

Characterizability in Belief Revision

György Turán*

University of Illinois at Chicago

MTA-SZTE Research Group on Artificial Intelligence, Szeged

and

Jon Yaggie*

University of Illinois at Chicago

Abstract

A formal framework is given for the postulate characterizability of a class of belief revision operators, obtained from a class of partial preorders using minimization. It is shown that for classes of posets characterizability is equivalent to a special kind of definability in monadic second-order logic, which turns out to be incomparable to first-order definability. Several examples are given of characterizable and non-characterizable classes. For example, it is shown that the class of revision operators obtained from posets which are not total is not characterizable.

1 Introduction

How to incorporate changes into a knowledge base is a fundamental problem of AI, referred to as the problem of belief change. The main approach to belief change is the AGM approach pioneered by [Alchourrón, Gärdenfors, and Makinson, 1985]. There is a large variety of results characterizing classes of belief change operators in terms of rationality postulates [Hansson, 1999].

Are there cases where no characterization can be given?

In view of the many positive results, it seems natural to ask this question. Non-axiomatizability and undefinability questions are much studied in mathematical logic and similar issues were also considered in modal logic [Blackburn, de Rijke, and Venema, 2001] and dynamic epistemic logic [van Ditmarsch, van der Hoek, and Kooi, 2007].

Having tools to prove non-characterizability could be useful when one tries to understand the properties of a class of revision operators. While the standard presentations of belief change start with a set of postulates and then find a matching class of revision operators, there may be situations when a class of revision operators comes first. This may happen, for example, when one considers a class of revision operators which is a natural variant of previously considered ones, such as belief bases instead of belief sets. Also, one may introduce a class of efficiently computable revision operators for practical purposes. In these cases one may inquire what are

the properties of the given class of revision operators? For instance, does it form a “nice” class, which can be characterized by postulates?

Proving non-characterizability presupposes a formal definition of a postulate. However, as noted in the survey [Fermé and Hansson, 2011]

“theories of belief change developed in the AGM tradition are not logics in a strict sense, but rather informal axiomatic theories of belief change. Instead of characterizing the models of belief and belief change in a formalized object language, the AGM approach uses a natural language (ordinary mathematical English) to characterize the mathematical structures under study.”

As far as we know, the first result on characterizability for belief revision was given by [Schlechta, 2004] in the framework of distance semantics [Lehmann, Magidor, and Schlechta, 2001]. This work was extended by [Ben-Naim, 2006].

In this paper, we provide a formal framework for studying characterizability based on the approach of [Katsuno and Mendelzon, 1991]. A revision operator $*$ is considered to assign a revised knowledge base $K * \varphi$ to every knowledge base K and every revising formula φ . However, the results remain valid if one considers $*$ to act on a fixed knowledge base K and an arbitrary revising formula φ . We are not discussing iterated revision here, so there is no interaction between the revisions of different knowledge bases. Katsuno and Mendelzon prove the following results:


- a) There is a finite set of postulates such that a revision operator satisfies these postulates iff there is a faithful *total preorder* representing it with minimization.
- b) There is a finite set of postulates such that a revision operator satisfies these postulates iff there is a faithful *partial preorder* representing it with minimization.
- c) There is a finite set of postulates such that a revision operator satisfies these postulates iff there is a faithful *poset* representing it with minimization.

Part a) is a finite version of Grove’s characterization of the AGM postulates in terms of systems of spheres [Grove, 1988]. Part b) is about a more general class of operators and

*Partially supported by NSF grant CCF-0916708

the postulates characterizing it are different from *a*). The postulates in *c*) turns out to be the same as in *b*), i.e., every operator satisfying the postulates in *b*) can already be represented by a faithful poset.

We consider the following general question. Let \mathcal{R} be a family of partial preorders. Is there a finite set of postulates such that a revision operator satisfies these postulates iff there is a faithful partial preorder from \mathcal{R} representing it with minimization? The general setup is illustrated below¹.



This problem corresponds to situations when the partial preorders over the possible worlds are known to satisfy certain properties, and the question is what additional properties of the revision operators are entailed by this extra information. For example, total orders mean that any two models can be compared. Bounded height means that the length of any chain of comparable models is bounded; thus, in a sense, our ability to rank models is limited. Another interesting class (not considered in this paper) is that of orders of bounded dimension [Trotter, 1992]. Here the assumption is that the comparison between two models is determined by a bounded number of underlying criteria, each represented by a total order, and a model is preferred to another if it is preferred according to each criterion.

Our goal is to “characterize postulate characterizability”. This is achieved for families of partial orders by showing that the answer to the question is positive iff the family is definable by a $\forall MSO_{\min}$ sentence, which is a special kind of monadic second-order sentence. Interestingly, $\forall MSO_{\min}$ definability turns out to be incomparable to first-order definability. We use the characterization to give several examples of characterizable and non-characterizable classes of revision operators. The negative results are proved using a “forgetful” version of the Ajtai-Fagin variant of Ehrenfeucht-Fraïssé games [Ajtai and Fagin, 1990; Libkin, 2004].

In Sections 2-7 of the paper we set up the framework for the discussion of characterizability and provide the tools used in the second part of the paper. Sections 8-10 define the classes of revision operators considered, and prove the results on characterizability, resp., non-characterizability. Section 11 contains a diagram summarizing the relationships between characterizability and logical definability. Due to lack of space, proofs are outlined only or omitted.

¹As discussed later on, the mapping from posets to operators is a bijection, unlike in the general case of partial preorders. Therefore we restrict ourselves to posets in this paper.

2 Preliminaries

We consider propositional logic knowledge bases K over a fixed finite set of variables. We write K_n to indicate that K is over n variables. A knowledge base is represented by a single formula². Truth assignments (or interpretations, models, possible worlds) are assignments of truth values to the variables. The set of truth assignments satisfying a formula φ is denoted by $|\varphi|$. Given a set A of truth assignments, $\langle A \rangle$ is some formula φ such that $|\varphi| = A$.

Given a knowledge base K , a belief revision operator $*$ assigns a formula $K * \varphi$ to every formula φ . Here φ is called the revising formula, and $K * \varphi$ is called the revised knowledge base.

A partial preorder is $R = (X, \leq)$, where X is a finite ground set and \leq is a reflexive, transitive binary relation. A poset (or partial order) is, in addition, antisymmetric. Elements a and b are comparable if $a \leq b$ or $b \leq a$ holds. Otherwise they are incomparable. The comparability graph of R is the undirected graph over X such that for any pair of vertices (a, b) is an edge iff $a \leq b$ or $b \leq a$. Elements a and b in a partial preorder are twins, denoted by $a \approx b$, if $a \leq b$ and $b \leq a$. An element a is minimal if there is no b such that $b < a$, where $b < a$ iff $b \leq a$ but $a \not\leq b$. If $X' \subseteq X$ then a is minimal in X' if $a \in X'$ and there is no $b \in X'$ such that $b < a$. The set of minimal elements of X' is denoted by $\min_{\leq} X'$, or simply $\min X'$ if \leq is clear from the context.

Partial preorders will be used to represent preferences over truth assignments, with truth assignments satisfying the knowledge base being most preferred³. This assumes that partial preorders considered have a special structure, referred to as regularity in this paper.

Definition 2.1. (Regular partial preorders) A partial preorder is regular if

1. every minimal element is smaller than any non-minimal element and
2. the number of elements is a power of 2.

The first assumption means that every truth assignment satisfying the knowledge base is preferred to every truth assignment not satisfying it. This is a standard assumption in belief change theory. The second assumption is made because the number of truth assignments is always a power of two, and we identify the elements of the partial preorder with truth assignments. From the point of view of partial preorders these are mild technical assumptions which do not have an essential effect on definability. An example of a non-regular partial preorder is the 4-element poset with $a < b, c < d$ and no other comparability. Condition 1 is satisfied, for example, if there is a unique minimal element.

The following is one of the standard representations of an epistemic state, i.e., a knowledge base with additional epistemic information used in the belief change operations.

²We note that because of finiteness this representation corresponds to the belief set framework. Computational complexity issues are not discussed here thus the details of the representation are irrelevant.

³Following standard usage, $a \leq b$ is taken to mean that a is preferred to b .

Definition 2.2. (*Faithful partial preorder*) A faithful partial preorder for a knowledge base K_n is a pair $F = (R, t)$, where $R = (X, \leq)$ is a regular partial preorder and $t : X \rightarrow \{0, 1\}^n$ is a bijection between the elements of X and truth assignments, such that $a \in X$ is minimal iff $t(a)$ satisfies K_n .

In the standard definition the partial preorder is defined over the set of truth assignments. For our discussion it is more convenient to separate the partial preorder and the labeling of its elements by truth assignments. Similar distinctions are made in modal logic as well [Blackburn, de Rijke, and Venema, 2001].

3 Revision Using Minimization

One of the basic approaches to belief change is to perform minimization using epistemic states represented by faithful partial preorders.


Definition 3.1. (*Revision using minimization*) The revision operator $*_F$ for K , determined by a faithful partial preorder F for K , using minimization is

$$K *_F \varphi = \langle \min_{\leq} t^{-1}(|\varphi|) \rangle.$$

Thus the revised knowledge base is satisfied by the minimal satisfying truth assignments of the revising formula. Faithfulness implies that if the revising formula is consistent with the knowledge base then the revised knowledge base is the conjunction of the knowledge base and the revising formula.

The definition is illustrated in the figure below. Let $n = 2$ and the knowledge base be $K = x_1 \wedge x_2$. Then F_1, F_2 and F_3 are faithful partial preorders for K . The double line in F_3 indicates that $t^{-1}(01)$ and $t^{-1}(00)$ are twins. Consider the revising formula $\varphi = \bar{x}_1$. Elements belonging to $t^{-1}(|\varphi|)$ are shown as black dots and the elements of $\min_{\leq} t^{-1}(|\varphi|)$ are in ovals. Then it holds that

$$K *_F \varphi = \bar{x}_1 \wedge x_2 \quad \text{and} \quad K *_F \varphi = K *_F \varphi = \bar{x}_1.$$



Note that the regular partial preorders underlying F_1 and F_2 are posets, the first one is total and the second one is not. The revision operator $*_{F_1}$ satisfies the postulate “if K and φ are inconsistent then $|K *_F \varphi|$ is a singleton”. The revision operator $*_{F_2}$ does not satisfy the postulate. Also, note that the regular partial preorders underlying F_2 and F_3 are not isomorphic, but the corresponding revision operators are identical. The difference between revision operators determined by partial preorders and posets, indicated by this example, can be formulated as follows.

Definition 3.2. A revision operator $*$ is poset-based if it is of the form $*_F$ for some faithful poset F .

Lemma 3.3. a) There are non-isomorphic regular partial pre-orders R_1 and R_2 such that there are faithful partial pre-orders $F_1 = (R_1, t_1)$ and $F_2 = (R_2, t_2)$ with $*_{F_1} = *_{F_2}$.

b) Let K be a knowledge base, and $*$ be a poset-based revision operator. Then there is a unique faithful poset F such that $* = *_F$.

4 Postulates and Characterizability

Consider the AGM postulate

$$\text{if } K \wedge \varphi \text{ is satisfiable, then } K *_F \varphi = K \wedge \varphi. \quad (1)$$

Here $K, K *_F \varphi$ and φ can be considered as unary predicates over the set of truth assignments, and thus the above postulate can be rewritten as

$$[\exists x(K(x) \wedge \varphi(x))] \rightarrow [\forall x((K *_F \varphi)(x) \leftrightarrow (K(x) \wedge \varphi(x)))]. \quad (2)$$

Postulates refer to a fixed knowledge base K , and are implicitly universally quantified over formula symbols such as φ . They express general requirements that are supposed to hold for all revising formulas. Our proposed definition generalizes this example. This definition is one possibility to formalize the notion of a postulate, It seems natural and covers most postulates considered for belief revision operators (but see the remark after the definition). In order to emphasize that this is just one, though hopefully basic, notion, we refer to such postulates as basic.

Definition 4.1. (*Basic postulate*) A basic postulate \mathcal{P} is a first-order sentence with unary predicate symbols $K, \varphi_1, \dots, \varphi_\ell$ and $K *_F \mu_1, \dots, K *_F \mu_m$, where μ_1, \dots, μ_m are Boolean combinations of $\varphi_1, \dots, \varphi_\ell$.

A revision operator satisfies a basic postulate for a knowledge base K if the basic postulate holds for all $\varphi_1, \dots, \varphi_\ell$, with the variables ranging over the set of sets of truth assignments.

Allowing for Boolean combinations as arguments to the belief revision operator is necessary as many postulates use such constructs. For example, the AGM postulates refer to $K *_F (\varphi \wedge \psi)$. Definition 4.1 covers most postulates in [Katsuno and Mendelzon, 1991] and in Section 7.3 of [Hansson, 1999]. The postulate of relevance seems to be an example which is not covered by this definition. Using terminology to be introduced later on, its standard definition is still monadic second-order but it contains a quantifier alternation.

The definition of characterizability applies to revision operators within the partial preorder minimization framework. Classes of revision operators can be defined by specifying a class of partial preorders.

Definition 4.2. (*\mathcal{R} -revision operator*) Let \mathcal{R} be a family of regular partial preorders. Let K be a knowledge base and $*$ be a revision operator for K . Then $*$ is an \mathcal{R} -revision operator iff there is a faithful partial preorder $F = (R, t)$ for K with $R \in \mathcal{R}$, representing $*$ using minimization.

Lemma 3.3 shows that there is a bijection between poset-based revision operators and the posets generating them,

while this is not always the case for revision operators generated by partial preorders. Therefore, in the rest of the paper we restrict the discussion to posets for simplicity.

Definition 4.3. (Characterization, characterizability) Let \mathcal{R} be a family of regular posets. A finite set of basic postulates \mathcal{P} characterizes \mathcal{R} -revision operators if for every knowledge base K and every poset-based revision operator $*$ for K the following holds: $*$ satisfies the basic postulates in \mathcal{P} iff $*$ is an \mathcal{R} -revision operator.

The family of \mathcal{R} -revision operators is characterizable if there is a finite set of basic postulates characterizing \mathcal{R} -revision operators.

5 min-formulas and Translation

We define a translation of basic postulates as defined in Definition 4.1 into sentences over an extension of the first-order language of posets. The language of posets contains the binary relation symbol \leq and equality.

The translated sentences also contain additional unary predicate symbols A_1, \dots, A_ℓ . These correspond to propositional formulas $\varphi_1, \dots, \varphi_\ell$ occurring in the basic postulates. We introduce some notation for the translations.

Definition 5.1. (Hat) Given a Boolean combination μ of $\varphi_1, \dots, \varphi_\ell$, we denote by $\hat{\mu}$ the first-order formula obtained by replacing the φ 's with A 's.

For instance, for $\mu(x) = \varphi_1(x)$, one has $\hat{\mu}(x) = A_1(x)$, and for $\mu(x) = \varphi_1(x) \wedge \varphi_2(x)$, one has $\hat{\mu}(x) = A_1(x) \wedge A_2(x)$.

Given a formula ν over the language \leq, A_1, \dots, A_ℓ with a single free variable x we write \min_{\leq}^{ν} for a formula expressing that x is a minimal element satisfying ν , i.e.,

$$\min_{\leq}^{\nu}(x) \equiv \nu(x) \wedge \forall y(\nu(y) \rightarrow \neg(y < x)). \quad (3)$$

When \leq is clear from the context it is omitted as a subscript. Minimal elements in the poset are defined by

$$\min(x) \equiv \forall y(\neg(y < x)).$$

We need the special cases when the formula ν is a Boolean combination of the unary predicates A_1, \dots, A_ℓ .

Definition 5.2. (min-formula) A min-formula over the unary predicate symbols A_1, \dots, A_ℓ is a first-order formula built from the A_i s and formulas of the form $\min_{\leq}^{\nu}(x)$, where the ν 's are arbitrary Boolean combinations of the A_i s.

Now we can define the translation of a basic postulate.

Definition 5.3. (Translation) The translation $\tau(P)$ of a basic postulate P is the min-sentence obtained from P by replacing

1. every occurrence of $K(x)$ with $\min(x)$,
2. every occurrence of $\varphi_i(x)$ and $\mu_i(x)$ with their "hat" versions and
3. every occurrence of $K * \mu_i$ with $\min^{\hat{\mu}_i}(x)$.

Note that Part 2 in the definition is redundant as the definition for φ_i is a special case of the definition for μ_i . The translation of a basic postulate is a first-order sentence over

the predicate symbols \leq, A_1, \dots, A_ℓ . It is based on the observation that the definition of revision by minimization (Definition 3.1) uses the underlying poset in a restricted manner, by simply taking the minimal elements of the revising formula.

Example 5.4. (Translation of basic postulate (2)) Applying Definition 5.3 we get the min-sentence

$$\begin{aligned} & [\exists x(\min(x) \wedge A_1(x))] \\ & \rightarrow [\forall x(\min^{A_1}(x) \leftrightarrow (\min(x) \wedge A_1(x)))]. \end{aligned} \quad (4)$$

The following is a direct consequence of the definitions, as τ is a simple syntactic transformation.

Proposition 5.5. The mapping τ is a bijection between basic postulates containing revising formulas $\varphi_1, \dots, \varphi_\ell$ and min-sentences over unary predicates A_1, \dots, A_ℓ .

In order to interpret min-formulas let us introduce the following.

Definition 5.6. (ℓ -extension) Let $R = (X, \leq)$ be a partial preorder. An ℓ -extension of R is a structure

$$R' = (X, \leq, A_1, \dots, A_\ell),$$

where A_1, \dots, A_ℓ are unary relations⁴.

Given $K, \varphi_1, \dots, \varphi_\ell$ and a faithful partial preorder F for K , the $(\varphi_1, \dots, \varphi_\ell)$ -extension of F is determined in the standard way, by interpreting the unary predicate symbols A_1, \dots, A_ℓ by $A_i(a) = \varphi_i(t(a))$. We recall that a distinction is made between an element a of the poset, and the truth assignment $t(a)$ assigned to that element. Again, the following proposition is a direct consequence of the definitions.

Proposition 5.7. Let K be a knowledge base, $F = (R, t)$ be a faithful partial preorder for K and let $*_F$ be the revision operator determined by F using minimization. Let $\varphi_1, \dots, \varphi_\ell$ be propositional formulas and \mathcal{P} be a basic postulate. Then \mathcal{P} is satisfied by $*_F$ for $\varphi_1, \dots, \varphi_\ell$ iff the $(\varphi_1, \dots, \varphi_\ell)$ -extension of F satisfies $\tau(\mathcal{P})$.

6 $\forall MSO_{\min}$ -definability

In the next two sections we develop the concepts and tools for the logical characterization of postulate characterizability using finite model theory [Ebbinghaus and Flum, 2006; Libkin, 2004].

As propositional logic formulas occurring in the basic postulates are implicitly universally quantified and such formulas are translated into subsets of the partial preorders, it is natural to consider universal monadic second-order logic.

A universal monadic second-order ($\forall MSO$) sentence is of the form

$$\Phi = \forall A_1, \dots, A_\ell \Psi, \quad (5)$$

where A_1, \dots, A_ℓ range over unary predicates (or subsets) of the universe, and Ψ is a first-order sentence using the unary

⁴With an abuse of notation, we use the same notation for a predicate symbol and its interpretation over a structure, assuming that the structure is clear from the context.

predicate symbols A_1, \dots, A_ℓ in addition to the original language (in our case \leq and equality). An existential second-order ($\exists MSO$) sentence is of the form $\Phi = \exists A_1, \dots, A_\ell \Psi$.

As noted in the previous section, translations of basic postulates, such as (4), have special structure. They only refer to the underlying order relation \leq in a *restricted manner*, only inside a \min_{\leq}^{μ} -formula. Thus we are using only a fragment of universal monadic second-order logic, defined as follows. The definition of $\exists MSO_{\min}$ sentences is analogous.

Definition 6.1. ($\forall MSO_{\min}$ sentence) A $\forall MSO_{\min}$ sentence is a universal second order sentence where Ψ is a min-sentence.

For example, universally quantifying in (4) over the second-order variable A_1 we get the $\forall MSO_{\min}$ sentence

$$\forall A_1([\exists x(\min(x) \wedge A_1(x)) \rightarrow [\forall x(\min^{A_1}(x) \leftrightarrow (\min(x) \wedge A_1(x)))]]). \quad (6)$$

Thus from the formal rewriting of the AGM postulate (1) as (2), we arrive at (6) through the intermediate step (4). Each step is a simple, invertible syntactic transformation.

Definition 6.2. ($\forall MSO_{\min}$ -definability) A family \mathcal{R} of regular posets is $\forall MSO_{\min}$ -definable if there is a $\forall MSO_{\min}$ sentence Φ such that for every regular poset R it holds that $R \in \mathcal{R}$ iff R satisfies Φ .

The logical characterization of basic postulate characterizability can now be stated.

Theorem 6.3. Let \mathcal{R} be a family of regular posets. The family of \mathcal{R} -revision operators is characterizable iff the family \mathcal{R} is $\forall MSO_{\min}$ -definable.

This theorem reduces questions about characterizability to questions about $\forall MSO_{\min}$ -definability. The next section develops tools from finite model theory for proving undefinability.

7 Games

The q -round first-order Ehrenfeucht - Fraïssé game over two relational structures is played by two players, Spoiler and Duplicator. In each round Spoiler picks one of the structures and an element of that structure. Duplicator responds by picking an element in the other structure. After q rounds Duplicator wins if the mapping, assigning elements picked in the same round to each other, yields isomorphic substructures. Otherwise Spoiler wins. A basic result is that a class of structures is first-order definable iff there is a q such that if the q -round game is played on two structures, one belonging to the class and the other not, then Spoiler has a winning strategy.

In the following definitions we use $\exists MSO_{\min}$ instead of $\forall MSO_{\min}$ for convenience; $\exists MSO_{\min}$ -definability of a class is the same as $\forall MSO_{\min}$ -definability of its complement. First we discuss $\exists MSO$, and then its min-version. The first-order Ehrenfeucht - Fraïssé game has a variant corresponding to $\exists MSO$ definability. A modified version, which is easier to use for proving undefinability, is defined in [Ajtai and Fagin, 1990]. Let \mathcal{R} be a class of structures.

Definition 7.1. ((\mathcal{R}, ℓ, q) - $\exists MSO$ Ajtai-Fagin game) Given a class \mathcal{R} of structures and parameters ℓ and q , the (\mathcal{R}, ℓ, q) - $\exists MSO$ Ajtai-Fagin game is played as follows:

1. Duplicator picks a structure $R_1 \in \mathcal{R}$,
2. Spoiler picks ℓ subsets A_1, \dots, A_ℓ of the universe of R_1 ,
3. Duplicator picks a structure $R_2 \notin \mathcal{R}$ and subsets B_1, \dots, B_ℓ of the universe of R_2 ,
4. Spoiler and Duplicator play a q -round first-order Ehrenfeucht - Fraïssé game on the structures extended with the subsets (i.e., unary relations) selected.

The connection between \exists -definability and the Ajtai-Fagin game is as follows.

Theorem 7.2. [Ajtai and Fagin, 1990] A class \mathcal{R} is $\exists MSO$ -definable iff there are ℓ, q such that Spoiler has a winning strategy in the (\mathcal{R}, ℓ, q) - $\exists MSO$ Ajtai-Fagin game.

We introduce now a further modification of the Ajtai-Fagin game, defined for partial orders only. We refer to this game as the Katsuno-Mendelzon (KM) game. First let us introduce the notion of a variant.

Let A_1, \dots, A_ℓ be unary predicate symbols. There are $L = 2^{2^\ell}$ logically inequivalent Boolean combinations μ of A_1, \dots, A_ℓ . Let us pick unary predicate symbols M_1, \dots, M_L representing them.

Definition 7.3. (ℓ -min-variant) Let $R = (X, \leq)$ be a partial order. An ℓ -min-variant of R is a structure

$$R'' = (X, A_1, \dots, A_\ell, M_1, \dots, M_L),$$

where $R' = (X, \leq, A_1, \dots, A_\ell)$ is an ℓ -extension of R , and M_1, \dots, M_L are the interpretations of the formulas $\min_{\leq}^{\nu}(x)$, for Boolean combinations ν of the A_i s.

Note that R'' is a structure with unary predicates only, the relation \leq is *not* included, it is “forgotten”⁵ (but \leq is used in order to determine the M_i s). So R'' is not an extension of R : therefore it is referred to as a variant.

Definition 7.4. ((\mathcal{R}, ℓ, q) - $\exists MSO_{\min}$ game, or Katsuno-Mendelzon game) Given a class \mathcal{R} of posets and parameters ℓ and q , the (\mathcal{R}, ℓ, q) - $\exists MSO_{\min}$ game is played by Spoiler and Duplicator as follows:

1. Duplicator picks a poset $R_1 = (X_1, \leq_1)$ in \mathcal{R} ,
2. Spoiler picks ℓ subsets A_1, \dots, A_ℓ of X_1 ,
3. Duplicator picks a poset $R_2 = (X_2, \leq_2) \notin \mathcal{R}$ and subsets B_1, \dots, B_ℓ of X_2 ,
4. Form the ℓ -min-variant R_1'' of R_1 determined by the ℓ -extension $R_1' = (X_1, \leq_1, A_1, \dots, A_\ell)$, and the ℓ -min-variant R_2'' of R_2 determined by the ℓ -extension $R_2' = (X_2, \leq_2, B_1, \dots, B_\ell)$,
5. Spoiler and Duplicator play a q -round first-order Ehrenfeucht - Fraïssé game on R_1'' and R_2'' .

Theorem 7.5. A class \mathcal{R} of posets is $\exists MSO_{\min}$ -definable iff there are ℓ and q such that Spoiler wins the (\mathcal{R}, ℓ, q) - $\exists MSO_{\min}$ game.

⁵This is the “forgetful property” of the game, mentioned in the introduction.

For non-characterizability we use the following corollary. The formulation takes into account that the theorem holds for general posets, but we are only interested in regular posets.

Corollary 7.6. *Let \mathcal{R} be a class of regular posets. Assume that for every ℓ and q , Duplicator has a winning strategy in the (\mathcal{R}, ℓ, q) - $\exists MSO_{\min}$ game such that the posets R_2 are also regular. Then \mathcal{R} is not $\exists MSO_{\min}$ -definable.*

8 Classes of Posets

A chain (resp., antichain) is a set of pairwise comparable (resp., incomparable) elements. The *height* (resp., *width*) of a poset is size of a largest chain (resp., antichain). A poset is total (aka linear) if it has width 1. A poset is connected (resp., disconnected) if its comparability graph is connected (resp., disconnected).

As for technical reasons we deal with regular posets in this paper, we use somewhat modified notions, by disregarding the minimal elements of the poset. The reason is that for our purposes the contribution of the minimal elements is not essential, as they relate to the other elements in a trivial way, but they may interfere with the structure of the remaining elements. For example, minimal elements are always incomparable, thus if minimal elements are included then a poset for a knowledge base of more than one truth assignment cannot be total.

The r -height (resp. r -width) of a regular poset is the height (resp., width) of the poset obtained by removing the minimal elements. A regular poset is r -total (resp., r -connected, r -disconnected) if the poset obtained by removing the minimal elements is total (resp., connected, disconnected).

We denote by $\mathcal{H}_{\leq k}$ the class of regular posets with r -height at most k . The classes $\mathcal{H}_{\geq k}, \mathcal{H}_{< k}, \mathcal{H}_{> k}, \mathcal{H}_{=k}$ are defined similarly. For width we use the notations $\mathcal{W}_{\leq k}, \mathcal{W}_{\geq k}, \mathcal{W}_{< k}, \mathcal{W}_{> k}$. The class of r -total (resp., r -connected, r -disconnected) regular posets is denoted by \mathcal{T} (resp., \mathcal{C}, \mathcal{D}).

9 Characterizable Classes

The characterizability results of this section are proved by showing that a given class of posets is $\forall MSO_{\min}$ -definable, or that its complement is $\exists MSO_{\min}$ -definable. It may happen that the standard definition of a class is not suitable for such a definition, but it can be replaced by an equivalent definition which is suitable. As a first example of characterizability we give a basic postulate characterizing revision operators obtained from r -total regular posets.

Theorem 9.1. *The class of \mathcal{T} -revision operators is characterized by the basic postulate*

$$(\forall x(\neg(K(x) \wedge \phi(x))) \wedge \exists x\phi(x)) \rightarrow \exists!x((K * \phi)(x)). \quad (7)$$

The standard definition of totality is that any two elements are comparable. With the modification to consider only non-minimal elements, this becomes the statement that any two non-minimal elements are comparable. This is not suitable for our purposes, as min-formulas cannot talk directly about comparability. An equivalent formulation in the language of min-formulas is that every nonempty set of non-minimal elements has a unique minimal element, expressed by (7).

When is it possible to find such a rewriting? We give examples of both kinds of answers. The positive results are summarized in the following two theorems.

Theorem 9.2. *1. For every k , the class of $\mathcal{H}_{\leq k}$ -revision operators is characterizable.*

2. For every k , the class of $\mathcal{H}_{\geq k}$ -revision operators is characterizable.

3. For every k , the class of $\mathcal{W}_{\leq k}$ -revision operators is characterizable.

We outline the proof of Part 2 by showing that $\mathcal{H}_{< k}$ is $\exists MSO_{\min}$ -definable. A regular poset has height less than k iff there are sets A_1, \dots, A_{k-1} such that the A_i s form a partition of the set of non-minimal elements, and each A_i is an antichain. This definition starts with existential second-order monadic quantifiers, thus the second-order quantifier structure is of the right form. The standard description of a set A_i being an antichain is that there are no comparable pairs in A_i . However, min-definability cannot directly express comparability. An attempt to circumvent this would be to say that a and b are comparable iff the set of minimal elements of the set $\{a, b\}$ is a singleton. This would require a second-order quantifier to follow a first-order one, which is also not possible. The solution is given by the observation that A_i is an antichain iff it is equal to the set of its minimal elements, i.e., $A_i = \min A_i$.

One case where a replacement of the standard definition by a $\forall MSO_{\min}$ -sentence cannot be found is having width at least k : the classes $\mathcal{W}_{\geq k}$ are turn out to be not $\forall MSO_{\min}$ -definable, even though they are first-order definable. Connectedness and disconnectedness are also not $\forall MSO_{\min}$ -definable. These results are presented in the next section.

On the other hand, connected posets of bounded height turn out to be $\forall MSO_{\min}$ -definable. This is an example of an $\forall MSO_{\min}$ -definable class which is *not* first-order.

Theorem 9.3. *For every k , the class of $\mathcal{C} \cap \mathcal{H}_{\leq k}$ -revision operators is characterizable.*

10 Non-characterizable Classes

The negative results are based on Theorem 6.3 and Corollary 7.6, by constructing winning strategies for the Duplicator. We first consider the class of $\mathcal{W}_{\geq 2}$ -revision operators, i.e., the class of revision operators obtained from regular posets which are *not* r -total. As non r -total regular posets form a simple and natural first-order definable class, the non-characterizability of the corresponding class of revision operators might be considered somewhat surprising.

Theorem 10.1. *The class of $\mathcal{W}_{\geq 2}$ -revision operators is not characterizable.*

We outline the proof. Given ℓ and q , we have to describe a winning strategy of the Duplicator in the (\mathcal{R}, ℓ, q) - $\exists MSO_{\min}$ game for the class of regular posets not in $\mathcal{W}_{\geq 2}$, i.e., for the class \mathcal{T} of r -total regular posets.

The poset $R_1 = (X_1, \leq_1)$ picked in step 1. will be a chain on N elements, for some power of 2 to be determined later. Assume that Spoiler picks ℓ subsets A_1, \dots, A_ℓ . For any element $a \in X_1$ we associate a bit-vector with components indicating which subsets A_i a belongs to, and for every Boolean

combination ν of the A_i s, whether a is minimal among elements belonging to ν . These bit-vectors form a coloring of X_1 with $L = 2^{\ell+2^{2^\ell}}$ colors. As there is at most one element which is minimal in a Boolean combination, there are at least $N - 2^{2^\ell}$ elements which are never minimal. If $N - 2^{2^\ell} > L$ then there are at least two never-minimal elements of the same color. Duplicator then forms R_2 by picking such an element a , splitting it into two incomparable elements a and b , and deleting another element c from the same color class. The sets B_i are the same as the A_i s, with the exception of b replacing c . The ℓ -min-variants R_1^* and R_2^* are isomorphic and thus Duplicator can win the first-order game in Step 5.

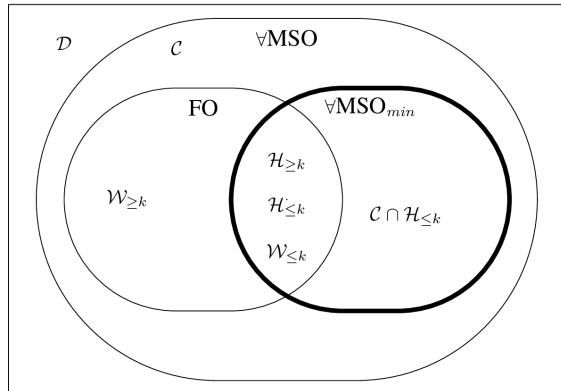
The non-characterizability of $\mathcal{W}_{\geq k}$ -revision operators follows similarly. We have seen in Theorem 9.3 that the class of revision operators generated by connected regular posets of bounded height is characterizable. On the other hand, the full class of revision operators generated by connected regular posets turns out to be non-characterizable. This shows another limitation of $\forall MSO_{\min}$ -definability, as connected posets are $\forall MSO$ -definable. Finally, the class of revision operators generated by disconnected regular posets is not even $\forall MSO$ -definable.

Theorem 10.2. *The classes of \mathcal{C} -revision operators and \mathcal{D} -revision operators are not characterizable.*

11 Summary

The results on classes of revision operators are summarized in the figure below. Characterizable classes (denoted by $\forall MSO_{\min}$ in the figure) turn out to be a proper subset of universal monadic second-order definable classes, incomparable with first-order definable classes. The discussion of some details of the diagram (such as $\mathcal{C} \cap \mathcal{H}_{\leq k}$ not being first-order definable) is omitted due to lack of space.

Characterizable	Non-characterizable
1. Height $\leq k$, height $\geq k$	4. Width $\geq k$
2. Width $\leq k$	5. Connected
3. Connected and height $\leq k$	6. Disconnected



In future work we plan to extend the results of this paper to partial preorders and also to contraction operators. Further questions include characterizability for other types of epistemic information, and for iterated revision.

References

- [Ajtai and Fagin, 1990] Ajtai, M., and Fagin, R. 1990. Reachability is Harder for Directed than for Undirected Finite Graphs. *J. Symbolic Logic* 55(1):113–150.
- [Alchourrón, Gärdenfors, and Makinson, 1985] Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *J. Symbolic Logic* 50(2):510–530.
- [Ben-Naim, 2006] Ben-Naim, J. 2006. Lack of Finite Characterizations for the Distance-based Revision. In *10th International Conference on Principles of Knowledge Representation and Reasoning (KR'06)*, 239–248.
- [Blackburn, de Rijke, and Venema, 2001] Blackburn, P.; de Rijke, M.; and Venema, Y. 2001. *Modal Logic*. New York, NY, USA: Cambridge University Press.
- [Ebbinghaus and Flum, 2006] Ebbinghaus, H.-D., and Flum, J. 2006. *Finite Model Theory*. Springer Monographs in Mathematics. Berlin: Springer-Verlag, enlarged edition.
- [Fermé and Hansson, 2011] Fermé, E., and Hansson, S. O. 2011. AGM 25 Years. Twenty-five Years of Research in Belief Change. *J. Philosophical Logic* 40(2):295–331.
- [Grove, 1988] Grove, A. 1988. Two Modellings for Theory Change. *J. Philosophical Logic* 17:157–170.
- [Hansson, 1999] Hansson, S. O. 1999. *A Textbook of Belief Dynamics*, volume 11 of *Applied Logic Series*. Dordrecht: Kluwer Academic Publishers.
- [Katsuno and Mendelzon, 1991] Katsuno, H., and Mendelzon, A. O. 1991. Propositional Knowledge Base Revision and Minimal Change. *Artificial Intelligence* 52(3):263–294.
- [Lehmann, Magidor, and Schlechta, 2001] Lehmann, D. J.; Magidor, M.; and Schlechta, K. 2001. Distance Semantics for Belief Revision. *J. Symbolic Logic* 66(1):295–317.
- [Libkin, 2004] Libkin, L. 2004. *Elements of Finite Model Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer-Verlag, Berlin.
- [Schlechta, 2004] Schlechta, K. 2004. *Coherent Systems*. Studies in Logic and Practical Reasoning. Elsevier Science.
- [Trotter, 1992] Trotter, W. T. 1992. *Combinatorics and Partially Ordered Sets*. Johns Hopkins Series in the Mathematical Sciences. Baltimore, MD: Johns Hopkins University Press.
- [van Ditmarsch, van der Hoek, and Kooi, 2007] van Ditmarsch, H.; van der Hoek, W.; and Kooi, B. 2007. *Dynamic Epistemic Logic*. Springer Verlag.