



UvA-DARE (Digital Academic Repository)

Logics of Preference when there is no Best

Grossi, D.; van der Hoek, W.; Kuijer, L.B.

DOI

[10.24963/kr.2020/46](https://doi.org/10.24963/kr.2020/46)

Publication date

2020

Document Version

Final published version

Published in

KR2020 : Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning

[Link to publication](#)

Citation for published version (APA):

Grossi, D., van der Hoek, W., & Kuijer, L. B. (2020). Logics of Preference when there is no Best. In D. Calvanese, E. Erdem, & M. Thielscher (Eds.), *KR2020 : Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning: Rhodes, Greece, September 12-18, 2020* (pp. 455-464). IJCAI. <https://doi.org/10.24963/kr.2020/46>

General rights

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

Disclaimer/Complaints regulations

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

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Logics of Preference when There Is No Best

Davide Grossi^{1,2}, Wiebe van der Hoek³, Louwe B. Kuijer³

¹University of Groningen

²University of Amsterdam

³University of Liverpool

d.grossi@rug.nl, {wiebe,lbkuijer}@liverpool.ac.uk

Abstract

Well-behaved preferences (e.g., total pre-orders) are a cornerstone of several areas in artificial intelligence, from knowledge representation, where preferences typically encode likelihood comparisons, to both game and decision theories, where preferences typically encode utility comparisons. Yet weaker (e.g., cyclical) structures of comparison have proven important in a number of areas, from argumentation theory to tournaments and social choice theory. In this paper we provide logical foundations for reasoning about this type of preference structures where no obvious best elements may exist. Concretely, we compare and axiomatize a number of ways in which the concepts of maximality and optimality can be lifted to this general class of preferences. In doing so we expand the scope of the long-standing tradition of the logical analysis of preference.

1 Introduction

Among the most fundamental mathematical notions in the formal representation of knowledge, inference and decision-making are arguably the notions of maximality and optimality. A maximal element is one to which nothing is strictly preferred; an optimal element is one that is weakly preferred to everything else. Given how a set of alternatives (options, states, strategies) compare with one another, one selects the maximal or optimal ones as the ones that are ‘best’.

Since at least the 60s (Von Wright 1963), formal logic has been used as a foundational tool to study reasoning about structures of pairwise comparisons, or preferences. Such a 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 (Kraus, Lehmann, and Magidor 1990; Boutilier 1990; Boutilier 1994); within epistemology, for the representation and analysis of processes of information acquisition by rational agents, in the AGM tradition (Alchourrón, Gärdenfors, and Makinson 1985), the conditionals tradition (Stalnaker 1968; Lewis 1973; Kratzer 1981; Burgess 1981; Halpern 2010), and more recently the dynamic-epistemic logic tradition (van Benthem 2004; Baltag and Smets 2008; Demey 2011; van Benthem 2011); within deontic logic, for the representation and analysis of conditional obligations (Hansson 1969;

Lewis 1974; Spohn 1975; Parent 2008; Parent 2013). 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 90s (van Benthem 1989; Nejdil 1991; Katsuno and Satoh 1991; Katsuno and Mendelzon 1991; Makinson 1993).

Aim & Focus All the above approaches to the characterization of a *logic of preference* make a fundamental assumption: the way in which alternatives (worlds, situations, options) compare with one another forms a structure where the set of maximal, or of optimal, elements is never empty. Often, this assumption builds on top of another classical one: preferences are transitive. Here we lift such assumptions and study logical foundations for possibly non-transitive pairwise comparison structures lacking maximals or optimals. Such structures underpin several areas within AI, eminent examples being abstract argumentation theory (Bench-Capon and Dunne 2007; Rahwan and Simari 2009) and the theory of tournaments (Laslier 1997; Brandt, Brill, and Harrenstein 2015). The aim of the paper is to extend the foundational tradition in the logic of preference mentioned above to bear on these areas.

We focus on the most general possible structures arising from pairwise comparisons of alternatives: unconstrained binary relations. In particular, such structures are not assumed to be transitive, anti-symmetric or total and may lack maximal and optimal elements. Besides the notions of maximality and optimality, we therefore consider two further definitions of ‘best’ applicable to this general setting: the notions of *unmatchedness* (inexistence of weakly preferred alternatives) and, in particular, of *acceptability* (membership to so-called *minimal retentive sets*, rooted in economic theory). We then develop sound and strongly complete axiom systems for these different notions of ‘best’, on general pairwise comparisons structures, based on a dyadic operator $B(\varphi \mid \psi)$, which can be read as “of all ψ states, the best ones are φ ”.

Related Work Conditional logics of preference structures based on dyadic $B(\varphi \mid \psi)$ operators have been investigated since (Hansson 1969), leading to a well-established literature especially, but not exclusively, in so-called deontic logic (e.g.: (Lewis 1974; Spohn 1975; Åqvist 1984; Parent 2008;

Parent 2013)). The logics we present here are rooted into this tradition but, crucially, they lift the assumptions on the existence of maximal or optimal elements. Such assumptions have taken two forms in the literature: imposing the set of ‘best’ elements (maximal or optimal) be non-empty (the so-called limitedness (Åqvist 1984));¹ or requiring that for any non-empty set X , a given element is either a ‘best’ element in X , or there is a ‘better’ element that is ‘best’ in X (so-called smoothness (Kraus, Lehmann, and Magidor 1990), or stopperedness (Makinson 2005)). Two sound and strongly complete axiom systems are known for dyadic operators $B(\varphi \mid \psi)$ (cf. (Parent 2013)) for reflexive and, respectively, reflexive and transitive relations in the presence of smoothness assumptions on maximality and optimality.

Contribution & Outline Our paper further extends these axiomatization results by providing a sound and strongly complete axiom system for: general binary relations under various notions of ‘best’ (maximality, optimality, unmatchedness) but without any smoothness assumptions; and general binary relations under the novel acceptability-based notion of ‘best’. The first system can be regarded as an axiomatization of the weakest possible dyadic modal logic of preference. The second system is, to the best of our knowledge, the first axiomatization of the dyadic modal logic of retentive sets. We believe such axiomatization to be also of technical interest as it requires a novel type of canonical model construction.

Section 2 introduces and compares the four concepts of maximality, optimality, unmatchedness and acceptability. Section 3 defines the semantics induced by them. We present then the two axiomatizations for $B(\varphi \mid \psi)$, one that is sound and strongly complete for the maximality, optimality and unmatchedness semantics (Section 5.1), and one for the acceptability semantics (Section 5.2).

2 Preliminaries

Let \mathcal{P} be a countable set of propositional atoms.

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
- $V : \mathcal{P} \rightarrow 2^S$ assigns to each atom a set of states.

We interpret $s_1 \preceq s_2$ as “ s_2 is (weakly) better than s_1 ” and define \prec, \succeq and \succ from \preceq in the usual way. Furthermore, we use $s_1 \perp s_2$ to denote $s_1 \not\preceq s_2$ and $s_2 \not\preceq s_1$. Finally, we use $s_1 \approx s_2$ to denote $s_1 \preceq s_2$ and $s_2 \preceq s_1$. Note that we do not place any conditions on \preceq . In particular, it is not required to be a partial order.²

¹Alternatively one can require the asymmetric part of the preference relation to be conversely well-founded (cf. (van Benthem, Grossi, and Liu 2013)).

²We also do not require \preceq to be reflexive. Requiring reflexivity would not influence the results in any way, however, as all notions of ‘best’ we consider are invariant under taking the reflexive closure of \preceq .

The meaning of S and \preceq depends on context. When representing belief, S is a set of possible worlds and \preceq an agent’s plausibility ordering, with $s_1 \preceq s_2$ meaning that s_2 is at least as plausible as s_1 . In that case, the agents believes the actual world to be among the “best” ones. When representing individual or societal preferences, S is again a set of possible worlds, but \preceq is a preference relation with $s_1 \preceq s_2$ meaning that s_2 is weakly preferred over s_1 . In that case, there is an obligation to make one of the “best” worlds become true. In game theory S might be a set of strategies and \preceq a dominance order, with $s_1 \preceq s_2$ meaning that s_2 weakly dominates s_1 . The “best” strategies are the ones that are rational. Finally, in argumentation theory S is a set of arguments and \preceq represents the attack relation, with $s_1 \preceq s_2$ meaning that either s_2 attacks s_1 or that a symmetric attack exists between s_2 and s_1 . The “best” arguments must be able to attack the “sub-best” ones.

It is important to note that for each interpretation there are situations where \preceq cannot be assumed to be transitive. If \preceq represents plausibility, non-transitivity can occur when an agent is irrational, or when a rational agent bases its plausibility on a majority judgment among multiple experts (cf. the well-known Condorcet paradox), or when an agent noisily estimates an underlying transitive order (cf. (Truchon 2008)). If \preceq represents preference, non-transitivity can also occur due to individual irrationality or through aggregation from multiple rational agents. If \preceq represents dominance, or an argumentative attack relation, it is clearly possible to have a cycle, such as *paper* \prec *scissor* \prec *rock* \prec *paper*. Likewise, there are scenarios where \preceq cannot be assumed to be total or anti-symmetric.

Before defining the four different ways in which a state can be among the “best” states, we first 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 $Y' \subset Y$ that is retentive in X .

Where X is understood we often write “minimal retentive” for “minimal retentive in X .”

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 there is no $s_2 \in X$ such that $s_1 \prec s_2$.
- *optimal* in X if for every $s_2 \in X \setminus \{s_1\}$, $s_2 \preceq s_1$.
- *unmatched* in X if for every $s_2 \in X \setminus \{s_1\}$, $s_1 \not\preceq s_2$.
- *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 in X w.r.t. M by $\max^M(X)$, the set of optimal elements in X w.r.t. M by $\text{opt}^M(X)$, the set of unmatched elements in X w.r.t. M by $\text{unm}^M(X)$ and the set of acceptable elements in X w.r.t. M by $\text{acc}^M(X)$. Where this should not cause confusion we will omit reference to M and speak of $\max(X)$, $\text{unm}(X)$, $\text{opt}(X)$ and $\text{acc}(X)$.




Figure 1: Four example models.

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.

Maximality and optimality are well studied, see almost any of the sources cited in the introduction. Unmatchedness in preference structures has not, to the best of our knowledge, been studied previously, but it is a straightforward dual of optimality.

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., (Laslier 1997). In the related field of voting theory, minimal retentive sets are also known as (*generalized*) *Condorcet sets* or sometimes *Smith sets*, after (Smith 1973). In game theory, minimal retentive sets are known as *sink equilibria* (Goemans, Mirrokni, and Vetta 2005).

The intuition behind a minimal retentive set is that it is collectively maximal. Thus, while there might not be a compelling argument to consider any individual member of the set to be “best”, the set as a whole should arguably be considered “best”.

Example 1. 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

We use the models M_1 – M_4 only to illustrate the technical aspects of the various notions of being the “best”. See Section 4 for a number of examples that focus more on how models, and the language defined in Section 3, can be interpreted.

Note that this example shows that the four notions are all different. They are not, however, unrelated.

Proposition 4. For every $M = (S, \preceq, V)$ and every $X \subseteq S$,

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

Furthermore, apart from $\text{opt}(X) \subseteq \text{acc}(X)$ and $\text{unm}(X) \subseteq \text{acc}(X)$ (which follow by transitivity of \subseteq) no other inclusions hold in general.

Proof. Inclusions 1., 2. and 3. follow immediately from Definition 3. 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 1. \square

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

Proposition 5. Let $M = (S, \preceq, V)$ be a model and $X \subseteq S$.

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

Before proving Proposition 5, let us consider an auxiliary lemma.

Lemma 6. 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 $Y_1 = \{s_1\} \cup \{t \in Y \mid s_1 \prec t\}$ and $Y_2 = \{s_2\} \cup \{t \in Y \mid s_2 \prec t\}$. 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.

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 is retentive. Similar reasoning shows that Y_2 is also 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 $s_1 \prec s_2$ 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. \square

We now continue to prove Proposition 5.

Proof of Proposition 5. Suppose that M is transitive and that $s \in \text{acc}(X)$. Then there is a minimal retentive $Y \subseteq X$ such that $s \in Y$. By Lemma 6, $Y = \{s\}$. As Y is retentive, we then have $s \not\prec x$ for all $x \in X \setminus \{s\}$. So $s \in \text{max}(X)$.

Suppose that M is anti-symmetric and that $s \in \text{max}(X)$. Then for every $x \in X$, $s \not\prec x$. Due to anti-symmetry this implies that $s \not\prec x$ for all $x \in X \setminus \{s\}$, so $s \in \text{unm}(X)$.

Suppose that M is total and that $s \in \text{max}(X)$. 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 $s \in \text{opt}(X)$. \square

The main lesson of Propositions 4 and 5 is that for a well-behaved relation \preceq 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.

Before continuing with the language and semantics, let us consider two more lemmas that will be useful later. The first lemma states that max depends only on the strict relation \prec , and follows immediately from the definition of max .

Lemma 7. 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 $X \subseteq S$, $\max^{M_1}(X) = \max^{M_2}(X)$.

The second lemma shows that \max , opt and unm satisfy a “pointwise monotonicity” property.

Lemma 8. Let $M = (S, \preceq, V)$ be any model and $\beta \in \{\max, \text{opt}, \text{unm}\}$. Take any $Y \subseteq X \subseteq S$ and let $s \in \beta(X) \cap Y$. Then $s \in \beta(Y)$.

Proof. We give the proof for $\beta = \max$, the proofs for $\beta = \text{opt}$ and $\beta = \text{unm}$ are analogous. Suppose that $s \in \max(X) \cap Y$. Then for every $t \in X$, $s \not\prec t$. Because $Y \subseteq X$, it follows that for every $t \in Y$, $s \not\prec t$, so $s \in \max(Y)$. \square

Note that this pointwise monotonicity 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.

3 Language and Semantics

The models developed in the preceding section allow us to represent belief, preference or dominance. 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.

The main operators in this language is $B(\varphi \mid \psi)$. The formula $B(\varphi \mid \psi)$ holds if and only if, among the states that satisfy ψ , all best states satisfy φ . The way to read $B(\varphi \mid \psi)$ depends on the meaning of \preceq . For example:

- 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.
- 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.
- If \preceq is a dominance relation among strategies, then $B(\varphi \mid \psi)$ represents 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.

In addition to $B(\varphi \mid \psi)$ we also use its dual $P(\varphi \mid \psi)$, which represents conditional plausibility, conditional permission and conditional rational permissibility. We also use the standard Boolean operators and a universal modality $\Box\varphi$ which holds if φ is true in every state.

Definition 9. 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 $p \in \mathcal{P}$.

The operators \wedge , \rightarrow , \leftrightarrow , \top , \perp , \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 $\neg\Box\neg$.

The formula $P(\varphi \mid \psi)$ can be read as φ being conditionally plausible, permissible or playable, depending on the meaning of \preceq .

Based on the four notions of “bestness”, we get four variants or our semantics.




Figure 2: The model M_A for Example 2.

Definition 10. Let $\beta \in \{\max, \text{unm}, \text{opt}, \text{acc}\}$, let $M = (S, \preceq, V)$ be a model and let $s \in S$. The satisfaction relation \models_β is given recursively by

$$\begin{aligned} M, s \models_\beta p &\Leftrightarrow s \in V(p) \\ M, s \models_\beta \neg\varphi &\Leftrightarrow M, s \not\models_\beta \varphi \\ M, s \models_\beta \varphi \vee \psi &\Leftrightarrow M, s \models_\beta \varphi \text{ or } M, s \models_\beta \psi \\ M, s \models_\beta B(\varphi \mid \psi) &\Leftrightarrow \forall s' \in \beta(\llbracket \psi \rrbracket_\beta) : M, s' \models_\beta \varphi \\ M, s \models_\beta \Box\varphi &\Leftrightarrow \forall s' \in S : M, s' \models_\beta \varphi \end{aligned}$$

where $\llbracket \psi \rrbracket_\beta := \{t \in S \mid M, t \models_\beta \psi\}$.

If $M, s \models_\beta \varphi$ for every state s of M we write $M \models_\beta \varphi$. If $M \models_\beta \varphi$ for every model M , 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 have $M, s \models_\beta \varphi$ we write $\Gamma \models_\beta \varphi$. When β is clear from context we write \models for \models_β and $\llbracket \psi \rrbracket$ for $\llbracket \psi \rrbracket_\beta$.

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 2. Alice is currently inside, in a position where she cannot directly observe the outside. She is reasoning about whether it is 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 2. Note that Alice considers s_2 and s_3 to be equally plausible.

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 $M_A \models P(\neg r \wedge \neg w \mid \top) \wedge B(\neg r \wedge \neg w)$. Her conditional beliefs do depend on the exact semantics that we use, however.

Consider $\llbracket w \rrbracket = \{s_2, s_3\}$. Both s_2 and s_3 are acceptable, maximal and optimal in $\{s_2, s_3\}$. As such, for $\beta \in \{\max, \text{opt}, \text{acc}\}$ 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 r and $\neg r$ plausible given 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, $M_A \models_{\text{unm}} \neg P(r \mid w) \wedge \neg P(\neg r \mid w)$.

The different semantics correspond to different standards of evidence for when something is to be considered plausible. The appropriate choice of semantics therefore depends on the standard used by Alice.

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

Example 3. 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.

In \max , opt and unm semantics, the cycle between p , c and f means that there is no best option. We then have $\neg P(\top \mid \top)$, 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 the clear best option is c , 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 4. *Heartstone* is a digital card game. Here we look at the pre-match strategy in *Hearthstone*, which consists of choosing 30 virtual cards to form a deck.³ 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 five such decks: “Control Warrior” (CtW), “Combo Warrior” (CbW), “Control Rogue” (CtR), “Tempo Rogue” (TR) and “Zoo Warlock” (ZL). In a match between different decks one is usually favoured over the other. For example, CtW usually beats, and therefore dominates, TR . There are also some cases where no side is favoured, however. For example, ZW is neither favoured nor disfavoured against CbW . See Figure 3 for the model M_C that shows the dominance relations between all five decks. Note that TR , CtW and CbW form a “rock-paper-scissor”-like cycle, where each strategy is dominated by the next one.

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: $M_C \models_{\text{acc}} P(TR \mid \top) \wedge P(CtW \mid \top) \wedge P(CbW \mid \top)$ and $M_C \models_{\text{acc}} B(TR \vee CtW \vee CbW \mid \top)$. 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 it made matches last too long. Consider therefore the restriction of M_C to

³There is also strategy involved in the moment to moment play during a match, but here we focus on the pre-match part.




Figure 3: The model M_C for Example 4.




Figure 4: The restriction of M_C to $\llbracket \neg el \rrbracket$.

$\llbracket \neg el \rrbracket$, shown in Figure 4. The reason CbW was previously playable was that it beats CtW . In the restricted strategy space there is no longer any reason to play CbW , so it is not 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 $M_C \models_{\text{acc}} \neg P(CbW \mid \neg el) \wedge \neg P(CtR \mid \neg el) \wedge P(TR \mid \neg el) \wedge P(ZW \mid \neg el)$. Note that we have at least one instance of each possible combination of unrestricted and restricted playability: TR remains playable, CtR remains unplayable, ZW becomes playable and CbW becomes unplayable.

5 Axiomatizations

We now move to the main technical results of the paper. Due to space constraints we cannot present all proofs here, an extended version with the additional proofs is available online.⁴

5.1 An Axiomatization for \max , opt and unm

We will introduce axiomatizations for \max , opt , unm and acc . The most interesting axiomatization, in our opinion, is the one for acc . The canonical model construction for that axiomatization is rather complex however, so we begin by considering an axiomatization **MOU** that is sound and complete for \max , opt and unm . We start by showing that it is sound and complete for opt .

Definition 11. The axiomatization **MOU** contains the following rules and axioms:

PL Substitution instances of propositional validities

⁴At https://sites.google.com/site/lbkuijer/maximality_KR.pdf.

\square S5	the S5 axioms and necessitation for \square
R-Ext	$\square(\varphi \leftrightarrow \psi) \rightarrow (B(\chi \mid \varphi) \rightarrow B(\chi \mid \psi))$
L-Ext	$\square(\varphi \leftrightarrow \psi) \rightarrow (B(\varphi \mid \chi) \rightarrow B(\psi \mid \chi))$
MP	from φ and $\varphi \rightarrow \psi$ infer ψ
Abs	$B(\varphi \mid \psi) \leftrightarrow \square B(\varphi \mid \psi)$
K	$B(\varphi \rightarrow \psi \mid \chi) \rightarrow (B(\varphi \mid \chi) \rightarrow B(\psi \mid \chi))$
Id	$B(\varphi \mid \varphi)$
Sh	$B(\varphi \mid \psi \wedge \chi) \rightarrow B(\psi \rightarrow \varphi \mid \chi)$

If a formula φ is derivable in **MOU** from a set Γ of premises, we write $\Gamma \vdash_{\text{MOU}} \varphi$. When the proof system is clear from context we write \vdash for \vdash_{MOU} .

The axioms are self-explanatory except perhaps for **Sh** which is named after (Shoham 1988) and expresses a form of conditionalization (cf. also (Kraus, Lehmann, and Magidor 1990)). **MOU** is closely related to the system known as **F** (Åqvist 1984) but without the axiom enforcing $\llbracket B(\top \mid \psi) \rrbracket$ to be non-empty whenever $\llbracket \psi \rrbracket$ is non-empty.

Lemma 12 (Soundness). *If $\Gamma \vdash_{\text{MOU}} \varphi$ then $\Gamma \models_{\text{opt}} \varphi$.*

Proof. Soundness of all rules and axioms other than **Sh** is immediately obvious from the semantics for **opt**. We therefore prove only the soundness of **Sh** in detail.

Take any model $M = (S, \preceq, V)$ and any $s \in S$. Suppose that $M, s \not\models B(\psi \rightarrow \varphi \mid \chi)$. Then there is some $t \in \text{opt}(\llbracket \chi \rrbracket)$ such that $M, t \not\models \psi \rightarrow \varphi$ and therefore $M, t \models \psi$ and $M, t \not\models \varphi$. This implies that $t \in \llbracket \psi \wedge \chi \rrbracket$. By pointwise monotonicity (Lemma 8) it follows that $t \in \text{opt}(\llbracket \psi \wedge \chi \rrbracket)$. 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, $M, s \models B(\varphi \mid \psi \wedge \chi) \rightarrow B(\psi \rightarrow \varphi \mid \chi)$. \square

Completeness of **MOU** for **opt** is shown through a reasonably standard canonical model construction, although there are a few small complications.

Definition 13. A set $\Gamma \subseteq \mathcal{L}$ is **MOU-consistent** if $\Gamma \not\vdash_{\text{MOU}} \perp$. It is **maximal MOU-consistent** if it is **MOU-consistent** and, furthermore, for every $\varphi \in \mathcal{L}$ either $\varphi \in \Gamma$ or $\neg\varphi \in \Gamma$. The set of maximal **MOU-consistent** sets is denoted MCS^{MOU} .

When this should not cause confusion we will omit reference to the proof system with respect to which a set is (maximal) consistent.

Definition 14. Let $\Gamma \subseteq \mathcal{L}$. The B -inverse of Γ with respect to ψ is given by $B_{\psi}^{-1}(\Gamma) := \{\varphi \mid B(\varphi \mid \psi) \in \Gamma\}$. The \square -inverse of Γ is given by $\square^{-1}\Gamma := \{\varphi \mid \square\varphi \in \Gamma\}$.

We can now begin to construct the canonical model. As usual, the states of the canonical model will be based on maximal consistent sets. However, it is not quite as simple as saying that the set of states equals the set of maximal consistent sets. Firstly, because we have a universal modality \square we need to keep the \square -inverse constant throughout the model. This means that the canonical model will be relative to some set Ξ , and we will only consider those maximal consistent sets Γ where $\square^{-1}\Gamma = \Xi$. Secondly, we will need

multiple copies of each maximal consistent set. Specifically, for every $\psi \in \mathcal{L}$ we will have a copy (Γ, ψ) . The formula ψ in this case is the “intended relativization”. Generally, it is the ψ -copy of an MCS that will be optimal in $\llbracket \psi \rrbracket$.⁵

Definition 15. A set $\Xi \subseteq \mathcal{L}$ is \square -maximal if there is a maximal consistent set Γ such that $\Xi = \square^{-1}\Gamma$.

For a \square -maximal set Ξ , take $MCS_{\Xi}^{\text{MOU}} := \{\Gamma \in MCS^{\text{MOU}} \mid \square^{-1}\Gamma = \Xi\}$.

Before defining the canonical model we need a few more lemmas. For most of these lemmas the proof is either standard and well known or very straightforward. We therefore omit those proofs.

Lemma 16 (Lindenbaum lemma). *If Γ is consistent then there is a maximal consistent set $\Delta \supseteq \Gamma$.*

Lemma 17. *Let Γ be a maximal consistent set and suppose that $\square\xi \in \Gamma$. Then for every ψ , we have $B(\top \mid \psi) \in \Gamma$, $B(\xi \mid \psi) \in \Gamma$ and $B(\square\xi \mid \psi) \in \Gamma$.*

Lemma 18. *If $\varphi \in x$ for all $x \in MCS_{\Xi}$, then $\varphi \in \Xi$.*

Lemma 19. *For every $\varphi, \psi \in \mathcal{L}$, every \square -maximal $\Xi \subseteq \mathcal{L}$ and every $x \in MCS_{\Xi}^{\text{MOU}}$, we have $B(\varphi \mid \psi) \in x$ if and only if $B(\varphi \mid \psi) \in \Xi$.*

The next lemma is significantly less straightforward. The proof is rather tedious and not very insightful however, so we include it only in the online version.

Lemma 20. *If $x \in MCS_{\Xi}$ and $\varphi \notin B_{\psi}^{-1}(x)$, then $\{\neg\varphi\} \cup B_{\psi}^{-1}(x) \cup \square\Xi \cup \{\neg\square\zeta \mid \zeta \in \mathcal{L} \setminus \Xi\}$ is consistent.*

Now, we can define the canonical model.

Definition 21. Let Ξ be \square -maximal. Then the canonical model $M_{\Xi}^{\text{MOU}} = (S, \preceq, V)$ is given by

- $S = \{(x, \varphi) \mid x \in MCS_{\Xi}^{\text{MOU}}, \varphi \in \mathcal{L}\}$,
- $V(p) = \{(x, \varphi) \in S \mid p \in x\}$,
- $\preceq \subseteq S \times S$ is the largest relation such that
 - if $B_{\psi}^{-1}(\Xi) \not\subseteq x$, then $(y, \psi) \not\preceq (x, \varphi)$ and
 - if $\psi \in x$ and $\varphi \notin y$ then $(y, \psi) \not\preceq (x, \varphi)$.

For $\delta \in \mathcal{L}$ we define $[\delta] := \{(x, \varphi) \mid \delta \in x\}$.

Lemma 22 (Truth lemma for **MOU** and **opt**). *For every \square -maximal Ξ and every $\varphi \in \mathcal{L}$, we have $[\varphi] = \llbracket \varphi \rrbracket$.*

See the online version for a proof of the truth lemma.

Lemma 23 (Completeness). *If $\Gamma \models_{\text{opt}} \varphi$ then $\Gamma \vdash_{\text{MOU}} \varphi$.*

We have now proved soundness (Lemma 12) and completeness (Lemma 23) of **MOU** for **opt**.

Theorem 1. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models_{\text{opt}} \varphi$ if and only if $\Gamma \vdash_{\text{MOU}} \varphi$.*

We now have a sound and complete axiomatization for \models_{opt} . Left to do is prove that **MOU** is sound and complete for \models_{max} and \models_{unm} as well.

Theorem 2. *Let $\beta_1, \beta_2 \in \{\text{max}, \text{opt}, \text{unm}\}$. 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 $\varphi \in \mathcal{L}$*

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

⁵This approach is also taken, e.g., in (Parent 2013).

Proof. We consider only two cases here. The remaining cases can be shown similarly, and are proven in the online version.

- Suppose $\beta_1 = \max$ and $\beta_2 = \text{opt}$. 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 7 we have $\max^{M_1}(X) = \max^{M_2}(X)$ for all $X \subseteq S$. Furthermore, \preceq_2 is total so by Proposition 5, $\max^{M_2}(X) = \text{opt}^{M_2}(X)$. Together, these two equalities imply that $\max^{M_1}(X) = \text{opt}^{M_2}(X)$, from which it follows easily that $M_1, s_1 \models_{\max} \varphi$ iff $M_2, s_1 \models_{\text{opt}} \varphi$.
- Suppose $\beta_1 = \text{opt}$ and $\beta_2 = \max$. 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
 - $(s, i) \approx_2 (s, j)$ 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 $s \in \text{opt}(X)$. 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 $(s, i) \in \max(X \times \{0, 1, 2\})$.
 Now, take any $X \subseteq S_1$ and any $s \notin \text{opt}(X)$. 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, $(s, i) \prec_2 (t, i + 1 \bmod 3)$ for all i . In either case, $(s, i) \notin \max(X \times \{0, 1, 2\})$.
 Together, this shows that $\max^{M_2}(X \times \{0, 1, 2\}) = \text{opt}^{M_1}(X) \times \{0, 1, 2\}$, from which it follows easily that $M_1, s_1 \models_{\text{opt}} \varphi$ iff $M_2, s_2 \models_{\max} \varphi$. □

Corollary 24. For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models_{\max} \varphi$ if and only if $\Gamma \vdash_{\text{MOU}} \varphi$, and $\Gamma \models_{\text{unm}} \varphi$ if and only if $\Gamma \vdash_{\text{MOU}} \varphi$.

The corollary does not hold for acc however. We will therefore introduce a different axiomatization for acc.

5.2 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 axiom is unsound for acc. Simply removing the axiom would result in an incomplete axiomatization, however. Instead, we must replace **Sh** by a weaker axiom. The axiom that we will use is **Cut**:

$$\mathbf{Cut} \quad (B(\varphi \mid \psi) \wedge B(\chi \mid \varphi \wedge \psi)) \rightarrow B(\chi \mid \psi)$$

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 axiom is also known as *cumulative transitivity* (cf. (Makinson 2005)).

Definition 25. The axiomatization **ACC** contains the following rules and axioms:

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

$$\mathbf{Cut} \quad (B(\varphi \mid \psi) \wedge B(\chi \mid \varphi \wedge \psi)) \rightarrow B(\chi \mid \psi)$$

Derivability, consistency and maximal consistent sets are defined as for **MOU**, mutatis mutandis.

Lemma 26 (Soundness). If $\Gamma \vdash_{\text{ACC}} \varphi$ then $\Gamma \models_{\text{acc}} \varphi$.

Proof. Soundness of all rules and axioms other than **Cut** is, again, obvious from the semantics. We therefore prove only the soundness of **Cut** in detail.

Suppose that $M, s \models B(\varphi \mid \psi) \wedge B(\chi \mid \varphi \wedge \psi)$. Let $X \subseteq \llbracket \psi \rrbracket$ be minimal retentive in $\llbracket \psi \rrbracket$. We will show that X is minimal retentive in $\llbracket \varphi \wedge \psi \rrbracket$.

First, note that $B(\varphi \mid \psi)$ implies that $X \subseteq \llbracket \varphi \rrbracket$. By assumption $X \subseteq \llbracket \psi \rrbracket$, so $X \subseteq \llbracket \varphi \wedge \psi \rrbracket$. Furthermore, since X is retentive in $\llbracket \psi \rrbracket$ it is also retentive in all subsets of $\llbracket \psi \rrbracket$, in particular, in $\llbracket \varphi \wedge \psi \rrbracket$. Finally, suppose towards a contradiction that some $X' \subset X$ is retentive in $\llbracket \varphi \wedge \psi \rrbracket$. Because X is retentive in $\llbracket \psi \rrbracket$, there can be no $x \in X'$ and $y \in \llbracket \psi \rrbracket \setminus X$ such that $x \prec y$. Furthermore, since X' is retentive in $\llbracket \varphi \wedge \psi \rrbracket$ and $X \subseteq \llbracket \varphi \wedge \psi \rrbracket$ there can be no $x \in X', y \in X \setminus X'$ such that $x \prec y$. Together, this shows that X' is retentive in $\llbracket \psi \rrbracket$. But that contradicts the minimality of X . This completes the proof that X is minimal retentive in $\llbracket \varphi \wedge \psi \rrbracket$.

Because $M, s \models B(\chi \mid \varphi \wedge \psi)$ and X is minimal retentive in $\llbracket \varphi \wedge \psi \rrbracket$, we have $X \subseteq \llbracket \chi \rrbracket$. This holds for any minimal retentive set in $\llbracket \psi \rrbracket$, so $M, s \models B(\chi \mid \psi)$. We therefore have $\models (B(\varphi \mid \psi) \wedge B(\chi \mid \varphi \wedge \psi)) \rightarrow B(\chi \mid \psi)$, i.e., **Cut** is sound. □

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_{Ξ}^{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 denote $\Gamma_{\psi}(\Xi) := \{x \in \text{MCS}_{\Xi}^{\text{ACC}} \mid B_{\psi}^{-1}(\Xi) \subseteq x\}$ and $\Delta_{\psi}(\Xi) := \{(x, y, \psi, b, i) \mid x, y \in \Gamma_{\psi}(\Xi)\}$.

Definition 27. Let Ξ be \Box -maximal. The model $M_{\Xi}^{\text{ACC}} = (S, \preceq, V)$ is given by

- $S = \{(x, y, \varphi, b, i) \mid x, y \in \text{MCS}_{\Xi}^{\text{ACC}}, \varphi \in \mathcal{L}, b \in \{0, 1\}, i \in \mathbb{N}\}$,
- $V(p) = \{(x, y, \varphi, b, i) \in S \mid p \in x\}$,
- \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 $x' \in MCS_{\Xi}^{\text{ACC}}$ such that $\psi \notin x'$: $(x, y, \psi, b, i) \preceq (x', x', \perp, 0, 0)$.

The set $\{(x, y, \psi, b, i) \in S \mid \varphi \in x\}$ is denoted $[\varphi]$.

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 $M_{\Xi}^{\text{ACC}}, (x, y, \psi, b, i) \models \varphi$ 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 are not. The bit b is used to create two copies of each ascending 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 figures to explain it. There are five kinds of arrows in M_{Ξ}^{ACC} : the ones that follows from conditions 1, 2a, 2b, 2c and 2d, respectively. Only arrows of type 2d are between states with different relativizing formulas ψ and φ , so Figure 5 shows a “slice” of the model where ψ is held constant.

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 5.

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. This is represented by the second chain from the left in Figure 5. 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 5. 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 5.

It is important to note that any two states $(x, y, \psi, b, i), (x', y', \psi, b', i') \in \Delta_\psi(\Xi)$ are reachable from each other by some number of \prec steps.

We now consider a few lemmas that will be useful in the truth lemma.

Lemma 28. *If $\varphi \neq \psi$ and $(x, y, \varphi, b, i) \preceq (x', y', \psi, b', j)$ then $\psi = \perp$, $(x, y, \varphi, b, i) \prec (x', y', \varphi, b', j)$ and $(x', y', \psi, b', j) \not\preceq^* (x, y, \varphi, b, i)$, where \preceq^* is the transitive closure of \preceq .*

See the online version for a proof.

Lemma 29. *If $(x, y, \psi, i, b) \notin \Delta_\psi(\Xi)$, then for every $\varphi \in \mathcal{L}$, $(x, y, \psi, i, b) \notin \text{acc}([\varphi])$.*

Proof sketch. This follows from the fact that an infinitely ascending chain without outgoing arrows is never acceptable. See the online version for details. \square




Figure 5: A schematic representation of a single ψ -slice of M_{Ξ}^{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.

Lemma 30. *$\Delta_\psi(\Xi)$ is minimal retentive in $[\psi]$.*

Proof. First, let us note that $\Delta_\psi(\Xi) \subseteq [\psi]$, because by **Id** every \square -maximal Ξ must contain $B(\psi \mid \psi)$.

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.

Minimality of $\Delta_\psi(\Xi)$ follows from the fact that, as observed earlier, any two states in $\Delta_\psi(\Xi)$ are reachable from one another. \square

Lemma 31. *If $[\chi] \not\subseteq [\psi]$ or $\Delta_\psi(\Xi) \not\subseteq [\chi]$ then $\Delta_\psi(\Xi) \cap \text{acc}([\chi]) = \emptyset$.*

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 $(x', x', \perp, 0, 0) \in [\chi] \setminus [\psi]$. Suppose towards a contradiction that $(x, y, \psi, b, i) \in \Delta_\psi(\Xi) \cap \text{acc}([\chi])$. Then $(x, y, \psi, b, i) \in X$ for some X that is minimal retentive in $[\chi]$.

Because $(x', x', \perp, 0, 0) \notin [\psi]$ we have $\psi \notin x'$. So, by clause 2d in the definition of \preceq , we have $(x, y, \psi, b, i) \preceq (x', x', \perp, 0, 0)$. By Lemma 28 we then have $(x, y, \psi, b, i) \prec (x', x', \perp, 0, 0)$ and $(x', x', \perp, 0, 0) \not\preceq^* (x, y, \psi, b, i)$.

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

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 $(x', x', \perp, 0, 0)$ where $\psi \notin x$, so $(x', x', \perp, 0, 0) \notin [\psi]$. From $[\chi] \subseteq [\psi]$ it then follows that $(x', x', \perp, 0, 0) \notin [\chi]$, 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 $(x, y, \psi, b, i) \in \Delta_\psi(\Xi) \cap \text{acc}([\chi])$, so there is a minimal retentive set X such that $(x, y, \psi, b, i) \in X$. We distinguish two cases. Firstly, suppose that $\chi \not\subseteq y$, 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. It follows that 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.

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

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

Lemma 32 (Truth lemma for ACC). *For every \square -maximal Ξ , every $\varphi, \psi \in \mathcal{L}$, every $x, y \in MCS_{\Xi}^{\text{ACC}}$, every $b \in \{0, 1\}$ and every $i \in \mathbb{B}$ we have $M_{\Xi}^{\text{ACC}}, (x, y, \psi, b, i) \models_{\text{acc}} \varphi$ 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 $\varphi = B(\gamma \mid \delta)$.

If $B(\gamma \mid \delta) \in x$ then $\square B(\gamma \mid \delta) \in x$ and therefore $\gamma \in B_{\delta}^{-1}(\Xi)$. Take any $s \in \text{acc}([\delta])$. By the induction hypothesis, $[\delta] = [\delta]$, so $s \in \text{acc}([\delta])$. By Lemmas 29–31 we must have one of the following: (i) $s \in \Delta_{\delta}(\Xi)$ or (ii) $s \in \Delta_{\xi}(\Xi)$ where $[\delta] \subseteq [\xi]$ and $\Delta_{\xi}(\Xi) \subseteq [\delta]$. In the first case we have $s \in [\gamma]$, due to $\gamma \in B_{\delta}^{-1}(\Xi)$.

Consider then the second case. Because $[\delta] \subseteq [\xi]$ we have that $\delta \in z$ implies $\xi \in z$ for all $z \in MCS_{\Xi}$. This implies that $\delta \leftrightarrow (\delta \wedge \xi) \in z$, for every z . By Lemma 18 this implies $\delta \leftrightarrow (\delta \wedge \xi) \in \Xi$. Because $x \in MCS_{\Xi}$, we obtain $\square(\xi \leftrightarrow (\xi \wedge \delta)) \in x$. By assumption we have $B(\gamma \mid \delta) \in x$, so by **R-Ext** we have $B(\gamma \mid \xi \wedge \delta) \in x$.

Furthermore, from $\Delta_{\xi}(\Xi) \subseteq [\delta]$ it follows that for every $z \in MCS_{\Xi}$ we have that if $\Delta_{\xi}(\Xi) \subseteq z$ then $\delta \in z$. So there is no maximal consistent set z such that $\square^{-1}z = \Xi$, $B_{\xi}^{-1}(\Xi) \subseteq z$ and $\neg\delta \in z$. By the Lindenbaum lemma (Lemma 16) every consistent set can be extended to a maximal consistent set, so $\{\neg\delta\} \cup B_{\xi}^{-1}(\Xi) \cup \square\Xi \cup \{\neg\square\zeta \mid \zeta \in \mathcal{L} \setminus \Xi\}$ is inconsistent.

We have $B_{\xi}^{-1}(\Xi) = B_{\xi}^{-1}(x)$, so $\{\neg\delta\} \cup B_{\xi}^{-1}(x) \cup \square\Xi \cup \{\neg\square\zeta \mid \zeta \in \mathcal{L} \setminus \Xi\}$ is inconsistent. By Lemma 20 this implies that $\delta \in B_{\xi}^{-1}(x)$. So we have $B(\delta \mid \xi) \in x$.

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

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

If $B(\gamma \mid \delta) \notin x$, then $\gamma \notin B_{\delta}^{-1}(x)$. By Lemma 20, the set $\{\neg\gamma\} \cup B_{\delta}^{-1}(x) \cup \square\Xi \cup \{\neg\square\zeta \mid \zeta \in \mathcal{L} \setminus \Xi\}$ is then consistent. It can therefore be extended to a maximal consistent set x' . By construction, $\square^{-1}x' = \Xi$, so $x' \in MCS_{\Xi}$. Furthermore, $B^{-1}(\Xi) \subseteq y$ so $(x', x', \delta, 0, 0) \in \Delta_{\delta}(\Xi)$.

By Lemma 30 we have $(x', x', \delta, 0, 0) \in \text{acc}([\delta])$. By the induction hypothesis this implies $(x', x', \delta, 0, 0) \in \text{acc}([\delta])$. Finally, $\neg\gamma \in x'$ so $(x', x', \delta, 0, 0) \notin [\gamma]$. By the induction hypothesis this implies $(x', x', \delta, 0, 0) \notin [\gamma]$. We therefore have $\text{acc}([\delta]) \not\subseteq [\gamma]$, so $M_{\Xi}^{\text{ACC}}, x \not\models B(\gamma \mid \delta)$, which was to be shown. \square

Completeness now follows immediately.

Lemma 33 (Completeness). *If $\Gamma \models_{\text{acc}} \varphi$ then $\Gamma \vdash_{\text{ACC}} \varphi$.*

We have now proven soundness (Lemma 26) and completeness (Lemma 33) of **ACC**.

Theorem 3. *For every $\Gamma \subseteq \mathcal{L}$ and every $\varphi \in \mathcal{L}$, we have $\Gamma \models_{\text{acc}} \varphi$ if and only if $\Gamma \vdash_{\text{ACC}} \varphi$.*

6 Conclusions

We have studied the logic of preference when preferences arise from structures of pairwise comparisons which may fail to contain ‘best’ elements according to maximality and optimality. The logical language of choice for our study has been that of dyadic operators of the type $B(\varphi \mid \psi)$. These operators have been investigated mostly in the deontic logic tradition and are object-level representations of propositional logic consequence relations based on preferential models. We showed that on general binary relations—and in the absence of axioms stipulating the existence of ‘best’ elements—maximality, optimality, and the related notion of unmatchedness all give rise to the same logic. We provided a sound and strongly complete axiom system for such logic. We then proposed the use of minimal retentive sets as a less-demanding notion of ‘best’ to support the semantics of $B(\varphi \mid \psi)$ operators, giving rise to a novel logic of preference for which we provided a sound and strongly complete axiom system.

In this paper we focused on axiomatizations, but several technical avenues for future work present themselves. Much remains to be understood about $B(\varphi \mid \psi)$ logics over general structures of comparison from—in particular—the model-theoretic perspective (but see (Demey 2011) for relevant ideas applied to the case of plausibility models) and the complexity theoretic perspective.

References

- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* 50:510–530.
- Åqvist, L. 1984. Deontic logic. In Gabbay, D., and Guenther, F., eds., *Handbook of Philosophical Logic. Volume II: Extensions of Classical Logic*. Reidel. 605–714.
- Baltag, A., and Smets, S. 2008. A qualitative theory of dynamic interactive belief revision. In Bonanno, G.; Van der Hoek, W.; and Wooldridge, M., eds., *Logic and the Foundation of Game and Decision Theory—LOFT 07*, volume 3 of *Texts in Logic and Games*. Amsterdam University Press. 11–58.
- Bench-Capon, T., and Dunne, P. E. 2007. Argumentation in artificial intelligence. *Artificial Intelligence* 171(10):619–641.
- Boutilier, C. 1990. Conditional logics of normality as modal systems. In *Proceedings of AAAI’90*, 594–599.
- Boutilier, C. 1994. Conditional logics of normality: A modal approach. *Artificial Intelligence* 68:87–154.
- Brandt, F.; Brill, M.; and Harrenstein, P. 2015. Tournament solutions. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A., eds., *Handbook of Computational Social Choice*. Cambridge University Press.
- Burgess, J. 1981. Quick completeness proofs for some logics of conditionals. *Notre Dame Journal of Formal Logic* 22(1):76–84.
- Demey, L. 2011. Some remarks on the model theory of epistemic plausibility models. *Journal of Applied Non-Classical Logic* 21:375–395.
- Goemans, M.; Mirrokni, V.; and Vetta, A. 2005. Sink equilibria and convergence. In *Proceedings of FOCS’05*, 142–151.
- Halpern, J. 2010. From causal models to counterfactual structures. In *Proceedings of the 12th International Conference on the Principles of Knowledge Representation and Reasoning (KR’10)*. AAAI.
- Hansson, B. 1969. An analysis of some deontic logics. *Nous* 3:373–398.
- Katsuno, H., and Mendelzon, A. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52:263–294.
- Katsuno, H., and Satoh, K. 1991. A unified view of consequence relations, belief revision, and conditional logic. In *Proceedings of IJCAI-91, 12th International Joint Conference on Artificial Intelligence*, 406–412.
- Kratzer, A. 1981. Partition and revision: The semantics of counterfactuals. *Journal of Philosophical Logic* 10(2).
- Kraus, S.; Lehmann, D. J.; and Magidor, M. 1990. Non-monotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44(1-2):167–207.
- Laslier, J.-F. 1997. *Tournament Solutions and Majority Voting*. Springer.
- Lewis, D. 1973. *Counterfactuals*. Oxford: Blackwell.
- Lewis, D. 1974. Semantic analysis for dyadic deontic logic. In Stenlund, S., ed., *Logical Theory and Semantical Analysis*. Reidel Publishing Company. 1–14.
- Makinson, D. 1993. Five faces of minimality. *Studia Logica* 52:339–379.
- Makinson, D. 2005. *Bridges from Classical to Nonmonotonic Logic*, volume 5 of *Texts in Computing*. King’s College Publications.
- Nejdl, W. 1991. The P-systems: A semantic classification of logics of nonmonotonicity. In *Proceedings of AAAI’91*, 366–372. AAAI Press.
- Parent, X. 2008. On the strong completeness of Åqvist’s dyadic deontic logic G. In *Proceedings of DEON’08*.
- Parent, X. 2013. Maximality vs. optimality in dyadic deontic logic. completeness results for systems in hansson’s tradition. *Journal of Philosophical Logic* 43:1101–1128.
- Rahwan, I., and Simari, G., eds. 2009. *Argumentation in Artificial Intelligence*. Springer.
- Shoham, Y. 1988. *Reasoning about Change*. MIT Press.
- Smith, J. H. 1973. Aggregation of preferences with variable electorate. *Econometrica* 41(6):1027–1041.
- Spohn, W. 1975. An analysis of hansson’s dyadic deontic logic. *Journal of Philosophical Logic* 4:237–252.
- Stalnaker, R. 1968. A theory of conditionals. In Rescher, N., ed., *Studies in Logical Theory*. Oxford. 98–112.
- Truchnon, M. 2008. Borda and the maximum likelihood approach to vote aggregation. *Mathematical Social Sciences* 55(1):96–102.
- van Benthem, J.; Grossi, D.; and Liu, F. 2013. Priority structures in deontic logic. *Theoria* 80(2):116–152.
- van Benthem, J. 1989. Semantic parallels in natural language and computation. In Ebbinghaus, H., ed., *Logic Colloquium ’87*. North-Holland.
- van Benthem, J. 2004. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logic* 14(2).
- van Benthem, J. 2011. *Logical Dynamics of Information and Interaction*. Cambridge University Press.
- Von Wright, G. H. 1963. *The Logic of Preference*. Edinburgh University Press.