we have $M_C \models_{\mathsf{acc}} \neg P(CbW \mid \neg el) \wedge \neg P(CtR \mid \neg el) \wedge P(TR \mid \neg el) \wedge P(ZW \mid \neg el)$. Note that we have at least one instance of each possible combination of unrestricted and restricted playability: TR was playable and remains so, CtR was and remains unplayable, ZW was unplayable and becomes playable, while CbW was playable and becomes unplayable.

Finally, we consider one example of meta-preferences, which we model by making the preference relation state-dependent.

Example 19. Dave likes sugary drinks but, for health reasons, would prefer not to like them. In order to represent this scenario, we use a model M_D with four states: s_{ld}, in which Dave likes sugary drinks and has one, s_l in which he likes sugary drinks but does not have one, s_d in which he has a sugary drink but does not like them and s_ϵ in which he neither likes sugary drinks nor has one. We assume that Dave does not currently have a drink, so s_l is the actual state.

Whether Dave has a drink is represented by the atom d, which holds in s_{ld} and s_d but not in s_l and s_ϵ. Whether Dave likes sugary drinks is not represented by an atom, but instead by the preference relation associated with a given state. This means that we need to define four relations: $\preceq_{s_{ld}}$, $\preceq_{s_l}$, $\preceq_{s_d}$ and $\preceq_{s_\epsilon}$. We have not yet fully specified Dave's preferences, so we make the following additional assumptions.

- Dave's preferences do not change depending on whether he has a drink, so $\preceq_{s_{ld}}$ is the same as $\preceq_{s_l}$ and $\preceq_{s_d}$ is the same as $\preceq_{s_\epsilon}$.
- While Dave's preference for drinks differs between s_{ld} and s_l on the one hand and s_d and s_ϵ on the other, his preference for preferring not to like sugary drinks is constant.

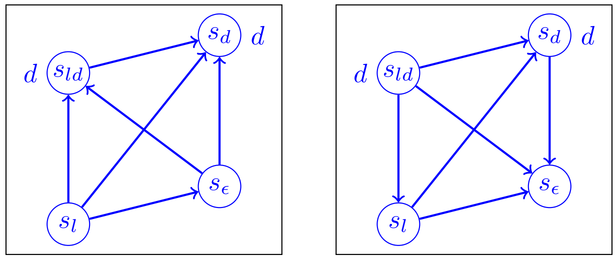

(a) Relation $\preceq_{s_l}$ (b) Relation $\preceq_{s_\epsilon}$

Figure 8: The model M_D for Example 19.

- In s_{ld} and s_l, Dave's preference for a sugary drink outweighs his preference for not liking them. So if he must choose between s_{ld} (having a drink and liking it) or s_ϵ (not having a drink and not liking them) he prefers s_{ld}.
- In s_d and s_ϵ, Dave's preference for not liking sugary drinks outweighs his preference for not having them. So if he must choose between s_d (having a drink despite not liking it) or s_l (not having a drink but wanting one), he prefers s_d.

We then have $s_l \prec_{s_l} s_{ld}$, $s_\epsilon \prec_{s_l} s_d$, $s_l \prec_{s_l} s_d$ and $s_\epsilon \prec_{s_l} s_{ld}$, because, in s_l, Dave prefers having a drink over not having one. Furthermore, we have $s_l \prec_{s_l} s_\epsilon$ and $s_{ld} \prec_{s_l} s_d$, because Dave prefers not to like sugary drinks. These preferences are shown in Figure 8a.

Similarly, $s_{ld} \prec_{s_\epsilon} s_l$ and $s_\epsilon \prec_{s_\epsilon} s_d$, because, in s_ϵ, Dave prefers not to have a drink. Furthermore, $s_l \prec_{s_\epsilon} s_\epsilon$, $s_l \prec_{s_\epsilon} s_d$, $s_{ld} \prec_{s_\epsilon} s_\epsilon$ and $s_{ld} \prec_{s_\epsilon} s_d$, because Dave likes that he does not like sugary drinks. These preferences are shown in Figure 8b.

The relations $\preceq_{s_l}$ and $\preceq_{s_\epsilon}$, and therefore also $\preceq_{s_{ld}}$ and $\preceq_{s_d}$, are transitive, total and anti-symmetric. So max, opt, unm and acc coincide on this model. The unique best state according to $\preceq_{s_d}$ is s_ϵ, and $s_\epsilon \in [\![\neg d]\!]$. We therefore have $M_D, s_d \models B(\neg d \mid \top)$. The unique best state according to $\preceq_{s_l}$ is s_d. We have $s_d \in [\![d]\!]$, so $M_D, s_l \models B(d \mid \top)$. Furthermore, we already established that $s_d \in [\![B(\neg d \mid \top)]\!]$, so we have $M_D, s_l \models B(B(\neg d \mid \top) \mid \top)$.

This shows that $M_D, s_l \models B(d \mid \top) \wedge B(B(\neg d \mid \top) \mid \top)$, as desired: in the actual state s_l, Dave likes sugary drinks but would prefer not to like them. This is entirely unsurprising, of course, since we constructed the model for that exact purpose. But it does illustrate how state-dependent relations can be used to model meta-preferences.

5. Shared axioms

The main technical contribution of this paper is the introduction of sound and strongly complete axiomatizations for $\mathcal{L}$ with respect to each of the four semantics, both on all models and on models satisfying any combination of transitivity, anti-symmetry and totality.

In total, we consider four types of semantics and 8 different combinations of transitivity, totality and anti-symmetry, for a total of 32 variants. For 31 of these variants we introduce an axiomatization and prove it to be sound and strongly complete. For the one remaining variant, acc semantics on total anti-symmetric models, we prove that no such axiomatization exists. Fortunately, while we provide axiomatizations for 31 variants, we do not need 31 different axiomatizations, since most axiomatizations are sound and complete for multiple variants.

As shown in Proposition 8, under certain assumptions the notions of maximality, optimality, unmatchedness and acceptability coincide. It follows immediately that the semantics based on these notions are, under those assumptions, equivalent. In particular, they then have the same axiomatizations. For example, on transitive models $\mathsf{max}(X) = \mathsf{acc}(X)$, so any axiomatization that is sound and complete for max semantics on transitive models will also be sound and complete for acc semantics on transitive models. Likewise, since all four semantics coincide on models that are transitive, total and anti-symmetric, they all have the same axiomatization on such models.

Some other variants are not equivalent, but equisatisfiable. For example, if Γ is a set of formulas and $M, s \models_{\mathsf{max}} \Gamma$ then we generally do not have $M, w \models_{\mathsf{unm}} \Gamma$, but there are M' and s' such that $M', s' \models_{\mathsf{unm}} \Gamma$, see Theorem 37. As such, max and unm have the same axiomatization for unconstrained models. Similarly, due to the duality between opt and unm and the duality between totality and anti-symmetry, see Lemma 10, opt semantics on total models is equisatisfiable to unm semantics on anti-symmetric models, and vice versa.

Overall, with a single assumption (transitivity, totality or anti-symmetry) two semantics are equivalent, with two assumptions three semantics are equivalent and with three assumptions all four semantics are equivalent. This results in 20 non-equivalent variants. One of those variants cannot be axiomatized, and because of equi-satisfiability the remaining 19 variants require only 8 different axiomatizations.

All of these axiomatizations consist of the same base set of axioms, plus a few more axioms associated with the exact semantics and constraints. In this section, we introduce these shared axioms, and then provide a number of definitions and lemmas that we will rely on later when we introduce the full axiomatizations for each variant.

The base set of axioms is similar to existing axiomatizations for preference logics, such as the ones given in [2, 26]. Some of the axioms from those proof systems are unsound in our context, however, since they depend on smoothness or transitivity assumptions. For obvious reasons we do not include these unsound axioms in our axiomatizations.

Definition 20. The axiomatization **BASE** contains the following rule and axioms:

PL	Substitution instances of propositional validities
□S5	the **S5** axioms and necessitation for □
R-Ext	$\Box(\varphi \leftrightarrow \psi) \to (B(\chi \mid \varphi) \to B(\chi \mid \psi))$
L-Ext	$\Box(\varphi \leftrightarrow \psi) \to (B(\varphi \mid \chi) \to B(\psi \mid \chi))$
MP	from φ and $\varphi \to \psi$ infer ψ
Abs	$B(\varphi \mid \psi) \leftrightarrow \Box B(\varphi \mid \psi)$
K	$B(\varphi \to \psi \mid \chi) \to (B(\varphi \mid \chi) \to B(\psi \mid \chi))$
Id	$B(\varphi \mid \varphi)$
Triv	$B(\top \mid \varphi)$

For any proof system **X** extending **BASE**, we say that formula φ is *derivable* in **X** from a set Γ of premises, denoted $\Gamma \vdash_{\mathbf{X}} \varphi$, if there is a finite list of formulas such that (1) φ occurs in the list and (2) every formula in the list is a premise, and instance of an axiom of **X** or derived from formulas occurring earlier in the list using a rule of **X**. If $\emptyset \vdash_{\mathbf{X}} \varphi$ we say that φ is *derivable* in **X** and write $\vdash_{\mathbf{X}} \varphi$.

Broadly speaking, one can think of **PL** and **MP** as representing the fact that $\mathcal{L}$ is based on propositional logic and **□S5** as representing the fact that □ is a universal modality. The axioms **R-Ext** and **L-Ext** represent the fact that $B(\cdot \mid \cdot)$ is extensional.

Abs holds because a model has only one preference relation $\preceq$. This axiom would therefore have to be dropped when using state-dependent preferences, as discussed in Section 3.

The axiom **K** represents the fact that $B(\cdot \mid \cdot)$ is a "box-like" modality, in the sense that $B(\varphi \mid \psi)$ holds if φ is true in *all* of the best ψ-states. **Id** and **Triv**, meanwhile, represent our use of relativization; obviously all of the best φ states are, in particular, φ states and $\top$ states so we have $B(\varphi \mid \varphi)$ and $B(\top \mid \varphi)$.

The axiom **Triv** is unnecessary in most of our axiom systems, in the sense that it can be derived from the other axioms and rules. We do need it in a few of the weaker systems, however, and even where it is derivable it still simplifies proofs to have **Triv** as an axiom.[7]

Given the above explanations we consider the soundness of the **BASE** axioms to be obvious, so we state it without proof.

Lemma 21 (Soundness of **BASE**). *If* $\Gamma \vdash_{\mathbf{BASE}} \varphi$, *then* $\Gamma \models_\beta \varphi$ *for every* $\beta \in \{\mathsf{max}, \mathsf{opt}, \mathsf{unm}, \mathsf{acc}\}$.

By itself, **BASE** is not complete for any of four the semantics (see Proposition 42). But because the complete axiomatizations all extend **BASE**, it is convenient to define a number of notions, such as maximal consistent sets, for any **X** extenting **BASE**.

[7] Specifically, the derivation of **Triv** uses the axiom **Sh**, which we introduce in Definition 32. This axiom is not included in the proof systems **BASE** and **ACC** (see Definition 40).

5.1. Maximal consistent sets

Definition 22. MCS Let **X** be any of the proof systems extending **BASE**. A set $\Gamma \subseteq \mathcal{L}$ is **X**-*consistent* if $\Gamma \not\vdash_{\mathbf{X}} \bot$. It is *maximal* **X**-*consistent* if it is **X**-consistent and, furthermore, for every $\varphi \in \mathcal{L}$ either $\varphi \in \Gamma$ or $\neg\varphi \in \Gamma$. The set of maximal **X**-consistent sets is denoted $MCS^{\mathbf{X}}$. When the proof system **X** is clear from context we omit reference to it.

As usual, the states in our canonical models will be based on maximal consistent sets. We cannot simply take the set of states to be equal to $MCS^{\mathbf{X}}$, however, for two reasons. Firstly, due to the presence of the universal operator $\Box$, not all consistent sets can occur in the same model. For example, $\{\Box p\}$ and $\{\Box\neg p\}$ are both consistent, but one requires p to hold in every state of the model and one requires $\neg p$ to hold, so they cannot occur in the same model. We will therefore need to relativize our models to some set Ξ of formulas that are universally true across the model. The set of states in the canonical model will then be based not on all maximal consistent sets, but only those that contain $\Box\Xi$.

Secondly, most of our canonical models will use multiple copies of each MCS. So instead of having states of the form Γ, states in those models will be of the form $(\Gamma, \vec{\lambda})$, where $\vec{\lambda}$ is a vector of indices. We say that a state of $(\Gamma, \vec{\lambda})$ is *based on* Γ. The set Γ takes the usual role of a maximal consistent set in a canonical model, i.e., φ will be true in the state $(\Gamma, \vec{\lambda})$ if and only if $\varphi \in \Gamma$. For ease of notation we write $[\varphi]$ for $\{(\Gamma, \vec{\lambda}) \mid \varphi \in \Gamma\}$, so the truth lemma for a given canonical model will show that $[\varphi] = [\![\varphi]\!]$.

Definition 23 ($\Box$-inverse)**.** Let $\Gamma \subseteq \mathcal{L}$ be a set of formulas. The $\Box$-*inverse* of Γ, denoted $\Box^{-1}\Gamma$ is given by $\Box^{-1}\Gamma = \{\varphi \mid \Box\varphi \in \Gamma\}$.

A set $\Gamma \subseteq \mathcal{L}$ is $\Box$-*maximal* if it is the $\Box$-inverse of a maximal consistent set.

Definition 24. Let $\Xi \subseteq \mathcal{L}$. Then $MCS^{\mathbf{X}}_{\Xi} = \{\Gamma \in MCS^{\mathbf{X}} \mid \Box^{-1}\Gamma = \Xi\}$.

As usual, consistent sets can be extended to maximal consistent sets. The proof of the lemma is completely standard, so we omit it.

Lemma 25 (Lindenbaum lemma)**.** *If* $\Gamma \subseteq \mathcal{L}$ *is consistent, then there is a maximal consistent* Δ *such that* $\Gamma \subseteq \Delta$.

Where the set of states of the canonical model is based on the set of maximal consistent sets, the preference relation between the states is based on the B-inverse of these MCS.

Definition 26 (B-inverse)**.** Let $\Gamma \subseteq \mathcal{L}$. The B-*inverse* of Γ with respect to ψ, denoted $B^{-1}_{\psi}\Gamma$, is given by $B^{-1}_{\psi}\Gamma = \{\varphi \mid B(\varphi \mid \psi) \in \Gamma\}$.

When Γ is clear from context, we omit it and write B^{-1}_{ψ} for $B^{-1}_{\psi}\Gamma$. In particular, because **BASE** contains the axioms **Abs**, for every $\Box$-maximal Ξ

and every $\Gamma, \Delta \in MCS^{\mathbf{X}}_{\Xi}$ we have

$$\begin{aligned} B(\varphi \mid \psi) \in \Gamma &\Leftrightarrow \Box B(\varphi \mid \psi) \in \Gamma \\ &\Leftrightarrow B(\varphi \mid \psi) \in \Xi \\ &\Leftrightarrow \Box B(\varphi \mid \psi) \in \Delta \\ &\Leftrightarrow B(\varphi \mid \psi) \in \Delta, \end{aligned}$$

which implies that $B^{-1}_{\psi}\Gamma = B^{-1}_{\psi}\Delta$, so we can unambiguously write B^{-1}_{ψ} for either.

5.2. Support and "good" states

The exact way in which the preference relation of the canonical model depends on B-inverses differs per axiomatization. Here, however, we introduce two notions that will be important for each canonical model. These are *support* and *"goodness"*.

Definition 27 (Good MCS)**.** Let $\Gamma \subseteq \mathcal{L}$ be maximal consistent. We say that Γ is φ-*good* if $B^{-1}_{\varphi} \subseteq \Gamma$. Γ is *good* if there is a $\varphi \in \mathcal{L}$ such that Γ is φ-good. If Γ is not φ-good then it is φ-*bad*, and if it is not good then it is *bad*.

A φ-good MCS is one that is a candidate for being in $\beta(\llbracket\varphi\rrbracket)$. After all, every $\psi \in B^{-1}_{\varphi}$ holds on every state $s \in \beta(\llbracket\varphi\rrbracket)$. If Γ is φ-good, then at least one of the states based on Γ will be in $\beta(\llbracket\varphi\rrbracket)$. Next, we have the notion of *support*.

Definition 28 (Support)**.** Let $\Gamma, \Delta \subseteq \mathcal{L}$. We say that Γ ψ-*supports* Δ, denoted $\Gamma \overset{\psi}{\rightsquigarrow} \Delta$ if $\psi \in \Gamma$ and $B^{-1}_{\psi} \subseteq \Delta$. If there is a ψ such that Γ ψ-supports Δ we say that Γ *supports* Δ, denoted $\Gamma \rightsquigarrow \Delta$. We take $supp(\Gamma) = \{\Delta \mid \Gamma \rightsquigarrow \Delta\}$.

If Γ ψ-supports Δ then, in particular, Δ is ψ-good. So there will be some state s based on Δ that should end up in $\beta(\llbracket\psi\rrbracket)$. Furthermore, it also implies that $\psi \in \Gamma$, so every state t based on Γ will be in $\llbracket\psi\rrbracket$. In other words, if Γ supports Δ then there is some extension of which t is a member in which s is best. Obviously, this has consequences for the relation between s and t. For example, if $\beta = \mathsf{opt}$ and $\Gamma \rightsquigarrow \Delta$ then there must be some s based on Δ such that for every t based on Γ we have $t \preceq s$.

Every canonical model will need to satisfy such a property connecting $\preceq$ to $\rightsquigarrow$. For some models we will ensure that this is the case by basing the definition of $\preceq$ directly on $\rightsquigarrow$. Other models will not directly use $\rightsquigarrow$ in their definition, but these models still satisfy the same property.

5.3. A few lemmas

Before we continue to define the axiomatizations extending **BASE**, we should consider a few more auxiliary lemmas that will be useful throughout.

Lemma 29. *If Γ is a maximal consistent set and $\Box\xi \in \Gamma$, then for every ψ we have $B(\xi \mid \psi) \in \Gamma$ and $B(\Box\xi \mid \psi) \in \Gamma$.*

Proof. We first show that $B(\xi \mid \psi) \in \Gamma$. We have $\vdash \xi \to (\top \leftrightarrow \xi)$, because of **PL**. By $\Box$-necessitation (included in $\Box$**S5**), this implies $\vdash \Box(\xi \to (\top \leftrightarrow \xi))$. Another use of $\Box$**S5** then yields $\vdash \Box\xi \to \Box(\xi \leftrightarrow \top)$.

By assumption, $\Box\xi \in \Gamma$. As Γ is maximal consistent, we therefore have $\Box(\xi \leftrightarrow \top) \in \Gamma$. Because of **Triv**, we also have $B(\top \mid \psi) \in \Gamma$. Using **L-Ext**, we can therefore conclude that $B(\xi \mid \psi) \in \Gamma$.

The second claim from the lemma, that $B(\Box\xi \mid \psi) \in \Gamma$, follows simply from the fact that, by $\Box$**S5**, $\Box\xi \in \Gamma$ implies $\Box\Box\xi \in \Gamma$, at which point $B(\Box\xi \mid \psi) \in \Gamma$ follows from the first claim in the lemma. $\Box$

Lemma 30. *If $\Gamma \in MCS_\Xi$ and $\varphi \notin B_\psi^{-1}(\Gamma)$, then the set*

$$\Delta = \{\neg\varphi\} \cup B_\psi^{-1}(\Gamma) \cup \Box\Xi \cup \{\neg\Box\zeta \mid \zeta \in \mathcal{L} \setminus \Xi\}$$

is consistent.

Proof. Suppose towards a contradiction that Δ is inconsistent, so $\Delta \vdash \bot$. Because every proof is, by definition, finite, this implies that there is a finite subset $\Delta' \subseteq \Delta$ such that $\Delta' \vdash \bot$. This implies that there are $\xi_1, \cdots, \xi_k \in \Xi$, $\zeta_1, \cdots, \zeta_n \in \mathcal{L} \setminus \Xi$ and $\delta_1, \cdots, \delta_m \in B_\psi^{-1}(\Gamma)$ such that $\vdash (\Box\xi_1 \wedge \cdots \wedge \Box\xi_k \wedge \neg\Box\zeta_1 \wedge \cdots \wedge \neg\Box\zeta_n \wedge \delta_1 \wedge \cdots \wedge \delta_m \wedge \neg\varphi) \to \bot$, which can be rewritten (using **PL**) as $\chi := \Box\xi_1 \to (\Box\xi_2 \to (\cdots \to (\delta_m \to \varphi)))$. Finally, using $\Box$**S5** this implies that $\vdash \Box\chi$.

Since Γ is maximal consistent, we then have $\Box\chi \in \Gamma$. By Lemma 29 this implies $B(\chi \mid \psi) \in \Gamma$. Furthermore, for every $1 \leq i \leq k$ we have $\Box\xi_i \in \Gamma$ and therefore, by Lemma 29, $B(\Box\xi_i \mid \psi) \in \Gamma$. For every $1 \leq i \leq n$ we have $\neg\Box\zeta_i \in \Gamma$ and therefore, because $\Box$ is S5, $\Box\neg\Box\zeta_i \in \Gamma$. By Lemma 29 this gives us $B(\neg\Box\zeta_i \mid \psi) \in \Gamma$. Finally, for every $1 \leq j \leq m$ we have $B(\delta_i \mid \psi) \in \Gamma$, since $\delta_i \in B_\psi^{-1}(\Gamma)$. Repeated application of **K** therefore yields $B(\varphi \mid \psi) \in \Gamma$. But that contradicts the assumption that $\varphi \notin B_\psi^{-1}(\Gamma)$.

From this contradiction, we conclude that Δ must be consistent, which was to be shown. $\Box$

The important thing to note is that, by construction, we have $\Box\xi \in \Delta$ for every $\xi \in \Xi$ and $\neg\Box\zeta \in \Delta$ for every $\zeta \notin \Xi$. This implies that for every maximal consistent set Δ' extending Δ, we have $\Box^{-1}\Delta' = \Xi$ and therefore $\Delta' \in MCS_\Xi$. So if $\varphi \notin B_\psi^{-1}(\Gamma)$, then not only can we find a maximal consistent set where φ is false, this set can be chosen inside MCS_Ξ.

We consider one more lemma. It follows immediately from Definition 24 that if $\xi \in \Xi$ then $\Box\xi \in \Gamma$ for each $\Gamma \in MCS_\Xi$. Because $\Box$ is an S5 operator and Γ is an MCS, this implies that $\xi \in \Gamma$. The following lemma states that the converse holds as well, which we need for the $\Box$-step in the truth lemma.

Lemma 31. *If $\varphi \in \Gamma$ for all $\Gamma \in MCS_\Xi$, then $\varphi \in \Xi$.*

Proof. From $\varphi \in \Gamma$ for all $\Gamma \in MCS_\Xi$ it follows that $\neg\varphi$ is not consistent with $\Box\Xi$. For some finite subset Δ of Ξ we therefore have $\vdash \bigwedge \Box\Delta \to \varphi$. Using $\Box$**S5** we obtain $\vdash \bigwedge \Box\Delta \to \Box\varphi$. Because Γ contains $\Box\Delta$ and is maximal consistent we have $\Box\varphi \in \Gamma$ and therefore $\varphi \in \Xi$. $\Box$

6. Axiomatizations for unconstrained preference relations

In this section, we introduce axiomatizations for all four semantics (max, opt, unm and acc) on models with unconstrained preference relation. In fact, it will turn out that we need only two axiomatizations: **MOU** which is sound and complete for max, opt and unm, and **ACC** which is sound and complete for acc.

For reasons of brevity we will not include the full soundness and completeness proofs for both axiomatizations in the main text. As **ACC** has, in our opinion, the more interesting proof, we include that in the main text and leave most of the details of the **MOU** proof for the appendix.

6.1. An axiomatization for max*,* opt *and* unm

In order to create a complete axiomatization for max, opt and unm semantics, we need to add one additional axiom to **BASE**. The axiom that we add is known as *Shoham's Rule*, so we denote it by **Sh**.

Definition 32. The axiomatization **MOU** consists of the rules and axioms of **BASE** plus the axiom

$$\textbf{Sh} \quad B(\varphi \mid \psi \wedge \chi) \rightarrow B(\psi \rightarrow \varphi \mid \chi)$$

We first show soundness and completeness of **MOU** for opt semantics. Soundness of **Sh** for opt follows from Lemma 12, which we recall states that if $Y \subseteq X$, $x \in Y$ and $x \in \mathsf{opt}(X)$ then $x \in \mathsf{opt}(Y)$.

Lemma 33 (Soundness of **MOU** for opt)**.** *If* $\Gamma \vdash_{\mathbf{MOU}} \varphi$ *then* $\Gamma \models_{\mathsf{opt}} \varphi$*.*

Proof. **BASE** is sound for every semantics (Lemma 21), so it suffices to show that **Sh** is also sound.

Take any model $M = (S, \preceq, V)$ and suppose that $M, s \not\models B(\psi \rightarrow \varphi \mid \chi)$. Then there is some $t \in \mathsf{opt}([\![\chi]\!])$ such that $M, t \not\models \psi \rightarrow \varphi$ and therefore $M, t \models \psi$ and $M, t \not\models \varphi$. This implies that $t \in [\![\psi \wedge \chi]\!]$. By Lemma 12 it follows that $t \in \mathsf{opt}([\![\psi \wedge \chi]\!])$. Because $M, t \not\models \varphi$ this implies that $M, s \not\models B(\varphi \mid \psi \wedge \chi)$.

We have shown that $M, s \models \neg B(\psi \rightarrow \varphi \mid \chi) \rightarrow \neg B(\varphi \mid \psi \wedge \chi)$ and therefore, by contraposition, $M, s \models B(\varphi \mid \psi \wedge \chi) \rightarrow B(\psi \rightarrow \varphi \mid \chi)$. □

For completeness, we define a canonical model. The states of this model are of the form (x, φ) where x is a maximal consistent set and φ is an "intended relativization". This means that if x is φ-good, then the copy (x, φ) is generally the one that will be optimal in $[\![\varphi]\!]$. For ease of notation, if $s = (x, \varphi)$ we write $\psi \in s$ for $\psi \in x$.

Definition 34. Let Ξ be $\Box$-maximal. Then $M^{\mathbf{MOU}}_{\Xi} = (S, \preceq, V)$ is given by

- $S = \{(x, \varphi) \mid x \in MCS^{\mathbf{MOU}}_{\Xi}, \varphi \in \mathcal{L}\}$,
- $V(p) = \{(x, \varphi) \in S \mid p \in x\}$,
- $\preceq \subseteq S \times S$ is the largest relation such that for all $\varphi, \psi \in \mathcal{L}$ and all $x, y \in MCS^{\mathbf{MOU}}_{\Xi}$,

– if x is φ-bad, then $(y, \psi) \not\preceq (x, \varphi)$,

– if $\psi \in x$ and $\varphi \notin y$ then $(y, \psi) \not\preceq (x, \varphi)$.

We can now state the truth lemma. We present only a proof sketch here, see the appendix for a detailed proof.

Lemma 35 (Truth lemma for **MOU** and opt)**.** *For every $\Box$-maximal Ξ, every $\varphi, \psi \in \mathcal{L}$ and every $x \in MCS^{\mathbf{MOU}}_{\Xi}$ we have $M^{\mathbf{MOU}}_{\Xi}, (x, \psi) \models_{\mathsf{opt}} \varphi$ if and only if $\varphi \in x$.*

Proof sketch. Recall that $[\delta] = \{(y, \varphi) \mid \delta \in y\}$. The main step in the proof is to show that (i) if $(y, \chi) \in \mathsf{opt}([\delta])$ then y is δ-good and (ii) if y is δ-good, then $(y, \delta) \in \mathsf{opt}([\delta])$.

For (i), first note that if $(y, \chi) \in \mathsf{opt}([\delta])$ then y is χ-good, since otherwise there would be no arrows towards (y, χ) due to the first exception clause in the definition of $\preceq$, and hence (y, χ) would not be optimal. Furthermore, we have $[\delta] \subseteq [\chi]$, since otherwise there would be some z such that $\delta \in z$ and $\chi \notin z$. By the second exception clause for $\preceq$ we would then have $(z, \delta) \not\preceq (y, \chi)$, contradicting the optimality of (y, χ). We can then use **R-Ext** and **Sh** to show that for every γ, if $\gamma \in B^{-1}_{\delta}$ then $\gamma \in B^{-1}_{\chi}$ (see appendix for details). From y being χ-good it therefore follows that y is δ-good.

For (ii), it suffices to note that in $M^{\mathbf{MOU}}_{\Xi}$, we have $(z, \epsilon) \preceq (y, \delta)$ unless one of two exceptions apply: either y is δ-bad, or $\epsilon \in y$ and $\delta \notin z$. The first possibility is excluded by y being δ-good. The second possibility only applies for $(z, \epsilon) \notin [\delta]$. So $(y, \delta) \in \mathsf{opt}([\delta])$. $\Box$

Now that we have a truth lemma, the rest of the completeness proof is entirely standard, so we omit it. Since we already proved soundness (Lemma 33), we have the desired soundness and completeness result.

Theorem 36. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models_{\mathsf{opt}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU}} \varphi$.*

We now have a sound and complete axiomatization for $\models_{\mathsf{opt}}$. Left to do is prove that **MOU** is sound and complete for $\models_{\mathsf{max}}$ and $\models_{\mathsf{unm}}$ as well. We do this not by creating further canonical models, but by showing that any model M_1 satisfying a set of formulas in one type of semantics can be transformed into a model M_2 satisfying the same formulas under different semantics.

Theorem 37. *Let $\beta_1, \beta_2 \in \{\mathsf{max}, \mathsf{opt}, \mathsf{unm}\}$. For every model $M_1 = (S_1, \preceq_1, V_1)$ and $s_1 \in S_1$ there are a model $M_2 = (S_2, \preceq_2, V_2)$ and $s_2 \in S_2$ such that for every $\varphi \in \mathcal{L}$*

$$M_1, s_1 \models_{\beta_1} \varphi \Leftrightarrow M_2, s_2 \models_{\beta_2} \varphi.$$

Proof sketch. If $\beta_1 = \mathsf{max}$, this follows from the fact that max semantics only depend on the strict relation $\prec$ (Lemma 11), and that max coincides with opt/unm on total/anti-symmetric models (Proposition 8).

If $\beta_1 = \mathsf{unm}$ and $\beta_2 = \mathsf{max}$, then for every $s \in S_1$ we create multiple copies $s^0, s^1, s^2 \in S_2$. We then make sure that if $s \approx_1 t$, then for every copy s^i of s

there is some copy t^j of t such that $s^i \prec_2 t^j$. This can be done, for example, by taking $s^0 \prec_2 t^1 \prec_2 s^2 \prec_2 t^0 \prec_2 s^1 \prec_2 t^2 \prec_2 s^0$.

If $s \prec_1 t$ then we take $s^i \prec_2 t^j$ for all i and j, and if $s \perp t$ then we take $s^i \approx t^j$. for all i and j. Finally, we take $s^i \approx_2 s^j$ for all i and j.

The model M_2 is constructed in such a way that $s \preceq_1 t$ if and only if for every i there is a j such that $s^i \prec_2 t^j$. As a result, it is relatively easy to verify that for any $X \subseteq S$, s is unmatched in X if and only if s^i is maximal in $X \times \{0,1,2\}$. It follows that $M_1, s \models_{\mathsf{unm}} \varphi$ iff $M_2, s^i \models_{\mathsf{max}} \varphi$.

For $\beta_1 = \mathsf{opt}$ and $\beta_2 = \mathsf{max}$ we use a similar construction. Finally, if $\beta_1 = \mathsf{opt}$ and $\beta_2 = \mathsf{unm}$ or vice versa, we go through max. For full details of the proof, see the appendix. □

Note that this implies that a formula is β_1-satisfiable if and only if it is β_2-satisfiable, but that the satisfying models may be different.

Corollary 38. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models_{\mathsf{max}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU}} \varphi$, and $\Gamma \models_{\mathsf{unm}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU}} \varphi$.*

The acc semantics do not, however, have the same set of validities. We will therefore need to introduce a different axiomatization for acc.

6.2. *An axiomatization for* acc

In this subsection we will introduce an axiomatization **ACC** that is sound and complete for acc semantics. This axiomatization does not contain the **Sh** axiom, as that is unsound for acc. This is a consequence of Lemma 12 not holding for acc.

Proposition 39 (Unsoundness of **Sh** for acc)**.** *There are M, s and φ, ψ, χ such that $M, s \not\models_{\mathsf{acc}} B(\varphi \mid \psi \wedge \chi) \to B(\psi \to \varphi \mid \chi)$.*

Proof. Consider the following model:

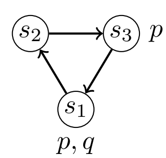

We have $\mathsf{acc}([\![p \wedge \top]\!]) = \{s_1\}$, and therefore $M \models_{\mathsf{acc}} B(q \mid p \wedge \top)$. Yet $\mathsf{acc}([\![\top]\!]) = \{s_1, s_2, s_3\}$, and therefore $M \not\models_{\mathsf{acc}} B(p \to q \mid \top)$. Hence $M \not\models_{\mathsf{acc}} B(q \mid p \wedge \top) \to B(p \to q \mid \top)$. □

Simply removing **Sh** would result in an incomplete axiomatization, however (see Proposition 42). Instead, we must replace it by a weaker axiom. The axiom that we will use is **Cut**:

$$\textbf{Cut} \quad (B(\varphi \mid \psi) \wedge B(\chi \mid \varphi \wedge \psi)) \to B(\chi \mid \psi)$$

As the name suggests, **Cut** is related to the cut-rule which is often used in sequent calculi. This is most clearly visible if we remember that $B(\varphi \mid \psi)$ is a type of conditional, so we could write alternative write it as $\psi \Rightarrow \varphi$. The **Cut** axiom then states that if $\psi \Rightarrow \varphi$ and $\varphi \wedge \psi \Rightarrow \chi$ then $\psi \Rightarrow \chi$.

The axiomatization **ACC** is obtained by adding **Cut** to **BASE**

Definition 40. The axiomatization **ACC** consists of the rules and axioms of **BASE** plus the axiom **Cut**.

Lemma 41 (Soundness of **ACC**). *If* $\Gamma \vdash_{\mathbf{ACC}} \varphi$ *then* $\Gamma \models_{\mathsf{acc}} \varphi$.

Proof. Because **BASE** is sound for all semantics (Lemma 21), it suffices to show that **Cut** is also sound. Suppose, therefore, that $M, s \models B(\varphi \mid \psi) \wedge B(\chi \mid \varphi \wedge \psi)$.

Let X be any minimal retentive set in $[\![\psi]\!]$. Then, because $M, s \models B(\varphi \mid \psi)$, we have $X \subseteq [\![\varphi]\!]$. It follows that $X \subseteq [\![\varphi \wedge \psi]\!]$. By Lemma 13, X is then minimal retentive in $[\![\varphi \wedge \psi]\!]$, so because $M, s \models B(\chi \mid \varphi \wedge \psi)$ we have $X \subseteq [\![\chi]\!]$. This holds for every minimal retentive $X \subseteq [\![\psi]\!]$, so we have $M, s \models B(\chi \mid \psi)$. Hence $M, s \models (B(\varphi \mid \psi) \wedge B(\chi \mid \varphi \wedge \psi)) \rightarrow B(\chi \mid \psi)$, as was to be shown. □

We mentioned before that **BASE** is not complete for any of the four semantics. Now that we have shown **Sh** to be sound for $\mathsf{max}, \mathsf{opt}$ and unm (Lemma 33) and **Cut** to be sound for acc (Lemma 41), we can prove this incompleteness.

Proposition 42. *For every* $\beta \in \{\mathsf{max}, \mathsf{opt}, \mathsf{unm}, \mathsf{acc}\}$*, there is a* $\varphi \in \mathcal{L}$ *such that* $\models_\beta \varphi$ *while* $\not\vdash_{\mathbf{BASE}} \varphi$.

Proof sketch. The proof is by showing that there are alternative semantics for $\mathcal{L}$ such that **BASE** is sound for the semantics while **Sh** and **Cut** are not. Let us denote these semantics by $\models_f$. Because of the unsoundness of **Sh** and **Cut** there are instances φ_1 of **Sh** and φ_2 of **Cut** such that $\not\models_f \varphi_1$ and $\not\models_f \varphi_2$. The soundness of **BASE** for $\models_f$ means that **BASE** can only derive formulas that are valid in $\models_f$, so $\not\vdash_{\mathbf{BASE}} \varphi_1$ and $\not\vdash_{\mathbf{BASE}} \varphi_2$.

Yet φ_1 is an instance of **Sh**, which is sound for $\mathsf{max}, \mathsf{opt}$ and unm semantics. Similarly, φ_2 is an instance of **Cut**, which is sound for acc semantics. So $\models_\beta \varphi_1$ for every $\beta \in \{\mathsf{max}, \mathsf{opt}, \mathsf{unm}\}$ and $\models_{\mathsf{acc}} \varphi_2$.

To create such semantics, we use the fact that the only axiom in **BASE** that relates two different relativizations is **R-Ext**. For example, the axiom **K**, which states that $B(\varphi \rightarrow \psi \mid \chi) \rightarrow (B(\varphi \mid \chi) \rightarrow B(\psi \mid \chi))$, only uses the relativization χ. The one axiom that does allow us to compare different relativizations, **R-Ext**, requires them to have the same extension. So in general **BASE** does not allow us to draw conclusions about what is best in $[\![\psi]\!]$ based on what is best in $[\![\delta]\!]$, unless $[\![\psi]\!] = [\![\delta]\!]$.

As a result, **BASE** is sound for semantics where we replace the relation $\preceq$ by a function $f : 2^S \rightarrow 2^S$ with the property that $f(X) \subseteq X$, and say that $M, s \models_f B(\varphi \mid \psi)$ if and only if $f([\![\psi]\!]) \subseteq [\![\varphi]\!]$. The axioms **Sh** and **Cut** are not sound for these semantics, so some instances of **Sh** and **Cut** are not derivable in **BASE**. It follows that **BASE** is not complete for any of the four semantics. See the appendix for a more detailed version of this proof. □

We continue by constructing a canonical model for **ACC**. Unfortunately, this canonical model is significantly more complex than the one for **MOU**. The states of $M_\Xi^{\mathbf{ACC}}$ are five-tuples (x, y, ψ, b, i) where x and y are maximal consistent sets, ψ is a formula, $b \in \{0, 1\}$ and $i \in \mathbb{N}$. We use $\Gamma_\psi(\Xi)$ to denote the ψ-good sets in $MCS_\Xi^{\mathbf{ACC}}$, and $\Delta_\psi(\Xi)$ to denote those tuples where both x and y are ψ-good, i.e., $\Delta_\psi(\Xi) := \{(x, y, \psi, b, i) \mid x, y \in \Gamma_\psi(\Xi)\}$.

Definition 43. Let Ξ be $\Box$-maximal. The model $M_\Xi^{\mathbf{ACC}} = (S, \preceq, V)$ is given by

- $S = \{(x, y, \varphi, b, i) \mid x, y \in MCS_\Xi^{\mathbf{ACC}}, \varphi \in \mathcal{L}, b \in \{0, 1\}, i \in \mathbb{N}\}$,
- $V(p) = \{(x, y, \varphi, b, i) \in S \mid p \in x\}$,
- $\preceq$ is the smallest relation such that
  1. for all $(x, y, \psi, b, i) \in S$: $(x, y, \psi, b, i) \preceq (x, y, \psi, b, i + 1)$,
  2. for all $(x, y, \psi, b, i) \in \Delta_\psi(\Xi)$:
     (a) if $x \neq y$ then for all $i > 0$: $(x, y, \psi, b, i) \preceq (y, y, \psi, 0, 0)$,
     (b) for all $i > 0$: $(x, x, \psi, 0, i) \preceq (x, x, \psi, 1, 0)$,
     (c) for all $i > 0$ and all $(x', y', \psi, b, i) \in \Delta_\psi(\Xi)$: if $b = 0$ or $x' \neq y'$ then $(x, x, \psi, 1, 1) \prec (x', y', \psi, b, 0)$ and
     (d) for all $x' \in MCS_\Xi^{\mathbf{ACC}}$ such that $\psi \notin x'$: $(x, y, \psi, b, i) \preceq (x', x', \bot, 0, 0)$.

Recall that $[\varphi]$ denotes the set of states that are based on an MCS containing φ, so in this case $[\varphi] = \{(x, y, \psi, b, i) \in S \mid \varphi \in x\}$.

In a five-tuple (x, y, ψ, b, i), x takes the usual role of denoting which formulas are supposed to be true at a state, i.e., we will have $M_\Xi^{\mathbf{ACC}}, (x, y, \psi, b, i) \models \varphi$ if and only if $\varphi \in x$. The formula ψ, as before, denotes the intended relativization. The index i is used to create infinitely ascending chains, which will be used to manipulate which states are acceptable and which ones are not. The bit b is used to create two copies of each ascending chain. This is because, under certain circumstances, acceptability of the chain with $b = 0$ requires a single state that is strictly preferred over all but finitely many states in that chain. The state we use for that is the first in the $b = 1$ chain. Finally, the set y is used to create certain connections between different chains. Specifically, (x, y, ψ, b, i) will connect to (y, y, ψ, b, i).

This canonical model is rather complex, so let us use a figure to explain it. There are five kinds of arrows in $M_\Xi^{\mathbf{ACC}}$: the ones that follow from conditions 1, 2a, 2b, 2c and 2d, respectively. Only arrows of type 2d are between states with different relativizing formulas ψ and φ. Figure 9 shows a "slice" of the model where ψ is held constant, so it does not contain arrows of type 2d.

For fixed x, y, ψ and b, the set $\{(x, y, \psi, b, i) \mid i \in \mathbb{N}\}$ forms an infinitely ascending chain, using type 1 arrows. If $(x, y, \psi, b, i) \notin \Delta_\psi(\Xi)$ then none of the type 2 arrows are applicable, so $\{(x, y, \psi, b, i) \mid i \in \mathbb{N}\}$ is isolated, as shown in the leftmost chain of states in Figure 9.

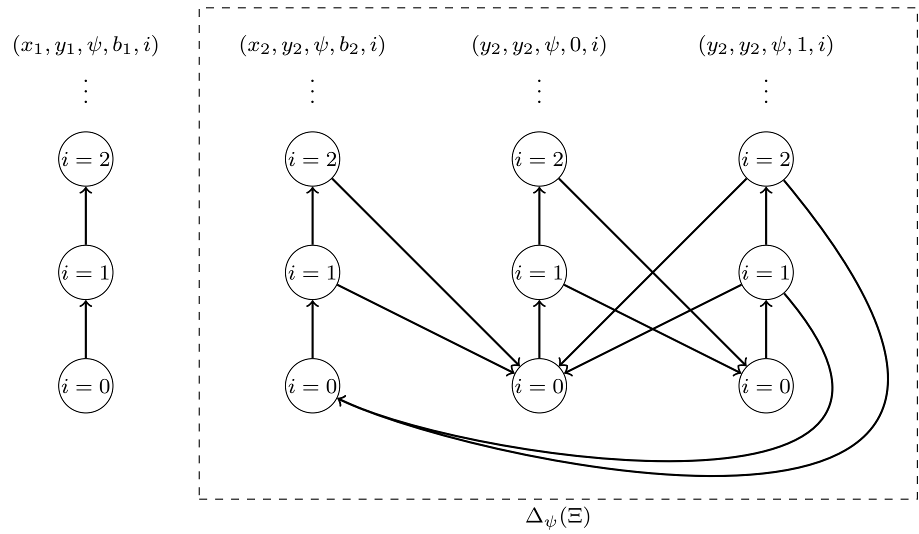

Figure 9: A schematic representation of a single ψ-slice of $M_{\Xi}^{\mathbf{ACC}}$. The dashed line indicates which chains are part of $\Delta_\psi(\Xi)$. Note that because only a single slice is drawn, no type 2d arrows are included, as those connect different slices.

If $(x, y, \psi, b, i) \in \Delta_\psi(\Xi)$, then it becomes important whether $x = y$. If $x \neq y$, then every state (x, y, ψ, b, i) with $i > 0$ is beaten by $(y, y, \psi, 0, 0)$, due to a type 2a arrow. Such states are represented by the second chain from the left in Figure 9., and the type 2a arrows are the one from the second chain to the third. States of the form $(x, x, \psi, 0, i)$ with $i > 0$ are beaten by $(x, x, \psi, 1, 0)$, due to type 2b arrows. This is represented by the third chain from the left in Figure 9, and the arrows to the fourth chain. Finally, states of the form $(x, x, \psi, 1, i)$ with $i > 0$ are beaten by all $(x', y', \psi, b, 0) \in \Delta_\psi(\Xi)$, where $b = 0$ or $x \neq y$ due to type 2c arrows. This is represented by the fourth chain in Figure 9, and the arrows to the second and third chains.

The purpose of the set $\Delta_\psi(\Xi)$ is to be minimal retentive in $[\psi]$, see Lemma 46. Every formula φ that is consistent with B_ψ^{-1} occurs in at least one x such that $(x, x, \psi, 0, 0) \in \Delta_\psi(\Xi)$, which provides the witness to show that $P(\varphi \mid \psi)$ holds in the canonical model. Yet $\Delta_\psi(\Xi)$ should generally[8] not be acceptable in $[\chi]$ when $\chi \neq \psi$, in order to prevent it from providing "false witness" to the truth of $P(\varphi \mid \chi)$ where $\neg\varphi \in B_\chi^{-1}$.

In order to give $\Delta_\psi(\Xi)$ these properties, the relation $\preceq$ of the canonical model satisfies three properties. Firstly, $\Delta_\psi(\Xi)$ has no outgoing arrows in $[\psi]$, so it is retentive. Secondly, for every $s_1, s_2 \in \Delta_\psi(\Xi)$, there is a path $s_1 \prec \cdots \prec s_2$ in $\Delta_\psi(\Xi)$. This ensures the minimality of $\Delta_\psi(\Xi)$, since any retentive set

[8]The exceptions are when $B_\chi^{-1} \subseteq B_\varphi^{-1}$, in which case it is unproblematic for $\Delta_\psi(\Xi)$ to be acceptable in $[\chi]$.

that contains s_1 must also contain s_2. Finally, except in some special cases, if $[\chi] \neq [\psi]$ then $\Delta_\psi(\Xi) \cap [\chi]$ ends in an infinitely ascending chain, and is therefore not minimally retentive.

We now prove a few lemmas that will be useful in the truth lemma. Recall that we use * to indicate the transitive closure of a relation.

Lemma 44. *If $\varphi \neq \psi$ and $(x, y, \varphi, b, i) \preceq (x', y', \psi, b', j)$ then $\psi = \bot$, $(x, y, \varphi, b, i) \prec (x', y', \varphi, b', j)$ and $(x', y', \psi, b', j) \not\preceq^* (x, y, \varphi, b, i)$.*

Proof. The relation $\preceq$ is the smallest one satisfying clauses 1, 2a, 2b, 2c and 2d of Definition 43. Hence $(x, y, \varphi, b, i) \preceq (x', y', \psi, b', j)$ implies that one of those clauses applies. The only clause that allows arrows between states with different third coordinate is 2d, so we must have $(x', y', \psi, b', j) = (x', x', \bot, 0, 0)$.

Now note that $B_\bot^{-1}$ is inconsistent for every $\Box$-maximal set Ξ, because $\vdash_{\mathbf{ACC}} B(\bot \mid \bot)$ by **Id**. This implies that $\Delta_\bot(\Xi) = \emptyset$. As such, no state with $\bot$ as third coordinate can be the origin of any arrow of type 2. It follows that $(x', x', \bot, 0, 0) \not\preceq^* (x, y, \varphi, b, i)$. Because we previously established that $(x, y, \varphi, b, i) \preceq (x', x', \bot, 0, 0)$, this also shows that $(x, y, \varphi, b, i) \prec (x', x', \bot, 0, 0)$ $\Box$

Lemma 45. *If $(x, y, \psi, i, b) \notin \Delta_\psi(\Xi)$, then for every $\varphi \in \mathcal{L}$, $(x, y, \psi, i, b) \notin \mathsf{acc}([\varphi])$.*

Proof. Because $(x, y, \psi, i, b) \notin \Delta_\psi(\Xi)$, the set $\{(x, y, \psi, j, b) \mid j \in \mathbb{N}\}$ forms an infinitely ascending chain without "outgoing" arrows to any state outside the chain. Furthermore, as membership in $[\varphi]$ depends only on the first coordinate, either $\{(x, y, \psi, j, b) \mid j \in \mathbb{N}\} \subseteq [\varphi]$ or $\{(x, y, \psi, j, b) \mid j \in \mathbb{N}\} \cap [\varphi] = \emptyset$.

Suppose towards a contradiction that $(x, y, \psi, b, i) \in \mathsf{acc}([\varphi])$. Then there is a minimal retentive set $X \subseteq [\varphi]$ such that $(x, y, \psi, b, i) \in X$. This implies that, in particular, $(x, y, \psi, b, i) \in [\varphi]$ and therefore $\{(x, y, \psi, b, j) \mid j \in \mathbb{N}\} \subseteq [\varphi]$. Because X is retentive, it follows that $\{(x, y, \psi, b, j) \mid j \geq i\} \subseteq X$. But then X cannot be minimal, as $\{(x, y, \psi, b, j) \mid j > i\} \subset X$ is also retentive. $\Box$

Lemma 46. *$\Delta_\psi(\Xi)$ is minimal retentive in $[\psi]$.*

Proof. First, let us show that $\Delta_\psi(\Xi) \subseteq [\psi]$. By **Id** we have $\vdash_{\mathbf{ACC}} B(\psi \mid \psi)$ and therefore, by $\Box$**S5**, $\vdash_{\mathbf{ACC}} \Box B(\psi \mid \psi)$. This means that every maximal consistent set contains $\Box B(\psi \mid \psi)$, so the $\Box$-maximal set Ξ must contain $B(\psi \mid \psi)$. It follows that $\psi \in B_\psi^{-1}$. By the definition of $\Delta_\psi(\Xi)$, if $(x, y, \psi, b, i) \in \Delta_\psi(\Xi)$ then x and y are ψ-good. In particular, this implies that $B_\psi^{-1} \subseteq x$, which together with $\psi \in B_\psi^{-1}$ yields $\psi \in x$, so $(x, y, \psi, b, i) \in [\psi]$. This holds for every $(x, y, \psi, b, i) \in \Delta_\psi(\Xi)$, so $\Delta_\psi(\Xi) \subseteq [\psi]$.

Next, let us show that $\Delta_\psi(\Xi)$ is retentive. Suppose towards a contradiction that there are $(x, y, \psi, b, i) \in \Delta_\psi(\Xi)$ and $(x', y', \chi, b', j) \in [\psi] \setminus \Delta_\psi(\Xi)$ such that $(x, y, \psi, b, i) \prec (x', y', \chi, b', j)$. Then, in particular, $(x, y, \psi, b, i) \preceq (x', y', \chi, b', j)$. Because $\preceq$ is the smallest relation satisfying conditions 1 and 2, it follows that one of the conditions 1, 2a, 2b, 2c or 2d must apply to (x, y, ψ, b, i) and (x', y', χ, b', j).

It cannot be condition 1, since that would require $x = x', y = y'$ and $\psi = \chi$, which would contradict $(x', y', \chi, b', j) \notin \Delta_\psi(\Xi)$. Similarly, it cannot be conditions 2a, 2b or 2c, since those also contradict $(x', y', \chi, b', j) \notin \Delta_\psi(\Xi)$. Finally, it cannot be condition 2d since that would require $\psi \notin x'$, which contradicts $(x', y', \chi, b', j) \in [\psi]$. We have arrived at a contradiction, and thereby conclude that $\Delta_\psi(\Xi)$ is retentive.

Left to show is that $\Delta_\psi(\Xi)$ is minimal. As observed above, every two states (x, y, ψ, b, i) and (x', y', ψ, b', j) in $\Delta_\psi(\Xi)$ are reachable from each other by some number of $\prec$ steps. It follows that no set that is retentive in $[\psi]$ can include one but not the other. So no non-empty strict subset of $\Delta_\psi(\Xi)$ is retentive, which means that $\Delta_\psi(\Xi)$ is mimimal retentive. □

Lemma 47. *If* $[\chi] \not\subseteq [\psi]$ *or* $\Delta_\psi(\Xi) \not\subseteq [\chi]$ *then* $\Delta_\psi(\Xi) \cap \mathsf{acc}([\chi]) = \emptyset$.

Proof. Suppose that $[\chi] \not\subseteq [\psi]$. Then there is an $(x', y', \delta, b, i) \in [\chi] \setminus [\psi]$. This implies that we also have $(x', x', \bot, 0, 0) \in [\chi] \setminus [\psi]$. Now, suppose towards a contradiction that $(x, y, \psi, b, i) \in \Delta_\psi(\Xi) \cap \mathsf{acc}([\chi])$. Then $(x, y, \psi, b, i) \in X$ for some set X that is minimal retentive in $[\chi]$.

Because $(x', x', \bot, 0, 0) \notin [\psi]$ we have $\psi \notin x'$. So, by clause 2d in the definition of $\preceq$, we have $(x, y, \psi, b, i) \preceq (x', x', \bot, 0, 0)$. By Lemma 44 we then have $(x, y, \psi, b, i) \prec (x', x', \bot, 0, 0)$ and $(x', x', \bot, 0, 0) \not\preceq^* (x, y, \psi, b, i)$.

We already showed that $(x', x', \bot, 0, 0) \in [\chi]$, so if $(x, y, \psi, b, i) \in X$ and X is retentive in $[\chi]$ then $(x', x', \bot, 0, 0) \in X$. Minimality of X would then require $(x', x', \bot, 0, 0) \prec^* (x, y, \psi, b, i)$, but that contradicts $(x', x', \bot, 0, 0) \not\preceq^* (x, y, \psi, b, i)$. We have arrived at a contradiction and thereby conclude that $\Delta_\psi(\Xi) \cap \mathsf{acc}([\chi]) = \emptyset$.

Suppose then that $[\chi] \subseteq [\psi]$ and that $\Delta_\psi(\Xi) \not\subseteq [\chi]$. First, note that $\Delta_\psi(\Xi)$ does not have any outgoing arrows in $[\chi]$. Such arrows would need to be of type 2d, but such arrows can only be to states $(x', x', \bot, 0, 0)$ where $\psi \notin x$, so $(x', x', \bot, 0, 0) \notin [\psi]$. From $[\chi] \subseteq [\psi]$ it then follows that $(x', x', \bot, 0, 0) \notin [\chi]$, so there are no type 2d arrows from $\Delta_\psi(\Xi)$ to states in $[\chi]$. It follows that if there is a minimal retentive set X such that $X \cap \Delta_\psi(\Xi) \neq \emptyset$ then $X \subseteq \Delta_\psi(\Xi)$.

Now, suppose towards a contradiction that $(x, y, \psi, b, i) \in \Delta_\psi(\Xi) \cap \mathsf{acc}([\chi])$, so there is a minimal retentive set X such that $(x, y, \psi, b, i) \in X$. We distinguish two cases. Firstly, suppose that $\chi \notin y$, so $(y, y, \psi, 0, 0) \notin [\chi]$. The infinitely ascending chain $\{(x, y, \psi, b, j) \mid j \in \mathbb{N}\}$ is beaten only by $(y, y, \psi, 0, 0)$, so $\{(x, y, \psi, b, j) \mid j \in \mathbb{N}\}$ is an infinitely ascending chain in $[\chi]$ without any outgoing arrows. This implies that it cannot be part of any minimal retentive set.

The second case is $\chi \in y$. Because $\Delta_\psi(\Xi) \not\subseteq [\chi]$ there must be a y' such that $(y', z', \psi, b', j) \in \Delta_\psi(\Xi) \setminus [\chi]$. But, because $[\chi]$-membership depends only on the first coordinate, $(x, y', \psi, b, 0) \in [\chi]$. Furthermore, $(x, y', \psi, b, 0)$ is reachable from (x, y, ψ, b, i) by a $\preceq$-chain that only contains $[\chi]$ states.[9] It follows that

[9]If $x \neq y$ this chain is $(x, y, \psi, b, i) \prec (x, y, \psi, b, i+1) \prec (y, y, \psi, 0, 0) \prec (y, y, \psi, 0, 1) \prec (y, y, \psi, 1, 0) \prec (y, y, \psi, 1, 1) \prec (x, y', \psi, b, 0)$. If $x = y$ and $b = 0$ it is $(x, x, \psi, 0, i) \prec$

any retentive set containing (x, y, ψ, b, i) must also contain $(x, y', \psi, b, 0)$. We have now reduced the second case to the first case. That case resulted in a contradiction, so (x, y, ψ, b, i) cannot be part of any minimal retentive set.

We have now shown that if $\Delta_\psi(\Xi) \not\subseteq [\chi]$ then $\Delta_\psi(\Xi) \cap \mathsf{acc}([\chi]) = \emptyset$. Together with the previous conclusion that if $[\chi] \not\subseteq [\psi]$ then $\Delta_\psi(\Xi) \cap \mathsf{acc}([\chi]) = \emptyset$, this proves the lemma. □

Using the preceding three lemmas, the proof of the truth lemma is relatively easy.

Lemma 48 (Truth lemma for **ACC**)**.** *For every $\Box$-maximal Ξ, every $\varphi, \psi \in \mathcal{L}$, every $x, y \in MCS^{\mathbf{ACC}}_\Xi$, every $b \in \{0, 1\}$ and every $i \in \mathbb{B}$ we have $M^{\mathbf{ACC}}_\Xi, (x, y, \psi, b, i) \models_{\mathsf{acc}} \varphi$ if and only if $\varphi \in x$.*

Proof. By induction on the complexity of φ and then by a case distinction on the main connective. The only interesting case is $\varphi = B(\gamma \mid \delta)$, so that is the case we consider in detail. Suppose therefore that $\varphi = B(\gamma \mid \delta)$.

- If $B(\gamma \mid \delta) \in x$ then $\gamma \in B^{-1}_\delta$. Take any $s \in \mathsf{acc}([\![\delta]\!])$. By the induction hypothesis, $[\![\delta]\!] = [\delta]$, so $s \in \mathsf{acc}([\delta])$. By Lemmas 45–47 we must have one of the following:

  1. $s \in \Delta_\delta(\Xi)$ or
  2. $s \in \Delta_\zeta(\Xi)$ where $[\delta] \subseteq [\zeta]$ and $\Delta_\zeta(\Xi) \subseteq [\delta]$.

  In the first case we have $s \in [\gamma]$, due to $\gamma \in B^{-1}_\delta$.

  Consider then the second case. Because $[\delta] \subseteq [\zeta]$ we have that $\delta \in z$ implies $\zeta \in z$ for all $z \in MCS_\Xi$. This implies that $\delta \leftrightarrow (\zeta \wedge \delta) \in z$, for every z. By Lemma 31 this implies $\delta \leftrightarrow (\zeta \wedge \delta) \in \Xi$. Because $x \in MCS_\Xi$, we obtain $\Box(\delta \leftrightarrow (\zeta \wedge \delta)) \in x$. By assumption we have $B(\gamma \mid \delta) \in x$, so by **R-Ext** we have $B(\gamma \mid \zeta \wedge \delta) \in x$.

  Furthermore, from $\Delta_\zeta(\Xi) \subseteq [\delta]$ it follows that for every $z \in MCS_\Xi$ we have that if $\Delta_\zeta(\Xi) \subseteq z$ then $\delta \in z$. So there is no maximal consistent set z such that $\Box^{-1}z = \Xi$, $B^{-1}_\zeta \subseteq z$ and $\neg\delta \in z$. By the Lindenbaum lemma (Lemma 25) every consistent set can be extended to a maximal consistent set, so $\{\neg\delta\} \cup B^{-1}_\zeta \cup \Box\Xi \cup \{\neg\Box\zeta \mid \zeta \in \mathcal{L} \setminus \Xi\}$ is inconsistent. By Lemma 30 this implies that $\delta \in B^{-1}_\zeta$. So we have $B(\delta \mid \zeta) \in x$.

  We have now shown that $B(\delta \mid \zeta) \in x$ and $B(\gamma \mid \delta \wedge \zeta) \in x$. Using **Cut** this implies that $B(\gamma \mid \zeta) \in x$, and therefore $B(\gamma \mid \zeta) \in \Xi$, which means that $\gamma \in B^{-1}_\zeta$. From $s \in \Delta_\zeta(\Xi)$ we therefore obtain $s \in [\gamma]$.

  In either case, we have shown $s \in [\gamma]$. By the induction hypothesis, this implies that $M^{\mathbf{ACC}}_\Xi, s \models \gamma$. This holds for every $s \in \mathsf{acc}([\![\delta]\!])$, so we have $M^{\mathbf{ACC}}_\Xi, x \models B(\gamma \mid \delta)$, which was to be shown.

$(x, x, \psi, 0, i+1) \prec (x, x, \psi, 1, 0) \prec (x, x, \psi, 1, 1) \prec (x, y', \psi, b, 0)$. If $x = y$ and $b = 1$ it is $(x, x, \psi, 1, i) \prec (x, x, \psi, 1, i+1) \prec (x, y', \psi, b, 0)$.

<table>
<tr><th rowspan="2">frame conditions</th><th colspan="4">Semantics</th></tr>
<tr><th>max</th><th>opt</th><th>unm</th><th>acc</th></tr>
<tr><td>none</td><td colspan="3">MOU</td><td>ACC</td></tr>
<tr><td>tr</td><td>MOU</td><td>MOU+DM</td><td colspan="2">MOU</td></tr>
<tr><td>to</td><td colspan="2">MOU</td><td>MOU+Un+WDM</td><td>ACC</td></tr>
<tr><td>as</td><td>MOU</td><td>MOU+Un+WDM</td><td>MOU</td><td>ACC</td></tr>
<tr><td>tr + to</td><td colspan="2">MOU+DR+DM</td><td>MOU+Un+SP+DM</td><td>MOU+DR+DM</td></tr>
<tr><td>tr + as</td><td>MOU</td><td>MOU+Un+DM</td><td colspan="2">MOU</td></tr>
<tr><td>to + as</td><td colspan="3">MOU+DR+DM</td><td>—</td></tr>
<tr><td>tr + to + as</td><td colspan="4">MOU+DR+Un+SP+DM</td></tr>
</table>

Figure 10: The axiomatizations for each combination of semantics and frame conditions.

- If $B(\gamma \mid \delta) \notin x$, then $\gamma \notin B_\delta^{-1}$. By Lemma 30, the set $\{\neg\gamma\} \cup B_\delta^{-1} \cup \Box\Xi \cup \{\neg\Box\zeta \mid \zeta \in \mathcal{L} \setminus \Xi\}$ is then consistent. It can therefore be extended to a maximal consistent set x'. By construction, $\Box^{-1}x' = \Xi$, so $x' \in MCS_\Xi$. Furthermore, $B_\delta^{-1} \subseteq y$ so $(x', x', \delta, 0, 0) \in \Delta_\delta(\Xi)$.

  By Lemma 46 we have $(x', x', \delta, 0, 0) \in \mathsf{acc}([\delta])$. By the induction hypothesis this implies $(x', x', \delta, 0, 0) \in \mathsf{acc}(\llbracket\delta\rrbracket)$. Finally, $\neg\gamma \in x'$ so $(x', x', \delta, 0, 0) \notin [\gamma]$. By the induction hypothesis this implies $(x', x', \delta, 0, 0) \notin \llbracket\gamma\rrbracket$. We therefore have $\mathsf{acc}(\llbracket\delta\rrbracket) \not\subseteq \llbracket\gamma\rrbracket$, so $M_\Xi^{\mathbf{ACC}}, x \not\models B(\gamma \mid \delta)$, which was to be shown.

□

Completeness follows immediately from the truth lemma, and soundness was proven in Lemma 41. We therefore have the desired soundness and completeness result.

Theorem 49. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models_{\mathsf{acc}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{ACC}} \varphi$.*

This concludes the discussion of the proof systems for unconstrained $\preceq$. We continue to consider cases where $\preceq$ satisfies further constraints.

7. Axiomatizations for constrained preference relations

We consider it to be unreasonable to demand that $\preceq$ be transitive, anti-symmetric or total in all situations. However, there are some contexts where it is reasonable to assume that $\preceq$ has such properties. For example, if $\preceq$ represents the preference relation of a rational agent then $\preceq$ should be transitive, and if $\preceq$ represents defeat among strategies then $\preceq$ should be anti-symmetric.

Making such assumptions regarding $\preceq$ may render a formula, or a set of formulas, unsatisfiable even if it is satisfiable in general. Or, putting it a different way, a belief about, say, the playability of certain strategies may be generally consistent, yet inconsistent under the assumption of, for example, transitivity.

There are 32 combinations of a semantics ($\mathsf{max}, \mathsf{opt}, \mathsf{unm}$ or acc) and a subset of {transitivity, totality, anti-symmetry} of additional constraints. Four of those, with the empty set of constraints, we discussed in the previous section. For 27 of the remaining combinations we give a sound and strongly complete axiomatization in this section. Several variants have the same sets of validities, so in total we only need 8 different axiomatizations, see Figure 10 for an overview. For the one remaining combination, acc semantics on total anti-symmetric models, we do not provide a sound and strongly complete axiomatization, and instead show that such an axiomatization does not exist.

Unfortunately, while we did find axiomatizations for most variants, we should emphasize that these are not correspondence results, in the sense of modal correspondence theory. With correspondence, once we have axioms for several properties, an axiomatization for models that satisfy all of these properties can be obtained simply by adding the axioms for each property.

Here, that is not the case. For example, **MOU** is complete for unm semantics on transitive models and **MOU** + **Un** + **WDM** is complete for unm semantics on total models, yet **MOU**+**Un**+**WDM** is not complete for unm on transitive total models.

7.1. The axioms

Each of the axioms is contained in multiple axiomatizations, so we begin by introducing the axioms, and proving a few things about their effects on a canonical model.

Definition 50. The axioms **Un**, **DR** and **SP** are as follows

- **Un** $P(\varphi \mid \psi) \to B(\varphi \mid \psi)$
- **DR** $B(\varphi \mid \psi \lor \chi) \to (B(\varphi \mid \psi) \lor B(\varphi \mid \chi))$
- **SP** $(P(\chi \mid \varphi) \land P(\xi \mid \psi) \land B(\neg\varphi \mid \varphi \lor \psi)) \to P(\xi \mid (\varphi \land \neg\chi) \lor \psi)$

The axiom **Un** is a *uniqueness* axiom, it implies that all of the best ψ-states are indistinguishable from one another.

The axiom **DR** is known in the literature as *disjunctive rationality* (see, e.g., [19]).

The axiom **SP** is, sadly, somewhat more complex. It can be seen as saying that (under certain conditions) there is a clear separation between states that are the best in some extension and those states that are not. We therefore refer to it as the *separation principle*. The axiom is about how the best states from two sets $[\![\psi]\!]$ and $[\![\varphi]\!]$ lie in the larger extension $[\![\varphi \lor \psi]\!]$. The antecedent of the axiom says that there are an $s_1 \in \beta([\![\psi]\!]) \cap [\![\xi]\!]$ and an $s_2 \in \beta([\![\varphi]\!]) \cap [\![\chi]\!]$, and that s_2 is *not* among the best states in $[\![\varphi \lor \psi]\!]$. Under these conditions, the consequent says that s_1 must be among the best states in $[\![(\varphi \land \neg\chi) \lor \psi]\!]$, i.e., s_1 must be best among all $\varphi \lor \psi$ states except s_2.

Intuitively, we can think of this as saying that, unless the φ states are simply better than the ψ ones (in which case s_2 would be among the best in $[\![\varphi \lor \psi]\!]$), the fact that s_1 is the best in its own extension $[\![\psi]\!]$ means that it must beat all of the non-best φ states, and therefore be best in $[\![\varphi \lor \psi]\!] \setminus \{s_2\}$.

To the best of our knowledge this axiom has not been considered before, nor has anything equivalent to it. It is not, by itself, particularly convincing: we see no philosophical reason why one would adopt the principle **SP**. Fortunately, **SP** does not itself need to be convincing; in the proof systems where it is used, we are justified in adopting the axiom because it is sound.

In addition to **Un**, **DR** and **SP**, we also need two *monotonicity* axioms, **DM** and **WDM**. For the purpose of comparison we also take a look at the axiom **RM** from the literature, although this axiom is not itself used in any of our axiomatizations.

The idea of a monotonicity axiom is that if a particular condition suffices to conclude ψ, then a stronger condition should suffice as well. So we go from $B(\psi \mid \alpha)$ to $B(\psi \mid \gamma)$ where γ implies α. Full monotonicity, i.e., $B(\psi \mid \alpha) \to B(\psi \mid \gamma)$ whenever $\gamma \to \alpha$, is unsound for any of the variants that we consider. More restricted versions of monotonicity, where $B(\psi \mid \alpha)$ implies $B(\psi \mid \gamma)$ under certain conditions, are, in some cases, sound though.

A version of monotonicity that is used in the literature (e.g., [26]) is called *rational monotonicity.*

$$\textbf{RM} \quad (P(\varphi \mid \chi) \wedge B(\psi \mid \chi)) \to B(\psi \mid \varphi \wedge \chi)$$

This axiom **RM** is sound for opt-semantics on transitive models. Unfortunately, while **RM** is sound for those models, **MOU** + **RM** is not complete. We will therefore use a different monotonicity axiom.

$$\textbf{DM} \quad (P(\chi \mid \varphi) \wedge B(\psi \mid \varphi \vee \chi)) \to B(\psi \mid \chi)$$

Later on we will also use a third monotonicity axiom.

$$\textbf{WDM} \quad (P(\chi \mid \varphi) \wedge P(\varphi \mid \chi) \wedge B(\psi \mid \varphi \vee \chi)) \to B(\psi \mid \chi)$$

Where **RM** strengthens by adding a conjunction (χ becomes $\varphi \wedge \chi$), **DM** and **WDM** strengthen by removing a disjunction ($\varphi \vee \chi$ becomes χ). We therefore refer to **DM** as *disjunctive monotonicity.* The axiom **WDM** is strictly weaker than **DM**, because it has an extra conjunct in the antecedent, so we call it *weak disjunctive monotonicity.*

7.2. Effect of the axioms

The presence of any of the above axioms influences the set of maximal consistent sets. In particular, they affect the support relation.

Lemma 51. *For the proof systems that include* **DM**, *the support relation* $\rightsquigarrow$ *is transitive.*

Proof. Suppose that $x \rightsquigarrow y$ and $y \rightsquigarrow z$. Then there are φ and ψ such that $x \overset{\varphi}{\rightsquigarrow} y$ and $y \overset{\psi}{\rightsquigarrow} z$. Now, take any χ such that $B(\chi \mid \varphi \vee \psi)$.

By $x \overset{\varphi}{\rightsquigarrow} y$ we have $B_\varphi^{-1} \subseteq y$ and by $y \overset{\psi}{\rightsquigarrow} z$ we have $\psi \in y$. Together, this implies $P(\psi \mid \varphi)$. Combined with $B(\chi \mid \varphi \vee \psi)$, this allows us to use **DM** to conclude $B(\chi \mid \psi)$. The latter, by $y \overset{\psi}{\rightsquigarrow} z$, implies that $\chi \in z$.

This holds for every $\chi \in B_{\varphi \vee \psi}^{-1}$, so $x \overset{\varphi \vee \psi}{\rightsquigarrow} z$ and therefore $x \rightsquigarrow z$. □

Lemma 52. *For the proof systems that include* **Un**, *if* $B^{-1}_\varphi \subseteq x$ *then* $B^{-1}_\varphi = x$.

Proof. For every $\psi \in x$ we have $P(\psi \mid \varphi)$ and therefore, by **Un**, $B(\psi \mid \varphi)$. □

Lemma 53. *For the proof systems that include* **DR**, *if* x *and* y *are good then* x *supports* y *or* y *supports* x.

Proof. Let φ and ψ be witnesses for x and y being good, respectively, so $B^{-1}_\varphi \subseteq x$ and $B^{-1}_\psi \subseteq y$. Now, suppose towards a contradiction that $B^{-1}_{\varphi\vee\psi}$ is not a subset of x or y. Then there are $\gamma, \delta \in B^{-1}_{\varphi\vee\psi}$ such that $\gamma \notin x$ and $\delta \notin y$, which implies that $\gamma \notin B^{-1}_\varphi$ and $\delta \notin B^{-1}_\psi$.

We then have $B(\gamma\wedge\delta \mid \varphi\vee\psi)$, $\neg B(\gamma\wedge\delta \mid \varphi)$ and $\neg B(\gamma\wedge\delta \mid \psi)$, contradicting **DR**. From this contradiction, we obtain that $B^{-1}_{\varphi\vee\psi}$ is a subset of x or y, so either $x \overset{\varphi\vee\psi}{\leadsto} y$ or $y \overset{\varphi\vee\psi}{\leadsto} x$. □

The effect of **SP**, like the axioms itself, is a bit more complex.

Lemma 54. *For the proof systems that include* **SP** *and* **DM**, *if* x, y *and* z *are distinct,* z *is good,* $x \leadsto y$ *and* $z \not\leadsto y$, *then* $x \leadsto z$.

Proof. Let φ be such that $x \overset{\varphi}{\leadsto} y$, let $\chi \in y \setminus x$, let ψ be such that $B^{-1}_\psi \subseteq z$, and take any $\xi \in z$.

Now, suppose towards a contradiction that we have $P(\varphi \mid \varphi \vee \psi)$, so there is some x' such that $B^{-1}_{\varphi\vee\psi} \subseteq x'$ and $\varphi \in x'$. Then $z \overset{\varphi\vee\psi}{\leadsto} x'$ and $x' \overset{\varphi}{\leadsto} x$, so by transitivity of support (due to inclusion of **DM**), $z \leadsto x$, contradicting our assumptions. From this contradiction, we obtain $B(\neg\varphi \mid \varphi \vee \psi)$.

We now have $P(\chi \mid \varphi)$, $P(\xi \mid \psi)$, $B(\neg\varphi \mid \varphi \vee \psi)$. By **SP**, this implies $P(\xi \mid (\varphi\wedge\neg\chi)\vee\psi)$. This holds for every $\xi \in z$, so $B^{-1}_{(\varphi\wedge\neg\chi)\vee\psi} \subseteq z$, and therefore $x \overset{(\varphi\wedge\neg\chi)\vee\psi}{\leadsto} z$. □

Note that this lemma more or less follows the intuition behind **SP** that we discussed in Section 7.1. Because $x \overset{\varphi}{\leadsto} y$, we can think of y as being among the best φ-states, and of x as being a non-best φ-state. The state z is good, so we can think of it as being among the best ψ-states for some ψ. The lemma then says that if z is not beaten by y (so $z \not\leadsto y$), then z must beat all the non-best φ-states, hence z must beat x (so $x \leadsto z$).

Finally, we consider proof systems that include both **Sh** and either **DM** or **WDM**. When proving the truth lemma for these proof systems, whenever $\varphi \in x$ and B^{-1}_φ is inconsistent we will need to find a witness y such that $\varphi \in y$ and $y \not\leadsto x$. The following two lemmas show that such a y exists.

Lemma 55. *Let* **X** *be a proof system that includes either* **WDM** *or* **DM**. *Furthermore, let* $x \in MCS^{\mathbf{X}}_{\Xi}$ *and take* $\Theta = \{\theta \mid B^{-1}_\theta \subseteq x\}$. *Then for every* $\theta \in \Theta$ *and every finite* $\Theta' \subseteq \Theta$, *we have* $P(\theta \mid \bigvee\Theta') \in x$ *and* $P(\bigvee\Theta' \mid \theta) \in x$.

Proof. Let $\Theta' = \{\theta_1, \cdots, \theta_n\}$. For every θ_i we have $B^{-1}_{\theta_i} \subseteq x$, and therefore in particular $\theta_i \in x$. As such, we also have $\bigvee \Theta' \in x$. Because $B^{-1}_{\theta} \subseteq x$, this implies that $\neg \bigvee \Theta' \notin B^{-1}_{\theta}$, and therefore $P(\bigvee \Theta' \mid \theta)$.

We now continue to show that $P(\theta \mid \bigvee \Theta')$, by induction on n. As base case, suppose that $n = 1$, so $\Theta' = \{\theta_1\}$. We have $\theta \in x$ and $B^{-1}_{\theta_1} \subseteq x$, so $\neg\theta \notin B^{-1}_{\theta_1}$, and therefore $P(\theta \mid \theta_1)$.

Suppose then as induction hypothesis that $n > 1$ and that the lemma holds for all strict subsets of Θ'. Suppose towards a contraction that $B(\neg\theta \mid \theta_1 \vee \cdots \vee \theta_n)$. By the induction hypothesis, we have $P(\theta_n \mid \theta_1 \vee \cdots \vee \theta_{n-1})$ and $P(\theta_1 \vee \cdots \vee \theta_{n-1} \mid \theta_n)$. Using **WDM** or **DM**, we then obtain $B(\neg\theta \mid \theta_1 \vee \cdots \vee \theta_{n-1})$. But that contradicts $P(\theta \mid \theta_1 \vee \cdots \vee \theta_{n-1})$, which follows from the induction hypothesis. From this contradiction we conclude that $P(\theta \mid \theta_1 \vee \cdots \vee \theta_n)$. This completes the induction step, and thereby the proof. □

Lemma 56. *Let* **X** *be a proof system that includes* **Sh** *and either* **WDM** *or* **DM**. *Furthermore, let $x \in MCS^{\mathbf{X}}_{\Xi}$ and take $\Theta = \{\theta \mid B^{-1}_{\theta} \subseteq x\}$. Finally, let φ be such that $\varphi \in x$ and B^{-1}_{φ} is inconsistent with Ξ. Then $\{\varphi\} \cup \Box\Xi \cup \neg\Theta$ is consistent.*

Proof. Suppose towards a contradiction that $\{\varphi\} \cup \Box\Xi \cup \neg\Theta$ is inconsistent. Then there are finite subsets $\Xi' \subseteq \Xi$ and $\Theta' = \{\theta_1, \cdots, \theta_n\} \subseteq \Theta$ such that $\{\varphi\} \cup \Box\Xi' \cup \neg\Theta'$ is inconsistent. It follows that $\vdash \bigwedge \Box\Xi' \to (\varphi \leftrightarrow (\varphi \wedge \bigvee \Theta'))$. For every $y \in MCS^{X}_{\Xi}$ we therefore have $\varphi \leftrightarrow (\varphi \wedge \bigvee \Theta') \in y$. By Lemma 31 we then have $\varphi \leftrightarrow (\varphi \wedge \bigvee \Theta') \in \Xi$, which means that $\Box(\varphi \leftrightarrow (\varphi \wedge \bigvee \Theta')) \in x$.

As B^{-1}_{φ} is inconsistent with Ξ, we have $B(\bot \mid \varphi)$. Using **R-Ext** we then obtain $B(\bot \mid \varphi \wedge \bigvee \Theta')$. Using **Sh**, this yields $B(\varphi \to \bot \mid \bigvee \Theta')$. By Lemma 55, we also have $P(\theta_n \mid \theta_1 \vee \cdots \theta_{n-1})$ and $P(\theta_1 \vee \cdots \vee \theta_{n-1} \mid \theta_n)$. Using either **DM** or **WDM**, we obtain $B(\varphi \to \bot \mid \theta_n)$.

As $B^{-1}_{\theta_n} \subseteq x$, this implies that $\varphi \to \bot \in x$. By assumption $\varphi \in x$, so $\bot \in x$, contradicting the consistency of x. From this contradiction, we conclude that $\{\varphi\} \cup \Box\Xi \cup \neg\Theta$ is consistent. □

8. Transitive models

Consider models where $\preceq$ is assumed to be transitive. Such an assumption can be justified if $\preceq$ represents the plausibility, preference or value to a single rational agent. It is generally not justified if $\preceq$ represents a defeat relation, an aggregated judgement from multiple agents, or a judgement from an irrational agent.

8.1. max*,* unm *and* acc *semantics*

It is generally not the case that a model and its transitive closure satisfy the same formulas under max semantics. Consider for example the model M with three states s_1, s_2 and s_3 such that $s_1 \prec s_2 \prec s_3 \prec s_1$, and let M^* be the transitive closure of this model. Then $M, s \models_{\mathsf{max}} B(\bot \mid \top)$, whereas

$M^*, s \not\models_{\mathsf{max}} B(\bot \mid \top)$. However, it turns out that **MOU** is nonetheless complete for $\models^{tr}_{\mathsf{max}}$.

Unfortunately, while this means that for every model M and state s such that $M, s \models \bigwedge \Gamma \wedge \neg\varphi$ there are a transitive model M' and state s' such that $M', s' \models \bigwedge \Gamma \wedge \neg\varphi$, this model M' is not easy to find. As such, we will prove completeness of **MOU** for $\models^{tr}_{\mathsf{max}}$ "from scratch", i.e., by constructing a new canonical model. In this model, states are of the form (x, φ, i), and we define $\Delta_\psi(\Xi) := \{(x, \psi, i) \mid B^{-1}_\psi \subseteq x\}$.

Definition 57. The model $M^{\mathsf{max}(tr)}_\Xi = (S, \preceq, V)$ is given by

- $S = \{(x, \varphi, i) \mid x \in MCS^{\mathbf{MOU}}_\Xi, \varphi \in \mathcal{L}, i \in \mathbb{N}\}$,
- $V(p) = \{(x, \varphi, i) \in s \mid p \in x\}$,
- $\preceq$ is the smallest relation such that
  1. if $(x, \varphi, i) \notin \Delta_\varphi(\Xi)$ then $(x, \varphi, i) \preceq (x, \varphi, j)$ for all $j > i$,
  2. if $(x, \varphi, i) \in \Delta_\varphi(\Xi)$ then for every y such that $\varphi \notin y$, $(x, \varphi, i) \preceq (y, \bot, j)$.

Lemma 58. *$M^{\mathsf{max}(tr)}_\Xi$ is transitive.*

Proof. Suppose that $(x, \varphi, i) \preceq (y, \psi, j)$ and $(y, \psi, j) \preceq (z, \chi, k)$. We distinguish two cases. First, suppose that $(x, \varphi, i) \notin \Delta_\varphi(\Xi)$. Then the only outgoing arrows from (x, φ, i) are of type 1, so we must have $x = y = z$, $\varphi = \psi = \chi$ and $i < j < k$. It follows that, by an arrow of type 1, $(x, \varphi, i) \preceq (z, \chi, k)$.

Secondly, suppose that $(x, \varphi, i) \in \Delta_\varphi(\Xi)$. Then we must have $\psi = \bot$ and $\varphi \notin y$. We have $(y, \bot, i) \notin \Delta_\bot(\Xi)$, since $\Delta_\bot(\Xi) = \emptyset$, so $z = y$, $\chi = \bot$ and $j < k$. Then $(x, \varphi, i) \preceq (z, \bot, k)$. □

We can now present the truth lemma for this canonical model. We omit the proof here, and include it in the appendix.

Lemma 59 (Truth lemma for $M^{\mathsf{max}(tr)}_\Xi$ with respect to max)**.** *$M^{\mathsf{max}(tr)}_\Xi, (x, \varphi, i) \models_{\mathsf{max}} \psi$ if and only if $\psi \in x$.*

Soundness of **MOU** for $\models^{tr}_{\mathsf{max}}$ is trivial, since it is sound for $\models_{\mathsf{max}}$. The truth lemma immediately implies that **MOU** is complete for $\models_{\mathsf{max}}$. We therefore have soundness and completeness.

Theorem 60. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models^{tr}_{\mathsf{max}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU}} \varphi$.*

We showed that $M^{\mathsf{max}(tr)}_\Xi$ is transitive. Furthermore, note that $M^{\mathsf{max}(tr)}_\Xi$ is anti-symmetric. Proposition 8 therefore allows us to conclude that max, unm and acc semantics coincide on the model. Hence $M^{\mathsf{max}(tr)}_\Xi$ is also a canonical model for acc and unm semantics.

Lemma 61 (Truth lemma for $M_\Xi^{\mathsf{max}(tr)}$ with respect to unm and acc)**.** *The following are equivalent: (i)* $\psi \in x$*, (ii)* $M_\Xi^{\mathsf{max}(tr)}, (x, \varphi, i) \models_{\mathsf{unm}} \psi$ *and (iii)* $M_\Xi^{\mathsf{max}(tr)}, (x, \varphi, i) \models_{\mathsf{acc}} \psi$*.*

This also results in the following soundness and completeness theorem.

Theorem 62. *For every* $\Gamma \subseteq \mathcal{L}$ *and every* $\varphi \in \mathcal{L}$*, the following are equivalent: (i)* $\Gamma \models_{\mathsf{unm}}^{tr} \varphi$*, (ii)* $\Gamma \models_{\mathsf{acc}}^{tr} \varphi$ *and (iii)* $\Gamma \vdash_{\mathbf{MOU}} \varphi$*.*

Corollary 63. *The proof system* **ACC** *is strictly weaker than* **MOU***.*

Proof. Suppose $\Gamma \vdash_{\mathbf{ACC}} \varphi$. Then, by soundness of **ACC**, $\Gamma \models_{\mathsf{acc}} \varphi$, and therefore also $\Gamma \models_{\mathsf{acc}}^{tr} \varphi$. By Theorem 62, this implies $\Gamma \vdash_{\mathbf{MOU}} \varphi$. So **MOU** is at least as strong as **ACC**. Furthermore, **MOU** contains the axiom **Sh**, which is unsound for acc. Because **ACC** is sound for acc, it follows that **ACC** is not at least as strong as **MOU**. □

8.2. opt *semantics*

Somewhat surprisingly, considering that max and unm validities are not affected by the restriction to transitive models, **MOU** is not complete for opt on transitive models. We need one additional axiom. Recall that **DM** is as follows,

$$\mathbf{DM} \quad (P(\chi \mid \varphi) \wedge B(\psi \mid \varphi \vee \chi)) \rightarrow B(\psi \mid \chi)$$

and that it makes support transitive. First, let us show that **DM** is sound.

Proposition 64. *The axiom* **DM** *is sound for* opt*-semantics on transitive models.*

Proof. Suppose that M is a transitive model and that $M, s \models P(\chi \mid \varphi) \wedge B(\psi \mid \varphi \vee \chi)$. Take any $s_1 \in \mathsf{opt}(\llbracket \chi \rrbracket)$. We will show that $s_1 \in \mathsf{opt}(\llbracket \varphi \vee \chi \rrbracket)$, so take any $s_2 \in \llbracket \varphi \vee \chi \rrbracket$. By $M, s \models P(\chi \mid \varphi)$, there is some $s_3 \in \mathsf{opt}(\llbracket \varphi \rrbracket)$ such that $M, s_3 \models \chi$.

If $s_2 \in \llbracket \chi \rrbracket$, then $s_2 \preceq s_1$ by the χ-optimality of s_1. If $s_2 \notin \llbracket \chi \rrbracket$, then $s_2 \in \llbracket \varphi \rrbracket$ and therefore, by φ-optimality of s_3, $s_2 \preceq s_3$. Furthermore, by χ-optimality of s_1, $s_3 \preceq s_1$. Because the model is transitive, this implies that $s_2 \preceq s_1$. This holds for every $s_2 \in \llbracket \varphi \vee \chi \rrbracket$, so $s_1 \in \mathsf{opt}(\llbracket \varphi \vee \chi \rrbracket)$.

From $B(\psi \mid \varphi \vee \chi)$ it then follows that $M, s_1 \models \psi$. This holds for every $s_1 \in \mathsf{opt}(\llbracket \chi \rrbracket)$, so $M, s \models B(\psi \mid \chi)$. This suffices to show that **DM** is sound on transitive models. □

Left to show is that adding the axiom **DM** to **MOU** yields a complete proof system for opt semantics on transitive models. We start by defining a canonical model.

Definition 65. The model $M_\Xi^{\mathsf{opt}(tr)} = (S, \preceq, V)$ is given by

- $S = MCS_\Xi^{\mathbf{MOU+DM}} \times \{0, 1\}$,
- $V(p) = \{(x, i) \in S \mid p \in x\}$,

- $(x,i) \preceq (y,j)$ if and only if $x \rightsquigarrow y$.

For $s = (x,i) \in S$ we abuse notation by identifying x and s wherever that should not cause confusion.

Lemma 66. *$M_\Xi^{\mathsf{opt}(tr)}$ is transitive.*

Proof. Follows immediately from transitivity of support (Lemma 51). ☐

We now consider an auxiliary lemma, followed by the truth lemma. For both lemmas, the proof is included in the appendix.

Lemma 67. *For every state x of $M_\Xi^{\mathsf{opt}(tr)}$ and every formula φ, if $x \in \mathsf{opt}([\varphi])$, then $B_\varphi^{-1} \subseteq x$.*

Lemma 68 (Truth lemma for $M_\Xi^{\mathsf{opt}(tr)}$)**.** *$M_\Xi^{\mathsf{opt}(tr)}, x \models \varphi$ if and only if $\varphi \in x$.*

Completeness follows immediately from the truth lemma and the fact that $M_\Xi^{\mathsf{opt}(tr)}$ is transitive. We already showed soundness, so the proof system is sound and complete.

Theorem 69. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models_{\mathsf{opt}}^{tr} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU+DM}} \varphi$.*

9. Total models

Consider models where $\preceq$ is assumed to be total. Such an assumption can often be justified if $\preceq$ represents an ordering based on a single characteristic, but it usually not justified if $\preceq$ is based on multiple incomparable characteristics.

9.1. max*,* opt *and* acc *semantics*

The max-semantics and acc-semantics are sensitive only to the strict relation $\prec$, see Lemma 11. It follows that a formula is max-valid on all models if and only if it is max-valid on total models, and similarly for acc-validity.

Lemma 70. *Let $M = (S, \preceq, V)$ be any model and let $M' = (S, \preceq', V)$ be the model given by $s_1 \preceq' s_2$ if and only if $s_1 \preceq s_2$ or $s_1 \perp s_2$. Then for every $\varphi \in \mathcal{L}$, we have $M, s \models_{\mathsf{max}} \varphi$ if and only if $M', s \models_{\mathsf{max}} \varphi$ and $M, s \models_{\mathsf{acc}} \varphi$ if and only if $M', s \models_{\mathsf{acc}} \varphi$.*

Furthermore, on total models maximality and optimality collapse, see Proposition 8, so $\Gamma \models_{\mathsf{max}}^{to} \varphi$ if and only if $\Gamma \models_{\mathsf{opt}}^{to} \varphi$. If follows that **MOU** is sound and complete for max and opt on total models and that **ACC** is sound and complete for acc on total models.

Corollary 71. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models_{\mathsf{max}}^{to} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU}} \varphi$.*

Corollary 72. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models_{\mathsf{opt}}^{to} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU}} \varphi$.*

Corollary 73. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models_{\mathsf{acc}}^{to} \varphi$ if and only if $\Gamma \vdash_{\mathbf{ACC}} \varphi$.*

Unlike max, opt and acc semantics, the validities with respect to unm semantics do change when we restrict ourselves to total models. Mainly, this is because the unmatched states in a total model must be unique.

Lemma 74. *Let $M = (S, \preceq, V)$ be a total model. Then for every $X \subseteq S$, if* $\mathsf{unm}(X) \neq \emptyset$ *then X is a singleton.*

Proof. Suppose towards a contradiction that $x_1, x_2 \in \mathsf{unm}(X)$ and $x_1 \neq x_2$. Then $x_1 \not\preceq x_2$, as x_1 is unmatched in X, and $x_2 \not\preceq x_1$, as x_2 is unmatched in X. So $x_1 \perp x_2$. But that contradicts M being total. □

As a result, **Un** becomes sound. Recall that this axiom is given by

$$\textbf{Un} \quad P(\varphi \mid \psi) \to B(\varphi \mid \psi).$$

Lemma 75. *The axiom* **Un** *is sound for* unm*-semantics on total models.*

Proof. Let $M = (S, \preceq, V)$ be a total model and $s \in S$, and suppose that $M, s \models P(\varphi \mid \psi)$. Then there is some $x \in \mathsf{unm}(\llbracket\psi\rrbracket)$ such that $M, x \models \varphi$. From Lemma 74 it then follows that $\mathsf{unm}(\llbracket\psi\rrbracket) = \{x\}$, so for every $y \in \mathsf{unm}(\llbracket\psi\rrbracket)$ we have $M, y \models \varphi$, and therefore $M, s \models B(\varphi \mid \psi)$. □

Adding only the **Un** axiom to **MOU** does not, however, result in a proof system that is complete for unm on total models. We also need a kind of monotonicity axiom. This is where we use the axiom **WDM** that we introduced previously:

$$\textbf{WDM} \quad (P(\chi \mid \varphi) \wedge P(\varphi \mid \chi) \wedge B(\psi \mid \varphi \vee \chi)) \to B(\psi \mid \chi).$$

Unlike the stronger **DM**, **WDM** is sound for unm-semantics on total models.

Lemma 76. *The axiom* **WDM** *is sound for* unm*-semantics on total models.*

Proof. Let $M = (S, \preceq, V)$ be a total model, and suppose that $M, s \models P(\chi \mid \varphi) \wedge P(\varphi \mid \chi) \wedge B(\psi \mid \varphi \vee \chi)$. By $M, s \models P(\chi \mid \varphi)$ and $M, s \models P(\varphi \mid \chi)$ there are $x \in \mathsf{unm}(\llbracket\chi\rrbracket) \cap \llbracket\varphi\rrbracket$ and $y \in \mathsf{unm}(\llbracket\varphi\rrbracket) \cap \llbracket\chi\rrbracket$. We will show that $x = y$. Suppose towards a contradiction that $x \neq y$. Then, because M is total, we have either $x \preceq y$, which contradicts x being unmatched in $\llbracket\chi\rrbracket$, or $y \preceq x$, which contradicts y being unmatched in $\llbracket\varphi\rrbracket$. By contradiction, it follows that $x = y$.

Now, take any $z \in \llbracket\varphi \vee \chi\rrbracket$. As we have $z \in \llbracket\varphi\rrbracket$ or $z \in \llbracket\chi\rrbracket$, unmatchedness of x in $\llbracket\varphi\rrbracket$ and $\llbracket\chi\rrbracket$ implies that $x \not\preceq z$. So $x \in \mathsf{unm}(\llbracket\varphi \vee \chi\rrbracket)$. From $M, s \models B(\psi \mid \varphi \vee \chi)$ it therefore follows that $x \in \llbracket\psi\rrbracket$. Furthermore, from $x \in \mathsf{unm}(\llbracket\chi\rrbracket)$ it follows by Lemma 74 that $\mathsf{unm}(\llbracket\chi\rrbracket) = \{x\}$. As $x \in \llbracket\psi\rrbracket$ this implies that $B(\psi \mid \chi)$. We have now shown that **WDM** is sound for unm on total models. □

Now that we have shown that **Un** and **WDM** are sound for unm-semantics on total models, all that is left to do is to show that **MOU** + **Un** + **WDM** is complete.

Definition 77. The model $M_{\Xi}^{\mathsf{unm}(to)} = (S, \preceq, V)$ is given by

- $S = S_{good} \cup S_{bad}$ where $S_{good} = \{x \in MCS_{\Xi}^{\mathbf{MOU+Un+WDM}} \mid x \text{ is good}\}$ and $S_{bad} = \{(x,i) \in MCS_{\Xi}^{\mathbf{MOU+Un+WDM}} \times \{0,1\} \mid x \text{ is bad}\}$
- $V(p) = \{x \in S_{good} \mid p \in x\} \cup \{(x,i) \in S_{bad} \mid p \in x\}$,
- $s_1 \preceq s_2$ if and only if $x_2 \not\rightsquigarrow x_1$, where $s_1 = x_1$ or $s_1 = (x_1, i)$ and $s_2 = x_2$ or $s_2 = (x_2, i)$.

As with the other proof systems, we consider a few auxiliary lemmas before proving the truth lemma.

Lemma 78. *If $x \overset{\varphi}{\rightsquigarrow} y$ and $y \overset{\psi}{\rightsquigarrow} x$ then $x = y$.*

Proof. Because the conditions on x and y are symmetric, it suffices to show that $x \subseteq y$. So take any $\chi \in x$, and suppose towards a contradiction that $\chi \notin y$.

First, note that by **Id** we have $B(\varphi \mid \varphi)$ and $B(\psi \mid \psi)$, so $\varphi \in y$ and $\psi \in x$. Because $\chi \notin y$ and $\varphi, \psi \in y$ we have $(\psi \wedge \varphi) \to \chi \notin y$. So $\neg B((\psi \wedge \varphi) \to \chi \mid \varphi)$.

By contraposition of **Sh** this implies $\neg B(\chi \mid (\psi \wedge \varphi) \wedge \varphi)$. Using **R-Ext** and **L-Ext** (plus some propositional reasoning) we obtain $\neg B(\neg\neg\chi \mid \psi \wedge \varphi)$, i.e., $P(\neg\chi \mid \psi \wedge \varphi)$.

Using **Un**, we then get $B(\neg\chi \mid \psi \wedge \varphi)$, which by **Sh** implies that $B(\varphi \to \neg\chi \mid \psi)$. Because $B_{\psi}^{-1} \subseteq x$, this implies that $\varphi \to \neg\chi \in x$. We have $\varphi \in x$, which implies that $\neg\chi \in x$. But we chose χ such that $\chi \in x$, so we have arrived at a contradiction. We therefore have $\chi \in y$. This holds for every $\chi \in x$, so $x \subseteq y$, as was to be shown.

□

Corollary 79. *$M_{\Xi}^{\mathsf{unm}(to)}$ is total.*

Proof. If $x \not\preceq y$ and $y \not\preceq x$, then $x \rightsquigarrow y$ and $y \rightsquigarrow x$. By the preceding lemma, this implies that $x = y$. □

Now, let us continue with the truth lemma for unm-semantics on total models. As with most other truth lemmas, see the appendix for the proof.

Lemma 80 (Truth lemma for $M_{\Xi}^{\mathsf{unm}(to)}$)**.** *$M_{\Xi}^{\mathsf{unm}(to)}, x \models \varphi$ if and only if $\varphi \in x$.*

Completeness follows immediately. We already showed soundness, so we now have the desired soundness and completeness.

Theorem 81. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models_{\mathsf{unm}}^{to} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU+DM+Un}} \varphi$.*

10. Anti-symmetric models

Consider models where $\preceq$ is assumed to be anti-symmetric. Such an assumption can be justified in the presence of a tie-breaking mechanism, but is typically unjustified otherwise.

10.1. max*,* unm *and* acc *semantics*

As in Section 9.1, recall that max and acc semantics are sensitive only to the strict relation $\prec$ (Lemma 11).

Lemma 82. *Let $M = (S, \preceq, V)$ be any model and let $M' = (S, \preceq', V)$ be the model given by $s_1 \preceq' s_2$ if and only if $s_1 \prec s_2$. Then for every $\varphi \in \mathcal{L}$, we have $M, s \models_{\mathsf{max}} \varphi$ if and only if $M', s \models_{\mathsf{max}} \varphi$ and $M, s \models_{\mathsf{acc}} \varphi$ if and only if $M', s \models_{\mathsf{acc}} \varphi$.*

Furthermore, on anti-symmetric models maximality and unmatchedness collapse, see Proposition 8, so $\Gamma \models^{as}_{\mathsf{max}} \varphi$ if and only if $\Gamma \models^{as}_{\mathsf{unm}} \varphi$. If follows that **MOU** is sound and complete for max and unm on anti-symmetric models and that **ACC** is sound and complete for acc on anti-symmetric models.

Corollary 83. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models^{as}_{\mathsf{max}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU}} \varphi$.*

Corollary 84. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models^{as}_{\mathsf{unm}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU}} \varphi$.*

Corollary 85. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models^{as}_{\mathsf{acc}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{ACC}} \varphi$.*

10.2. opt *semantics*

Recall from Lemma 10 that if M_1 and M_2 satisfy (i) $s \prec_1 t$ iff $s \prec_2 t$, (ii) $s \perp_1 t$ iff $s \approx_2 t$ and (iii) $s \approx_1 t$ iff $s \perp_2 t$, for all s and t, then $\mathsf{opt}^{M_1}(X) = \mathsf{unm}^{M_2}(X)$ for all X. Now, note that if M_1 and M_2 satisfy these three conditions then M_1 is total if and only if M_2 is anti-symmetric, and vice versa. As the proof system $\mathbf{MOU} + \mathbf{WDM} + \mathbf{Un}$ is sound and complete for unm semantics on total models, it follows that the same proof system is sound and complete for opt semantics on anti-symmetric models.

Corollary 86. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models^{as}_{\mathsf{opt}}$ if and only if $\Gamma \vdash_{\mathbf{MOU}+\mathbf{WDM}+\mathbf{Un}} \varphi$.*

11. Transitive and total models

Consider models where $\preceq$ is assumed to be both transitive and total.

11.1. max*,* opt *and* acc *semantics*

On transitive models max and acc semantics coincide, and on total model max and opt semantics coincide, see Proposition 8. It follows that, in particular, any axiomatization that is sound and complete for one of these three semantics on transitive and total models is sound and complete for the other two semantics as well. Left to do is to find such an axiomatization.

The proof system $\mathbf{MOU} + \mathbf{DM}$ is sound for opt on transitive models. As such, our proof system for max, opt and acc on transitive and total models will need to be at least as strong as $\mathbf{MOU} + \mathbf{DM}$. In fact, we will need one further axiom, namely **DR**.

$$\textbf{DR} \quad B(\varphi \mid \psi \vee \chi) \to (B(\varphi \mid \psi) \vee B(\varphi \mid \chi))$$

First, let us show that **DR** is sound with respect to opt (and therefore also with respect to max and acc) on transitive total models.

Lemma 87. *The axiom* **DR** *is sound for* opt*-semantics on transitive total models.*

Proof. Let $M = (S, \preceq, V)$ be any transitive total model. We will show that **DR** holds by contraposition, so suppose that $M, s \not\models B(\varphi \mid \psi) \vee B(\varphi \mid \chi)$. Then $M, s \models P(\neg\varphi \mid \psi)$ and $M, s \models P(\neg\varphi \mid \chi)$. Let s_1 and s_2 be the witnesses for $P(\neg\varphi \mid \psi)$ and $P(\neg\varphi \mid \chi)$, respectively, so $s_1 \in \mathsf{opt}([\![\psi]\!]) \cap [\![\neg\varphi]\!]$ and $s_2 \in \mathsf{opt}([\![\chi]\!]) \cap [\![\neg\varphi]\!]$. Note that, in particular, for every $t \in [\![\psi \vee \chi]\!]$ we have either $t \in [\![\psi]\!]$ or $t \in [\![\chi]\!]$, and therefore $t \preceq s_1$ or $t \preceq s_2$.

We distinguish two cases. First, suppose that $s_1 = s_2$. Then $t \preceq s_1$ for all $t \in [\![\psi \vee \chi]\!]$, so $s_1 \in \mathsf{opt}([\![\psi \vee \chi]\!]) \cap [\![\varphi]\!]$, and therefore $M, s \models P(\neg\varphi \mid \psi \vee \chi)$.

Secondly, suppose that $s_1 \neq s_2$. As M is total, we then have either $s_1 \preceq s_2$ or $s_2 \preceq s_1$. Suppose, without loss of generality, that $s_2 \preceq s_1$. For every $t \in [\![\psi]\!]$, we have $t \preceq s_1$ because $t \in \mathsf{opt}([\![\psi]\!])$. For every $t \in [\![\chi]\!]$, we have $t \preceq s_1$ because $t \preceq s_2$ (since $s_2 \in \mathsf{opt}([\![\chi]\!])$), $s_2 \preceq s_1$ and the model is transitive. It follows that $t \preceq s_1$ for each $t \in [\![\psi \vee \chi]\!]$, so $s_1 \in \mathsf{opt}([\![\psi \vee \chi]\!])$. We have $s_1 \in [\![\neg\varphi]\!]$, so this implies $M, s \models P(\neg\varphi \mid \psi \vee \chi)$.

In either case, $M, s \models P(\neg\varphi \mid \psi \vee \chi)$, and therefore $M, s \not\models B(\varphi \mid \psi \vee \chi)$. We have now shown that if $M, s \not\models B(\varphi \mid \psi)$ and $M, s \not\models B(\varphi \mid \chi)$ then $M, s \not\models B(\varphi \mid \psi \vee \chi)$, so it follows that $M, s \models B(\varphi \mid \psi \vee \chi) \to (B(\varphi \mid \psi) \vee B(\varphi \mid \chi))$. This holds for every transitive total model M and every state s, so **DR** is sound on transitive total models. □

The next step is to show that **MOU** + **DM** + **DR** is complete for opt on transitive total models. As with the other completeness proofs, we start by defining the canonical model, then prove a number of auxiliary lemmas and after that prove the truth lemma.

Definition 88. The model $M_\Xi^{\mathsf{opt}(tr,to)} = (S, \preceq, V)$ is given by

- $S = S_g \cup S_b \cup S_u$, where
  - $S_g = \{(x,i) \in MCS_\Xi^{\mathbf{MOU}+\mathbf{DM}+\mathbf{DR}} \times \mathbb{N} \mid x \text{ is good}\}$
  - $S_b = \{(x,i) \in MCS_\Xi^{\mathbf{MOU}+\mathbf{DM}+\mathbf{DR}} \times \mathbb{N} \mid x \text{ is bad and } supp(x) \neq \emptyset\}$
  - $S_u = \{(x,i) \in MCS_\Xi^{\mathbf{MOU}+\mathbf{DM}+\mathbf{DR}} \times \mathbb{N} \mid x \text{ is bad and } supp(x) = \emptyset\}$
- $(x,i) \preceq (y,j)$ if and only if
  - $(y,j) \in S_g$ and $supp(y) \subseteq supp(x)$,
  - $(y,j) \in S_b$ and $supp(y) \subset supp(x)$,
  - $(x,i), (y,j) \in S_b$ and $supp(y) \subseteq supp(x)$
  - $(x,i) \notin S_u$ and $(y,j) \in S_u$ or

  - $(x,i) \in S_u$ and $(y,j) \in S_u$ and $i \leq j$.

- $(x,i) \in V(p)$ if and only if $p \in x$.

Recall that, because of **DM**, support is transitive (Lemma 51). Furthermore, because of **DR**, for any two states one must support the other (Lemma 53). Finally, recall that $[\varphi] = \{(x,i) \in S \mid \varphi \in x\}$.

Lemma 89. *For every $(x,i),(y,j) \in S$, either $supp(x) \subseteq supp(y)$ or $supp(y) \subseteq supp(x)$.*

Proof. Suppose towards a contradiction that $supp(x) \not\subseteq supp(y)$ and $supp(y) \not\subseteq supp(x)$. Then there are $u \in supp(x) \setminus supp(y)$ and $v \in supp(y) \setminus supp(x)$. Because $u \in supp(x)$ and $v \in supp(y)$, both u and v are good. Therefore, either u supports v or v supports u. As support is transitive, this implies that x supports v or y supports u, contradicting either $u \in supp(x) \setminus supp(y)$ or $v \in supp(y) \setminus supp(x)$. ☐

Lemma 90. *For every $(x,i),(y,j) \in S$, if x supports y then $supp(y) \subseteq supp(x)$.*

Proof. This follows immediately from the transitivity of support. ☐

Now we can begin to show that $M_\Xi^{\mathsf{opt}(tr,to)}$ satisfies the required properties. The proofs can be found in the appendix.

Lemma 91. *The model $M_\Xi^{\mathsf{opt}(tr,to)}$ is transitive.*

Lemma 92. *The model $M_\Xi^{\mathsf{opt}(tr,to)}$ is total.*

Lemma 93. *If $B_\varphi^{-1} \subseteq x$, then $(x,i) \in \mathsf{opt}([\varphi])$.*

Lemma 94. *If $B_\varphi^{-1} \not\subseteq x$, then $(x,i) \notin \mathsf{opt}([\varphi])$.*

The truth lemma now follows immediately.

Lemma 95 (Truth lemma for $M_\Xi^{\mathsf{opt}(tr,to)}$)**.** *$M_\Xi^{\mathsf{opt}(tr,to)}, (x,i) \models \varphi$ if and only if $\varphi \in x$.*

Proof. By induction and a case distinction on the main connective of φ. As usual, only case $\varphi = B(\gamma \mid \delta)$ is interesting.

- Suppose $B(\gamma \mid \delta) \notin x$. Then $B_\delta^{-1} \cup \{\neg\gamma\}$ is consistent and can therefore be extended to a $y \in MCS_\Xi^{\mathbf{MOU+DM+DR}}$. By Lemma 93 we have $y \in \mathsf{opt}([\delta])$ and therefore, by the induction hypothesis, $(y,j) \in \mathsf{opt}([\![\delta]\!])$. Furthermore, $\gamma \notin y$ and therefore, by the induction hypothesis, $M_\Xi^{\mathsf{opt}(tr,to)}, (y,j) \not\models \gamma$. It follows that $M_\Xi^{\mathsf{opt}(tr,to)}, (x,i) \not\models B(\gamma \mid \delta)$.

- Suppose $B(\gamma \mid \delta) \in x$. Then we have $\gamma \in B_\delta^{-1}$ and therefore, by Lemma 94, $\mathsf{opt}([\delta]) \subseteq [\gamma]$. By the induction hypothesis, this implies that $\mathsf{opt}([\![\delta]\!]) \subseteq [\![\gamma]\!]$, so $M_\Xi^{\mathsf{opt}(tr,to)}, (x,i) \models B(\gamma \mid \delta)$.

□

If follows that $\mathbf{MOU} + \mathbf{DM} + \mathbf{DR}$ is sound and complete for opt-semantics on transitive total models.

Theorem 96. *For every* $\Gamma \subseteq \mathcal{L}$ *and every* $\varphi \in \mathcal{L}$*, we have* $\Gamma \models^{tr,to}_{\mathsf{opt}} \varphi$ *if and only if* $\vdash_{\mathbf{MOU+DM+DR}} \varphi$.

Because opt, max and acc semantics coincide on transitive total models we also have the following two corollaries.

Corollary 97. *For every* $\Gamma \subseteq \mathcal{L}$ *and every* $\varphi \in \mathcal{L}$*, we have* $\Gamma \models^{tr,to}_{\mathsf{max}} \varphi$ *if and only if* $\vdash_{\mathbf{MOU+DM+DR}} \varphi$.

Corollary 98. *For every* $\Gamma \subseteq \mathcal{L}$ *and every* $\varphi \in \mathcal{L}$*, we have* $\Gamma \models^{tr,to}_{\mathsf{acc}} \varphi$ *if and only if* $\vdash_{\mathbf{MOU+DM+DR}} \varphi$.

11.2. unm *semantics*

Previously, we showed that $\mathbf{MOU} + \mathbf{Un} + \mathbf{WDM}$ is sound and complete for unm semantics on total models. So our proof system for unm on transitive total models will need to be at least as strong. In fact, our proof system here will need to be strictly stronger than $\mathbf{MOU}+\mathbf{Un}+\mathbf{WDM}$. For one thing, the axiom $\mathbf{DM}$ is sound for these models.

Lemma 99. *The axiom* $\mathbf{DM}$ *is sound for* unm *semantics on transitive total models.*

Proof. Let $M = (S, \preceq, V)$ be a transitive total model, let $s \in S$ and suppose that $M, s \models P(\chi \mid \varphi) \wedge B(\psi \mid \varphi \vee \chi)$. To show is that $M, s \models B(\psi \mid \chi)$. So take any $s_1 \in \mathsf{unm}([\![\chi]\!])$.

From $P(\chi \mid \varphi)$ it follows that there is some $s_2 \in \mathsf{unm}([\![\varphi]\!]) \cap [\![\chi]\!]$. We will now show that $s_1 \in \mathsf{unm}([\![\varphi \vee \chi]\!])$. So take any $s_3 \in [\![\varphi \vee \chi]\!] \setminus \{s_1\}$. We first consider the case where s_1, s_2 and s_3 are all distinct.

By totality of the model, we have $s_2 \preceq s_1$ or $s_1 \preceq s_2$. Because $s_1 \in \mathsf{unm}([\![\chi]\!])$, we have $s_1 \not\preceq s_2$ and therefore $s_2 \prec s_1$. From $s_3 \in [\![\varphi \vee \chi]\!]$ it follows that $s_3 \in [\![\varphi]\!]$ or $s_3 \in [\![\chi]\!]$. If $s_3 \in [\![\chi]\!]$, then $s_1 \not\preceq s_3$ because s_1 is unmatched in $[\![\chi]\!]$. If $s_3 \in [\![\varphi]\!]$ then $s_2 \not\preceq s_1$ because s_2 in unmatched in $[\![\varphi]\!]$. By totality of the model this implies that $s_3 \prec s_2$. If we would have $s_1 \preceq s_3$ then it would follow, by transitivity, that $s_1 \preceq s_2$, contradicting $s_1 \in \mathsf{unm}([\![\chi]\!])$. So we have $s_1 \not\preceq s_3$.

If $s_1 = s_2$ then s_1 is unmatched in both $[\![\varphi]\!]$ and $[\![\chi]\!]$, so it is unmatched in $[\![\varphi \vee \chi]\!]$. If $s_2 = s_3$, then $s_3 \in [\![\chi]\!]$ and thereofre, by unmatchedness of s_1, $s_1 \not\preceq s_3$. In every case, $s_1 \not\preceq s_3$. This holds for every $s_3 \in [\![\varphi \vee \chi]\!]$, so $s_1 \in \mathsf{unm}([\![\varphi \vee \chi]\!])$.

We have now shown that $\mathsf{unm}([\![\chi]\!]) \subseteq \mathsf{unm}([\![\varphi \vee \chi]\!])$. From $B(\psi \mid \varphi \vee \chi)$ it therefore follows that $B(\psi \mid \chi)$, as was to be shown. □

We also need the axiom $\mathbf{SP}$, which we recall is given by

$$\mathbf{SP} \quad (P(\chi \mid \varphi) \wedge P(\xi \mid \psi) \wedge B(\neg\varphi \mid \varphi \vee \psi)) \rightarrow P(\xi \mid (\varphi \wedge \neg\chi) \vee \psi)$$

Lemma 100. *The axiom* **SP** *is sound for* unm *semantics on transitive total models.*

Proof. Suppose that M is total and that $M, s \models P(\chi \mid \varphi) \wedge P(\xi \mid \psi) \wedge B(\neg\varphi \mid \varphi \vee \psi)$. Then there are states $s_1 \in \mathsf{unm}([\![\varphi]\!]) \cap [\![\chi]\!]$ and $s_2 \in \mathsf{unm}([\![\psi]\!]) \cap [\![\xi]\!]$.

Suppose now towards a contradiction that $s_1 = s_2$. Then s_1 is unmatched in both $[\![\varphi]\!]$ and $[\![\psi]\!]$, so it is unmatched in $[\![\varphi \vee \psi]\!]$. This contradicts $M, s \models B(\neg\varphi \mid \varphi \vee \psi)$. From this contradiction we conclude that $s_1 \neq s_2$.

Suppose then, again towards a contradiction, that $s_1 \not\preceq s_2$. Because M is total it follows that $s_2 \prec s_1$. For every $s_3 \in [\![\psi]\!]$, if we were to have $s_1 \preceq s_3$ then, because M is transitive, we would have $s_2 \preceq s_3$. But that is impossible, as s_2 is unmatched in $[\![\psi]\!]$. So we have $s_1 \not\preceq s_3$. It follows that $s_1 \in \mathsf{unm}([\![\psi]\!])$. Together with $s_1 \in \mathsf{unm}([\![\varphi]\!])$, this implies $s_1 \in \mathsf{unm}([\![\varphi \vee \psi]\!])$, contradicting $B(\neg\varphi \mid \varphi \vee \psi)$. From this contradiction we conclude that $s_1 \preceq s_2$.

Finally, suppose towards a contradiction that for some $s_4 \in [\![\varphi]\!] \setminus \{s_1\}$ we have $s_2 \preceq s_4$. Together with $s_1 \preceq s_2$, using the transitivity of M, this implies that $s_1 \preceq s_4$, contradicting $s_1 \in \mathsf{unm}([\![\varphi]\!])$. From this contradiction we conclude that $s_2 \not\preceq s_4$ for every $s_4 \in [\![\varphi]\!] \setminus \{s_1\}$.

Now note that, because $s_1 \in [\![\chi]\!]$, for every $s_5 \in [\![(\varphi \wedge \neg\chi) \vee \psi]\!]$ we have either $s_5 \in [\![\varphi]\!] \setminus \{s_1\}$ or $s_5 \in [\![\psi]\!]$. In either case, $s_2 \not\preceq s_5$. So $s_2 \in \mathsf{unm}([\![(\varphi \wedge \neg\chi) \vee \psi]\!])$. Because $s_2 \in [\![\xi]\!]$, this implies that $M, s \models P(\xi \mid (\varphi \wedge \neg\chi) \vee \psi)$, as was to be shown. □

Left to show is the completeness of **MOU** + **Un** + **DM** + **SP**.

Definition 101. The model $M_\Xi^{\mathsf{unm}(tr,to)} = (S, \preceq, V)$ is given by

- $S = \{x \in MCS_\Xi^{\mathbf{MOU+Un+DM+SP}} \mid x \text{ is good}\} \cup$ $(\{x \in MCS_\Xi^{\mathbf{MOU+Un+DM+SP}} \mid x \text{ is bad}\} \times \{0, 1\})$,
- $x \in V(p)$ or $(x, i) \in V(p)$ if and only if $p \in x$.
- if x is good, then $x \preceq x$ and if x is bad, then $(x, i) \preceq (x, i)$,
- if $x \neq y$ and x and y are good, then $x \preceq y$ iff $y \not\rightsquigarrow x$,
- if x is bad and y is good, then $(x, i) \prec y$ if $x \rightsquigarrow y$ and $y \prec (x, i)$ otherwise,
- if $x \neq y$ and x and y are bad, then $(x, i) \preceq (y, j)$ iff $\mathit{supp}(y) \subseteq \mathit{supp}(x)$,
- $(x, i) \preceq (x, j)$.

Note that we have $(x, 0) \preceq s$ if and only if $(x, 1) \preceq s$, so we can abuse notation by writing $x \preceq s$.

In order to show that $M_\Xi^{\mathsf{unm}(tr,to)}$ is a canonical model for unm on transitive total models, we first need to show that the model is indeed transitive and total. The proofs for these lemmas are included in the appendix.

Lemma 102. $M_\Xi^{\mathsf{unm}(tr,to)}$ *is total.*

Lemma 103. $M_{\Xi}^{\mathsf{unm}(tr,to)}$ *is transitive.*

Now that we have shown that the model is transitive and total, what remains is to prove the truth lemma. We first prove two auxiliary lemmas.

Lemma 104. *If* $B_{\varphi}^{-1} \subseteq x$ *then* x *is unmatched in* $[\varphi]$.

Proof. Let $y \neq x$ be any state such that $x \preceq y$. From $B_{\varphi}^{-1} \subseteq x$ we know that x is good. By Definition 101, we have the following: if y were bad and $y \rightsquigarrow x$ then $y \prec x$, and if y were good and $y \rightsquigarrow x$ then $x \not\preceq y$. Either would contradict $x \preceq y$, so $y \not\rightsquigarrow x$. In particular, this means that y does not φ-support x. Together with $B_{\varphi}^{-1} \subseteq x$, this implies that $\varphi \notin y$.

So x is unmatched in $[\varphi]$. $\square$

Lemma 105. *If* $B_{\varphi}^{-1} \not\subseteq x$, *then* x *is not unmatched in* $[\varphi]$.

Proof. If $\varphi \notin x$ the lemma is trivial, so assume that $\varphi \in x$.

First, suppose that B_{φ}^{-1} is consistent, so there is some y such that $B_{\varphi}^{-1} \subseteq y$, which implies that x supports y. If x is bad, this implies that $x \prec y$. If x is good, then we have $y \not\preceq x$, so by the previously shown totality $x \prec y$.

Suppose then that B_{φ}^{-1} is inconsistent but x is good. Then take $\Theta = \{\theta \mid B_{\theta}^{-1} \subseteq x\}$. By Lemma 56, $\varphi \wedge \neg\Theta$ is consistent with $\Box\Xi$ and can therefore be extended to some y. By construction, this y does not support x. If x and y are both good, or if x is good and y is bad, this immediately implies $x \preceq y$.

Finally, if x is bad recall that a bad state is represented by two states $(x, 0)$ and $(x, 1)$, with the property that $(x, 0) \approx (x, 1)$. So no bad state can be unmatched. $\square$

The truth lemma now follows immediately, since $B(\varphi \mid \psi)$ is the only hard part in it.

Lemma 106. $M_{\Xi}^{\mathsf{unm}(tr,to)}, x \models_{\mathsf{unm}} \varphi$ *if and only if* $\varphi \in x$.

Furthermore, since we already proved soundness, we now have the desired result for the axiomatization.

Theorem 107. *For every* $\Gamma \subseteq \mathcal{L}$ *and every* $\varphi \in \mathcal{L}$, *we have* $\Gamma \models_{\mathsf{unm}}^{tr,to} \varphi$ *if and only if* $\Gamma \vdash_{\mathbf{MOU+WDM+Un+SP}} \varphi$.

12. Transitive and anti-symmetric models

Consider models where $\preceq$ is assumed to be both transitive and anti-symmetric, i.e., a partial order.

On transitive models max and acc semantics coincide, and on anti-symmetric models max and opt semantics coincide, see Proposition 8. It follows that, in particular, any axiomatization that is sound and complete for one of these three semantics on transitive and anti-symmetric models is sound and complete for the other two semantics as well. Left to do is to find such an axiomatization.

Here we once again use the fact that max semantics are not sensitive to addition or removal of bi-directional arrows.

Lemma 108. *Let $M = (S, \preceq, V)$ be a transitive model, and let $M' = (S, \preceq', V)$ be such that $s_1 \preceq' s_2$ if and only if $s_1 \prec s_2$. Then M' is transitive and anti-symmetric.*

Proof. Anti-symmetry of M' is trivial: we can have at most one of $s_1 \prec s_2$ and $s_2 \prec s_1$, and therefore at most one of $s_1 \preceq' s_2$ and $s_2 \preceq' s_1$. Left to show is that M' is transitive. So suppose that $s_1 \preceq' s_2 \preceq' s_3$. Then $s_1 \prec s_2 \prec s_3$, so by transitivity of M it follows that $s_1 \preceq s_3$. Now, suppose towards a contradiction that $s_3 \preceq s_1$. Then we would have $s_2 \prec s_3 \preceq s_1$ and therefore, by transitivity, $s_3 \preceq s_2$, contradicting $s_2 \prec s_3$. From this contradiction, we conclude that $s_3 \not\preceq s_1$. Together with $s_1 \preceq s_3$ this implies that $s_1 \prec s_3$ and therefore $s_1 \preceq' s_3$. □

Lemma 109. *If $\Gamma \not\models^{tr}_{\mathsf{max}} \varphi$ then $\Gamma \not\models^{tr,as}_{\mathsf{max}} \varphi$.*

Proof. Suppose that $\Gamma \not\models^{tr}_{\mathsf{max}} \varphi$. Then there are a transitive model $M = (S, \preceq, V)$ and a state $s \in S$ such that $M, s \models \gamma$ for all $\gamma \in \Gamma$ and $M, s \not\models \varphi$. Let $M' = (S, \preceq', V)$ be the model such that $s_1 \preceq' s_2$ if and only if $s_1 \prec s_2$.

By Lemma 11, for every $X \subseteq S$ we have $\mathsf{max}^M(X) = \mathsf{max}^{M'}(X)$. It follows that $M', s \models \gamma$ for all $\gamma \in \Gamma$ and $M', s \not\models \varphi$. Furthermore, by Lemma 108, the model M' is transitive and anti-symmetric. So $\Gamma \not\models^{tr,as}_{\mathsf{max}} \varphi$. □

The other direction, that $\Gamma \models^{tr}_{\mathsf{max}} \varphi$ implies $\Gamma \models^{tr,as}_{\mathsf{max}} \varphi$, is trivial. So we have $\Gamma \models^{tr}_{\mathsf{max}} \varphi$ if and only if $\Gamma \models^{tr,as}_{\mathsf{max}} \varphi$. As **MOU** is sound and complete for max semantics on transitive models, it is also sound and complete for max semantics on transitive total models.

Corollary 110. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models^{tr,as}_{\mathsf{max}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU}} \varphi$.*

Furthermore, as remarked above, max, unm and acc semantics coincide on transitive anti-symmetric models, so **MOU** is sound and complete for those semantics as well.

Corollary 111. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models^{tr,as}_{\mathsf{unm}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU}} \varphi$.*

Corollary 112. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models^{tr,as}_{\mathsf{acc}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU}} \varphi$.*

The axiom **DM** is sound for opt semantics on transitive models. The axioms **WDM** and **Un** are sound for opt semantics on anti-symmetric models. It follows that **DM** and **Un** are sound for opt semantics on transitive anti-symmetric models. Furthermore, we will show that $\mathbf{MOU} + \mathbf{DM} + \mathbf{Un}$ is in fact complete for opt on these models.

Definition 113. The model $M_\Xi^{\mathsf{opt}(tr,as)} = (S, \preceq, V)$ is given by

- $S = MCS^{\mathbf{MOU+DM+Un}} \times \mathbb{N}$,
- $V(p) = \{(x,i) \in S \mid p \in x\}$,
- $(x,i) \preceq (y,j)$ if and only if
  1. $x \neq y$ and $x \rightsquigarrow y$,
  2. $x = y$, x is good and $j < i$ or
  3. $x = y$, x is bad and $i < j$.

First, we need to show that this model is indeed anti-symmetric and transitive.

Lemma 114. *$M_\Xi^{\mathsf{opt}(tr,as)}$ is anti-symmetric.*

Proof. Suppose towards a contradiction that $(x,i) \preceq (y,j)$ and $(y,j) \preceq (x,i)$. We distinguish two cases.

- Suppose that $x \neq y$. Then from $(x,i) \preceq (y,j)$ and $(y,j) \preceq (x,i)$ it follows that $y \in supp(x)$ and $x \in supp(y)$. But in the presence of **Un** this implies that $x = y$, contradicting our assumption that $x \neq y$.
- Suppose that $x = y$. Then from $(x,i) \preceq (y,j)$ and $(y,j) \preceq (x,i)$ it follows that $i < j$ and $j < i$, which is a contradiction.

□

Lemma 115. *$M_\Xi^{\mathsf{opt}(tr,as)}$ is transitive.*

Proof. This follows from the transitivity of support (in the presence of **DM**) and of $<$. □

Now, we can work on the truth lemma. See the appendix for the proof.

Lemma 116 (Truth lemma for $M_\Xi^{\mathsf{opt}(tr,as)}$)**.** *$M_\Xi^{\mathsf{opt}(tr,as)}, (x,i) \models \varphi$ if and only if $\varphi \in x$.*

Completeness follows immediately from the truth lemma, and we had already established soundness. We therefore have the following theorem.

Theorem 117. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models_{\mathsf{opt}}^{tr,as} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU+DM+Un}} \varphi$.*

13. Totality and anti-symmetry

Consider models where $\preceq$ is assumed to be both total and anti-symmetric. These models are also sometimes referred to as *tournaments*.

13.1. max, opt *and* unm *semantics*

On total models, maximality and optimality coincide. On anti-symmetric models, maximality and unmatchedness coincide. On total anti-symmetric models, therefore, all three of these semantics are the same.

The system **MOU+Un+WDM** is sound for unm-semantics on total models and for opt-semantics on anti-symmetric models. It is therefore also sound for unm and opt semantics on total anti-symmetric models. We will show that it is also complete for unm on these models, which implies that it is complete for opt and max as well.

Let $M = (S, \preceq, V)$ be any total model. We create a total, anti-symmetric model $M' = (S', \preceq', V')$ from M. For certain states $s'_1, s'_2 \in S'$ we need to have either $s'_1 \prec' s'_2$ or $s'_2 \prec' s'_1$ but it does not matter which. In order to make a choice between the two options, let $\leq$ be any linear order on S. Now, M' is constructed in the following way:

- $S' = S \times \{0, 1\}$,
- for every p, $V'(p) = V(p) \times \{0, 1\}$ and
  - $(s, 0) \prec' (s, 1)$,
  - if $s_1 \prec s_2$, then $(s_1, i) \prec' (s_2, j)$ for every $i, j \in \{0, 1\}$,
  - if $s_1 \approx s_2$ and $s_1 \neq s_2$, then
    * $(s_1, 1) \prec' (s_2, 0)$,
    * $(s_2, 1) \prec' (s_1, 0)$,
    * if $s_1 \leq s_2$ then $(s_1, 1) \prec' (s_2, 1)$ and $(s_1, 0) \prec' (s_2, 0)$.

By construction, M' is total and anti-symmetric. Furthermore, for any $X \subset S$ we have $s \in \mathsf{unm}(X)$ if and only if $(s, 1) \in \mathsf{unm}(X \times \{0, 1\})$.

It follows that, for every φ and every s, we have $M, s \models \varphi$ if and only if $M', (s, 1) \models \varphi$. Since this process can be done for every model M, we have $\Gamma \models^{to,as}_{\mathsf{unm}} \varphi$ iff $\Gamma \models^{to}_{\mathsf{unm}} \varphi$.

Corollary 118. *For every $\Gamma \subset \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models^{to,as}_{\mathsf{unm}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU+WDM+Un}} \varphi$.*

Because max, opt and unm coincide, we also get the following.

Corollary 119. *For every $\Gamma \subset \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models^{to,as}_{\mathsf{max}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU+WDM+Un}} \varphi$.*

Corollary 120. *For every $\Gamma \subset \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models^{to,as}_{\mathsf{opt}} \varphi$ if and only if $\Gamma \vdash_{\mathbf{MOU+WDM+Un}} \varphi$.*

13.2. acc *semantics*

Finally, we come to acc semantics on total, anti-symmetric models. This is a special case; for each other combination of semantics and frame conditions we have a sound and strongly complete axiomatization. For this variant, we instead show that such an axiomatization does not exist.

First, however, we introduce a definition and a lemma that we will use later.

Definition 121 (Upset)**.** Let $M = (S, \preceq, V)$, $X \subseteq S$ and $x \in X$. The *upset* of x in X, denoted $x^{\uparrow(X)}$, is the set of states y such that there is a path $x \prec y_1 \prec \cdots \prec y_n \prec y$ in X.

Lemma 122. *Let $M = (S, \preceq, V)$ be total and anti-symmetric. Then* $\mathsf{acc}(X) = \{y \mid \forall x \in X : y \in x^{\uparrow(X)}\}$.

Proof. Take any $y \in \{y \mid \forall x \in X : y \in x^{\uparrow(X)}\}$ and let $z \in X$ be such that $y \prec z$. Then for every $x \in X$ we have $x \prec \cdots \prec y$ and therefore $x \prec \cdots \prec y \prec z$. So $z \in \{y \mid \forall x \in X : y \in x^{\uparrow(X)}\}$. So $\{y \mid \forall x \in X : y \in x^{\uparrow(X)}\}$ is retentive.

Now, suppose towards a contradiction that there is a $Y \subset \{y \mid \forall x \in X : y \in x^{\uparrow(X)}\}$ that is retentive. Then there are $y \in Y$ and $z \in \{y \mid \forall x \in X : y \in x^{\uparrow(X)}\} \setminus Y$. Due to the retentiveness of Y, we have $y^{\uparrow(X)} \subseteq Y$. But $z \in y^{\uparrow(X)}$, contradicting the choice of z. It follows that $\{y \mid \forall x \in X : y \in x^{\uparrow(X)}\}$ is not just retentive but minimally so. $\square$

Now, we can show that no sound and strongly complete axiomatization exists for acc on total anti-symmetric models. We do this by showing that acc semantics are not compact on total anti-symmetric models, i.e., there is an inconsistent set $\Gamma \subseteq \mathcal{L}$ such that every finite subset of Γ is consistent. A proof system demanding finite proofs, which all the proof systems we consider in this paper do, therefore cannot detect such inconsistency.

Let $q, p_1, p_2, \cdots \in \mathbf{P}$ be distinct, and let Γ consist of exactly the following formulas:

- $\Box(p_i \to \neg p_j)$ for all $i \neq j$,
- $\Box(q \to \neg p_i)$ for all i,
- $P(\top \mid q)$,
- $P(q \mid \neg p_i)$ for all i,
- $B(\neg p_j \mid \neg p_i)$ for all $j < i$,
- $P(p_i \mid \top)$ for all i.

Lemma 123. $\Gamma \models^{to,anti}_{\mathsf{acc}} \bot$.

Proof. Suppose towards a contradiction that Γ is satisfied in some total anti-symmetric model $M = (S, \preceq, V)$. We define $T = \{s \in S \mid \forall i : s \not\models p_i\}$ and $S_i = \{s \in S \mid s \models p_i\}$. Furthermore, we define $T^+ = T \cap \bigcap_i \mathsf{acc}(\llbracket \neg p_i \rrbracket)$.

We structure the remainder of the proof using 8 claims.

Claim 1: $\{T, S_0, S_1, \cdots\}$ is a partition of S.
Proof of Claim 1: It is immediate from the definition of the sets that $T \cup \bigcup_i S_i = S$. It is also immediate that $T \cap S_i = \emptyset$. Left to show is that $S_i \cap S_j = \emptyset$ for $i \neq j$, which follows from the fact that $M \models \Gamma$ and therefore, in particular, $M \models \Box(p_i \to \neg p_j)$.

Claim 2: For every i, $S_i \neq \emptyset$. Furthermore, $\llbracket q \rrbracket \subseteq T$ and $\llbracket q \rrbracket \neq \emptyset$.
Proof of Claim 2: Non-emptyness of S_i follows from $M \models P(p_i \mid \top)$. That $\llbracket q \rrbracket \subseteq T$ follows from $M \models \Box(q \to \neg p_i)$ for all i. Finally, non-emptyness of $\llbracket q \rrbracket$ follows from $M \models P(\top \mid q)$.

Claim 3: $\mathsf{acc}(\llbracket q \rrbracket) \neq \emptyset$ and $\mathsf{acc}(\llbracket q \rrbracket) \subseteq \mathsf{acc}(\llbracket \neg p_i \rrbracket)$, for every i.
Proof of Claim 3: Non-emptyness of $\mathsf{acc}(\llbracket q \rrbracket)$ follows from $M \models P(\top \mid q)$. Furthermore, from $M \models P(q \mid \neg p_i)$ it follows that $\mathsf{acc}(\llbracket \neg p_i \rrbracket) \cap \llbracket q \rrbracket$ is nonempty.

So let $y \in \mathsf{acc}(\llbracket \neg p_i \rrbracket) \cap \llbracket q \rrbracket$ and $z \in \mathsf{acc}(\llbracket q \rrbracket)$. Then for any $x \in \llbracket \neg p_i \rrbracket$ we have $y \in x^{\uparrow(\llbracket \neg p_i \rrbracket)}$ and $z \in y^{\uparrow(\llbracket q \rrbracket)}$, by Lemma 122. We also have $\llbracket q \rrbracket \subseteq \llbracket \neg p_i \rrbracket$, so the path $x \prec \cdots \prec y \prec \cdots \prec z$ lies entirely in $\llbracket \neg p_i \rrbracket$, which implies that $z \in x^{\uparrow(\llbracket \neg p_i \rrbracket)}$. This is true for any $x \in \llbracket \neg p_i \rrbracket$, so by Lemma 122 we have $z \in \mathsf{acc}(\llbracket \neg p_i \rrbracket)$. Therefore, $\mathsf{acc}(\llbracket q \rrbracket) \subseteq \mathsf{acc}(\llbracket \neg p_i \rrbracket)$.

Claim 4: T^+ is non-empty.
Proof of Claim 4: By Claim 3, $\mathsf{acc}(\llbracket q \rrbracket) \neq \emptyset$ and $\mathsf{acc}(\llbracket q \rrbracket) \subseteq T^+$.

Claim 5: For every $x \in T^+$ and every $y \in T \setminus T^+$, $y \prec x$.
Proof of Claim 5: Since the model is total and anti-symmetric, we have $y \prec x$ or $x \prec y$. So suppose towards a contradiction that $x \prec y$. Because $y \in T$, we have $y \in \llbracket \neg p_i \rrbracket$ for each i. Because $x \in T^+$, $x \in \mathsf{acc}(\llbracket \neg p_i \rrbracket)$ for each i. As such, by retentiveness of $\mathsf{acc}(\llbracket \neg p_i \rrbracket)$ we have $y \in \mathsf{acc}(\llbracket \neg p_i \rrbracket)$ for each i, and therefore $y \in T^+$, contradicting $y \in T \setminus T^+$.

Claim 6: For every $x \in T^+$ and every $y \in S \setminus T^+$, $y \prec x$.
Proof of Claim 6: For $y \in T$ we already showed the claim, so assume without loss of generality that $y \in S_i$ for some i. Suppose towards a contradiction that $y \not\prec x$, and therefore $x \prec y$. We have $x \in \mathsf{acc}(\llbracket \neg p_{i+1} \rrbracket)$ and $y \in \llbracket \neg p_{i+1} \rrbracket)$ so by retentiveness we also have $y \in \mathsf{acc}(\llbracket \neg p_{i+1} \rrbracket)$. As $y \in S_i$ this implies that $P(p_i \mid \neg p_{i+1})$. But that contradicts $B(\neg p_i \mid \neg p_{i+1})$. As we have arrived at a contradiction we must have $y \prec x$.

Claim 7: $\mathsf{acc}(\llbracket \top \rrbracket) \subseteq T^+$.
Proof of Claim 7: Take any $x \in T^+$. By Claim 6, we have $x^{\uparrow(\llbracket \top \rrbracket)} \subseteq T^+$. By Lemma 122 this implies that $\mathsf{acc}(\llbracket \top \rrbracket) \subseteq T^+$.

Claim 8: For every i, $M \not\models P(p_i \mid \top)$.
Proof of Claim 8: This follows immediately from Claim 7.

Note that Claim 8 contradicts our assumption that Γ is satisfied in M. It therefore follows that Γ is unsatisfiable, as was to be shown. $\square$

Left to show is that any finite subset of Γ is satisfiable.

Lemma 124. *For any finite $\Gamma' \subseteq \Gamma$, there is a total anti-symmetric model $M = (S, \preceq, V)$ such that $M \models \Gamma'$.*

Proof. Because Γ' is finite, there is some $n \in \mathbb{N}_{\geq 3}$, such that Γ' is a subset of

- $\Box(p_i \to \neg p_j)$ for all $i \neq j$,
- $\Box(q \to \neg p_i)$ for all i,
- $P(\top \mid q)$,
- $P(q \mid \neg p_i)$ for all $i \leq n$,
- $B(\neg p_j \mid \neg p_i)$ for al $j < i \leq n$,
- $P(p_i \mid \top)$ for all $i \leq n$.

We will show that Γ' is satisfied in the following model $M = (S, \preceq, V)$:

- $S = \{t, s_0, s_1, \cdots, s_n + 2\}$,
- $s_i \succ s_{i+1}$ for all $0 \leq i \leq n+1$,
- $s_i \succ s_j$ for all i, j where $j \geq i+2$
- $t \prec s_{n+2}$,
- $t \succ s_i$ for all $i < n+2$,
- $V(p_i) = \{s_i\}$ if $i \leq n+2$ and $V(p_i) = \emptyset$ if $i > n+2$,
- $V(q) = \{t\}$.

So in M, the state s_{n+2} is preferred over the single q-state t, while t is preferred over all other s_i states. Generally, s_j is preferred over s_i if $j > i$, with the one exception being if $j = i+1$, in which case s_i is preferred over s_j. A schematic drawing of this model is shown in Figure 11.

In M, all atoms are mutually exclusive, so we have $M \models \Box(p_i \to \neg p_j)$ for $i \neq j$ and $M \models \Box(q \to \neg p_i)$. Additionally, we have $[\![q]\!] = \mathsf{acc}([\![q]\!]) = \{t\}$, so $M \models P(\top \mid q)$.

Now, consider any $s_j \in [\![\neg p_i]\!]$, with $i \leq n$. If $j < n+2$, then $s_j < t$. If $j = n+2$ then $s_j = s_{n+2} < s_{n+1} < t$. So for every s_j, we have $t \in s_j^{\uparrow([\![\neg p_i]\!])}$. By Lemma 122, this implies $t \in \mathsf{acc}([\![\neg p_i]\!])$ and therefore $M \models P(q \mid \neg p_i)$.

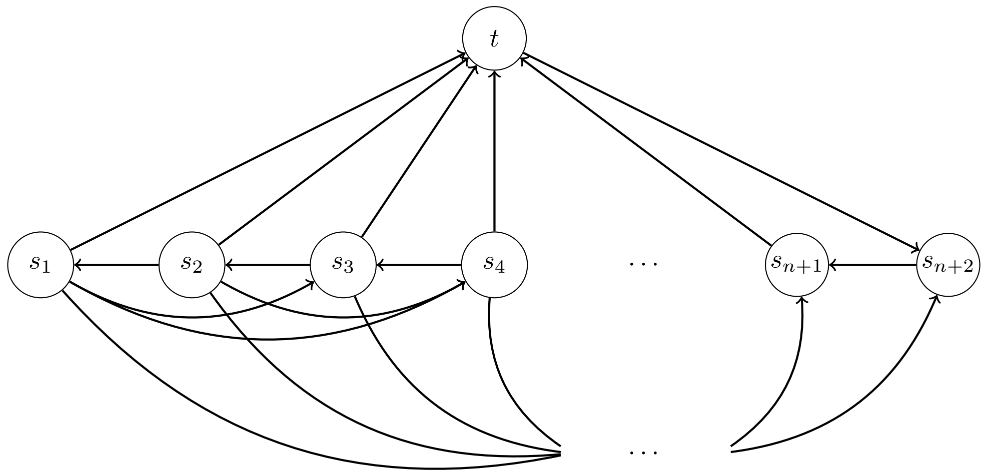

Figure 11: The model used in the proof of Lemma 124.

Next, consider any $s_j \in [\![\neg p_i]\!]$, with $j < i \leq n$. The only path from s_{n+2} to s_j in M is $s_{n+2} < s_{n+1} < s_n < \cdots < s_{j+1} < s_j$. Because $j < i \leq n$, at least one of those states is not in $[\![\neg p_i]\!]$, so $s_j \notin s_{n+2}^{\uparrow([\![\neg p_i]\!]}$. By Lemma 122 this implies $s_j \notin \mathsf{acc}([\![\neg p_i]\!])$, and therefore $M \models B(\neg p_j \mid \neg p_i)$.

Finally, consider $[\![\top]\!] = S$. We have $t < s_{n+2} < s_{n+2} < \cdots < s_2 < s_1 < t$, so for any two states x, y we have $x \in y^{\uparrow([\![\top]\!]}$ and therefore, by Lemma 122, $\mathsf{acc}([\![\top]\!]) = [\![\top]\!]$. This implies that $M \models P(p_i \mid \top)$.

We have now shown that $M \models \Gamma'$, so that set is satisfiable. □

Theorem 125. *On total, anti-symmetric models,* acc *is not compact.*

Corollary 126. *There is no sound and strongly complete axiomatization for* acc *on total anti-symmetric models.*

14. Transitive, total and anti-symmetric models

Finally, consider the models that are transitive, total and anti-symmetric, i.e., $\preceq$ is a linear order. In this case, all four semantics coincide. Furthermore, we inherit soundness for all of the axioms we considered so far, i.e., **MOU** + **DM** + **DR** + **Un** + **SP** is sound for any of the semantics on these models. We will show that this proof system is also complete.

We begin with the construction of the canonical model. Because we only care about the best states in each extension, a maximal consistent set does not fix an ordering between any two bad states (which are not best in any extension). In order to create a linear order we therefore have to introduce some form of tie-breaking. We do this using a function $f : 2^{\mathcal{L}} \to \mathbb{R}$.

Definition 127. Let $f : 2^{\mathcal{L}} \to \mathbb{R}$ be an injection. The model $M_{\Xi}^{tr,to,anti} = (S, \preceq, V)$ is given by

- $S = MCS_{\Xi}^{\mathbf{MOU+DM+DR+Un+SP}} \times \mathbb{N}$,
- $V(p) = \{(x,i) \mid p \in x\}$,
- $(x,i) \preceq (y,j)$ if
  - $x = y$, x is good and $j < i$,
  - $x = y$, x is bad and $i < j$,
  - $supp(y) \subset supp(x)$,
  - $supp(y) = supp(x)$ and y is good or
  - $supp(y) = supp(x)$, x and y are both bad and $f(x) < f(y)$.

First, let us show that this model is indeed transitive, total and anti-symmetric. Note that in the presence of these axioms,

- support is transitive (Lemma 51),
- for every x, y either $supp(x) \subseteq supp(y)$ or $supp(y) \subseteq supp(x)$ (Lemma 89),
- no two good states support each other (Lemma 78).

Lemma 128. *$M_{\Xi}^{tr,to,anti}$ is total, transitive and anti-symmetric.*

Proof. If $i \neq j$ then $i < j$ or $i > j$, so $(x,i) \preceq (x,j)$ or $(x,i) \succeq (x,j)$. If x and y are different, we have $supp(x) \subseteq supp(y)$ or $supp(y) \subseteq supp(x)$, so the remaining three clauses of the definition of $\preceq$ guarantee that $(x,i) \preceq (y,j)$ or $(x,i) \succeq (y,j)$. So the model is total.

Transitivity is inherited from the transitivity of $<$ (in $\mathbb{N}$ as well as $\mathbb{R}$) and $\subseteq$.

For anti-symmetry, note that no two good states support each other, so from $supp(x) = supp(y)$ and y being good it follows that x is bad. This implies that $(x,i) \prec (y,j)$. In all other cases, anti-symmetry is inherited from anti-symmetry of $<$ and $\subset$. □

Lemma 129. *If $B_{\varphi}^{-1} \subseteq x$, then $(x,0) \in \mathsf{max}([\varphi])$.*

Proof. Because $B_{\varphi}^{-1} \subseteq x$, we know that x is good. Therefore, $(x,0) \succ (x,i)$ for all $i > 0$.

Consider then any $y \in [\varphi] \setminus \{x\}$. Then y supports x, so by transitivity of support we have $supp(x) \subseteq supp(y)$. If $supp(x) \subset supp(y)$ then $(x,0) \succ (y,j)$. If $supp(x) = supp(y)$ then, because x is good, we still have $(x,0) \succ (y,j)$. □

Lemma 130. *If $B_{\varphi}^{-1} \not\subseteq x$, then $(x,i) \notin \mathsf{max}([\varphi])$.*

Proof. If $(x, i) \notin [\varphi]$, then trivially $(x, i) \notin \mathsf{max}([\varphi])$. Furthermore, if x is bad then $(x, i) \prec (x, i + 1)$ and therefore $(x, i) \notin \mathsf{max}([\varphi])$. Assume therefore that $x \in [\varphi]$ and x is good.

If B_φ^{-1} is consistent, then there is some y such that $B_\varphi^{-1} \subseteq y$. Then x supports y, so $supp(y) \subseteq supp(x)$. No two good states support each other and x supports itself, as it is good. So $supp(y) \subset supp(x)$, and therefore $(x, i) \prec (y, j)$.

If B_φ^{-1} is not consistent, then $\varphi \cup \Box\Xi \cup \neg\Theta$ is consistent, where $\Theta = \{\theta \mid B_\theta^{-1} \subseteq x\}$, see Lemma 56. So it can be extended to some y that, by construction, does not support x. Because x is good it supports itself, so $supp(x) \not\subseteq supp(y)$. This yields $supp(y) \subset supp(x)$, and therefore $(x, i) \prec (y, j)$. □

The truth lemma, and then completeness, follow immediately.

Lemma 131. $M_\Xi^{\mathsf{unm}(tr,to,as)}, (x, i) \models_{\mathsf{max}} \varphi$ *if and only if* $\varphi \in x$.

Furthermore, since we already proved soundness, we now have the desired result for the axiomatization.

Theorem 132. *For every* $\Gamma \subseteq \mathcal{L}$ *and every* $\varphi \in \mathcal{L}$*, we have* $\Gamma \models_{\mathsf{max}}^{tr,to,as} \varphi$ *if and only if* $\Gamma \vdash_{\mathbf{MOU+WDM+DR+Un+SP}} \varphi$.

The other semantics coincide with max on these models.

Corollary 133. *For every* $\Gamma \subseteq \mathcal{L}$ *and every* $\varphi \in \mathcal{L}$*, we have* $\Gamma \models_{\mathsf{opt}}^{tr,to,as} \varphi$ *if and only if* $\Gamma \vdash_{\mathbf{MOU+WDM+DR+Un+SP}} \varphi$.

Corollary 134. *For every* $\Gamma \subseteq \mathcal{L}$ *and every* $\varphi \in \mathcal{L}$*, we have* $\Gamma \models_{\mathsf{unm}}^{tr,to,as} \varphi$ *if and only if* $\Gamma \vdash_{\mathbf{MOU+WDM+DR+Un+SP}} \varphi$.

Corollary 135. *For every* $\Gamma \subseteq \mathcal{L}$ *and every* $\varphi \in \mathcal{L}$*, we have* $\Gamma \models_{\mathsf{acc}}^{tr,to,as} \varphi$ *if and only if* $\Gamma \vdash_{\mathbf{MOU+WDM+DR+Un+SP}} \varphi$.

15. Conclusion

We have studied four different semantics for conditional logics based on a preference relation: $\mathsf{max}, \mathsf{opt}, \mathsf{unm}$ and acc. Of these, max and opt are relatively well studied, at least on transitive models that satisfy some kind of well-foundedness property. Semantics based on unm are a kind of dual form of opt, and acc is new in this context.

Importantly, we did not place the usual constraints on the preference relation: $\preceq$ was not necessarily assumed to be transitive (although we did consider the transitive case as well), limited or smooth. We then also looked at cases where $\preceq$ is assumed to satisfy some constraints, namely any combination of transitivity, totality and anti-symmetry.

With one exception, for each combination of a semantics ($\mathsf{max}, \mathsf{opt}, \mathsf{unm}, \mathsf{acc}$) and a set $U \subseteq \{tr, to, as\}$ of constraints, we have given a sound and strongly complete axiomatization. For the one exception, acc semantics on total anti-symmetric models, we showed that such an axiomatization does not exist.

Some clear directions for future research remain. Firstly, while we showed that no sound and *strongly* complete axiomatization for acc on total anti-symmetric models exist, it may be possible to find a sound and *weakly* complete axiomatization.

Secondly, some further semantics, in addition to max, opt, unm and acc, seem potentially interesting. In particular, one could consider *unique best* semantics, which requires a best state to be both optimal and unmatched, or *weakly acceptable* semantics, which considers a state to be best if it is acceptable or if there are no acceptable states.

It may also be interesting to study other constraints on $\preceq$, such as well-foundedness. Unfortunately, the connection between axioms on the one hand and constraints on the other is not as straightforward as the correspondence results for standard modal logic [31]. In particular, we cannot treat each constraint in isolation, so each extra constraint under consideration increases the number of variants that need to be considered exponentially. In general, the correspondence theory of these languages remain an open area of research with only a few results available so far.

Finally, we could consider constraints not on $\preceq$ but on other parts of the model. For example, we could look at models with a *unique valuation*, where any two different states are distinguishable from one another. This prevents the existence of two states that are indistinguishable by any formula, but where one is still preferred over the other. As such, a unique valuation is a reasonable assumption if, and only if, the set of atoms encodes all the information that is required to determine whether one state is better than another.

Acknowledgments

We are indebted to Frederik van de Putte and Xavier Parent for helpful feedback on an earlier version of this paper that was presented at the Workshop on Logics of Agency, Counterfactuals and Norms (4 December 2020, University of Amsterdam). We would also like to thank Eric Pacuit and Rahul Savani for their helpful suggestions about how to present and contextualize the results. Finally, we are grateful to three anonymous reviewers for their thorough and constructive feedback.

References

[1] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

[2] L. Åqvist. Deontic logic. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic. Volume II: Extensions of Classical Logic*, pages 605–714. Reidel, 1984.

[3] A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. Van der Hoek, and M. Wooldridge, editors,

Logic and the Foundation of Game and Decision Theory—LOFT 07, volume 3 of *Texts in Logic and Games*, pages 11–58. Amsterdam University Press, 2008.

[4] C. Boutilier. Conditional logics of normality as modal systems. In *Proceedings of AAAI'90*, pages 594–599, 1990.

[5] C. Boutilier. Conditional logics of normality: A modal approach. *Artificial Intelligence*, 68:87–154, 1994.

[6] J. Burgess. Quick completeness proofs for some logics of conditionals. *Notre Dame Journal of Formal Logic*, 22(1):76–84, 1981.

[7] L. F. del Cerro, A. Herzig, and J. Lang. From ordering-based nonmonotonic reasoning to conditional logics. *Artificial Intelligence*, 66(2):375–393, 1994.

[8] L. Demey. Some remarks on the model theory of epistemic plausibility models. *Journal of Applied Non-Classical Logic*, 21:375–395, 2011.

[9] C. Domshlak, E. Hüllermeier, S. Kaci, and H. Prade, editors. *Representing, Processing, and Learning Preferences: Theoretical and Practical Challenges*, volume 175, numbers 7–8 of *Artificial Intelligence*, 2011.

[10] M. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proceedings of FOCS'05*, pages 142–151, 2005.

[11] D. Grossi, W. van der Hoek, and L. B. Kuijer. Logics of Preference when There Is No Best. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning*, pages 455–464, 9 2020.

[12] J. Halpern. From causal models to counterfactual structures. In *Proceedings of the 12th International Conference on the Principles of Knowledge Representation and Reasoning (KR'10)*, pages 153–160. AAAI, 2010.

[13] B. Hansson. An analysis of some deontic logics. *Nous*, 3:373–398, 1969.

[14] J. F. Horty. Agency and obligation. *Synthese*, 108(2):269–307, 1996.

[15] H. Katsuno and A. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.

[16] H. Katsuno and K. Satoh. A unified view of consequence relations, belief revision, and conditional logic. In *Proceedings of IJCAI-91, 12th International Joint Conference on Artificial Intelligence*, pages 406–412, 1991.

[17] Barteld Kooi and Allard Tamminga. Moral conflicts between groups of agents. *Journal of Philosophical Logic*, 37:1–21, 2008.

[18] A. Kratzer. Partition and revision: The semantics of counterfactuals. *Journal of Philosophical Logic*, 10(2), 1981.

[19] S. Kraus, D. J. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1-2):167–207, 1990.

[20] J.-F. Laslier. *Tournament Solutions and Majority Voting*. Springer, 1997.

[21] D. Lewis. *Counterfactuals*. Blackwell, Oxford, 1973.

[22] D. Lewis. Semantic analysis for dyadic deontic logic. In S. Stenlund, editor, *Logical Theory and Semantical Analysis*, pages 1–14. Reidel Publishing Company, 1974.

[23] D. Makinson. Five faces of minimality. *Studia Logica*, 52:339–379, 1993.

[24] W. Nejdl. The P-systems: A semantic classification of logics of nonmonotonicity. In *Proceedings of AAAI'91*, pages 366–372. AAAI Press, 1991.

[25] X. Parent. On the strong completeness of Åqvist's dyadic deontic logic G. In *Proceedings of DEON'08*, pages 189–202, 2008.

[26] X. Parent. Maximality vs. optimality in dyadic deontic logic. Completeness results for systems in Hansson's tradition. *Journal of Philosophical Logic*, 43:1101–1128, 2013.

[27] T. Schwartz. On the possibility of rational policy evaluation. *Theory and Decision*, 1:89–106, 1970.

[28] A. K. Sen. Choice functions and revealed preference. *The Review of Economic Studies*, 38(3), 1971.

[29] W. Spohn. An analysis of Hansson's dyadic deontic logic. *Journal of Philosophical Logic*, 4:237–252, 1975.

[30] R. Stalnaker. A theory of conditionals. In N. Rescher, editor, *Studies in Logical Theory*, pages 98–112. Oxford, 1968.

[31] J. van Benthem. Correspondence theory. In *Handbook of Philosophical Logic*, pages 325–408. Springer, 1984.

[32] J. van Benthem. Semantic parallels in natural language and computation. In H. Ebbinghaus, editor, *Logic Colloquium '87*. North-Holland, 1989.

[33] J. van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logic*, 17(2), 2007.

[34] J. van Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, 2011.

[35] F. van de Putte and C. Strasser. Preferential semantics using non-smooth preference relations. *Journal of Philosophical Logic*, 43:903–942, 2013.

[36] G. H. Von Wright. *The Logic of Preference*. Edinburgh University Press, 1963.

Appendix A. Proofs

Lemma 35 (Truth lemma for **MOU** and opt)**.** *For every* $\Box$*-maximal* Ξ*, every* $\varphi, \psi \in \mathcal{L}$ *and every* $x \in MCS^{\mathbf{MOU}}_{\Xi}$ *we have* $M^{\mathbf{MOU}}_{\Xi}, (x, \psi) \models_{\mathsf{opt}} \varphi$ *if and only if* $\varphi \in x$*.*

Proof. By induction on the complexity of φ. If φ is atomic then the lemma follows immediately from the choice of V in $M^{\mathbf{MOU}}_{\Xi}$. Suppose therefore as induction hypothesis that φ is not atomic and that the lemma holds for all subformulas of φ. We continue by case distinction on the main connective of φ.

- Suppose the main connective of φ is Boolean. Then the lemma follows immediately from **PL**, **MP** and the induction hypothesis.
- Suppose $\varphi = \Box\chi$. Then

  $$\begin{aligned} \varphi \in x &\Leftrightarrow \chi \in \Xi \\ &\Leftrightarrow \forall y \in MCS_{\Xi} : \chi \in y \\ &\Leftrightarrow \forall s \in S : M^{\mathbf{MOU}}_{\Xi}, s \models \chi \\ &\Leftrightarrow M^{\mathbf{MOU}}_{\Xi}, (x, \psi) \models \Box\chi, \end{aligned}$$

  where the third equivalence uses the induction hypothesis.
- Suppose $\varphi = B(\gamma \mid \delta)$.
  - Suppose $\varphi \in x$, and take any $(y, \chi) \in \mathsf{opt}(\llbracket\delta\rrbracket)$. By the induction hypothesis, $\llbracket\delta\rrbracket = [\delta]$. From $(y, \chi) \in \mathsf{opt}([\delta])$ it follows that $(y, \chi) \in [\delta]$ and therefore that $(y, \theta) \in [\delta]$ for every θ. By optimality of (y, χ) it follows that $(y, \theta) \preceq (y, \chi)$. By the first clause for $\preceq$ in the construction of $M^{\mathbf{MOU}}_{\Xi}$, this implies that y is *not* χ-bad, so $B^{-1}_{\chi} \subseteq y$.

    Additionally, suppose towards a contradiction that $\delta \to \chi \notin \Xi$. Then there is a $z \in MCS_{\Xi}$ such that $\delta \in z$ and $\chi \notin z$. By the assumption that $(y, \chi) \in [\delta]$ we also have $\delta \in y$. From $\delta \in y$ and $\chi \notin z$ it follows that $(z, \delta) \not\preceq (y, \chi)$, due to the second exception clause in the construction of $M^{\mathbf{MOU}}_{\Xi}$. Yet $\delta \in z$, so that contradicts the optimality of (y, χ) in $[\delta]$. From the contradiction, we conclude that $\delta \to \chi \in \Xi$.

    Because $\delta \to \chi \in \Xi$, we also have $\delta \leftrightarrow (\delta \wedge \chi) \in \Xi$ and therefore $\Box(\delta \leftrightarrow (\delta \wedge \chi)) \in x$. Using **R-Ext** and the fact that we assumed $B(\gamma \mid \delta) \in x$ we obtain $B(\gamma \mid \delta \wedge \chi) \in x$. Using **Sh** this yields $B(\delta \to \gamma \mid \chi) \in x$.

    We showed that $B^{-1}_{\chi} \subseteq y$, so we get $\delta \to \gamma \in y$. Together with $\delta \in y$, this implies that $\gamma \in y$. By the induction hypothesis we obtain $M^{\mathbf{MOU}}_{\Xi}, (y, \chi) \models \gamma$. This holds for every $(y, \chi) \in \mathsf{opt}(\llbracket\delta\rrbracket)$, so we have $M^{\mathbf{MOU}}_{\Xi}, (x, \psi) \models B(\gamma \mid \delta)$, which was to be shown.
  - Suppose $\varphi \notin x$, so $B(\gamma \mid \delta) \notin x$. By Lemma 30 it follows that $B^{-1}_{\delta} \cup \Box\Xi \cup \{\neg\gamma\} \cup \{\neg\Box\zeta \mid \zeta \in \mathcal{L} \setminus \Xi\}$ is consistent, and can therefore be

extended to a maximal consistent set y. By construction, $\Box^{-1}y = \Xi$, so $y \in MCS_\Xi$.

We claim that $(y, \delta) \in \mathsf{opt}(\llbracket\delta\rrbracket)$. So take any $(z, \epsilon) \in \llbracket\delta\rrbracket$. Then $(z, \epsilon) \preceq (y, \delta)$ unless one of the two clauses from Definition 34 applies. By construction, $B_\delta^{-1} \subseteq y$. The first clause from Definition 34 therefore doesn't apply. For the second clause to apply we would need to have $\delta \notin z$, which by the induction hypothesis implies that $(z, \epsilon) \notin \llbracket\delta\rrbracket$, contradicting our choice of (z, ϵ). The second clause therefore doesn't apply either. It follows that $(z, \epsilon) \preceq (y, \delta)$ for every $(z, \epsilon) \in \llbracket\delta\rrbracket$, so $(y, \delta) \in \mathsf{opt}(\llbracket\delta\rrbracket)$.

We have $\gamma \notin y$, which by the induction hypothesis implies that $M_\Xi^{\mathbf{MOU}}, (y, \delta) \not\models \gamma$. This implies that $M_\Xi^{\mathbf{MOU}}, (x, \psi) \not\models B(\gamma \mid \delta)$, as was to be shown.

In each case we have shown that $M_\Xi^{\mathbf{MOU}}, (x, \psi) \models \varphi$ if and only if $\varphi \in x$, completing the induction step and thereby the proof. $\Box$

Theorem 37. *Let $\beta_1, \beta_2 \in \{\mathsf{max}, \mathsf{opt}, \mathsf{unm}\}$. For every model $M_1 = (S_1, \preceq_1, V_1)$ and $s_1 \in S_1$ there are a model $M_2 = (S_2, \preceq_2, V_2)$ and $s_2 \in S_2$ such that for every $\varphi \in \mathcal{L}$*

$$M_1, s_1 \models_{\beta_1} \varphi \Leftrightarrow M_2, s_2 \models_{\beta_2} \varphi.$$

Proof. It suffices to prove four cases: $\beta_1 = \mathsf{max}$ and $\beta_2 = \mathsf{opt}$, $\beta_1 = \mathsf{max}$ and $\beta_2 = \mathsf{unm}$, $\beta_1 = \mathsf{opt}$ and $\beta_2 = \mathsf{max}$, and $\beta_1 = \mathsf{unm}$ and $\beta_2 = \mathsf{max}$.

- Suppose $\beta_1 = \mathsf{max}$ and $\beta_2 = \mathsf{opt}$. Let $M_2 = (S_1, \preceq_2, V_1)$, where $s \preceq_2 t$ if and only if (i) $s \preceq_1 t$ or (ii) $s \perp_1 t$.

  We have $\prec_1 = \prec_2$, so by Lemma 11 we have $\mathsf{max}^{M_1}(X) = \mathsf{max}^{M_2}(X)$ for all $X \subseteq S$. Furthermore, $\preceq_2$ is total so by Proposition 8, $\mathsf{max}^{M_2}(X) = \mathsf{opt}^{M_2}(X)$. Together, these two equalities imply that $\mathsf{max}^{M_1}(X) = \mathsf{opt}^{M_2}(X)$, from which it follows easily that $M_1, s_1 \models_{\mathsf{max}} \varphi$ iff $M_2, s_1 \models_{\mathsf{opt}} \varphi$.

- Suppose $\beta_1 = \mathsf{max}$ and $\beta_2 = \mathsf{unm}$. Let $M_2 = (S_1, \preceq_2, V_1)$, where $s \preceq_2 t$ if and only if $s \prec_1 t$.

  We have $\prec_1 = \prec_2$, so by Lemma 11 we have $\mathsf{max}^{M_1}(X) = \mathsf{max}^{M_2}(X)$ for all $X \subseteq S$. Furthermore, $\preceq_2$ is anti-symmetric so by Proposition 8, $\mathsf{max}^{M_2}(X) = \mathsf{unm}^{M_2}(X)$. Together, these two equalities imply that $\mathsf{max}^{M_1}(X) = \mathsf{unm}^{M_2}(X)$, from which it follows easily that $M_1, s_1 \models_{\mathsf{max}} \varphi$ iff $M_2, s_1 \models_{\mathsf{unm}} \varphi$.

- Suppose $\beta_1 = \mathsf{opt}$ and $\beta_2 = \mathsf{max}$. Let $M_2 = (S_2, \preceq_2, V_2)$ be given by $S_2 = S_1 \times \{0, 1, 2\}$, $V_2(p) = V_1(p) \times \{0, 1, 2\}$ and for all $s \neq t$

  - $(s, i) \approx_2 (s, j)$ for all i and j,
  - if $s \prec_1 t$ then $(s, i) \prec_2 (t, j)$ for all i and j,
  - if $s \approx_1 t$ then $(s, i) \approx_2 (t, j)$ for all i and j and

- if $s \perp_1 t$ then $(s,0) \prec_2 (t,1)$, $(s,1) \prec_2 (t,2)$ and $(s,2) \prec_2 (t,0)$.

Take any $X \subseteq S_1$ and any $s \in \mathsf{opt}(X)$. Then for every $t \in X \setminus \{s\}$, $t \preceq_1 s$. For every (s,i) and every (t,j) we then have $(t,j) \preceq_2 (s,i)$ and therefore, in particular, $(s,i) \not\prec_2 (t,j)$. It follows that $(s,i) \in \mathsf{max}(X \times \{0,1,2\})$.

Now, take any $X \subseteq S_1$ and any $s \notin \mathsf{opt}(X)$. Then there is a $t \in X \setminus \{s\}$ such that $s \prec_1 t$ or $s \perp_1 t$. In the first case, $(s,i) \prec_2 (t,j)$ for all i and j. In the second case, $(s,i) \prec_2 (t, i+1 \mod 3)$ for all i. In either case, $(s,i) \notin \mathsf{max}(X \times \{0,1,2\})$.

Together, this shows that $\mathsf{max}^{M_2}(X \times \{0,1,2\}) = \mathsf{opt}^{M_1}(X) \times \{0,1,2\}$, from which it follows easily that $M_1, s_1 \models_{\mathsf{opt}} \varphi$ iff $M_2, s_2 \models_{\mathsf{max}} \varphi$.

- Suppose $\beta_1 = \mathsf{unm}$ and $\beta_2 = \mathsf{max}$. Let $M_2 = (S_2, \preceq_2, V_2)$ be given by $S_2 = S_1 \times \{0,1,2\}$, $V_2(p) = V_1(p) \times \{0,1,2\}$ and for all $s \neq t$

  - $(s,i) \approx_2 (s,j)$ for all i and j,
  - if $s \prec_1 t$ then $(s,i) \prec_2 (t,j)$ for all i and j,
  - if $s \perp_1 t$ then $(s,i) \approx_2 (t,j)$ for all i and j and
  - if $s \approx_1 t$ then $(s,0) \prec_2 (t,1)$, $(s,1) \prec_2 (t,2)$ and $(s,2) \prec_2 (t,0)$.

  Take any $X \subseteq S_1$ and any $s \in \mathsf{unm}(X)$. Then for every $t \in X \setminus \{s\}$ we have $s \not\preceq_1 t$, and therefore $t \prec_1 s$ or $t \perp_1 s$. In the first case, $(t,i) \prec_2 (s,j)$, in the second case $(t,i) \approx_2 (s,j)$ for all i and j. In either case, $(s,j) \not\prec_2 (t,i)$. It follows that $(s,j) \in \mathsf{max}(X \times \{0,1,2\})$.

  Now, take any $X \subseteq S_1$ and any $s \notin \mathsf{unm}(X)$. Then there is a $t \in X \setminus \{s\}$ such that $s \preceq_1 t$, and therefore $s \prec_1 t$ or $s \approx_1 t$. In the first case, $(s,i) \prec_2 (t,j)$ for all i and j, in the second case $(s,i) \prec_2 (t, i+1 \mod 3)$. In either case $(s,i) \notin \mathsf{max}(X \times \{0,1,2\})$.

  Together, this shows that $\mathsf{max}^{M_2}(X \times \{0,1,2\}) = \mathsf{unm}^{M_1}(X) \times \{0,1,2\}$, from which it follows easily that $M_1, s_1 \models_{\mathsf{unm}} \varphi$ iff $M_2, s_2 \models_{\mathsf{max}} \varphi$.

□

Proposition 42. *For every $\beta \in \{\mathsf{max}, \mathsf{opt}, \mathsf{unm}, \mathsf{acc}\}$, there is a $\varphi \in \mathcal{L}$ such that $\models_\beta \varphi$ while $\not\vdash_{\mathbf{BASE}} \varphi$.*

Proof. Consider alternative semantics for $\mathcal{L}$ based on models of the form $M = (S, f, V)$, where $f : 2^S \to 2^S$ satisfies $f(X) \subseteq X$ for every $X \subseteq S$. We then say that

$$M, s \models B(\varphi \mid \psi) \Leftrightarrow f([\![\psi]\!]) \subseteq [\![\varphi]\!]$$

We have not changed the semantics of propositional logic or the $\Box$ operator, so **PL**, $\Box$**S5** and **MP** are sound for these alternative semantics. Because the semantics are extensional, **R-Ext** and **L-Ext** are sound as well.

Soundness of **Abs** follows from the fact that whether $M, s \models B(\varphi \mid \psi)$ depends only on M, not on s. The constraint that $f(X) \subseteq X$ implies that

Id and **Triv** are sound. Finally, if $B(\varphi \to \psi \mid \chi)$ and $B(\varphi \mid \chi)$ hold then $f([\![\psi]\!]) \subseteq [\![\varphi \to \psi]\!] \cap [\![\varphi]\!] \subseteq [\![\psi]\!]$, so $B(\psi \mid \chi)$ holds as well. So **K** is also sound. We have now shown that all the rules and axioms of **BASE** are sound for these semantics.

But now consider a model $M = (S, f, V)$ given by

- $S = \{s_1, s_2\}$,
- $f(\emptyset) = f(\{s_1\}) = f(\{s_2\}) = \emptyset$,
- $f(\{s_1, s_2\}) = \{s_1, s_2\}$ and
- $V(q) = \{s_1\}$.

Then $f([\![q \wedge \top]\!]) = \emptyset \subseteq [\![\bot]\!]$, and therefore $M \models B(\bot \mid q \wedge \top)$. Yet $f([\![\top]\!]) = \{s_1, s_2\} \not\subseteq \{s_2\} = [\![q \to \bot]\!]$, and therefore $M \not\models B(q \to \bot \mid \top)$. So $M \not\models B(\bot \mid q \wedge \top) \to B(q \to \bot \mid \top)$.

Because **BASE** is sound for this logic, $M \not\models B(\bot \mid q \wedge \top) \to B(q \to \bot \mid \top)$ implies that $\not\vdash_{\mathbf{BASE}} B(\bot \mid q \wedge \top) \to B(q \to \bot \mid \top)$. But this formula is an instance of **Sh**, so $\models_\beta B(\bot \mid q \wedge \top) \to B(q \to \bot \mid \top)$ for each $\beta \in \{\mathsf{max}, \mathsf{opt}, \mathsf{unm}\}$.

Similarly, consider the model $M = (S, f, V)$ given by

- $S = \{s_1, s_2\}$,
- $f(\emptyset) = f(\{s_1\}) = f(\{s_2\}) = \emptyset$,
- $f(\{s_1, s_2\}) = \{s_1\}$ and
- $V(q) = \{s_1\}$.

Then $f([\![\top]\!]) = \{s_1\} \subseteq [\![q]\!]$, so $M \models B(q \mid \top)$. Furthermore, $f([\![q \wedge \top]\!]) = \emptyset \subseteq [\![\bot]\!]$, so $M \models B(\bot \mid q \wedge \top)$. Yet $f([\![\top]\!]) = \{s_1\} \not\subseteq \emptyset = [\![\bot]\!]$, so $M \not\models B(\bot \mid \top)$. So we have $M \not\models (B(q \mid \top) \wedge B(\bot \mid q \wedge \top)) \to B(\bot \mid \top)$.

Because of the soundness of **BASE**, this implies that $\not\vdash_{\mathbf{BASE}} (B(q \mid \top) \wedge B(\bot \mid q \wedge \top)) \to B(\bot \mid \top)$. But this formula is an instance of **Cut**, so $\models_{\mathsf{acc}} (B(q \mid \top) \wedge B(\bot \mid q \wedge \top)) \to B(\bot \mid \top)$. ☐

Lemma 59 (Truth lemma for $M_\Xi^{\mathsf{max}(tr)}$ with respect to max)**.** *$M_\Xi^{\mathsf{max}(tr)}, (x, \varphi, i) \models_{\mathsf{max}} \psi$ if and only if $\psi \in x$.*

Proof. By induction and a case distinction on the main connective of ψ. As usual, only the case $\psi = B(\gamma \mid \delta)$ is interesting.

- Suppose $B(\gamma \mid \delta) \in x$. Take any $(y, \chi, j) \in \mathsf{max}([\![\delta]\!])$. By the induction hypothesis this implies $(y, \chi, j) \in \mathsf{max}([\delta])$. Then two things hold:
    1. $(y, \chi, j) \in \Delta_\chi(\Xi)$, since otherwise $(y, \chi, j) \prec (y, \chi, j+1)$ and $(y, \chi, j+1) \in [\delta]$,
    2. $[\delta] \subseteq [\chi]$, since otherwise there would be a z such that $\delta \in z$ and $\chi \notin z$, and we would have $(y, \chi, j) \prec (z, \bot, k)$ and $(z, \bot, k) \in [\delta]$.

Because $[\delta] \subseteq [\chi]$ we have $\delta \leftrightarrow (\delta \wedge \chi) \in \Xi$ and therefore $\Box(\delta \leftrightarrow (\delta \wedge \chi))$. By assumption we also have $B(\gamma \mid \delta)$, so by **R-Ext** we obtain $B(\gamma \mid \delta \wedge \chi)$. By **Sh** it follows that $B(\delta \rightarrow \gamma \mid \chi)$.

Because $(y, \chi, j) \in \Delta_\chi(\Xi)$ we have $B_\chi^{-1} \subseteq y$. In particular, this implies that $\delta \rightarrow \gamma \in y$. Furthermore, by assumption $(y, \chi, j) \in \mathsf{max}([\delta])$ and thereforein particular $(y, \chi, j) \in [\delta]$. So we have $\delta \in y$. Together with $\delta \rightarrow \gamma \in y$ this yields $\gamma \in y$ and therefore, by the induction hypothesis, $(y, \chi, j) \in [\![\gamma]\!]$. This holds for every $(y, \chi, j) \in \mathsf{max}([\![\delta]\!])$, so $M_\Xi^{\mathsf{max}(tr)}, (x, \varphi, i) \models B(\gamma \mid \delta)$.

- Suppose $B(\gamma \mid \delta) \notin x$. Then by Lemma 30 there is a $y \in MCS_\Xi^{\mathbf{MOU}}$ such that $B_\delta^{-1} \subseteq y$ and $\gamma \notin y$. Then $(y, \delta, i) \in \Delta_\delta(\Xi)$, so any outgoing arrow from (y, δ, i) must be of type 2. So any such arrow must go to a state $(z, \bot, k)$ with $\delta \notin z$, implying that $(z, \chi, k) \notin [\delta]$. So $(y, \delta, i) \in \mathsf{max}([\delta])$.

  By the induction hypothesis, $\gamma \notin y$ implies $(y, \delta, i) \notin [\![\gamma]\!]$ and $(y, \delta, i) \in \mathsf{max}([\delta])$ implies $(y, \delta, i) \in \mathsf{max}([\![\delta]\!])$. So $M_\Xi^{\mathsf{max}(tr)}, (x, \varphi, i) \not\models B(\gamma \mid \delta)$.

□

Lemma 67. *For every state x of $M_\Xi^{\mathsf{opt}(tr)}$ and every formula φ, if $x \in \mathsf{opt}([\varphi])$, then $B_\varphi^{-1} \subseteq x$.*

Proof. Let $x \in \mathsf{opt}([\varphi])$, and suppose towards a contradiction that $B_\varphi^{-1} \not\subseteq x$. Note that, in particular, $x \in \mathsf{opt}([\varphi])$ implies that $x \in [\varphi]$. We distinguish between two cases.

1. Suppose that B_φ^{-1} is consistent with $\Box\Xi$, so there is a $y \in MCS_\Xi^{\mathbf{MOU+DM}}$ such that $B_\varphi^{-1} \subseteq y$. Because $x \in \mathsf{opt}([\varphi])$, we have $y \preceq x$, and therefore $y \overset{\chi}{\rightsquigarrow} x$ for some χ. Now, take any $\psi \in B_\varphi^{-1}$.

   Then we have $B(\psi \mid \varphi)$ and therefore, by **R-Ext**, $B(\psi \mid (\varphi \vee \chi) \wedge \varphi)$. By **Sh**, this implies that $B(\varphi \rightarrow \psi \mid \varphi \vee \chi)$. Note furthermore that, because $\chi \in y$, we have $\neg\chi \notin y$. The state y was chosen such that $B_\varphi^{-1} \subseteq y$, so it follows that we do not have $B(\neg\chi \mid \varphi)$, implying that $P(\chi \mid \varphi)$. Together with $B(\varphi \rightarrow \psi \mid \varphi \vee \chi)$, this implies, by **DM**, that $B(\varphi \rightarrow \psi \mid \chi)$.

   Because $y \overset{\chi}{\rightsquigarrow} x$ we have $B_\chi^{-1} \subseteq x$, so $\varphi \rightarrow \psi \in x$. From $x \in [\varphi]$, it then follows that $\psi \in x$. This holds for every $\psi \in B_\varphi^{-1}$, so we have $B_\varphi^{-1} \subseteq x$, contradicting the assumption that $B_\varphi^{-1} \not\subseteq x$.

2. Suppose that B_φ^{-1} is inconsistent with $\Box\Xi$. Then let $\Theta = \{\theta \mid B_\theta^{-1} \subseteq x\}$. By Lemma 56, the set $\{\varphi\} \cup \neg\Theta$ is consistent with $\Box\Xi$, so there is some y such that $\varphi \in y$ and $y \not\rightsquigarrow x$.

   Then $y \not\preceq x$, and $y \in [\varphi]$, so $x \notin \mathsf{opt}([\varphi])$.[10] This contradicts the assumption that $x \in \mathsf{opt}([\varphi])$.

[10] If $x = y$ we use that there are two copies $(x, 0)$ and $(x, 1)$ of x, and if $x \not\rightsquigarrow x$ then $(x, 0) \perp (x, 1)$.

In either case, we encounter a contradiction. So we conclude that $B_\varphi^{-1} \subseteq x$, as was to be shown. $\square$

Lemma 68 (Truth lemma for $M_\Xi^{\mathsf{opt}(tr)}$)**.** *$M_\Xi^{\mathsf{opt}(tr)}, x \models \varphi$ if and only if $\varphi \in x$.*

Proof. By induction and a case distinction on the main connective of φ. As usual, only the case $\varphi = B(\gamma \mid \delta)$ is interesting.

- Suppose $B(\gamma \mid \delta) \in x$, and let $y \in \mathsf{opt}([\![\delta]\!])$. By the induction hypothesis we then have $y \in \mathsf{opt}([\delta])$, so by Lemma 67 we have $B_\delta^{-1} \subseteq y$ and therefore, in particular, $\gamma \in y$. We therefore have $M_\Xi^{\mathsf{opt}(tr)}, x \models B(\gamma \mid \delta)$.
- Suppose $B(\gamma \mid \delta) \notin x$. Then by Lemma 30, the set $B_\delta^{-1} \cup \{\neg\gamma\}$ is consistent with Ξ, so it can be extended to a maximal consistent set $y \in MCS_\Xi^{\mathsf{opt}(tr)}$. Furthermore, for every $z \in [\delta]$ we have $z \overset{\delta}{\rightsquigarrow} y$, and therefore $z \preceq y$. As $[\varphi] = [\![\varphi]\!]$ by the induction hypothesis this implies that $y \in \mathsf{opt}([\![\varphi]\!])$.

  We have $\gamma \notin y$ and therefore, by the induction hypothesis, $M_\Xi^{\mathsf{opt}(tr)}, y \not\models \gamma$. It follows that $M_\Xi^{\mathsf{opt}(tr)}, x \not\models B(\gamma \mid \delta)$.

$\square$

Lemma 80 (Truth lemma for $M_\Xi^{\mathsf{unm}(to)}$)**.** *$M_\Xi^{\mathsf{unm}(to)}, x \models \varphi$ if and only if $\varphi \in x$.*

Proof. By induction and a case distinction on the main connective of φ. As usual, only the case $\varphi = B(\gamma \mid \delta)$ is interesting.

- Suppose $B(\gamma \mid \delta) \notin x$. Then B_δ^{-1} is consistent and therefore, by Lemma 52, it is maximal consistent. Furthermore, it contains $\Box\Xi$, so $B_\delta^{-1} = y \in MCS_\Xi^{\mathbf{MOU+Un+WDM}}$. For every $z \in [\delta]$ we have $z \overset{\delta}{\rightsquigarrow} y$, so $y \not\preceq z$. Furthermore, $\gamma \notin B_\delta^{-1} = y$, so $y \notin [\gamma]$. By the induction hypothesis $[\gamma] = [\![\gamma]\!]$ and $\mathsf{unm}([\delta]) = \mathsf{unm}([\![\delta]\!])$, so $y \in \mathsf{unm}([\![\delta]\!])$ while $y \notin [\![\gamma]\!]$ and therefore $M_\Xi^{\mathsf{unm}(to)}, x \not\models B(\gamma \mid \delta)$.
- Suppose $B(\gamma \mid \delta) \in x$. We distinguish two cases.
  - Suppose B_δ^{-1} is consistent. Then it is maximal consistent, so there is a $y \in MCS_\Xi^{\mathbf{MOU+Un+WDM}}$ such that $y = B_\delta^{-1}$. For every $z \in [\delta]$ we have $y \not\preceq z$, as $z \overset{\delta}{\rightsquigarrow} y$. We therefore have $y \in \mathsf{unm}([\delta])$. By the induction hypothesis, this implies that $y \in \mathsf{unm}([\![\delta]\!])$. As $M_\Xi^{\mathsf{unm}(to)}$ is total, $\mathsf{unm}([\![\delta]\!])$ contains at most one element, so $\mathsf{unm}([\![\delta]\!]) = \{y\}$.

    So $y = B_\delta^{-1}(x)$. As $B(\gamma \mid \delta) \in x$, this implies that $\gamma \in y$ and therefore, by the induction hypothesis, $y \in [\![\gamma]\!]$. It follows that $\mathsf{unm}([\![\delta]\!]) \subseteq [\![\gamma]\!]$, so $M_\Xi^{\mathsf{unm}(to)}, x \models B(\gamma \mid \delta)$.
  - Suppose B_δ^{-1} is inconsistent. We will show that $\mathsf{unm}([\delta]) = \emptyset$. Suppose towards a contradiction that there is a $y \in \mathsf{unm}([\delta])$. Note that this implies that y is good, since a bad state $(u, 0)$ is matched by

$(u, 1)$ and vice versa. Now, let $\Theta = \{\theta \in \mathcal{L} \mid B_\theta^{-1} \subseteq y\}$. Because y is good, we know that $\Theta \neq \emptyset$. By Lemma 56, $\{\delta\} \cup \neg\Theta$ is consistent with $\Box\Xi$, and can therefore be extended to some z.

By construction, there is no $\zeta \in z$ such that $B_\zeta^{-1} \subseteq y$, so $z \not\rightarrow y$, and therefore $y \preceq z$, contradicting the unmatchedness of y in $[\delta]$. As we have arrived at a contradiction, we conclude that $\mathsf{unm}([\delta]) = \emptyset$.

By the induction hypothesis, $\mathsf{unm}(\llbracket\delta\rrbracket) = \mathsf{unm}([\delta])$, so $\mathsf{unm}(\llbracket\delta\rrbracket) = \emptyset$, from which it follows that $M_\Xi^{\mathsf{unm}(to)}, x \models B(\gamma \mid \delta)$.

We have shown that in both possible cases $M_\Xi^{\mathsf{unm}(to)}, x \models B(\gamma \mid \delta)$, which was to be shown.

□

Lemma 91. *The model $M_\Xi^{\mathsf{opt}(tr,to)}$ is transitive.*

Proof. Suppose that $(x, i) \preceq (y, j) \preceq (z, k)$. We distinguish between a number of cases based upon the subsets of S that the states belong to.

First, suppose that $(x, i) \in S_u$. Any arrow from a state in S_u goes to another state in S_u, so we must also have $(y, j) \in S_u$ and $(z, k) \in S_u$. From $(x, i) \preceq (y, j)$ and $(y, j) \preceq (z, k)$ it therefore follows that $i \leq j$ and $j \leq k$, and therefore that $i \leq k$. We therefore have $(x, i) \preceq (z, k)$.

Secondly, suppose that $(x, i) \notin S_u$ while $(y, j) \in S_u$ or $(z, k) \in S_u$. Because S_u has no outgoing arrows it follows that $(z, k) \in S_u$. Together with $(x, i) \notin S_u$ this implies that $(x, i) \preceq (z, k)$.

Note that these first two cases cover all possibilities where $(x, i) \in S_u$, $(y, j) \in S_u$ or $(z, k) \in S_u$. In the remaining cases we therefore have $(x, i), (y, j), (z, k) \in S_g \cup S_b$. By the construction of the model this implies that, in particular, $supp(z) \subseteq supp(y)$ and $supp(y) \subseteq supp(x)$.

Suppose then, as third case, that $(z, k) \in S_g$. Then $supp(z) \subseteq supp(x)$ suffices to conclude that $(x, i) \preceq (z, k)$.

As fourth case, suppose that $(z, k) \in S_b$ and $(x, i) \notin S_b$. Then there are two possibilities. Either $(y, j) \in S_g$, in which case $(y, j) \preceq (z, k)$ implies that $supp(z) \subset supp(y)$, or $(y, j) \in S_b$, in which case $(x, i) \preceq (z, j)$ implies that $supp(y) \subset supp(x)$. In either case, $supp(z) \subset supp(x)$, so $(x, i) \preceq (z, k)$.

As fifth and final case, suppose that $(z, k) \in S_b$ and $(x, i) \in S_b$. Then $supp(z) \subseteq supp(x)$ suffices to conclude that $(x, i) \preceq (z, k)$.

These five cases are exhaustive, and in each case $(x, i) \preceq (z, k)$. So the relation is transitive. □

Lemma 92. *The model $M_\Xi^{\mathsf{opt}(tr,to)}$ is total.*

Proof. Take any $(x, i), (y, j) \in S$. By Lemma 89, either $supp(x) \subseteq supp(y)$, or $supp(y) \subseteq supp(x)$. Note that this also implies that $supp(x) \subseteq supp(y)$ or $supp(y) \subset supp(x)$.

- If $(x, i) \in S_u$ and $(y, j) \in S_u$, then it follows from $i \leq j$ or $j \leq i$ that $(x, i) \preceq (y, j)$ or $(y, j) \preceq (x, i)$.

- If $(x, i) \in S_u$ and $(y, j) \notin S_u$, then $(x, i) \preceq (y, j)$.
- If $(x, i) \in S_g$ and $(y, j) \in S_g$, or if $(x, i) \in S_b$ and $(y, j) \in S_b$, then it follows from $supp(x) \subseteq supp(y)$ or $supp(y) \subseteq supp(x)$ that $(x, i) \preceq (y, j)$ or $(y, j) \preceq (x, i)$.
- If $(x, i) \in S_g$ and $(y, j) \in S_b$, then from $supp(x) \subseteq supp(y)$ or $supp(y) \subset supp(x)$ it follows that $(x, i) \preceq (y, j)$ or $(y, j) \preceq (x, i)$.

Without loss of generality, we can assume that one of these four cases applies. In each case, $(x, i) \preceq (y, j)$ or $(y, j) \preceq (x, i)$. So the model is total. □

Lemma 93. *If $B_\varphi^{-1} \subseteq x$, then $(x, i) \in \mathsf{opt}([\varphi])$.*

Proof. First, note that x is good, since $B_\varphi^{-1} \subseteq x$. Now, take any $(y, j) \in [\varphi]$. Note that $x \in supp(y)$. This implies that $supp(y) \neq \emptyset$, and therefore $y \in S_g \cup S_b$. Furthermore, by Lemma 90 it implies that $supp(x) \subseteq supp(y)$. Together, this suffices to conclude that $(y, j) \preceq (s, i)$. As this holds for every $(y, j) \in [\varphi]$, we have $(x, i) \in \mathsf{opt}([\varphi])$. □

Lemma 94. *If $B_\varphi^{-1} \not\subseteq x$, then $(x, i) \notin \mathsf{opt}([\varphi])$.*

Proof. If $(x, i) \notin [\varphi]$ then trivially $(x, i) \notin \mathsf{opt}([\varphi])$. So assume that $(x, i) \in [\varphi]$. We distinguish five cases.

- Suppose $supp(x) = \emptyset$. Then $(x, i) \in S_u$, and therefore $(x, i) \prec (x, i + 1)$. So $(x, i) \notin \mathsf{opt}([\varphi])$.
- Suppose $supp(x) \neq \emptyset$, B_φ^{-1} is inconsistent and $x \in S_g$. Let $\Theta = \{\theta \mid B_\theta^{-1} \subseteq x\}$. By Lemma 56, $\{\varphi\} \cup \neg\Theta$ is consistent with $\Box\Xi$, so it can be extended to some $y \in MCS_\Xi$. Because $x \in S_g$ it is good, and therefore supports itself. By construction y does not support x, so $supp(x) \not\subseteq supp(y)$. This implies that $(y, j) \not\preceq (x, i)$, so $(x, i) \notin \mathsf{opt}([\varphi])$.
- Suppose $supp(x) \neq \emptyset$, B_φ^{-1} is inconsistent and $x \in S_b$. Because $supp(x) \neq \emptyset$, there is at least one $y \in supp(x)$. We have $x \rightsquigarrow y$, so y is good, and therefore $y \in S_g$. Furthermore, by transitivity of support, $supp(y) \subseteq supp(x)$. In particular, this implies that $supp(x) \not\subset supp(y)$, and therefore $(y, j) \not\preceq (x, i)$. So $(x, i) \notin \mathsf{opt}([\varphi])$.
- Suppose $supp(x) \neq \emptyset$, B_φ^{-1} is consistent and $x \in S_b$. Then B_φ^{-1} can be extended to some y. We have $x \overset{\varphi}{\rightsquigarrow} y$, so by Lemma 90 we have $supp(y) \subseteq supp(x)$. In particular, this means that $supp(x) \not\subset supp(y)$. Because $x \in S_b$ and $y \in S_g$, this implies that $y \not\preceq x$, so $(x, i) \notin \mathsf{opt}([\varphi])$.
- Finally, suppose $supp(x) \neq \emptyset$, B_φ^{-1} is consistent and $x \in S_g$. Then B_φ^{-1} can be extended to some y. We have $x \overset{\varphi}{\rightsquigarrow} y$, so by Lemma 90 we have $supp(y) \subseteq supp(x)$. Suppose towards a contradiction that $supp(x) \subseteq supp(y)$. Because x is good, it supports itself, so $supp(x) \subseteq supp(y)$

implies that $y \overset{\psi}{\rightsquigarrow} x$ for some ψ. We have $\psi \in y$ and $B^{-1}_\varphi \subseteq y$, so $P(\psi \mid \varphi)$. By **DM**, it the follows that for every χ, if $B(\chi \mid \varphi \vee \psi)$ then $B(\chi \mid \psi)$. So $B^{-1}_{\varphi \vee \psi} \subseteq B^{-1}_\psi \subseteq x$.

Furthermore, we have $\Box(\varphi \leftrightarrow (\varphi \wedge (\varphi \vee \psi))$. Using **R-Ext**, this yields $B(\chi \mid \varphi) \leftrightarrow B(\chi \mid \varphi \wedge (\varphi \vee \psi))$ for every χ. Using **Sh**, $B(\chi \mid \varphi \wedge (\varphi \vee \psi))$ implies $B(\varphi \to \chi \mid \varphi \vee \psi)$. So for every χ, if $B(\chi \mid \varphi)$ then $B(\varphi \to \chi \mid \varphi \vee \psi)$. As $B^{-1}_{\varphi \vee \psi} \subseteq x$, this implies that $\varphi \to \chi \in x$ for each such χ. Because $\varphi \in x$, that in turn implies $\chi \in x$. So $B^{-1}_\varphi \subseteq x$, contradicting the assumption from the lemma that $B^{-1}_\varphi \not\subseteq x$.

From this contradiction, we conclude that $supp(x) \not\subseteq supp(y)$. This implies that $(y, j) \not\preceq (x, i)$, so $x \notin \mathsf{opt}([\varphi])$.

In each of the five possible cases we concluded that $x \notin \mathsf{opt}([\varphi])$, so we have $x \notin \mathsf{opt}([\varphi])$, as was to be shown. □

Lemma 102. *$M^{\mathsf{unm}(tr,to)}_\Xi$ is total.*

Proof. Let $x, y \in S$. If x and y are both good, we would have $x \perp y$ only if x and y support each other. So there are $\varphi \in x, \psi \in y$ such that $B^{-1}_\varphi \subseteq y$ and $B^{-1}_\psi \subseteq x$. As in Lemma 78, we then have that $x = y$.

If x is bad and y is good, or vice versa, it is immediate from the definition that $x \prec y$ or $y \prec x$.

If x and y are both bad, suppose towards a contradiction that $x \perp y$. Then $supp(x) \not\subseteq supp(y)$ and $supp(y) \not\subseteq supp(x)$, so there are $u \in supp(x) \setminus supp(y)$ and $v \in supp(y) \setminus supp(x)$. Then there are φ and ψ such that $x \overset{\varphi}{\rightsquigarrow} u$ and $y \overset{\psi}{\rightsquigarrow} v$.

As the proof system contains **DM**, support is transitive. From the fact that $x \rightsquigarrow u$ but $x \not\rightsquigarrow v$, it therefore follows that $u \not\rightsquigarrow v$. Likewise, because $y \rightsquigarrow v$ but $y \not\rightsquigarrow u$, we have $v \not\rightsquigarrow u$.

Now, suppose towards a contradiction that $P(\varphi \mid \varphi \vee \psi)$. Then $B^{-1}_{\varphi \vee \psi}$ is consistent, and can therefore be extended to some z, with the property that $B^{-1}_{\varphi \vee \psi} \subseteq z$ and $\varphi \in z$. We have $v \overset{\varphi \vee \psi}{\rightsquigarrow} z$ and $z \overset{\varphi}{\rightsquigarrow} u$, so by transitivity of support $v \rightsquigarrow u$, contradicting our earlier conclusion. From this contradiction, we conclude that $B(\neg\varphi \mid \varphi \vee \psi)$.

The state x is bad while u is good, so $x \neq u$, implying that there is a $\chi \in u \setminus x$. Because $B^{-1}_\varphi \subseteq u$ and $\chi \in u$, we have $P(\chi \mid \varphi)$. Finally, take any $\xi \in v$, so $P(\xi \mid \psi)$. We now have

$$P(\chi \mid \varphi) \wedge P(\xi \mid \psi) \wedge B(\neg\varphi \mid \varphi \vee \psi)$$

By **SP** this implies that $P(\xi \mid (\varphi \wedge \neg\chi) \vee \psi)$. This holds for every $\xi \in v$, so $B^{-1}_{(\varphi \wedge \neg\chi) \vee \psi} \subseteq v$. But φ and χ were chosen in such a way that $\varphi \wedge \neg\chi \in x$, so x supports v with witness $(\varphi \wedge \neg\chi) \vee \psi$. We have arrived at a contradiction from the assumption that $x \perp y$, which completes the proof that $M^{\mathsf{unm}(tr,to)}_\Xi$ is total. □

Lemma 103. *$M^{\mathsf{unm}(tr,to)}_\Xi$ is transitive.*

Proof. Suppose that $x \preceq y \preceq z$. If $x = y$ or $y = z$, it trivially follows that $x \preceq z$. If $x = z$, then $x \preceq z$ because the model is reflexive. Assume then that x, y and z are all distinct. We show that $x \preceq z$ by case distinction on which of x, y and z are good.

- If x, y and z are bad, then $x \preceq y \preceq z$ implies $\mathit{supp}(y) \subseteq \mathit{supp}(x)$ and $\mathit{supp}(z) \subseteq \mathit{supp}(y)$. So $\mathit{supp}(z) \subseteq \mathit{supp}(x)$, and therefore $x \preceq z$.
- If x is good and y, z are bad, then $x \preceq y \preceq z$ implies that y doesn't support x, and $\mathit{supp}(z) \subseteq \mathit{supp}(y)$. So z doesn't support x, and therefore $x \preceq z$.
- If y is good and x, z are bad, then $x \preceq y \preceq z$ implies $x \rightsquigarrow y$ and $z \not\rightsquigarrow y$. So $\mathit{supp}(x) \not\subseteq \mathit{supp}(z)$, and therefore $z \not\preceq x$. By totality, it follows that $x \preceq z$.
- If z is good and x, y are bad, then $x \preceq y \preceq z$ implies that $\mathit{supp}(y) \subseteq \mathit{supp}(x)$ and $y \rightsquigarrow z$. So $x \rightsquigarrow z$, and therefore $x \preceq z$.
- If x, y are good and z is bad, then $x \preceq y \preceq z$ implies that $y \not\rightsquigarrow x$, and $z \not\rightsquigarrow y$.

  Suppose then, towards a contradiction, that $z \rightsquigarrow x$. Then we have $z \rightsquigarrow x$, $y \not\rightsquigarrow x$ and y is good. Hence, by Lemma 54, $z \rightsquigarrow y$, contradicting $z \not\rightsquigarrow y$.

  From this contradiction we obtain $z \not\rightsquigarrow x$, and therefore $x \preceq z$.
- If x, z are good and y is bad, then $x \preceq y \preceq z$ implies that $y \rightsquigarrow z$ but $y \not\rightsquigarrow x$.

  If we were to have $x \not\preceq z$ that would imply that $z \rightsquigarrow x$, which by transitivity of support would imply that $y \rightsquigarrow x$, contradicting $x \preceq y$. From the contradiction we obtain $x \preceq z$.
- If y, z are good and x is bad, then $x \preceq y \preceq z$ implies that $x \rightsquigarrow y$ and $z \not\rightsquigarrow y$. By Lemma 54 we then have $x \rightsquigarrow z$, which implies that $x \preceq z$.
- Finally, if x, y and z are good, then $x \preceq y \preceq z$ implies that $y \not\rightsquigarrow x$ and $z \not\rightsquigarrow y$.

  If we were to have $x \not\preceq z$ that would imply $z \rightsquigarrow x$. Using Lemma 54, $z \rightsquigarrow x$, $y \not\rightsquigarrow x$ and x being good imply that $z \rightsquigarrow y$, contradicting $z \not\rightsquigarrow y$. From the contradiction we obtain $x \preceq z$.

□

Lemma 116 (Truth lemma for $M_{\Xi}^{\mathsf{opt}(tr,as)}$)**.** *$M_{\Xi}^{\mathsf{opt}(tr,as)}, (x, i) \models \varphi$ if and only if $\varphi \in x$.*

Proof. By induction and a case distinction on the main connective of φ. As usual, only the case $\varphi = B(\gamma \mid \delta)$ is interesting.

- Suppose $B(\gamma \mid \delta) \notin x$. Then $y = B_{\delta}^{-1}$ is consistent. In the presence of **Un** this implies that it is maximal consistent. Furthermore, $\Box\Xi \subseteq B_{\delta}^{-1}$. So $(y, 0) \in S$. Now, note that for every $(z, j) \in [\delta] \setminus \{(y, 0)\}$ we have $z \overset{\delta}{\rightsquigarrow} y$ and either $z \neq B_\delta^{-1}$ or $z = B_\delta^{-1}$ and $0 < j$. Therefore, by the construction of the model, $(z, j) \preceq (y, 0)$. So we have $(y, 0) \in \mathsf{opt}([\delta])$. By the induction hypothesis, this implies that $(y, 0) \in \mathsf{opt}([\![\delta]\!])$.

Furthermore, $\gamma \notin y$, so by the induction hypothesis $(y, 0) \notin [\![\gamma]\!]$. We therefore have $M_\Xi^{\mathsf{opt}(tr,as)}, (x, i) \not\models B(\gamma \mid \delta)$.

- Suppose $B(\gamma \mid \delta) \in x$. We distinguish two sub-cases.
  - Suppose B_δ^{-1} is consistent. Then $(B_\delta^{-1}, 0) \in S$ and, furthermore, for every $(y, j) \in [\delta] \setminus \{(B_\delta^{-1}, 0)\}$ we have $(y, j) \preceq (B_\delta^{-1}, 0)$. So $(B_\delta^{-1}, 0) \in \mathsf{opt}([\delta])$. An optimal set in an anti-symmetric model contains at most one state, so $\mathsf{opt}([\delta]) = \{(B_\delta^{-1}, 0)\}$. We have $(B_\delta^{-1}, 0) \in [\gamma]$, so $\mathsf{opt}([\delta]) \subseteq [\gamma]$. By the induction hypothesis this implies that $\mathsf{opt}([\![\delta]\!]) \subseteq [\![\gamma]\!]$, and therefore $M_\Xi^{\mathsf{opt}(tr,as)}, (x, i) \models B(\gamma \mid \delta)$.
  - Suppose B_δ^{-1} is inconsistent. We will show that $\mathsf{opt}([\delta]) = \emptyset$. Suppose towards a contradiction that there is some $(y, j) \in \mathsf{opt}([\delta])$. Note that this implies that y is good, since otherwise we would have $(y, j) \prec (y, j+1)$. Now, let $\Theta = \{\theta \in \mathcal{L} \mid B_\theta^{-1} \subseteq y\}$. Because y is good, we know that $\Theta \neq \emptyset$.

    By Lemma 56, the set $\Box\Xi \cup \{\delta\} \cup \neg\Theta$ is consistent and can therefore be extended to a $z \in MCS_\Xi^{\mathbf{MOU+DM+Un}}$. So we have $(z, 0) \in S$. By construction, $y \notin supp(z)$. Furthermore, y is good, so in particular it supports itself. So we must have $y \neq z$. By the definition of the canonical model, it follows that $(z, 0) \not\preceq (y, j)$. But $(z, 0) \in [\delta]$, so this contradicts the optimality of (y, j). We have arrived at a contradiction, and conclude that $\mathsf{opt}([\delta]) = \emptyset$.

    By the induction hypothesis, this implies that $\mathsf{opt}([\![\delta]\!]) = \emptyset$, so we have $M_\Xi^{\mathsf{opt}(tr,as)}, (x, i) \models B(\gamma \mid \delta)$.

□