Journal of Artificial Intelligence Research 52 (2015) 287-329 Submitted 10/14; published 02/15

Revision by History

Paolo Liberatore LIBERATO@DIS.UNIROMA1.IT
Sapienza University of Rome, DIAG
Via Ariosto 25, 00185 Rome, Italy

Abstract

This article proposes a solution to the problem of obtaining plausibility information, which is necessary to perform belief revision: given a sequence of revisions, together with their results, derive a possible initial order that has generated them; this is different from the usual assumption of starting from an all-equal initial order and modifying it by a sequence of revisions. Four semantics for iterated revision are considered: natural, restrained, lexicographic and reinforcement. For each, a necessary and sufficient condition to the existence of an order generating a given history of revisions and results is proved. Complexity is proved coNP complete in all cases but one (reinforcement revision with unbounded sequence length).

1. Introduction

Many belief revision operators are based on some sort of plausibility order (Spohn, 1988; Boutilier, 1996; Nayak, 1994; Williams, 1994; Areces & Becher, 2001; Zhang, 2004; Benferhat, Kaci, Le Berre, & Williams, 2004; Hild & Spohn, 2008; Fermé & Hansson, 2011). Whenever revising can be done in two or more different ways, the result is the disjunction of either all of them (Alchourròn & Makinson, 1982; Fagin, Ullman, & Vardi, 1983; Winslett, 1988) or the most plausible ones only according to the order (Gärdenfors, 1988; Katsuno & Mendelzon, 1991; Peppas, 2008; Nebel, 1992; Fermé & Hansson, 2011). Fewer disjuncts imply more formulae; therefore, the more discriminating the order, the more informative the result. A fine-grained order is central to the usefulness of the revised knowledge base.

Iterated revision provides itself a way for obtaining a plausibility order. Even starting from an all-equal plausibility order (the least discriminating one), each revision changes it by making some possibilities more plausible than others (Spohn, 1988; Boutilier, 1996; Nayak, 1994; Williams, 1994; Booth & Meyer, 2006; Jin & Thielscher, 2007). A sequence of revisions produces an order that, depending on the revising formulae, is more or less informative. In some cases, this is a solution to the problem of obtaining a plausibility order: by a sequence of previous revisions (Konieczny & Pino Pérez, 2000; Baltag, Gierasimczuk, & Smets, 2011). However, even a long history of revisions may not produce a fine discrimination. In a limit case, after revising by a, $\neg a$, a, $\neg a$, etc., the final order only discriminates models of a from models of $\neg a$.

Is there any other way to obtain an initial plausibility order? One possibility is to derive it from knowledge of the previous results (this can also be done for merging, Liberatore, 2014b, 2014a). In other words, not only the previous revising formulae are given, but also the results they produced. If K_0 is the initial knowledge base and P_1 the first revising formula, the result is another knowledge base K_1, which can be further revised by P_2. In this article, all K_i's and P_i's are assumed known up to a certain point:

$$K_0 \xrightarrow{P_1} K_1 \xrightarrow{P_2} K_2 \xrightarrow{P_3} K_3 \ldots K_{n-1} \xrightarrow{P_n} K_n$$

Such a sequence of consistent formulae $[K_0, P_1, K_1, \ldots, P_n, K_n]$ is called a *revision sequence*. It gives information about the initial plausibility order over models, like the following example shows.

Example 1 *Let $[K_0, P_1, K_1]$ be the revision sequence where:*

$$\begin{aligned} K_0 &= a \\ P_1 &= \neg a \\ K_1 &= \neg a \wedge b \wedge c \end{aligned}$$

This specific $K_1 = \neg a \wedge b \wedge c$ is not the only possible result of revising $K_0 = a$ by $P_1 = \neg a$. While for example $\neg a \wedge \neg b \wedge c$ was also possible, $\neg a \wedge b \wedge c$ is the actual revision result. This means that the model $\{\neg a, b, c\}$ was considered more plausible than $\{\neg a, \neg b, c\}$, another model of P_1. This information can be useful in subsequent revisions.

A revision sequence may be seen as a form of training: the first n revisions are manually performed by human operators, the following others are done automatically using the initial plausibility order obtained from the training. Technically, from the revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ containing the revising formulae and their resulting knowledge bases up to K_n, one derives an initial plausibility order, which is revised by $P_1, \ldots, P_n, P_{n+1}, P_{n+2}, \ldots$ to obtain $K_{n+1}, K_{n+2}, \ldots$ A similar mechanism has been studied by Nittka and Booth (2008); a comparison with their approach is in the conclusions.

Example 2 *As a part of a research project, PhD student Polyxena is tasked with automating the incorporation of new data in a database, a process so far manually performed by the database maintainers, a specialized group of people. She soon realizes that the information to be incorporated may be disjunctive, so that revisions may be done in multiple ways. Introducing data in the form "either $\neg a$ or $\neg b$" when both a and b are in the database requires erasing either a, b or both. Studying the relevant literature on the topic, she finds out that such choices require plausibility information. Since the previous revisions were performed by the database maintainers, they possess this information. She therefore asks them about the rank of the models, getting answers ranging from "what do you mean?" to "I rank Brazilian redhead models at the top, but Swiss models are not bad". After spending half an hour trying to explain to them the concepts of plausibility, consistency, propositional models (and how they differ from "fashion models") and ranking, she gives up.*

When she is about to leave the room, one of the maintainers suggests to her to have a look at the logs, since these files record everything that happened to the database. She does, and indeed all previous revisions are stored in the logs: both the new information and how the database maintainers incorporated it. The problem now shifts from eliciting a rank from the database maintainers, which proved difficult, to determining it from the previous history of revisions. She finds out that the sequence of revisions and results is natural-compatible (Definition 4) but not compatible with the other revisions. This means *that not only an initial plausibility order can be calculated using Lemma 2, but also that the people who performed the previous revisions had (unknowingly) adopted a policy of minimal plausibility change.*

Apart from the remarks of the operators, admittedly surreal for comedic purposes, the example shows a not-so-uncommon scenario: a process performed by hand is to be automated, and eliciting information from the people who have performed it so far is difficult. First, this information may have never been expressed in an explicit form; second, it may be hard to formalize for people lacking a background in formal logic.

In the case of belief revision, the information needed to perform the subsequent revisions is the initial order of the models. However, eliciting an order is not as easy as it may seem, as shown by research on the similar concept of preference (Sandholm & Conen, 2010; Carson & Louviere, 2011), not to mention the experimental results in cognitive psychology (Tversky & Kahneman, 1983): after being given some background information, the majority of the participants in a test reckoned "Linda is a bank teller and is active in the feminist movement" to be more likely than "Linda is a bank teller", while probability theory forbids $a \wedge b$ to be more likely than a.

Furthermore, providing plausibility information is an additional work for the people who have manually performed the process so far. Instead, as Example 1 shows, some information may be derived from the previous history of revisions.

Another example is in data synchronization: the SyncML protocol (OMA, 2002) allows synchronizing data (phonebook, calendar notes, etc.) between a mobile phone and a computer, but conflicts may arise; an implementation may then ask the user what to do, or take a decision (like "phone always wins") which may however be later manually reversed by the user. Either way, the result tells how the conflict should have been solved. Again, both the revising formula and the resulting knowledge base are given, and can be used to derive information about the unknown plausibility order.

As Example 1 shows, from the knowledge of a revision sequence $[K_0, P_1, K_1, P_2, K_2, \ldots, P_n, K_n]$ some information about the initial plausibility of models can be derived. Such an information depends on the revision semantics, and some sequences are not generated at all by some semantics.

Example 3 *Let $[K_0, P_1, K_1, P_2, K_2, P_3, K_3]$ be the revision sequence defined as follows.*

$$
\begin{aligned}
K_0 &= a \\
P_1 &= b \\
K_1 &= a \wedge b \\
P_2 &= c \\
K_2 &= a \wedge b \wedge c \\
P_3 &= \neg a \\
K_3 &= \neg a
\end{aligned}
$$

It will be shown that there exists an initial order of the models that generates this sequence using the natural revision semantics. By contrast, no order generates it with the restrained and lexicographic semantics.

The technical results provided by this article are: first, equivalent formulations for the problem of establishing the existence of an order generating a revision sequence using the natural, lexicographical, restrained and reinforcement revision; second, how an initial order can be built if one exists; third, a complexity characterization.

Since the number of models is exponential in the number of variables, "there exists an order over models such that ..." is a quantification over a data structure that may be exponential in size. As a result, brute-force search takes double exponential time. The equivalent formulations avoid this high computational cost by recasting the problem in terms of polynomial-size data structures.

The problem of establishing the existence of an initial order generating a given sequence is coNP-complete in all cases but one (reinforcement revision with unbounded sequence length), therefore showing that the problem can be expressed as the validity of a ∀QBF. This proves that the problem can be recast in a form that does not contain the existential quantification over the initial order.

2. Preliminaries

Belief revision as considered in this article is on propositional formulae built on a finite alphabet of variables. A truth evaluation over such an alphabet is called a *model*: a function from the variables to either true or false. Following a common terminology of propositional logic, if a model satisfies a formula then it is a model *of* the formula, and the formula *has* that model. The set of models of a formula F is denoted $Mod(F)$. A *QBF* is a propositional formula where the variables are quantified, either universally (like in $\forall a.\neg a \wedge \neg b$), existentially (like in $\exists a \exists b.a \vee b$) or both (like in $\exists a \forall b.a \rightarrow b$). When all variables are universally quantified the formula is a $\forall QBF$.

A revision sequence represents the evolution of beliefs over time, including both the revising formulae and their results.

Definition 1 *A* revision sequence *is an odd sequence of consistent propositional formulae* $[K_0, P_1, K_1, \ldots, P_n, K_n]$ *over a finite set of variables.*

The semantics for belief revision considered in this article work on an ordering of the models, representing their relative plausibility, which is modified when new information arrives. Such orderings can be defined as follows.

Definition 2 *A* total preorder C *is a partition of the models into a finite sequence of classes* $[C(0), C(1), C(2), \ldots, C(m)]$ *with* $C(0) \neq \emptyset$.

Intuitively, such a partition represents a way to compare models: I and J compare the same if they are in the same class, I compares greater than J if it is in a class of higher index. The use of partitions instead of the usual notation $I \leq J$ simplifies definitions and proofs. Since classes can be empty (except the first), several partitions may represent the same way of comparing models. This is not a problem as total preorders are never checked for equality in this article.

A total preorder can be depicted as a shelf, as in Figure 1. The bottom drawer $C(0)$ contains the most plausible models. These represent the situations currently believed possible: $C(0) = Mod(K_0)$. Revising C by P_1 changes it into a new preorder C_{P_1} that takes into account the new information. The class $C_{P_1}(0)$ contains the models that are now considered the most plausible; therefore, $C_{P_1}(0) = Mod(K_1)$.

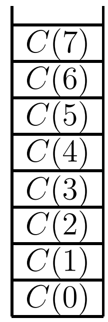

Figure 1: The graphical representation of a total preorder C

Such a partition formalizes the plausibility of models: models of $C(i)$ are more plausible than models of $C(i+1)$. The lower the class, the more plausible the model; for this reason, a total preorder is often seen as representing implausibility rather than plausibility. It is inverse of an ordinal conditional function (Spohn, 1988): $\kappa(I) = n$ if and only if $I \in C(n)$.

The study of two of the semantics considered in this article involves the prefixes and the maxsets of a sequence. Sequences are denoted using brackets $[\ldots]$. Given a sequence of formulae $[P_1, \ldots, P_n]$, its h-prefix is the sequence containing only its first $h-1$ formulae in the sequence. The maxset of a sequence extends the concept of maximal consistent subsets from sets to sequences.

$$\begin{array}{rcl}
\text{maxset}([P_1, \ldots, P_n]) & = & \text{maxset}(\epsilon; [P_1, \ldots, P_n]) \\
\text{maxset}([Q_1, \ldots, Q_i]; [P_1, P_2, \ldots, P_n]) & = & \begin{cases} \text{maxset}([Q_1, \ldots, Q_i, P_1]; [P_2, \ldots, P_n]) \\ \quad \text{if } Q_1 \wedge \cdots \wedge Q_i \wedge P_1 \text{ is consistent} \\ \text{maxset}([Q_1, \ldots, Q_i, \mathsf{true}], [P_2, \ldots, P_n]) \\ \quad \text{otherwise} \end{cases} \\
\text{maxset}([Q_1, \ldots, Q_n]; \epsilon) & = & [Q_1, \ldots, Q_n]
\end{array}$$

A sequence of formulae used in a context where a propositional formula is expected implicitly represents the conjunctions of its formulae. For example, $a \vee [b, \neg c, c \vee d]$ means $a \vee (b \wedge \neg c \wedge (c \vee d))$. According to this notation, $Q_1 \wedge \cdots \wedge Q_i \wedge P_1$ is inconsistent if and only if $[Q_1, \ldots, Q_i] \models \neg P_1$. As a result, $[Q_1, \ldots, Q_i, \mathsf{true}]$ can be replaced by $[Q_1, \ldots, Q_i, \neg P_1]$ in the definition of a maxset.

By definition of prefixes and maxsets, they commute: the h-prefix of the maxset of a sequence is the same as the maxset of the h-prefix of the same sequence. As a result, if P is the h-th element of the sequence S then $P \in \text{maxset}(S)$ if and only if $P \wedge \text{maxset}(\text{prefix}_h(S))$ is consistent.

The maxset is often written $\text{maxset}(P_1, \ldots, P_n)$ as a shorthand of $\text{maxset}([P_1, \ldots, P_n])$. A number of properties of maxsets are now shown. All proofs are in the appendix.

Lemma 1 *If F is consistent and $F \models \text{maxset}(P_1, \ldots, P_n)$, then $\text{maxset}(P_1, \ldots, P_n) \equiv \{P_i \mid 1 \leq i \leq n \text{ and } F \models P_i\}$.*

Lemma 2 *If F is consistent and $F \models \text{maxset}(P_1, \ldots, P_n)$, every consistent subset of $\{P_1, \ldots, P_n\}$ that contains all formulae entailed by F is equivalent to $\text{maxset}(P_1, \ldots, P_n)$.*

In this article, a sequence of formulae where some of them are replaced by `true` is called a subsequence. This is similar to the usual definition, with the difference that formulae maintain their position in the sequence.

Lemma 3 *If F is consistent then $F \not\models \text{maxset}(P_1, \ldots, P_n)$ if and only if there exists a subsequence R of $[P_1, \ldots, P_n]$ such that:*

1. *R is consistent;*
2. *if $F \models P_i$ then $P_i \in R$;*
3. *for some i, $P_i \notin R$ and $P_i \wedge \text{prefix}_i(R)$ is consistent.*

The conditions in the lemma are all of existential type: there exists R, there exists a model for R, either $F \not\models P_i$ or $P_i \in R$ and $P_i \wedge \text{prefix}_i(R)$ is consistent. This proves that checking $F \not\models \text{maxset}(P_1, \ldots, P_n)$ can be expressed as the validity of a $\exists$QBF, and is therefore in NP.

Corollary 1 *If F is consistent, checking $F \not\models \text{maxset}(P_1, \ldots, P_n)$ is in* NP.

The lemma avoids constructing the maxset one formula at time by replacing the test of satisfiability of $P_i \wedge \text{prefix}_i(\text{maxset}(P_1, \ldots, P_n))$ with $F \models P_i$, which is the same if F is consistent and entails the maxset. This way, the sequence of satisfiability checks required to build the maxset are parallelized, that is, turned into a number of validity checks that can be performed in parallel.

In order to check $F \equiv \text{maxset}(P_1, \ldots, P_n)$, one first checks whether $F \models \text{maxset}(P_1, \ldots, P_n)$, and then $\text{maxset}(P_1, \ldots, P_n) \models F$. Assuming that the first condition is true, the second can be shown in coNP.

Theorem 1 *If F is consistent, checking $F \equiv \text{maxset}(P_1, \ldots, P_n)$ is in* coNP.

In this article, only revisions satisfying the AGM postulate 4 are considered: $K * P \equiv K \wedge P$ if $K \wedge P$ is consistent. Also, all formulae P_i and K_i in the sequences are assumed consistent. When checking a sequence for being generated by some total preorder, $K_{i-1} \wedge P_i$ is consistent if and only if $K_i \equiv K_{i-1} \wedge P_i$. In particular, if the sequence is generated by some total preorder and $K_{i-1} \wedge P_i$ is consistent then by the AGM postulate $K_{i-1} \wedge P_i$ is equivalent to K_i; conversely, if $K_i \equiv K_{i-1} \wedge P_i$ is consistent then the consistency of K_i implies that of $K_{i-1} \wedge P_i$. This property is important because it allows replacing a satisfiability test with an unsatisfiability test.

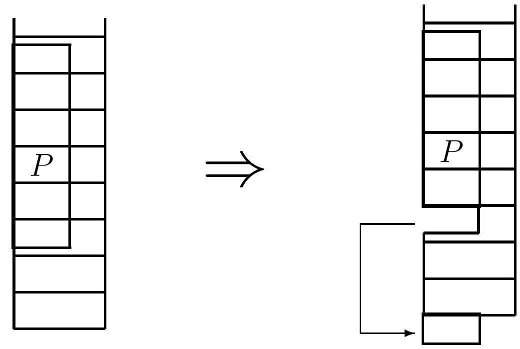

Figure 2: Natural revision

3. Natural Revision

Natural revision (Boutilier, 1996) modifies a total preorder of plausibility of models C in light of a new piece of information P into a new total preorder C_P that is as close as possible to the original one. In the new preorder P has to be true in all most plausible models, these of $C_P(0)$. A minimal change of C ensuring this is setting $C_P(0)$ to be the minimal models of P according to C, leaving the rest of the preorder unaltered.

Definition 3 *The natural revision of the total preorder C by formula P is defined as the total preorder C_P that follows, where i is the minimal index such that $C(i) \cap Mod(P) \neq \emptyset$:*

$$C_P(j) = \begin{cases} C(i) \cap Mod(P) & \text{if } j = 0 \\ C(j-1) \backslash C_P(0) & \text{otherwise} \end{cases}$$

For example, $C_P(0) = C(i) \cap Mod(P)$, while $C_P(1) = C(1-1) \backslash C_P(0) = C(0) \backslash C_P(0)$. Graphically, the change that P produces to the preorder in natural revision can be depicted as "cutting out" the lowest models of P and placing them below the others, as shown in Figure 2.

If $C_{P_1,\ldots,P_i}$ is the result of revising C by P_1, then P_2, etc. using natural revision, then $Mod(K_i) = C_{P_1,\ldots,P_i}(0)$ in the revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$.

Example 4 *Let $C = [C(0), C(1)]$ be the total preorder such that:*

$$\begin{aligned} C(0) &= Mod(a) \\ C(1) &= Mod(\neg a) \end{aligned}$$

Revising it by $P_1 = b$, $P_2 = c$ and $P_3 = \neg a$ using natural revision generates the revision sequence in Example 3. Indeed, revising C by $P_1 = b$ makes the minimal models of P_1 to form the new class zero. Since $C(0) \cap Mod(P) \neq \emptyset$, the index i in the definition of natural revision is zero. The resulting preorder is therefore:

$$\begin{array}{rcl} C_{P_1}(0) & = & C(0) \cap Mod(P_1) = Mod(a \wedge b) \\ C_{P_1}(1) & = & C(0) \backslash C_{P_1}(0) = Mod(a \wedge \neg b) \\ C_{P_1}(2) & = & C(1) \backslash C_{P_1}(0) = Mod(\neg a) \end{array}$$

Since $Mod(K_1) = C_{P_1}(0)$, it follows that $K_1 \equiv a \wedge b$. A similar change happens when revising by $P_2 = c$, since $C_P(0) \cap Mod(P_2) \neq \emptyset$, which implies $i = 0$.

$$\begin{array}{rcl} C_{P_1P_2}(0) & = & Mod(a \wedge b \wedge c) \\ C_{P_1P_2}(1) & = & Mod(a \wedge b \wedge \neg c) \\ C_{P_1P_2}(2) & = & Mod(a \wedge \neg b) \\ C_{P_1P_2}(3) & = & Mod(\neg a) \end{array}$$

Again, $Mod(K_2) = C_{P_1P_2}(0)$, which implies $K_2 \equiv a \wedge b \wedge c$. The minimal models of $P_3 = \neg a$ are the entire class $C_{P_1P_2}(3)$. Therefore, $i = 3$ and the preorder becomes:

$$\begin{array}{rcl} C_{P_1P_2P_3}(0) & = & Mod(\neg a) \\ C_{P_1P_2P_3}(1) & = & Mod(a \wedge b \wedge c) \\ C_{P_1P_2P_3}(2) & = & Mod(a \wedge b \wedge \neg c) \\ C_{P_1P_2P_3}(3) & = & Mod(a \wedge \neg b) \end{array}$$

This proves that $K_3 \equiv \neg a$. The revision sequence coincides with that of Example 3. Looking at this example in the other direction, it shows that the revision sequence $[a, b, a \wedge b, c, a \wedge b \wedge c, \neg a, \neg a]$ is generated by natural revision from some preorder. It will be proved that this is not the case for restrained and lexicographic revisions.

The aim of this article is to establish whether a sequence is generated by some preorder, and finding it. Unfortunately, a direct search in the space of total preorders is unfeasible: the number of models is exponential in the number of the variables, and the number of total preorders is therefore a double exponential. Fortunately, for natural revision this difficulty can be overcome thanks to a necessary and sufficient condition for a sequence to be generated by some total preorder. A number of lemmas are needed to prove it. The first shows that the revising by a formula does not alter the relative order of models that are not in the resulting knowledge base.

Lemma 4 *If $C_P(0) \cap Mod(F) = \emptyset$ then C_P compares the models of F as C does, where C_P is the natural revision of the total preorder C with formula P.*

This result can be iterated over a number of revising formulae: if the resulting knowledge bases K_i are all inconsistent with a formula F, the relative order of models of F is not changed. The result of a final revision by F can therefore be calculated from the original ordering, in this case. The following lemma is formulated over a fragment of a revision sequence because this is how it is later applied.

Lemma 5 *Let* $[K_j, P_{j+1}, \ldots, P_i, K_i]$ *be a revision sequence generated by natural revision from a total preorder* C*. If* $K_j \wedge P_i$ *is consistent while none of* $K_{j+1} \wedge P_i, \ldots, K_{i-1} \wedge P_i$ *is, then* $C_{P_{j+1},\ldots,P_i}(0) = Mod(K_j \wedge P_i)$*.*

This lemma is similar to a result by Boutilier (1996, Thm. 17), but lifts the assumption that all conjunctions $K_j \wedge P_{j+1}, \ldots, K_{i-2} \wedge P_{i-1}$ are consistent. It shows that if P_i is consistent with a previous K_j, then natural revision by P_i produces a result that can be determined from K_j only, independent of the initial preorder. The following lemma covers the other case, where P_i is inconsistent with all previous K_j.

Lemma 6 *If the revision sequence* $[K_0, P_1, K_1, \ldots, P_i, K_i]$ *is generated by natural revision from the total preorder* C *and* P_i *is inconsistent with each of* $K_0, \ldots, K_{i-1}$*, then the models of* K_i *are the minimal models of* P_i *according to* C*.*

The last two lemmas prove that, for natural revision, K_i is equivalent to $K_j \wedge P_i$ for the maximal j for which this conjunction is consistent if one exists, otherwise is determined from the initial preorder. The first is a necessary condition for the existence of a total preorder generating the sequence.

Definition 4 *A revision sequence* $[K_0, P_1, \ldots, P_n, K_n]$ *is* natural-compatible *if, for every* $i \in \{1, \ldots, n\}$*, it holds:*

1. $K_i \models P_i$*;*
2. *if* j *is the maximal index such that* $j < i$ *and* $K_j \wedge P_i$ *is consistent (if any), then* $K_i \equiv K_j \wedge P_i$*.*

If a sequence is natural-compatible then it is generated by natural revision from some initial total preorder.

Theorem 2 *If* $[K_0, P_1, K_1, \ldots, P_n, K_n]$ *is natural-compatible then it is generated by natural revision from the initial preorder* $C = [C(0), \ldots, C(n+1)]$*.*

$$C(i) = \begin{cases} Mod(K_i) & \text{if } i \leq n \text{ and } \forall l < i \;.\; K_l \wedge P_i \models \bot \\ \emptyset & \text{otherwise, if } i \leq n \\ \bigcup\{Mod(K_j) \mid \exists l < j \;.\; K_l \wedge P_j \not\models \bot\} & \text{if } i = n+1 \end{cases}$$

Natural-compatibility is not only a sufficient condition for a sequence being generated by natural revision from some initial preorder. The following theorem proves that it is also necessary. Therefore, it characterizes exactly the revision sequences that natural revision generates.

Theorem 3 *A revision sequence* $[K_0, P_1, K_1, \ldots, P_n, K_n]$ *is generated by natural revision from some initial total preorder if and only if it is natural-compatible.*

The following example illustrates natural compatibility and its application to natural revision.

Example 5 *The revision sequence in the previous example is natural-compatible. The first condition of natural compatibility is satisfied: $K_1 = a \wedge b$ implies $P_1 = b$; $K_2 = a \wedge b \wedge c$ implies $P_2 = c$; and $K_3 = \neg a$ implies $P_3 = \neg a$.*

The last preceding formula K_j consistent with $P_1 = b$ is $K_0 = a$, and indeed $K_1 = K_0 \wedge P_1 = a \wedge b$. The last preceding formula K_j consistent with $P_2 = c$ is $K_1 = a \wedge b$, and indeed $K_2 = K_1 \wedge P_2 = a \wedge b \wedge c$. Finally, $P_3 = \neg a$ is consistent with none of K_0, K_1, K_2. Therefore, the second condition of natural compatibility places no constraint on it.

Theorem 2 not only proves that the revision sequence is generated by natural revision from some preorder, but it also provides one: $C(0) = Mod(a)$, $C(1) = Mod(\neg a)$. This was indeed the preorder used in the previous example to generate the sequence.

Natural compatibility can be rewritten as a number of satisfiability and unsatisfiability tests. In particular, that j is the maximal index having the property can be written as: $K_j \wedge P_i$ is consistent and $K_h \wedge P_i$ is not, for $j < h < i$. This way, no satisfiability check is dependent on another, so the problem can be solved with two parallel calls to an NP oracle, one positive and one negative. The assumption that formulae in revision sequences are all consistent allows rewriting the first.

Lemma 7 *Checking the existence of a total preorder C generating a revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ using natural revision is in* coNP.

The problem is also hard for coNP. Therefore, it is coNP complete.

Theorem 4 *The problem of establishing the existence of a preorder generating a revision sequence using natural revision is* coNP *complete.*

4. Restrained Revision

Restrained revision (Booth & Meyer, 2006) has in common to natural revision that when revising a total preorder C by a formula P, the minimal models of P becomes the new class zero. In addition, every other class is split in two according to satisfaction of P: the models of P go in the lower class, the others in the higher.

Equivalently, every class i is refined (Papini, 2001) into classes $2i$ and $2i + 1$, where the first class contains the models of class i satisfying P and the second the models of class i not satisfying P; then, natural revision is applied.

Definition 5 *The restrained revision of the total preorder C by formula P is defined as the total preorder C_P that follows, where i is the minimal index such that $C(i) \cap Mod(P) \neq \emptyset$ and / denotes quotient (integer division, truncated):*

$$C_P(j) = \begin{cases} C(i) \cap Mod(P) & \text{if } j = 0 \\ (C((j-1)/2) \backslash C_P(0)) \cap Mod(P) & \text{if } j > 0 \text{ odd} \\ (C((j-1)/2) \backslash C_P(0)) \backslash Mod(P) & \text{otherwise} \end{cases}$$

For example, $C_P(0) = C(i) \cap Mod(P)$, $C_P(1) = (C((1-1)/2) \backslash C_P(0)) \cap Mod(P) = (C(0) \backslash C_P(0)) \cap Mod(P)$ and $C_P(2) = (C((2-1)/2) \backslash C_P(0)) \backslash Mod(P) = (C(0) \backslash C_P(0)) \backslash Mod(P)$ since $(2-1)/2 = 1/2 = 0$ using integer division.

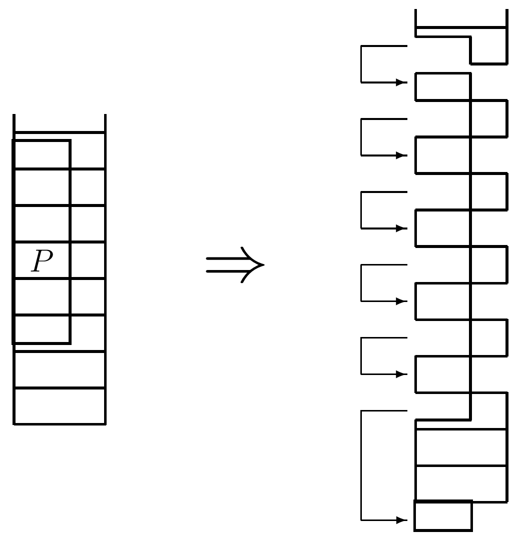

Figure 3: Restrained revision

A graphical example of the application of restrained revision to a total preorder is in Figure 3.

The revision sequence in Example 3 is not generated by any preorder using restrained revision. This will be proved using a necessary and sufficient condition to the existence of a preorder generating the sequence. For now, just to illustrate how restrained revision works, the preorder shown for natural revision is used.

Example 6 *Let C be the following total preorder:*

$$\begin{array}{rcl} C(0) & = & Mod(a) \\ C(1) & = & Mod(\neg a) \end{array}$$

Restrained revision by $P_1 = b$, $P_2 = c$ and $P_3 = \neg a$ generates a revision sequence different than that in Example 3. Since $K_0 = Mod(C(0))$, it follows $K_0 \equiv a$. Revising C by $P_1 = b$ makes the minimal models of P_1 to be the new class zero and splits every other class by $b/\neg b$. The resulting total preorder after removing the empty classes is therefore:

$$\begin{array}{rcl} C_{P_1}(0) & = & Mod(a \wedge b) \\ C_{P_1}(1) & = & Mod(a \wedge \neg b) \\ C_{P_1}(2) & = & Mod(\neg a \wedge b) \\ C_{P_1}(3) & = & Mod(\neg a \wedge \neg b) \end{array}$$

Since $Mod(K_1) = C_{P_1}(0)$, it follows that $K_1 \equiv a \wedge b$. A similar change happens when revising by $P_2 = c$:

$$
\begin{array}{lcl}
C_{P_1 P_2}(0) & = & Mod(a \wedge b \wedge c) \\
C_{P_1 P_2}(1) & = & Mod(a \wedge b \wedge \neg c) \\
C_{P_1 P_2}(2) & = & Mod(a \wedge \neg b \wedge c) \\
C_{P_1 P_2}(3) & = & Mod(a \wedge \neg b \wedge \neg c) \\
C_{P_1 P_2}(4) & = & Mod(\neg a \wedge b \wedge c) \\
C_{P_1 P_2}(5) & = & Mod(\neg a \wedge b \wedge \neg c) \\
C_{P_1 P_2}(6) & = & Mod(\neg a \wedge \neg b \wedge c) \\
C_{P_1 P_2}(7) & = & Mod(\neg a \wedge \neg b \wedge \neg c)
\end{array}
$$

Again, $Mod(K_2) = C_{P_1 P_2}(0)$, which implies $K_2 \equiv a \wedge b \wedge c$. The minimal models of $P_3 = \neg a$ are the whole class $C_{P_1 P_2}(4)$. The preorder therefore becomes:

$$
\begin{array}{lcl}
C_{P_1 P_2 P_3}(0) & = & Mod(\neg a \wedge b \wedge c) \\
C_{P_1 P_2 P_3}(1) & = & Mod(a \wedge b \wedge c) \\
C_{P_1 P_2 P_3}(2) & = & Mod(a \wedge b \wedge \neg c) \\
C_{P_1 P_2 P_3}(3) & = & Mod(a \wedge \neg b \wedge c) \\
C_{P_1 P_2 P_3}(4) & = & Mod(a \wedge \neg b \wedge \neg c) \\
C_{P_1 P_2 P_3}(5) & = & Mod(\neg a \wedge b \wedge \neg c) \\
C_{P_1 P_2 P_3}(6) & = & Mod(\neg a \wedge \neg b \wedge c) \\
C_{P_1 P_2 P_3}(7) & = & Mod(\neg a \wedge \neg b \wedge \neg c)
\end{array}
$$

As a result, $K_3 \equiv \neg a \wedge b \wedge c$. This revision sequence coincides with that of Example 3 up to P_3 but K_3 is different, as $K_3 = \neg a$ in the previous example. It will be shown that no preorder generates that sequence using restrained revision.

The following property is similar to Lemma 5 of natural revision, with the difference that a maxset is introduces to account for the class split.

Lemma 8 *Let $[K_j, P_{j+1}, \ldots, P_i, K_i]$ be a revision sequence generated by restrained revision from the initial total preorder C. If $K_j \wedge P_i$ is consistent while none of $K_{j+1} \wedge P_i$, $\ldots$, $K_{i-1} \wedge P_i$ is, then $C_{P_{j+1},\ldots,P_i}(0) = Mod(\mathrm{maxset}(K_j \wedge P_i, P_{j+1}, \ldots, P_{i-1}))$.*

This result is the first half of a necessary and sufficient condition for a sequence being generated by restrained revision from some preorder, which involves the following definition.

Definition 6 *A revision sequence $[K_0, P_1, K_1 \ldots, P_n, K_n]$ is* restrained-compatible *if, for every $i \in \{1, \ldots, n\}$, it holds:*

1. *$K_i \models P_i$;*
2. *$K_i \equiv \mathrm{maxset}(K_j \wedge P_i, P_{j+1}, \ldots, P_{i-i})$ if j is the maximal index such that $j < i$ and $K_j \wedge P_i$ is consistent, if any;*

3. *either $K_i \models P_l$ or $K_i \models \neg P_l$ for every $l < i$ if no such j exists.*

As a side remark, the condition that either $K_i \models P_l$ or $K_i \models \neg P_l$ in the third point of the definition refers to all indexes $l < i$, including the ones for which K_l is consistent with a previous K_j.

If a revision sequence is restrained-compatible it is generated by restrained revision from some initial preorder. The following lemma specifies which one that is.

Lemma 9 *If $[K_0, P_1, K_1, \ldots, P_n, K_n]$ is restrained-compatible then it is generated by restrained revision from the total preorder $C = [C(0), \ldots, C(n+1)]$, where:*

$$C(i) = \begin{cases} Mod(K_i) & \text{if } i \leq n \text{ and } \forall l < i \ . \ P_i \wedge K_l \models \bot \\ \emptyset & \text{otherwise, if } i \leq n \\ \bigcup\{Mod(K_j) \mid \exists l < j \ P_j \wedge K_l \not\models \bot\} & \text{if } i = n+1 \end{cases}$$

The results proved so far can be collected into an equivalent formulation of the existence of a preorder generating a sequence.

Theorem 5 *A revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ is generated by restrained revision from some initial total preorder if and only if it is restrained-compatible.*

The sequence in Example 3 can be shown not to be restrained-compatible. Therefore, not only it is not generated by restrained revision from the preorder that worked for natural revision, it is not generated from any other.

Example 7 *The revision sequence $[K_0, P_1, K_1, P_2, K_2, P_3, K_3]$ with $K_0 = a$, $P_1 = b$, $K_1 = a \wedge b$, $P_2 = c$, $K_2 = a \wedge b \wedge c$, $P_3 = \neg a$, $K_3 = \neg a$ is not restrained-compatible (this is the sequence of Example 3). Indeed, $P_3 = \neg a$ is inconsistent with each of K_0, K_1 and K_2, yet $K_3 = \neg a$ entails neither $P_1 = b$ nor $\neg P_1 = \neg b$, thereby violating the third condition of restrained compatibility.*

The definition of restrained compatibility involves consistency and entailment checks. The following lemma rewrites it in a form that can be shown to use only inconsistencies.

Lemma 10 *A revision sequence $[K_0, P_1, K_1 \ldots, P_n, K_n]$ is restrained-compatible if and only if, for every index i:*

1. *$K_i \models P_i$;*

2. *for every $0 \leq j < i$, either $K_j \wedge P_i$ is inconsistent or $K_i \models K_l \wedge P_i$ for some $j < l < i$ or $K_i \equiv \mathrm{maxset}(K_j \wedge P_i, P_{j+1}, \ldots, P_{i-1})$;*

3. *for every $0 \leq j < i$, $K_i \models P_j$, $K_i \models \neg P_j$, or $K_i \models K_l \wedge P_i$ for some $0 \leq l < i$.*

Since checking equivalence of a consistent formula and a maxset is a problem in coNP by Theorem 1, it can be expressed as a universally quantified formula. As a result, all quantifiers in the conditions of the lemma are universal, and the whole problem is in coNP. Hardness for the same class is easy to prove.

Theorem 6 *Establishing the existence of a total preorder generating a restrained revision sequence is* coNP*-complete.*

5. Intermezzo: Multiple Preorders

The same revision sequence may be generated by more than one total preorder. Examples exist even with two variables only, such as the sequence $[K_0, P_1, K_1, P_2, K_2]$ with $K_0 = a \wedge b$, $P_1 = a \wedge \neg b$, $K_1 = a \wedge \neg b$, $P_2 = \neg a \wedge b$ and $K_2 = \neg a \wedge b$.

$$a \wedge b \stackrel{a \wedge \neg b}{\longrightarrow} a \wedge \neg b \stackrel{\neg a \wedge b}{\longrightarrow} \neg a \wedge b$$

According to Lemma 9, the sequence is generated by restrained revision from the initial preorder $C = [C(0), C(1), C(2), C(3)]$.

$$\begin{array}{rcl} C(0) & = & Mod(a \wedge b) \\ C(1) & = & Mod(a \wedge \neg b) \\ C(2) & = & Mod(\neg a \wedge b) \\ C(3) & = & Mod(\neg a \wedge \neg b) \end{array}$$

However, this is not the only preorder generating the sequence. Classes $C(1)$ and $C(2)$ can be swapped, still leading to the same result.

$$\begin{array}{rcl} C(0) & = & Mod(a \wedge b) \\ C(1) & = & Mod(\neg a \wedge b) \\ C(2) & = & Mod(a \wedge \neg b) \\ C(4) & = & Mod(\neg a \wedge \neg b) \end{array}$$

Intuitively, revising by a single-model formula has always a single possible outcome: the formula itself. As a result, even if $Mod(K_1)$ is now contained in a class greater than $Mod(K_2)$, still the minimal models of P_1 are $Mod(K_1)$.

In a way, this is expected: when revision can only be performed in a single possible way, the initial preorder is irrelevant. This intuition holds for all considered revisions, and is confirmed by further simplifying the preorder: indeed, apart from $C(0)$, the other classes can be shuffled in every possible way, or even merged:

$$\begin{array}{rcl} C(0) & = & Mod(a \wedge b) \\ C(1) & = & Mod(\neg a \wedge b) \cup Mod(a \wedge \neg b) \cup Mod(\neg a \wedge \neg b) \end{array}$$

This kind of preorder works whenever each P_i has a single model, but not in general. Even restricting to the P_i's such that $P_i \wedge K_j \models \bot$ for all $j < i$, if $l < i$ then $P_i \wedge K_l$ is guaranteed to be inconsistent, but $P_l \wedge K_i$ is not. In other words, $Mod(K_i)$ cannot always be swapped or merged with $Mod(K_l)$.

When more than one preorder is possible, a sensible principle is to choose the least discriminating one. Such a preorder would compare I different from J only if strictly necessary to obtain the revision sequence. In other words, it does not carry plausibility information that does not follow from the revision sequence. Such a minimization is similar in spirit to the way rational closure in conditional logic is the most rationally rooted consequence relation (Lehmann & Magidor, 1992; Booth & Nittka, 2008). The last preorder shown obeys this principle, but the question of whether a least discriminating preorder exists for all revision sequences and all considered revision semantics is an open problem.

A different possible solution is to proceed by refutation: if a sequence $[K_0, P_1, K_2, \ldots, P_n, K_n]$ is generated by a revision from some preorder but $[K_0, P_1, K_2, \ldots, P_n, K_n, P_{n+1}, K_{n+1}]$ is not, then $\neg K_{n+1}$ might be considered true after revising by P_{n+1}.

As correctly pointed out by one of the reviewers, a revision sequence may even be generated by different revisions from different initial preorder, adding a second dimension to the problem: not only the preorder, but also the semantics of revision has to be chosen. The sequence shown in this example is one of this kind: it is generated by all four revision semantics considered in this article from every initial preorder C with $C(0) = Mod(a \wedge b)$.

6. Lexicographic Revision

In a seminal work on iterated revision, Spohn (1988) defined a tentative semantics based on the principle that newer formulae are more plausible than older ones at all levels of plausibility: even the most unlikely models of P are to be preferred over the most likely of $\neg P$. In spite of some apparent drawbacks pointed out by its author, this semantics was later recognized as a principled way to perform iterated revision (Nayak, 1994; Darwiche & Pearl, 1997; Booth & Meyer, 2006; Jin & Thielscher, 2007; Konieczny & Pino Pérez, 2000).

Like the other revisions used in this article, lexicographic revision works on a total preorder of plausibility of models C. In particular, revision by P changes it by moving all models of P to classes of index lower than the others.

Definition 7 *The lexicographic revision of a total preorder C by a formula P is defined as the following total preorder, where i and j are respectively the indexes of the lowest and highest classes containing models of P:*

$$C_P(k) = \begin{cases} C(k+i) \cap Mod(P) & \text{if } k \leq j - i \\ C(k-j+i-1) \backslash Mod(P) & \text{otherwise} \end{cases}$$

The new class zero is $C_P(0) = C(0+i) \cap Mod(P)$; as expected, it comprises the minimal models of P, since $C(i)$ is the lowest class containing models of P. The class $C_P(j-i) = C(j-i+i) \cap Mod(P) = C(j) \cap Mod(P)$ contains the highest-class models of P, since by assumption these are in $C(j)$. The models of $C(0)$ that satisfy $\neg P$, if any, are moved to class $C_P(j-i+1) = C(j-i+1-j+i-1) \backslash Mod(P) = C(0) \backslash Mod(P)$. The index is $j-i+1$ because the lower classes $C_P(0), \ldots, C_P(j-i)$ contain the models of P coming from $C(i), \ldots, C(j)$.

Figure 4 shows how a total preorder C is changed by a formula P using lexicographic revision.

Graphically, lexicographic revision "cuts out" the models of P from their classes and wedges them under the shelf. This way, every model of P belongs to a lower class than all models of $\neg P$. At the same time, the relative position of two models of P is not changed, and the same holds for every two models of $\neg P$.

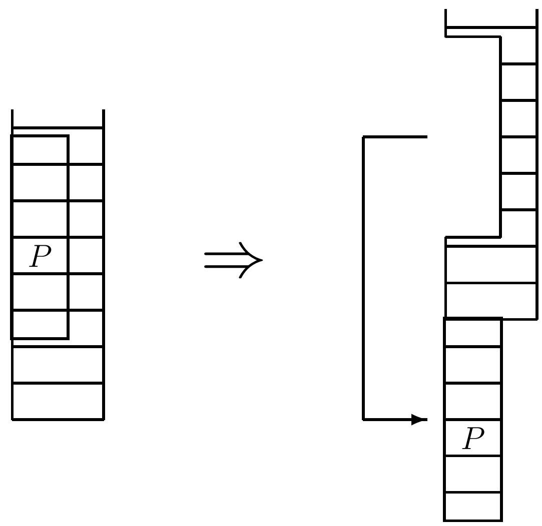

Figure 4: Lexicographic revision

Example 8 *It will be shown that the sequence in Example 3 is not generated by lexicographic revision from any total preorder. Meanwhile, to illustrate the definition of lexicographic revision the preorder used in the example for natural revision is revised by $P_1 = b$, $P_2 = c$ and $P_3 = \neg a$.*

$$\begin{array}{rcl} C(0) & = & Mod(a) \\ C(1) & = & Mod(\neg a) \end{array}$$

Revising C by $P_1 = b$ using lexicographic revision removes all models of P_1 from their classes and creates new classes for them at the bottom:

$$\begin{array}{rcl} C_{P_1}(0) & = & Mod(a \wedge b) \\ C_{P_1}(1) & = & Mod(\neg a \wedge b) \\ C_{P_1}(2) & = & Mod(a \wedge \neg b) \\ C_{P_1}(3) & = & Mod(\neg a \wedge \neg b) \end{array}$$

As a result, $K_1 = a \wedge b$. Revising by $P_2 = c$ has a similar effect:

$$\begin{array}{rcl} C_{P_1P_2}(0) & = & Mod(a \wedge b \wedge c) \\ C_{P_1P_2}(1) & = & Mod(\neg a \wedge b \wedge c) \\ C_{P_1P_2}(2) & = & Mod(a \wedge \neg b \wedge c) \\ C_{P_1P_2}(3) & = & Mod(\neg a \wedge \neg b \wedge c) \\ C_{P_1P_2}(4) & = & Mod(a \wedge b \wedge \neg c) \end{array}$$

$$
\begin{array}{rcl}
C_{P_1P_2}(5) & = & Mod(\neg a \wedge b \wedge \neg c) \\
C_{P_1P_2}(6) & = & Mod(a \wedge \neg b \wedge \neg c) \\
C_{P_1P_2}(7) & = & Mod(\neg a \wedge \neg b \wedge \neg c)
\end{array}
$$

This preorder produces $K_2 = a \wedge b \wedge c$. Finally, revising by $P_3 = \neg a$ and removing the empty classes makes the classes $1, 3, 5, 7$ to become the new classes $0, 1, 2, 3$.

$$
\begin{array}{rcl}
C_{P_1P_2P_3}(0) & = & Mod(\neg a \wedge b \wedge c) \\
C_{P_1P_2P_3}(1) & = & Mod(\neg a \wedge \neg b \wedge c) \\
C_{P_1P_2P_3}(2) & = & Mod(\neg a \wedge b \wedge \neg c) \\
C_{P_1P_2P_3}(3) & = & Mod(\neg a \wedge \neg b \wedge \neg c) \\
C_{P_1P_2P_3}(4) & = & Mod(a \wedge b \wedge c) \\
C_{P_1P_2P_3}(5) & = & Mod(a \wedge \neg b \wedge c) \\
C_{P_1P_2P_3}(6) & = & Mod(a \wedge b \wedge \neg c) \\
C_{P_1P_2P_3}(7) & = & Mod(a \wedge \neg b \wedge \neg c)
\end{array}
$$

Since $K_3 = \neg a \wedge b \wedge c$ is not equivalent to $\neg a$, the revision sequence in Example 3 is not generated by lexicographic revision from the initial preorder C. It will be shown that no preorder at all generates that sequence using lexicographic revision.

For every preorder C and consistent formula P, the revised preorder C_P using lexicographic revision has three properties:

1. $C_P(0)$ is the set of minimal models of P in the preorder C;
2. there exists an index h such $\bigcup_{i=0,\dots,h} C_P(i) = Mod(P)$;
3. if two models both satisfy P or both falsify it, C_P compares them as C does.

Lexicographic revision can be recast in terms of the maxsets of the reversed sequence. Booth and Nittka (2008) proved the following property; more precisely, what they proved implies the following property because an arbitrary preorder $[C(0), \dots, C(n)]$ is obtained by revising the ordering with all models in class zero with a sequence of formulae having $C(n), \dots, C(0)$ as their sets of models.

Property 1 *If $[K_0, P_1, \dots, K_n, P_n]$ is a revision sequence generated by lexicographic revision from the total preorder C, then $Mod(K_i)$ is the set of minimal models of* $\mathrm{maxset}(P_i, \dots, P_1)$ *according to C.*

This property has the following consequences:

1. all models of K_i are in the same class of C;

2. all other models of $\mathrm{maxset}(P_i, \ldots, P_1)$ are in greater classes.

Two further properties follow. First, if $K_i \wedge K_j$ is consistent then all models in $Mod(K_i) \cup Mod(K_j)$ are in the same class of C. Indeed, since $K_i \wedge K_j$ is consistent, it has a model I; since all models of K_i are in the same class of I, and the same for K_j, all models of these two formulae are in the same class. Second, if K_j has some models in common with $\neg K_i \wedge \mathrm{maxset}(P_i, \ldots, P_1)$, then all its models are in the same class of C, greater than that of the models of K_i.

This allows shifting from a total preorder among models to a total preorder among formulae K_i. Since this is a preorder over the formulae K_i, which are the results of the revision process, it is called a result preorder.

Definition 8 *A* result preorder *for the revision sequence* $[K_0, P_1, \ldots, P_n, K_n]$ *is a total preorder among its formulae* K_i *such that:*

1. *if* $K_i \wedge K_j$ *is consistent then* K_i *and* K_j *are in the same class;*
2. *if* $\neg K_i \wedge K_j \wedge \mathrm{maxset}(P_i, \ldots, P_1)$ *is consistent then* K_i *is in a lower class than* K_j.

The advantage of result preorders is that they can be built from the revision sequence, as it will be shown. Before, it is proved that the existence of a result preorder is the same as the existence of an initial total preorder over models generating the sequence by lexicographic revision. Since two preorders are involved (one among models, one among formulae), a distinction is made between "preorder among models C" and "preorder among formulae R"; result preorders are of the second kind.

Lemma 11 *If a revision sequence* $[K_0, P_1, K_1, \ldots, P_n, K_n]$ *is generated by lexicographic revision from the total preorder among models* C *then it has the result preorder* R *defined by:*

$$R(i) = \{K_i \mid Mod(K_i) \subseteq C(i)\}$$

The converse also holds: from a result preorder for a revision sequence one can derive a preorder among models that generates the sequence using lexicographic revision.

Lemma 12 *If* R *is a result preorder for a revision sequence* $[K_0, P_1, K_1, \ldots, P_n, K_n]$ *such that* $K_i \models \mathrm{maxset}(P_i, \ldots, P_1)$ *for every* i*, then lexicographic revision from the following total preorder* C *among models generates the revision sequence, where* z *is the index of the greatest class of* R*:*

$$C(h) = \begin{cases} \bigcup\{Mod(K_i) \mid K_i \in R(h)\} & \text{if } h \leq z \\ \{I \mid I \not\in K_1 \cup \cdots \cup K_n\} & \text{if } h = z+1 \end{cases}$$

The two lemmas together imply the following corollary.

Corollary 2 *A revision sequence* $[K_0, P_1, \ldots, P_n, K_n]$ *is generated by lexicographic revision from some initial preorder among models if and only if a result preorder for the sequence exists and* $K_i \models \mathrm{maxset}(P_i, \ldots, P_1)$ *for every* i.

The condition of existence of a result preorder is already simpler than that of existence of an initial total preorder over models, since the number of formulae K_i is linear in the size of the revision sequence while the number of models may be exponential. The existence of a result preorder can be further recast in terms of a condition over the revision sequence.

Definition 9 *A revision sequence $[K_0, P_1, \ldots, K_n, P_n]$ is* lexicographic compatible *if $K_i \models \text{maxset}(P_i, \ldots, P_1)$ for every i and the relations $\simeq$ and $<$ defined as follows do not form cycles containing some $<$ links:*

- *$K_i \simeq K_j$ if and only if $K_i \wedge K_j$ is consistent;*
- *$K_i < K_j$ if and only if $\neg K_i \wedge K_j \wedge \text{maxset}(P_i, \ldots, P_1)$ is consistent.*

The rationale of the definition is that $K_i \simeq K_j$ is the same as K_i and K_j being in the same class of R, and $K_i < K_j$ is the same as K_i being in a lower class. This intuition is confirmed by the following lemma.

Lemma 13 *A revision sequence is generated by lexicographic revision from some total preorder if and only if it is lexicographic compatible.*

This result can be applied to the running example.

Example 9 *The revision sequence presented in Example 3 is not lexicographic compatible. Indeed, $\text{maxset}(P_3, P_2, P_1) = \text{maxset}(\neg a, c, b) = \neg a \wedge c \wedge b$, which is not entailed by $K_3 = \neg a$. As a result, the revision sequence is not generated by lexicographic revision from any preorder.*

Given a revision sequence, one can determine the consistency of $K_i \wedge K_j$ and $\neg K_i \wedge K_j \wedge \text{maxset}(P_i, \ldots, P_1)$ for every pair of formulae K_i and K_j. The problem of non-existence of a preorder generating the sequence is then turned into that of existence of cycles, which is computationally easy (polynomial in the size of the revision sequence). The hard part is checking consistency. Since the problem is polynomial if an NP oracle is available (which turns all consistency checks into constant-time operations), the problem is in Δ_2^p. However, it can be proved to be even computationally easier than that.

Lemma 14 *A revision sequence $[K_0, P_1, K_1 \ldots, P_n, K_n]$ is not lexicographic compatible if and only if either $K_i \not\models \text{maxset}(P_i, \ldots, P_1)$ for some i or consistent sets $R_1, \ldots, R_n$ exist such that:*

1. *$\{P_j \mid 1 \leq j \leq i \text{ and } K_i \models P_j\} \subseteq R_i$ for every i; and*
2. *there exists a cycle $K_{i_1}, \ldots, K_{i_m} = K_{i_1}$ such that either $K_{i_j} \wedge K_{i_{j+1}}$ or $\neg K_{i_j} \wedge K_{i_{j+1}} \wedge R_i$ is consistent for all $i_j \in \{i_1, \ldots, i_{m-1}\}$, and the second is consistent for at least one such index.*

The advantage of this reformulation of lexicographic incompatibility is that all entailment tests it contains can be reformulated in terms of consistency. This means that incompatibility is in NP. Therefore, compatibility is in coNP. It can also be shown hard for the same class.

Theorem 7 *The problem of checking the existence of a total preorder generating a revision sequence using lexicographic revision is* coNP*-complete.*

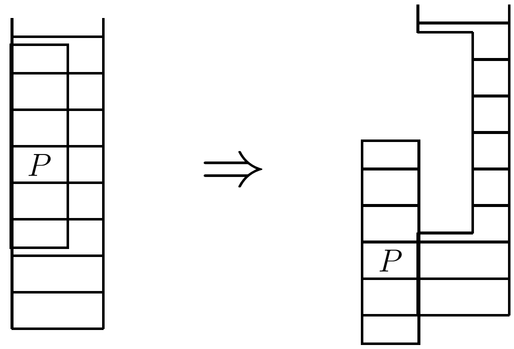

Figure 5: Reinforcement revision

7. Reinforcement Revision

Reinforcement revision (Jin & Thielscher, 2007) takes as input not only a total preorder to revise C and a revising formula P, but also a parameter m that encodes the degree of belief in P (more precisely, m is the degree of disbelief in $\neg P$). For the sake of simplicity, the restriction to the the case $m = 1$ is analyzed.

Definition 10 *The reinforcement revision of the total preorder C by formula P and parameter $m = 1$ is the following total preorder C_P, where i is the minimal index such that $C(i) \cap Mod(P) \neq \emptyset$.*

$$C_P(j) = \begin{cases} C(i) \cap Mod(P) & \text{if } j = 0 \\ C(j-1) \backslash Mod(P) \cup C(j+i) \cap Mod(P) & \text{if } j > 0 \end{cases}$$

The general definition has $j - m$ instead of $j - 1$; in this article, m is always 1. The two cases can be merged into the single one $C_P(j) = C(j-1)\backslash Mod(P) \cup C(j+i) \cap Mod(P)$ for all $j \geq 0$ by assuming that $C(-1) = \emptyset$. For example, $C_P(0) = C(-1)\backslash Mod(P) \cup C(i) \cap Mod(P) = C(i) \cap Mod(P)$, while $C_P(1) = C(0)\backslash Mod(P) \cup C(i+1) \cap Mod(P)$.

A graphical example of a revision is in Figure 5.

The behavior of this revision is shown on the preorder and formulae of Example 3.

Example 10 *Revising the preorder $C = [C(0), C(1)]$ where $C(0) = Mod(a)$ and $C(1) = Mod(\neg a)$ by $P_1 = b$ using reinforcement revision has the effect of increasing the class of every model of $\neg P_1$ by one; the models of P_1 do not decrease of class since some of them are already in class zero, which means that $i = 0$ in the definition of C_P.*

$$\begin{aligned} C_{P_1}(0) &= Mod(a \wedge b) \\ C_{P_1}(1) &= Mod((a \wedge \neg b) \vee (\neg a \wedge b)) \\ C_{P_1}(2) &= Mod(\neg a \wedge \neg b) \end{aligned}$$

The same happens when revising by $P_2 = c$: every class is the union of the original class conjoined with c and of the previous class conjoined with $\neg c$:

$$\begin{array}{rcl} C_{P_1P_2}(0) & = & \mathit{Mod}(a \wedge b \wedge c) \\ C_{P_1P_2}(1) & = & \mathit{Mod}((a \wedge b \wedge \neg c) \vee (a \wedge \neg b \wedge c) \vee (\neg a \wedge b \wedge c)) \\ C_{P_1P_2}(2) & = & \mathit{Mod}((a \wedge \neg b \wedge \neg c) \vee (\neg a \wedge b \wedge \neg c) \vee (\neg a \wedge \neg b \wedge c)) \\ C_{P_1P_2}(3) & = & \mathit{Mod}(\neg a \wedge \neg b \wedge \neg c) \end{array}$$

The minimal class containing models of $P_3 = \neg a$ is $C_{P_1P_2}(1)$. Therefore, models of $\neg a$ are decreased of one class. Models of $\neg P_3 = a$ are increased of one class, as usual:

$$\begin{array}{rcl} C_{P_1P_2P_3}(0) & = & \mathit{Mod}((\neg a \wedge b \wedge c)) \\ C_{P_1P_2P_3}(1) & = & \mathit{Mod}((a \wedge b \wedge c) \vee (\neg a \wedge \neg b \wedge c) \vee (\neg a \wedge b \wedge \neg c)) \\ C_{P_1P_2P_3}(2) & = & \mathit{Mod}((\neg a \wedge \neg b \wedge \neg c) \vee (a \wedge b \wedge \neg c) \vee (a \wedge \neg b \wedge c)) \\ C_{P_1P_2P_3}(3) & = & \mathit{Mod}(a \wedge \neg b \wedge \neg c) \end{array}$$

As a result, $K_3 = \neg a \wedge b \wedge c$.

As for the other revision semantics, every K_i is assumed consistent; the case of inconsistent K_i is degenerated, as no preorder produces such a knowledge base.

Given a fixed revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ generated by reinforcement revision from the initial total preorder C, the lowest class index of the models of P_i in the ordering $C_{P_1,\ldots,P_i}$ is denoted $D_C(i)$. Formally:

$$D_C(i) = \min\{j \mid C_{P_1,\ldots,P_{i-1}}(j) \cap \mathit{Mod}(P_i) \neq \emptyset\}$$

Revising by P_i shifts down all models of P_i of $D_C(i)$ classes and raises all other models of one class. As a result, $D_C(1), \ldots, D_C(n)$ tell how a model is moved up or down at every step, allowing to determine how a model is moved between step i and step j.

Definition 11 *The* movement *of I from i to j according to D_C in the revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ is $M_{D_C}(I, i, j)$ where:*

- $M_{D_C}(I, i, i) = 0$*;*
- $M_{D_C}(I, i, i+1) = -D_C(i+1)$ *if* $I \models P_{i+1}$ *and* $M_{D_C}(I, i, i+1) = 1$ *otherwise;*
- *if* $j > i$ *then* $M_{D_C}(I, i, j) = \sum_{l=i,\ldots,j-1} M_{D_C}(I, l, l+1)$*;*
- *if* $j < i$ *then* $M_{D_C}(I, i, j) = -M_{D_C}(I, j, i)$*.*

Since $D_C(i)$ is the minimal class of models of P_i at step $i-1$ according to an initial preorder, $M_{D_C}(I, i, j)$ is the change of classes between i and j when using the same preorder. The initial preorder C affects the definition of $M_{D_C}(I, i, j)$ only indirectly, via the sequence $D_C = [D_C(1), \ldots, D_C(n)]$. As a result, $M_V(I, i, j)$ can be defined from an arbitrary sequence of n numbers $V = [V(1), \ldots, V(n)]$.

Lemma 15 *For every sequence of n numbers $V = [V(1), \ldots, V(n)]$, it holds $M_V(I, i, j) = M_V(I, i, h) + M_V(I, h, j)$ for every three indexes i, j and h.*

This lemma holds even if h is not between i and j.

Since $M_{D_C}(I, i, j)$ has been defined so that it is the change of class of model I from step i to step j, in the particular case where $I \models K_i$ it is the class of I at step j, since that at step i is zero.

Lemma 16 *If a revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ is generated by reinforcement revision from the total preorder C and $D_C = [D_C(1), \ldots, D_C(n)]$ with $D_C(i) = \min\{j \mid C_{P_1,\ldots,P_{i-1}}(j) \cap Mod(P_i) \neq \emptyset\}$ and $I \models K_i$ then, for every j:*

- *$M_{D_C}(I, i, j) = 0$ if $I \models K_j$;*
- *$M_{D_C}(I, i, j) > 0$ otherwise.*

This lemma can be reversed, in the sense that a sequence of values having this property allows to determine a preorder generating the sequence.

Definition 12 *A sequence of nonnegative integer values $V = [V(1), \ldots, V(n)]$ is a* reinforcement mover *of the revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ if, for every i and j, if $I \models K_i$ then:*

- *$M_V(I, i, j) = 0$ if $I \models K_j$;*
- *$M_V(I, i, j) > 0$ otherwise.*

The previous lemma can therefore be recast as: if a sequence is generated by reinforcement revision from the total preorder C, then it has a reinforcement mover: D_C. The converse also holds: from a reinforcement mover one can determine a total preorder generating the revision sequence.

Lemma 17 *If $V = [V(1), \ldots, V(n)]$ is a reinforcement mover for the revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$, the following initial preorder $C = [C(0), \ldots, C(V(1) + \cdots + V(n+1))]$ generates the revision sequence by reinforcement revision and $D_C = V$.*

$$C(j) = \begin{cases} \{I \mid I \models K_i \text{ and } M_V(I, i, 0) = j\} & \text{if } j < V(1) + \ldots + V(n) + 1 \\ \{I \mid \forall i \;.\; I \not\models K_i\} & \text{if } j = V(1) + \ldots + V(n) + 1 \end{cases}$$

In contrast to the condition of compatibility for the other revision semantics, this one does not explicitly require $K_i \models P_i$. It is however implied: if $I \not\models P_i$ then $M_V(I, i-1, i) = 1$ by the definition of movement (Definition 11) and if $I \models K_i$ then $M_V(I, i, j) \geq 0$ for every j by definition of reinforcement mover (Definition 12). But in the particular case of $j = i - 1$ this is $M_V(I, i, j) = M_V(I, i, i-1) = -M_V(I, i-1, i) = -1$, which is not greater than or equal to zero. Therefore, no reinforcement mover exists if $K_i \not\models P_i$.

This lemma allows checking the existence of a preorder generating the sequence by guessing $V(1), \ldots, V(n)$ and then checking the class of every model I that satisfies at least a K_i. However, membership to the polynomial hierarchy follows only if all such values $V(i)$ are representable in polynomial space, which is only possible if their values are bounded by some exponential in the size of the sequence.

Lemma 18 *If reinforcement revision generates a revision sequence* $[K_0, P_1, K_1, \ldots, P_n, K_n]$ *from some total preorder, then it also generates the same sequence from a preorder in which the minimal initial class of models of* P_1 *is* 0 *or* 1*.*

This lemma proves that if a sequence is generated by some preorder is also generated by a preorder C such that $D_C(1)$ is either 0 or 1. In particular, if $K_0 \wedge P_1$ is consistent then $D_C(1) = 0$, otherwise $D_C(1) = 1$. This is the base case of a recursive proof giving a bound on the size of the other $D_C(i)$. Intuitively, this is done by lowering all models of P_i of the same number of classes in the initial preorder; after revising by P_i they will moved together so that their minimal ones are in class zero, resulting in the same ordering as obtained from the original one. This lowering cannot however be so large that some of these models enter class zero at a previous step j if they do not satisfy K_j.

Lemma 19 *If reinforcement revision generates the sequence* $[K_0, P_1, K_1, \ldots, P_n, K_n]$ *from a total preorder, it is also generates the same sequence from a total preorder* C *such that* $D_C(i+1) \leq D_C(1) + \ldots + D_C(i) + i + 1$*.*

This lemma can be seen as the inductive part of a proof, where the previous one was the base case. They lead to the following conclusion.

Lemma 20 *Reinforcement revision generates the revision sequence* $[K_0, P_1, K_1, \ldots, P_n, K_n]$ *from a total preorder if and only if it generates the same sequence from a total preorder such that each* $D_C(i)$ *is bounded by* $2^i - 1$*.*

The values of $D_C(1), \ldots, D_C(n)$ are bounded by an exponential in the value of n, which is also a lower bound on the size of the sequence $[K_0, P_1, \ldots, P_n, K_n]$. Therefore, each $D_C(i)$ can be represented in space polynomial in the size of the sequence.

Theorem 8 *Establishing the existence of a total preorder generating the revision sequence* $[K_0, P_1, \ldots, P_n, K_n]$ *by reinforcement revision is in* Σ_2^p*, and is in* coNP *if* n *is a constant.*

For the case of constant-length sequences, hardness is easy to prove.

Theorem 9 *Checking the existence of a preorder generating a reinforcement revision sequence* $[K_0, P_1, K_1, \ldots, P_n, K_n]$ *is* coNP*-complete, if* n *is a constant.*

8. Conclusions

Belief revision employs plausibility orders to revise a knowledge base, but how to obtain such orders has largely been neglected. The solution proposed and studied in this article is to assume knowledge of previous revisions, "reversing" them to obtain the initial order, which can be then used in further revisions. This method is similar to deriving the order from hypothetical revisions like $K * P \models Q$, $K * R \models T$, etc., which assumes knowledge of what would happen if the revising formula is P, if it is R, etc. In the approach considered in this article the results of iteratively incorporating a series of new formulae are given, like in nested counterfactuals (Eiter & Gottlob, 1996). Booth and Nittka (2008) considered the problem of deriving facts holding at some time points from partial information expressed in terms of positive and negative conditions (i.e., some formulae hold and some others do not hold at certain time points), and this is done by constructing an initial ordering. Their study focuses on lexicographic revision only, but allows for partial knowledge of both revisions and results, and includes the choice of an initial ordering among the possible ones. Rather than entailment information like $K * P * R \models Q$, in this article the history of revisions and resulting knowledge bases is assumed fully known. These could be the actual results of manually performed changes, and may already be available.

The analysis has shown simple equivalent conditions to the existence of an ordering generating a given series of revisions and results for natural (Boutilier, 1996), restrained (Booth & Meyer, 2006), lexicographic (Spohn, 1988; Nayak, 1994) and reinforcement revisions (Jin & Thielscher, 2007). The conditions allow to construct such an initial ordering if one exists.

Using these equivalent conditions, the complexity of establishing the existence of orderings generating a sequence has been established for the considered semantics. Surprisingly, they turned out to be relatively simple: coNP complete in all cases but one (reinforcement revision with unbounded sequence length). This is the same as checking propositional entailment, which means that checking generability does not increase complexity over that of the base language of propositional logic.

As shown in Section 5, the same revision sequence may be generated by more than one ordering. This leads to the question of whether one of them may be considered the “most natural” for the sequence. If for example I is less than J in an ordering but not in another, the second may be seen as more cautious, and in the end more rational: the sequence can be generated without assuming that I represents a more plausible world than J, so there is no reason to draw such a conclusion. The question of whether a single least informative ordering exists for every sequence, in each of the considered revision semantics, is an open problem.

Still open is the comparison of generable sequences in the various revisions: as shown by the running example, a sequence that is generated by natural revision is generated by no ordering with the restrained and lexicographic revision. Is there any sequence with the opposite property? If not, natural revision may be seen as more suited at explaining why a revision sequence has been generated. On the other hand, it may also give less information about the initial ordering used to generate them; this is the case if the orderings generating the same sequence are more numerous than in the other semantics.

The four semantics for iterated belief revisions are not the only ones defined in the literature (Williams, 1994; Darwiche & Pearl, 1997; Areces & Becher, 2001; Benferhat et al., 2004; Konieczny & Pino Pérez, 2000; Zhang, 2004); a recent survey counted at least twenty seven revision-related operators (Rott, 2009). Natural and lexicographic semantics are regarded as extreme forms of revision satisfying the Darwiche-Pearl postulates (Darwiche & Pearl, 1997), where minimal and maximal hearing is respectively given to the new information. Restrained and reinforcement revision can be considered to be in the middle, as they also obey some other conditions (Booth & Meyer, 2006; Jin & Thielscher, 2007). The four considered semantics therefore constitute a reasonable spectrum of possibilities, but others exist.

Some require additional information (like the strength of every revision, Spohn, 1988; Williams, 1994; Benferhat et al., 2004), others are families of revisions rather than single ones (Darwiche & Pearl, 1997; Zhang, 2004). These, in particular, open an interesting line of research: whether a revision sequence can be generated by some ordering and some revision semantics satisfying a given set of conditions, like the Darwiche and Pearl (1997) postulates. In other words, both the ordering and the semantics are the solution of the problem, while the given is only the revision sequence.

The assumption of reliability strictly increasing with time also gives other directions to study. Indeed, this principle has been realized as not true in general (Peppas, 2008) and iterated revision recognized as a form of prioritized merging (Delgrande, Dubois, & Lang, 2006). In this perspective, giving preference to the last formula is just a particular case. It is a case of interest, yet searching for the initial plausibility order is also possible in the general case.

Yet another open problem is how to combine the approach in this article with results about how people actually perform revision. Indeed, it has been experimentally proved that human revision suffers from a number of biases (Tversky & Kahneman, 1983; See, Morrison, Rothman, & Soll, 2011; Wang, Zhang, & Johnson, 2000), such as the anchoring or order effect, the excessive preference of knowledge acquired early. These studies show that revision performed by people is not fully rational, contrary to what belief revision formal semantics attempt to be. Such psychological, extra-logical biases should be kept into account when working on sequences of manually-performed revisions.

8.1 Acknowledgements

The author thanks the anonymous referees for their useful suggestions on the previous versions of this article.

Appendix A. Proofs

The following sections contain the proofs of all lemmas and theorems in the article.

A.1 Preliminaries: Proofs

Lemma 1 *If F is consistent and $F \models \text{maxset}(P_1, \ldots, P_n)$, then $\text{maxset}(P_1, \ldots, P_n) \equiv \{P_i \mid 1 \leq i \leq n \text{ and } F \models P_i\}$.*

Proof. Since $F \models \text{maxset}(P_1, \ldots, P_n)$ then F implies every P_i that is in the maxset. The only possibility for the claim not to hold is that F also implies some P_i not in the maxset. Since every element of $\text{maxset}(P_1, \ldots, P_n)$ is implied by F, also every element of $\text{prefix}_i(\text{maxset}(P_1, \ldots, P_n))$ is. Since F also entails P_i, every model of F satisfies both $\text{prefix}_i(\text{maxset}(P_1, \ldots, P_n))$ and P_i. The consistency of $P_i \wedge \text{prefix}_i(\text{maxset}(P_1, \ldots, P_n))$ contradicts the assumption that P_i is not in the maxset. □

Lemma 2 *If F is consistent and $F \models \text{maxset}(P_1, \ldots, P_n)$, every consistent subset of $\{P_1, \ldots, P_n\}$ that contains all formulae entailed by F is equivalent to $\text{maxset}(P_1, \ldots, P_n)$.*

Proof. By Lemma 1, since F is consistent and $F \models \text{maxset}(P_1, \ldots, P_n)$ then $\text{maxset}(P_1, \ldots, P_n) \equiv \{P_i \mid 1 \leq i \leq n \text{ and } F \models P_i\}$. This proves that $\text{maxset}(P_1, \ldots, P_n)$ contains all formulae entailed by F. Let R be a consistent proper superset of

$\text{maxset}(P_1, \ldots, P_n)$. Let P_i be the lowest-index formula that is in R but not in $\text{maxset}(P_1, \ldots, P_n)$. Since this is the lowest index, R and $\text{maxset}(P_1, \ldots, P_n)$ have the same formulae among $\{P_1, \ldots, P_{i-1}\}$. Since R is consistent, its intersection with this set is consistent as well. Since R also contains P_i, it follows that $\text{prefix}_i(P_1, \ldots, P_n) \cup \{P_i\}$ is consistent, contradicting the assumption that P_i is not in $\text{maxset}(P_1, \ldots, P_n)$. $\square$

Lemma 3 *If F is consistent then $F \not\models \text{maxset}(P_1, \ldots, P_n)$ if and only if there exists a subsequence R of $[P_1, \ldots, P_n]$ such that:*

1. *R is consistent;*
2. *if $F \models P_i$ then $P_i \in R$;*
3. *for some i, $P_i \not\in R$ and $P_i \wedge \text{prefix}_i(R)$ is consistent.*

Proof. Two cases are considered: first, F entails the maxset; second, F does not. In the first, no such R is proved to exists; in the second, one such R is shown.

If $F \models \text{maxset}(P_1, \ldots, P_n)$, then F implies all elements of the maxset. As a result, the second condition can only be true if R contains all formulae in the maxset. If R also contains some formulae not in the maxset, since a formula not in the maxset is inconsistent with the maxset, then R is inconsistent. As a result, R can be consistent only if it coincides with the maxset. This contradicts the third point, showing that for every R, if the first and the second conditions are true the third is false.

If $F \not\models \text{maxset}(P_1, \ldots, P_n)$, the three conditions are satisfied by R containing precisely the formulae P_i entailed by F. This choice meets the first condition because F is consistent and the second by construction. The third condition is now proved to hold as well.

Since F does not entail the maxset then F does not entail some formulae in the maxset. Let i be the least index of a formula in $\text{maxset}(P_1, \ldots, P_n)$ that is not entailed by F. By construction, every formula P_j in the maxset is in R if $j < i$. It can be shown that the converse also holds. To the contrary, let $j < i$ be the lowest index of a formula in R that is not in the maxset.

The assumptions that i and j are the least indexes on which R and the maxset differ (in a way or the other) imply $\text{prefix}_j(R) = \text{prefix}_j(\text{maxset}(P_1, \ldots, P_n))$. Since by assumption P_j is not in the maxset, it is inconsistent with $\text{prefix}_j(\text{maxset}(P_1, \ldots, P_n))$. As a result, it is also inconsistent with $\text{prefix}_j(R)$. Since $P_j \in R$, that implies the inconsistency of R, which was proved consistent.

This contradiction proves that R and the maxset are not only equal up to some index $j < i$, but up to i. In other words, $\text{prefix}_i(R) = \text{prefix}_i(\text{maxset}(P_1, \ldots, P_n))$.

By assumption P_i is a formula in the maxset that is not in R. Being in the maxset means that $P_i \wedge \text{prefix}_i(\text{maxset}(P_1, \ldots, P_n))$ is consistent. As it has just been proved, the latter is equivalent to $P_i \wedge \text{prefix}_i(R)$. This concludes the proof of the third condition. $\square$

Theorem 1 *If F is consistent, checking $F \equiv \text{maxset}(P_1, \ldots, P_n)$ is in* coNP*.*

Proof. $F \equiv \text{maxset}(P_1, \ldots, P_n)$ holds if $F \models \text{maxset}(P_1, \ldots, P_n)$ and $\text{maxset}(P_1, \ldots, P_n) \models F$. If the first condition is true, by Lemma 1 the maxset comprises exactly the formulae entailed by F. Therefore, the converse $\text{maxset}(P_1, \ldots, P_n) \not\models F$ can only happen if there exists a model I satisfying all formulae P_i entailed by F but not F itself. In other words:

$$\begin{array}{ll} F \equiv \text{maxset}(P_1, \ldots, P_n) & \\ \quad \text{iff} \quad F \models \text{maxset}(P_1, \ldots, P_n) \text{ and } \text{maxset}(P_1, \ldots, P_n) \models F \\ \quad \text{iff} \quad F \models \text{maxset}(P_1, \ldots, P_n) \text{ and } \neg(\exists I \;.\; I \not\models F \text{ and } \forall i(F \models P_i \rightarrow I \models P_i)) \\ \quad \text{iff} \quad F \models \text{maxset}(P_1, \ldots, P_n) \text{ and } \neg(\exists I \;.\; I \not\models F \text{ and } \forall i(F \not\models P_i \text{ or } I \models P_i)) \\ \quad \text{iff} \quad F \models \text{maxset}(P_1, \ldots, P_n) \text{ and } \forall I \;.\; I \models F \text{ or } \exists i(F \models P_i \text{ and } I \not\models P_i)) \end{array}$$

Lemma 3 reformulates the converse of the first condition $F \models \text{maxset}(P_1, \ldots, P_n)$ using only existential quantifiers (the second point is equivalent to either $F \not\models P_i$ or $P_i \in R$). As a result, the first condition can be expressed using only universal quantifiers. The only existential quantifier in the second condition is that over i, which can be replaced by a disjunction. Since all quantifiers are universal, the problem is in coNP. □

A.2 Natural Revision: Proofs

Lemma 4 *If $C_P(0) \cap Mod(F) = \emptyset$ then C_P compares the models of F as C does, where C_P is the natural revision of the total preorder C with formula P.*

Proof. Let I and J be two models of F, m and l their classes. Since these are models of F and $C_P(0) \cap Mod(F) = \emptyset$, they do not belong to $C_P(0)$. As a result, $C_P(m+1) = C(m) \backslash C_P(0)$ contains I and $C_P(l+1) = C(l) \backslash C_P(0)$ contains J. This proves that revising C by F increases the classes of I and J by one each. Therefore, I is greater or equal than J according to C_P if and only if it is according to C. □

Lemma 5 *Let $[K_j, P_{j+1}, \ldots, P_i, K_i]$ be a revision sequence generated by natural revision from a total preorder C. If $K_j \wedge P_i$ is consistent while none of $K_{j+1} \wedge P_i, \ldots, K_{i-1} \wedge P_i$ is, then $C_{P_{j+1},\ldots,P_i}(0) = Mod(K_j \wedge P_i)$.*

Proof. By Lemma 4, $C_{P_{j+1},\ldots,P_{i-1}}$ compares models of P_i in the same way C does. As a result, $Mod(K_i)$ are the minimal models of P_i in C. Since $C(0) = Mod(K_j)$ and $K_j \wedge P_i$ is consistent, these minimal models are $C(0) \cap Mod(P_i) = Mod(K_j \wedge P_i)$. □

Lemma 6 *If the revision sequence $[K_0, P_1, K_1, \ldots, P_i, K_i]$ is generated by natural revision from the total preorder C and P_i is inconsistent with each of $K_0, \ldots, K_{i-1}$, then the models of K_i are the minimal models of P_i according to C.*

Proof. By Lemma 4, revising C by $P_1, \ldots, P_{i-1}$ does not affect the order between the models of P_i. As a result, the minimal models of P_i according to $C_{P_1,\ldots,P_{i-1}}$ are the minimal models of P_i according to C. □

Theorem 2 *If $[K_0, P_1, K_1, \ldots, P_n, K_n]$ is natural-compatible then it is generated by natural revision from the initial preorder $C = [C(0), \ldots, C(n+1)]$.*

$$C(i) = \begin{cases} Mod(K_i) & \text{if } i \leq n \text{ and } \forall l < i \;.\; K_l \wedge P_i \models \bot \\ \emptyset & \text{otherwise, if } i \leq n \\ \bigcup\{Mod(K_j) \mid \exists l < j \;.\; K_l \wedge P_j \not\models \bot\} & \text{if } i = n+1 \end{cases}$$

Proof. Since no formula has index lower than zero, $C(0)$ is equal to $Mod(K_0)$ and is therefore not empty. The classes of C contain all models because $C(n+1)$ comprises every model not in $C(0), \ldots, C(n)$. To prove that C is a total preorder, it is to be proved that the sets $C(i)$ are disjoint. A model is in $C(i)$ with $i \leq n$ only if it is in $Mod(K_i)$. Therefore, it is also a model of P_i. Since P_i is inconsistent with each of $K_0, \ldots, K_{i-1}$, the model is not in any of $C(0), \ldots, C(i-1)$. The class $C(n+1)$ contains exactly the models that are not in $C(0) \cup \cdots \cup C(n)$. This proves that C is a total preorder.

The formulae K_i are generated in different ways depending on whether $K_j \wedge P_i$ is consistent for some $j < i$:

- if such an index j exists, then by the assumption of natural compatibility $K_i \equiv K_j \wedge P_i$ for the maximal such j; by Lemma 5, this is exactly the result of revising C by $P_1, \ldots, P_i$.
- otherwise, P_i is inconsistent with each of $K_0, \ldots, K_{i-1}$; by Lemma 6, $Mod(K_i)$ are the minimal models of P_i in the initial preorder C. Since P_i is inconsistent with $K_0, \ldots, K_{i-1}$, then: first, since $C(0), \ldots, C(i-1)$ are subsets of $Mod(K_0), \ldots, Mod(K_{i-1})$, these classes do not contain models of P_i; second, $C(i)$ is not empty but equal to $Mod(K_i)$. As a result, the minimal models of P_i are exactly $Mod(K_i)$.

□

Theorem 3 *A revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ is generated by natural revision from some initial total preorder if and only if it is natural-compatible.*

Proof. The previous theorem shows that every natural-compatible sequence is generated by natural revision from some total preorder. The converse is now proved: if a sequence is not natural-compatible then it is not generated by natural revision from any preorder.

Since $Mod(K_i) = C_{P_1,\ldots,P_{i-1},P_i}(0) = C_{P_1,\ldots,P_{i-1}}(l) \cap Mod(P_i)$ for some l, it holds $Mod(K_i) \subseteq Mod(P_i)$, which is the same as $K_i \models P_i$. Therefore, if $K_i \not\models P_i$ no preorder can generate the revision sequence. Otherwise, natural-compatibility is violated if, for some j and i:

1. $K_j \wedge P_i$ is consistent;
2. formulae $K_{j+1} \wedge P_i, \ldots, K_{i-1} \wedge P_i$ are inconsistent;
3. $K_j \wedge P_i$ is not equivalent to K_i.

By Lemma 5, the first two points imply that revising C_j by $P_{j+1}, \ldots, P_i$ generates a formula equivalent to $K_j \wedge P_i$, contradicting the third point. This proves that a revision sequence that is not natural-compatible is not generated by natural revision from any preorder. $\square$

Lemma 7 *Checking the existence of a total preorder C generating a revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ using natural revision is in* coNP.

Proof. According to Theorem 3, such a preorder C exists if and only if, for every $i \in \{1, \ldots, n\}$, it holds that $K_i \models P_i$ and, if j is the maximal index such that $j < i$ and $K_j \wedge P_i$ is consistent (if any), then $K_i \equiv K_j \wedge P_i$. The first part, $K_i \models P_i$ for all i, can be verified by a linear number of independent unsatisfiability tests.

The second part can be rewritten as: if $K_j \wedge P_i$ is consistent and $K_h \wedge P_i$ is inconsistent for every h between j and i, then $K_i \equiv K_j \wedge P_i$. Rewriting implication as disjunction, for every $j = 0, \ldots, i-1$:

1. either $K_0 \wedge P_i$ is inconsistent or $K_h \wedge P_i$ is consistent for some $0 < h < i$ or $K_i \equiv K_0 \wedge P_i$; and
2. either $K_1 \wedge P_i$ is inconsistent or $K_h \wedge P_i$ is consistent for some $1 < h < i$, or $K_i \equiv K_1 \wedge P_i$; and
3. $\ldots$;
4. either $K_{i-1} \wedge P_i$ is inconsistent or $K_i \equiv K_0 \wedge P_i$.

These are the conditions for index i. They have to hold for every $i \in \{1, \ldots, n\}$. Since the formulae K_i are consistent, the conditions can be simplified: if $K_j \wedge P_i$ is consistent, then either $K_h \wedge P_i$ is consistent for $j < h < i$ or $K_i \equiv K_j \wedge P_i$; in the latter case, $K_i \equiv K_h \wedge P_i$ for one such index h (the last). Therefore, the condition can be recast as:

$$K_j \wedge P_i \models \bot \text{ or } K_i \equiv K_{j+1} \wedge P_i \text{ or } \ldots K_i \equiv K_{i-1} \wedge P_i \text{ or } K_i \equiv K_j \wedge P_i$$

This condition can be checked with a number of independent unsatisfiability tests, and is therefore in coNP. $\square$

Theorem 4 *The problem of establishing the existence of a preorder generating a revision sequence using natural revision is* coNP *complete.*

Proof. Membership is proved in the previous lemma. Hardness is proved by reduction from the problem of checking whether a formula G is unsatisfiable. The instance is $[K_0, P_1, K_2]$ with $K_0 = a$, $K_1 = P_1 = a \rightarrow G$, where a is a new variable, not in G. If G is satisfiable, then $K_0 \wedge P_1$ is consistent; therefore, it should be equivalent to K_1. Instead, it is $a \wedge G$. If G is unsatisfiable, then $K_1 = P_1 = \neg a$ is inconsistent with K_0. By Theorem 3, a preorder generating the sequence using natural revision exists. $\square$

A.3 Restrained Revision: Proofs

Lemma 8 *Let $[K_j, P_{j+1}, \ldots, P_i, K_i]$ be a revision sequence generated by restrained revision from the initial total preorder C. If $K_j \wedge P_i$ is consistent while none of $K_{j+1} \wedge P_i$, $\ldots$, $K_{i-1} \wedge P_i$ is, then $C_{P_{j+1},\ldots,P_i}(0) = Mod(\text{maxset}(K_j \wedge P_i, P_{j+1}, \ldots, P_{i-1}))$.*

Proof. Proof is by induction on $i-j$. When $j = i+1$ the claim $C_{P_i}(0) = Mod(\text{maxset}(K_j \wedge P_i))$ holds because $K_j \wedge P_i$ is consistent by assumption.

The inductive claim is that $C_{P_{j+1},\ldots,P_{i-2},P_{i-1},P_i}(0) = Mod(\text{maxset}(K_j \wedge P_i, P_{j+1}, \ldots, P_{i-2}, P_{i-1}))$, the inductive assumption is the same without P_{i-1}:

$$C_{P_{j+1},\ldots,P_{i-2},P_i}(0) = Mod(\text{maxset}(K_j \wedge P_i, P_{j+1}, \ldots, P_{i-2}))$$

By definition, $C_{P_{j+1},\ldots,P_{i-2},P_i}(0)$ is $C_{P_{j+1},\ldots,P_{i-2}}(k) \cap Mod(P_i)$ where k is the minimal integer making this intersection non-empty. As a result, $C_{P_{j+1},\ldots,P_{i-2}}(l) \cap Mod(P_i) = \emptyset$ for all indexes l such that $0 \leq l < k$. Revising the preorder by P_{i-1} changes the classes of:

- the models of K_{i-1}, which become class zero;
- the models of the same class, which are split according to whether they satisfy P_{i-1}.

Since no model of K_{i-1} satisfies P_i by assumption, $C_{P_{j+1},\ldots,P_{i-2},P_{i-1}}(0) \cap Mod(P_i) = \emptyset$ holds. None of the classes of $C_{P_{j+1},\ldots,P_{i-2}}$ of index $0 \leq l < k$ intersect $Mod(P_i)$; therefore, neither do the ones resulting from splitting them. As a result, the minimal-index class intersecting $Mod(P_i)$ is one of the two resulting from splitting $C_{P_{j+1},\ldots,P_{i-2}}(k)$, which are:

$$\begin{array}{l} C_{P_{j+1},\ldots,P_{i-2}}(k) \cap Mod(P_{i-1}) \\ C_{P_{j+1},\ldots,P_{i-2}}(k) \backslash Mod(P_{i-1}) \end{array}$$

If the first intersects $Mod(P_i)$, this is $Mod(K_i)$; otherwise, $Mod(K_i)$ is the second. In formulae:

$$Mod(K_i) = \left\{ \begin{array}{ll} C_{P_{j+1},\ldots,P_{i-2}}(k) \cap Mod(P_{i-1}) \cap Mod(P_i) & \text{if not empty} \\ (C_{P_{j+1},\ldots,P_{i-2}}(k) \backslash Mod(P_{i-1})) \cap Mod(P_i) & \text{otherwise} \end{array} \right.$$

By the properties of set operators, this equation can be rewritten as:

$$Mod(K_i) = \left\{ \begin{array}{ll} C_{P_{j+1},\ldots,P_{i-2}}(k) \cap Mod(P_i) \cap Mod(P_{i-1}) & \text{if not empty} \\ (C_{P_{j+1},\ldots,P_{i-2}}(k) \cap Mod(P_i)) \backslash Mod(P_{i-1}) & \text{otherwise} \end{array} \right.$$

By the way k was defined, $C_{P_{j+1},\ldots,P_{i-2}}(k) \cap Mod(P_i)$ is equal to $C_{P_{j+1},\ldots,P_{i-2},P_i}(0)$. The latter is by the inductive assumption $Mod(\text{maxset}(K_j \wedge P_i, P_{j+1}, \ldots, P_{i-2}))$. Intersecting this set with $Mod(P_{i-1})$ if the result is not empty, and subtracting $Mod(P_{i-1})$ otherwise is the same as adding P_{i-1} to the end of the sequence, by the definition of maxset of a sequence. Since $Mod(K_i) = C_{P_{j+1},\ldots,P_{i-2},P_{i-1},P_i}(0)$, this proves the inductive claim that this set is equal to $Mod(\text{maxset}(K_j \wedge P_i, P_{j+1}, \ldots, P_{i-2}, P_{i-1}))$. $\square$

Lemma 9 *If $[K_0, P_1, K_1, \ldots, P_n, K_n]$ is restrained-compatible then it is generated by restrained revision from the total preorder $C = [C(0), \ldots, C(n+1)]$, where:*

$$C(i) = \begin{cases} Mod(K_i) & \text{if } i \leq n \text{ and } \forall l < i \;.\; P_i \wedge K_l \models \bot \\ \emptyset & \text{otherwise, if } i \leq n \\ \bigcup\{Mod(K_j) \mid \exists l < j \; P_j \wedge K_l \not\models \bot\} & \text{if } i = n+1 \end{cases}$$

Proof. No formula has index less than zero; therefore, $C(0) = Mod(K_0)$, which is not empty because all formulae in the sequences are consistent by assumption. Since $C(n+1)$ contains all models not in $C(0) \cup \cdots \cup C(n)$, the union of the classes include all models, and no model in $C(n+1)$ is also in another classes. In order to prove that C is a partition, it is shown that no model of $C(i)$ is also in $C(l)$ with $l < i \leq n$. If $I \in C(i)$ with $i \leq n$ then $I \in Mod(K_i)$ with $P_i \wedge K_l$ inconsistent for all $l < i$. Since $K_i \models P_i$, also $K_i \wedge K_l$ is inconsistent. This proves that a model I cannot be in $C(i)$ and also in $C(l)$ with $l < i$.

The previous lemma shows that, regardless of the initial total preorder, if a revision sequence is generated by restrained revision then $K_i \equiv \text{maxset}(K_j \wedge P_i, P_{j+1}, \ldots, P_{i-i})$ if j is the maximal index such that $j < i$ and $K_j \wedge P_i$ is consistent, if any. Since restrained compatibility ensures that this condition holds, all K_i for which such a j exists are obtained by revision regardless of C.

Remains to show that C generates K_i even if $P_i \wedge K_l$ is inconsistent for all $l < i$. The models of P_i are in $C(i) = Mod(K_i)$ and classes of greater index. Indeed, if a model of P_i were in $C(l)$ with $l < k$ then $P_i \wedge K_l$ would be consistent. Revising C by P_l, with $l = 1, \ldots, i-1$, changes the preorder in two ways:

1. the models of K_l are moved to class zero;
2. all other classes are split according to satisfaction of P_l.

Since $P_i \wedge K_l$ is inconsistent, no model of P_i is moved to class zero. As a result, the relative position of the models of P_i is only modified by the second change, the splitting of the classes. This could break the class $Mod(K_i)$ in two, but not in this case. Indeed, since either $K_i \models P_l$ or $K_i \models \neg P_l$, either all models of K_i satisfy P_l or all falsify P_l. In other words, the change may alter the comparison of two models of P_i, but does not if they are both in $Mod(K_i)$.

Two claims are therefore proved: that the models of P_i are in $C(i) = Mod(K_i)$ and greater classes; and that after the first $i-1$ revisions the minimal models of P_i are still $Mod(K_i)$. This proves that the result of the i-th revision is K_i. □

Theorem 5 *A revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ is generated by restrained revision from some initial total preorder if and only if it is restrained-compatible.*

Proof. The previous lemma shows that every restrained-compatible revision sequence is generated by a certain initial preorder. Remains therefore to prove the converse: if a revision sequence is not restrained-compatible then it is generated by no initial preorder.

Lemma 8 proves that every sequence generated by restrained revision satisfies the condition that $K_i \equiv \text{maxset}(P_i \wedge K_j, P_{j+1}, \ldots, P_{i-1})$ where j is the maximal $j < i$ such that

$P_i \wedge K_j$ is consistent. That $K_i \models P_i$ holds is a consequence of $Mod(K_i)$ being the set of minimal models satisfying P_i. Remains therefore to prove the necessity of the third condition of restrained compatibility: if P_i is consistent with all K_j with $j < k$ then either $K_i \models P_l$ or $K_i \models \neg P_l$ for every $l < i$.

To the contrary, let $l < i$ be such that $K_i \not\models P_l$ and $K_i \not\models \neg P_l$. These two conditions imply that $Mod(K_i)$ contains some models of P_l and some models of $\neg P_l$. Even if the models of K_i are in the same class after $l-1$ revisions, the l-th one separates the ones satisfying P_l from the ones satisfying $\neg P_l$. These may end in two consecutive classes, or in the new class zero and in another class, but either way the models of K_i are placed in two separate classes. Since constrained revision never merges classes, these models are still in separate classes after the $i-1$-th revision. As a result, revising by P_i can only select a part of $Mod(K_i)$ and not all of it. □

Lemma 10 *A revision sequence* $[K_0, P_1, K_1 \ldots, P_n, K_n]$ *is restrained-compatible if and only if, for every index* i*:*

1. $K_i \models P_i$*;*

2. *for every* $0 \leq j < i$*, either* $K_j \wedge P_i$ *is inconsistent or* $K_i \models K_l \wedge P_i$ *for some* $j < l < i$ *or* $K_i \equiv \mathrm{maxset}(K_j \wedge P_i, P_{j+1}, \ldots, P_{i-1})$*;*

3. *for every* $0 \leq j < i$*,* $K_i \models P_j$*,* $K_i \models \neg P_j$*, or* $K_i \models K_l \wedge P_i$ *for some* $0 \leq l < i$*.*

Proof. The three conditions of restrained compatibility can be rewritten as follows, for every index i.

1. $K_i \models P_i$;

2. for every j such that $0 \leq j < i$, either $K_j \wedge P_i$ is inconsistent or $K_l \wedge P_i$ is consistent for some $j < l < i$ or $K_i \equiv \mathrm{maxset}(K_j \wedge P_i, P_{j+1}, \ldots, P_{i-1})$;

3. for every j such that $0 \leq j < i$, either $K_i \models P_j$ or $K_i \models \neg P_j$ or $K_l \wedge P_i$ is consistent for some $0 \leq l < i$.

Both the second and the third point include "$K_l \wedge P_i$ is consistent for some $h \leq l < i$": in the first, $h = j + 1$; in the second, $h = 0$. In the particular conditions of the lemma, these two conditions will be shown to be equivalent:

1. $K_l \wedge P_i$ is satisfiable for some l such that $h \leq l < i$;

2. $K_i \models K_l \wedge P_i$ for some l such that $h \leq l < i$.

It is proved that: first, if the sequence is restrained-compatible then Condition 1 implies Condition 2; second, if the three conditions of the lemma are true then Condition 2 implies Condition 1.

Condition 1 is that $K_l \wedge P_i$ is consistent for some $h \leq l < i$. Either l is the maximal index with this property or some other index between h and i is. Let g be such a maximal index. By restrained compatibility, $K_i \equiv \mathrm{maxset}(K_g \wedge P_i, P_{i-1}, \ldots, P_{g+1})$. The maxset of a sequence implies its first element, if consistent. Since this is the case, $K_i \models K_g \wedge P_i$. This means that Condition 2 holds for index g.

Condition 2 is that $K_i \models K_h \wedge P_i$ for some $h \leq l < i$. Since K_i is consistent, $K_h \wedge P_i$ is consistent as well, proving Condition 1 for the same index l. □

Theorem 6 *Establishing the existence of a total preorder generating a restrained revision sequence is* coNP*-complete.*

Proof. Membership follows from the previous lemma. Its conditions can be all reformulated as entailments (like $K_i \models P_i$), inconsistencies (like that of $K_j \wedge P_i$) and equivalences with maxsets ($K_i \equiv \mathrm{maxset}(K_j \wedge P_i, P_{j+1}, \ldots, P_{i-1})$). All these problems are in coNP, and can therefore be rewritten as universal quantified formulae. The whole problem is a combination of these; by renaming all variables and taking out the quantifiers, a single universally quantified formula results. Since $\forall$QBF is in coNP, the problem is in coNP.

Hardness is proved by reduction from the problem of propositional unsatisfiability. Given a formula F, the associated revision sequence is $[a, b \vee F, a \wedge b]$, where a and b are fresh variables, not occurring in F. If F is unsatisfiable then $b \vee F$ is equivalent to b, and the sequence $[a, b, a \wedge b]$ is generated by the preorder $C = [C(0), C(1)]$ with $C(0) = \mathit{Mod}(a)$ and $C(1) = \mathit{Mod}(\neg a)$. If F is satisfiable then $a \wedge (b \vee F)$ is satisfiable but is not equivalent to $a \wedge b$, which makes the sequence not generated by restrained revision from any preorder. The sequence is therefore generated by restrained revision from some preorder if and only if F is unsatisfiable. □

A.4 Lexicographic Revision: Proofs

Lemma 11 *If a revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ is generated by lexicographic revision from the total preorder among models C then it has the result preorder R defined by:*

$$R(i) = \{K_i \mid \mathit{Mod}(K_i) \subseteq C(i)\}$$

Proof. R is a total preorder because: first, $R(0)$ is not empty because $C(0) = \mathit{Mod}(K_0)$, which implies $\mathit{Mod}(K_0) \subseteq C(0)$ and $K_0 \in R(0)$; second, no formula is in two classes of R because C is a partition. Third, every K_i is in a class of R because of Property 1: since $\mathit{Mod}(K_i)$ is a set of minimal models of C, all its models are in the same class $C(j)$; therefore, $K_i \in R(j)$. Remains to prove that this preorder satisfies the two conditions for being a result preorder for the sequence.

If K_i and K_j are consistent with each other, they have a common model I. Let h be its class. By construction, since $I \models K_i$ then $K_i \in R(h)$. The same holds for K_j, since $I \models K_j$. As a result, $K_j \in R(h)$.

Let K_i and K_j be such that $\neg K_i \wedge K_j \wedge \mathrm{maxset}(P_i, \ldots, P_1)$ is consistent. By assumption, the sequence is generated by lexicographic revision from C. As a result, the models of K_i are exactly the minimal models of $\mathrm{maxset}(P_i, \ldots, P_1)$ according to C. Let h be the class of such models. By construction of R, it holds $K_i \in R(h)$. All models of $\neg K_i \wedge \mathrm{maxset}(P_i, \ldots, P_1)$ belong to a greater class $l > h$. Since $\neg K_i \wedge K_j \wedge \mathrm{maxset}(P_i, \ldots, P_1)$ is consistent, K_j has some models in $C(l)$. Since all models of K_j are in the same class, it follows $Mod(K_j) \subseteq C(l)$, which implies $K_j \in R(l)$. Since $l > h$, the claim is proved. □

Lemma 12 *If R is a result preorder for a revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ such that $K_i \models \text{maxset}(P_i, \ldots, P_1)$ for every i, then lexicographic revision from the following total preorder C among models generates the revision sequence, where z is the index of the greatest class of R:*

$$C(h) = \begin{cases} \bigcup\{Mod(K_i) \mid K_i \in R(h)\} & \text{if } h \leq z \\ \{I \mid I \not\in K_1 \cup \cdots \cup K_n\} & \text{if } h = z+1 \end{cases}$$

Proof. C is proved to be a total preorder. First, $C(0)$ is not empty since $R(0)$ is not empty. Second, no model belongs to more than one class: if I is in both $C(h)$ and $C(l)$ then by definition there exists $K_i \in R(h)$ and $K_j \in R(l)$ that are both satisfied by I; this implies that $K_i \wedge K_j$ is consistent; since R is a result preorder for the sequence, it follows that $l = h$. Models of $C(z+1)$ are exactly the ones not contained in the other classes.

Remains to be proved that lexicographic revision generates the revision sequence from C: revising C by $P_1, \ldots, P_i$ generates K_i. Since $K_i \models \text{maxset}(P_i, \ldots, P_1)$ by assumption, all models of K_i are also models of $\text{maxset}(P_i, \ldots, P_1)$. Remains to be proved that they are the minimal ones: all other models are in greater classes. Let h be the class of the models of K_i, and I be a model of $\neg K_i \wedge \text{maxset}(P_i, \ldots, P_1)$. If I does not satisfy any other K_j then $I \in R(z+1)$, and $z+1 > h$ since z is the greatest index of the classes of R. If I satisfies K_j then it satisfies $\neg K_i \wedge K_j \wedge \text{maxset}(P_i, \ldots, P_1)$, which is therefore satisfiable. Since R is a result preorder for the sequence, K_j is in a class $R(l)$ with $l > h$. This implies that $I \in C(l)$. This proves that every model of $\text{maxset}(P_i, \ldots, P_1)$ that is not in $Mod(K_i)$ is in a class greater than h. □

Lemma 13 *A revision sequence is generated by lexicographic revision from some total preorder if and only if it is lexicographic compatible.*

Proof. If the revision sequence is generated by some total preorder then $K_i \models \text{maxset}(P_i, \ldots, P_1)$ holds for every i. Furthermore, a result preorder R for it exists by Lemma 11. From it, define the relations $\simeq$ and $<$ as: $K_i \simeq K_j$ if and only if K_i and K_j are in the same class $R(h)$; $K_i < K_j$ if $K_i \in R(h)$, $K_j \in R(l)$ and $h < l$. Since R is a revision preorder for the sequence, if $K_i \wedge K_j$ is consistent then K_i and K_j are in the same class, which implies $K_i \simeq K_j$. In the same way, if $\neg K_i \wedge K_j \wedge \text{maxset}(P_i, \ldots, P_1)$ is consistent then K_i is in a lower class than K_j, implying $K_i < K_j$. A cycle containing $<$ is impossible because $K_i < K_j$ means that the class of K_i is lower than the class of K_j.

This does not prove lexicographic compatibility, which requires $K_i \simeq K_j$ and $K_i < K_j$ to be respectively equivalent to the consistency of $K_i \wedge K_j$ and $\neg K_i \wedge K_i \wedge \text{maxset}(P_i, \ldots, P_1)$, not merely implied by them. However, by removing from these relations all pairs K_i and K_j not satisfying the respective condition makes these relations weaker. Therefore, no new cycle is created.

To prove the other direction, assume that $K_i \models \text{maxset}(P_i, \ldots, P_1)$ for every i and that $\simeq$ and $<$ are two relations defined from the sequence as specified in the definition of lexicographic compatibility. A result preorder for the sequence is shown. Together with $K_i \models \text{maxset}(P_i, \ldots, P_1)$, this implies that the revision sequence is generated by some initial preorder by Lemma 12.

From $\simeq$ and $<$, a relation $\leq$ is defined: $K_i \leq K_j$ if either $K_i \simeq K_j$ or $K_i < K_j$. This relation is not necessarily transitive, but has a number of other properties:

1. if $K_i \simeq K_j$ then $K_i \leq K_j$ and $K_j \leq K_i$; indeed, $K_i \simeq K_j$ holds only if $K_i \wedge K_j$ is consistent, which implies that $K_j \wedge K_i$ is consistent; therefore, $K_j \simeq K_i$, which implies $K_j \leq K_i$;

2. if $K_i < K_j$ then $K_i \leq K_j$ and $K_j \not\leq K_i$; to the converse, if $K_j \leq K_i$, then either $K_j \simeq K_i$ or $K_j < K_i$; in both cases, this is a cycle of $\simeq$ and $<$ containing one $<$ link, which contradicts the assumption that no such cycle exists;

3. $\leq$ is reflexive; indeed, all formulae K_i are consistent by assumption; therefore, $K_i \wedge K_i$ is consistent, which implies $K_i \simeq K_i$ and $K_i \leq K_i$;

4. $\leq$ is Suzumura consistent (Suzumura, 1976): it does not form cycles $K_{i_1}, \ldots, K_{i_m} = K_{i_1}$ such that $K_{i_j} \leq K_{i_{j+1}}$ for all j and for some j also $K_{i_{j+1}} \not\leq K_{i_j}$; this is proved below.

Properties 1 and 2 mean that $\simeq$ and $<$ are the equivalence and strict part of $\leq$, respectively. Indeed, $K_i \leq K_j$ holds only if either $K_i \simeq K_j$ or $K_i < K_j$; the former implies $K_j \leq K_i$, the latter $K_j \not\leq K_i$. Property 4 is a consequence of this fact and of the assumption of nonexistence of cycles of $\simeq$ and $<$ containing at least one link $<$.

Since $\leq$ is reflexive (Property 3) and Suzumura consistent (Property 4), by the Suzumura extension theorem (Suzumura, 1976) a total preorder R extending $\leq$ exists. Extending means that both the equivalence and strict parts of $\leq$ are preserved in R. Since these have been proved to be $\simeq$ and $<$, the total preorder R has K_i and K_j in the same class if $K_i \simeq K_j$ and K_i in a lower class if $K_i < K_j$.

If $K_i \wedge K_j$ is consistent then $K_i \simeq K_j$, which implies that K_i and K_j are in the same class of R. If $\neg K_i \wedge K_j \wedge \text{maxset}(P_i, \ldots, P_1)$ is consistent then $K_i < K_j$, which implies that the class of K_i in R is less than the class of K_j. This proves that R is a result preorder for the revision sequence. $\square$

Lemma 14 *A revision sequence $[K_0, P_1, K_1 \ldots, P_n, K_n]$ is not lexicographic compatible if and only if either $K_i \not\models \text{maxset}(P_i, \ldots, P_1)$ for some i or consistent sets $R_1, \ldots, R_n$ exist such that:*

1. *$\{P_j \mid 1 \leq j \leq i$ and $K_i \models P_j\} \subseteq R_i$ for every i; and*

2. *there exists a cycle $K_{i_1}, \ldots, K_{i_m} = K_{i_1}$ such that either $K_{i_j} \wedge K_{i_{j+1}}$ or $\neg K_{i_j} \wedge K_{i_{j+1}} \wedge R_i$ is consistent for all $i_j \in \{i_1, \ldots, i_{m-1}\}$, and the second is consistent for at least one such index.*

Proof. A revision sequence is lexicographic compatible if and only if $K_i \models \text{maxset}(P_i, \ldots, P_1)$ for all i and there is no cycle as specified by the definition. Inverting this condition, a sequence is not lexicographic compatible if either $K_i \not\models \text{maxset}(P_i, \ldots, P_1)$ or a cycle exists. As a result, one can check whether $K_i \not\models \text{maxset}(P_i, \ldots, P_1)$ for some i; if this is true, no further check is needed: the sequence is not lexicographic compatible. The presence of cycles is irrelevant in this case. The other case is that $K_i \models \text{maxset}(P_i, \ldots, P_1)$ for all i, and then the sequence is not lexicographic compatible if and only if it contains cycles.

The point is that the condition for the presence of cycles can be written under the assumption that $K_i \models \text{maxset}(P_i, \ldots, P_1)$, since this is the only case where this condition matters. By Lemma 2, the only consistent subset of $\{P_1, \ldots, P_i\}$ containing all formulae entailed by K_i is $\text{maxset}(P_i, \ldots, P_1)$. This means that R_i can be used in place of $\text{maxset}(P_i, \ldots, P_1)$, since the only R_i satisfying the first condition in the statement of the lemma is $\text{maxset}(P_i, \ldots, P_1)$. □

Theorem 7 *The problem of checking the existence of a total preorder generating a revision sequence using lexicographic revision is* coNP*-complete.*

Proof. Membership follows from the previous lemma: the sequence is not generated by lexicographic revision from any preorder if and only if it is not lexicographic compatible, which in turns can be checked by existential quantifiers only:

1. $K_i \not\models \text{maxset}(P_i, \ldots, P_1)$;
2. there exists R_i;
3. R_i is consistent;
4. either $K_i \not\models P_j$ or $P_j \in R_i$;
5. there exists a cycle $K_{i_1}, \ldots, K_{i_m} = K_{i_1}$;
6. $K_{i_j} \wedge K_{i_{j+1}}$ is consistent;
7. $\neg K_{i_j} \wedge K_{i_{j+1}} \wedge R_i$ is consistent.

The first condition can be expressed in terms of existential quantifiers only as shown by Lemma 3. The same holds for the other conditions as well. As a result, incompatibility is in NP, which means that the existence of preorder generating the sequence is in coNP.

Hardness is proved by reduction from propositional unsatisfiability. A formula F is satisfiable if and only if $[K_0, P_1, K_1]$ is generated by no preorder, where $K_0 = a$, $P_1 = a$ and $K_1 = a \vee F$ and a is a new variable not contained in F. Indeed, if F is unsatisfiable, then $K_1 = a$, and the sequence is generated by the preorder $C = [C(0), C(1)]$ with $C(0) = Mod(a)$ and $C(1) = Mod(\neg a)$. If F is satisfiable, then K_1 has some models that do not satisfy P_1: the ones of F. As a result, $K_1 \not\models P_1$, and the sequence is not generated by any preorder. □

A.5 Reinforcement Revision: Proofs

Lemma 15 *For every sequence of n numbers $V = [V(1), \ldots, V(n)]$, it holds $M_V(I, i, j) = M_V(I, i, h) + M_V(I, h, j)$ for every three indexes i, j and h.*

Proof. If $i < j$ then $M_V(I, i, j)$ is the sum of $M_V(I, l, l+1)$ for $i < l \leq j$. Otherwise, it is the sum of $-M_V(I, l, l+1) = M_V(I, l+1, l)$. Also $M_V(I, i, h)$ and $M_V(I, h, j)$ can be expressed in the same way. If h is between i and j the result follows immediately. Otherwise, if $i < j$ and $h > j$ then $M_V(I, i, h)$ includes $M_V(I, l, l+1)$ for $l > j$, but then $M_V(I, h, j)$ includes $M_V(I, l+1, l) = -M_V(I, l, l+1)$, which subtracts the same amount from the sum. The case $h < i$ is similar. □

Lemma 16 *If a revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ is generated by reinforcement revision from the total preorder C and $D_C = [D_C(1), \ldots, D_C(n)]$ with $D_C(i) = \min\{j \mid C_{P_1,\ldots,P_{i-1}}(j) \cap Mod(P_i) \neq \emptyset\}$ and $I \models K_i$ then, for every j:*

- $M_{D_C}(I, i, j) = 0$ *if* $I \models K_j$;
- $M_{D_C}(I, i, j) > 0$ *otherwise.*

Proof. If a model I is in $C_{P_1 \ldots P_{i-1}}(c)$ then it is in $C_{P_1 \ldots P_i}(c - D_C(i))$ if $I \models P_i$ and in $C_{P_1 \ldots P_{i-1}}(c+1)$ otherwise. As a result, $M_{D_C}(I, i, i+1)$ is the difference between the class of I in the preorder at step $i+1$ and the preorder at step i. Since $M_{D_C}(I, i, j)$ is the sum of these amounts from index i to index j, it is the difference between the class of I at step j and i. If $I \models K_i$ then its class at step i is zero. As a result, $M_{D_C}(I, i, j)$ is the class of I at step j. This is zero if $I \models K_j$ and greater otherwise. □

Lemma 17 *If $V = [V(1), \ldots, V(n)]$ is a reinforcement mover for the revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$, the following initial preorder $C = [C(0), \ldots, C(V(1) + \cdots + V(n+1))]$ generates the revision sequence by reinforcement revision and $D_C = V$.*

$$C(j) = \begin{cases} \{I \mid I \models K_i \text{ and } M_V(I, i, 0) = j\} & \text{if } j < V(1) + \ldots + V(n) + 1 \\ \{I \mid \forall i \,.\, I \not\models K_i\} & \text{if } j = V(1) + \ldots + V(n) + 1 \end{cases}$$

Proof. Since K_0 is consistent, it has at least a model I. Since $I \models K_0$ and $M_V(I, 0, 0) = 0$ by definition, it follows $I \in C(0)$, proving that $C(0)$ is not empty. Every model is in some class because if $I \models K_i$ then I is in $C(M_V(I, i, 0))$, otherwise it is in $C(V(1) + \ldots + V(n) + 1)$. To prove that C is a total preorder, remains to prove that no model belongs to two classes. If $I \models K_i$ and $I \models K_j$ then $M_V(I, j, i) = 0$ because V is a reinforcement mover; therefore, $M_V(I, i, 0) = 0 + M_V(I, i, 0) = M_V(I, j, i) + M_V(I, i, 0) = M_V(I, j, 0)$. A model of $C(V(1) + \ldots + V(n) + 1)$ does not satisfy any K_i; therefore, it cannot be in any other class $C(j)$. This proves that C is a total preorder.

The longest part of the proof is to show that reinforcement revision generates the sequence from C. Inductively, it is assumed that $D_C(1) = V(1)$, $D_C(2) = V(2)$, $\ldots$, $D_C(i) = V(i)$, where $D_C(l) = \min\{j \mid C_{P_1,\ldots,P_{l-1}}(j) \cap Mod(P_l) \neq \emptyset\}$. It is then proved that all models of K_i are in $C_{P_1,\ldots,P_i}(0)$ and all other models are in classes of greater index of the same total preorder. Furthermore, the minimal class of models of P_{i+1} in that preorder is $V(i+1)$, proving that $D_C(i+1) = V(i+1)$, which allows to iterate the proof.

The base case is with $i = 0$: what is to be proved is that $Mod(K_0) = C(0)$ and $D_C(1) = V(1)$. By construction of C, a model I is in $C(0)$ if $M_V(I, i, 0) = 0$ and $I \models K_i$ for some i; both hold for $i = 0$. Vice versa, if $I \in C(0)$ then $I \models K_i$ and $M_V(I, i, 0) = 0$ for some i. If $I \not\models K_0$, the definition of reinforced mover with $j = 0$ implies that $M_V(I, i, 0) > 0$, which in turn implies $I \notin C(0)$, a contradiction.

In order to prove that $D_C(1) = V(1)$, let I be a model of P_1. By construction of C, if $I \models K_1$ then I is in $C(V(1))$. If $I \not\models K_1$ two cases are possible. In the first, $I \not\models K_j$ for every j. This implies that $I \in C(V(1)+\ldots+V(n)+1)$, and $V(1)+\ldots+V(n)+1 > V(1)$. In the second case, $I \models K_j$ for some other j. By definition of reinforcement mover, $M_V(I, j, 1) > 0$. The class of C containing I has index $M_V(I, j, 0) = M_V(I, j, 1)+M_V(I, 1, 0) = M_V(I, j, 1)+V(1)$, the last step being a consequence of $I \models P_i$. Since $M_V(I, j, 1) > 0$, it follows that this amount is greater than $V(1)$. This proves that $V(1)$ is the index of the minimal class of models of P_1 in C: $D_C(1) = V(1)$. This concludes the base case.

It is now assumed that $D_C(1) = V(1)$, $\ldots$, $D_C(i) = V(i)$, and is proved that $C_{P_1\ldots P_i}$ has the class zero equal to $Mod(K_i)$ and the minimal models of P_{i+1} in class $V(i+1)$, which proves that $D_C(i+1) = V(i+1)$. By induction, this proves that the sequence is generated by reinforcement revision from the total preorder C.

Let I be a model of K_i. By construction, $I \in C(M_V(I, i, 0))$. The first i revisions increase the class of I by $M_{D_C}(I, 0, i)$. Since $D_C(1), \ldots, D_C(i)$ are equal to $V(1), \ldots, V(i)$ by the inductive assumption, this is the same as $M_V(I, 0, i)$. By definition of M_V, it holds $M_V(I, i, 0) = -M_V(I, 0, i)$. As a result, $M_V(I, i, 0) + M_{D_C}(I, 0, i) = M_V(I, i, 0) - M_V(I, i, 0) = 0$: the model I is in $C_{P_1\ldots P_i}(0)$.

If I is not a model of K_i, then it may be a model of some other K_j or not. In the second case, it is in $C(V(1) + \cdots + V(n) + 1)$. The first i steps reduce its class number by at most $D_C(1) + \ldots + D_C(i)$, leading to $V(1) + \cdots + V(n) + 1 + D_C(1) + \ldots + D_C(i)$. Since the first i values of V and D_C coincide, this is $V(1) + \cdots + V(n) + 1 - V(1) - \cdots - V(i) = 1 + V(i+1) + \cdots + V(n)$, which is greater than zero.

If I is a model of some K_j then $I \in C(M_V(I, j, 0))$ by the definition of C. Its class at step i is therefore $M_V(I, j, 0)+M_{D_C}(I, 0, i)$. Since $D_C(1), \ldots, D_C(i)$ coincide with $V(1), \ldots, V(i)$ by the induction hypothesis, the second term is equal to $M_V(I, 0, i)$. The sum is therefore equal to $M_V(I, j, 0) + M_V(I, 0, i) = M_V(I, j, 0) - M_V(I, i, 0) = M_V(I, j, i) + M_V(I, i, 0) - M_V(I, i, 0) = M_V(I, j, i)$. By the definition of reinforced mover with i and j reversed, since $I \models K_j$ and $I \not\models K_i$ then $M_V(I, j, i) > 0$. This proves that the class of I in the order $C_{P_1\ldots P_i}$ has index larger than zero.

The last step of the proof is to show that the models of P_{i+1} are in classes of index $V(i+1)$ and greater according to the ordering at step i, which is $C_{P_1\ldots P_i}$. Let I be a model of P_{i+1}. If it does not satisfy any K_j then it is in $C(V(1)+\cdots+V(n)+1)$. At step i its class is at least $V(1)+\cdots+V(n)+1-D_C(1)-\cdots-D_C(i)$, since $D_C(j)$ is the maximal decrease of classes at step j. Since $D_C(1), \ldots, D_C(i)$ coincide with $V(1), \cdots, V(i)$ by the induction assumption, this is equal to $V(1)+\cdots+V(n)+1-V(1)-\cdots-V(i) = V(i+1)+\cdots+V(n)+1$, which is larger than $V(i+1)$.

If I satisfies some K_j, possibly with $j = i+1$, then $I \in C(M_V(I,j,0))$ by the definition of C. At step i, its class index is $M_V(I,j,0) + M_{D_C}(I,0,i)$. Since $D_C(1), \ldots, D_C(i)$ are assumed equal to $V(1), \ldots, V(i)$, the second term is equal to $M_V(I,0,i)$ and the sum to $M_V(I,j,0) + M_V(I,0,i) = M_V(I,j,i)$. Since $I \models P_{i+1}$, $M_V(I,i,i+1) = -V(i+1)$ by the definition of M_V (Definition 11 with V in place of D_C). The definition of reinforced mover ensures that $M_V(I,j,i+1) = M_V(I,j,i) + M_V(I,i,i+1) = M_V(I,j,i) - V(i+1)$ is equal to 0 if $I \models K_{i+1}$ and greater otherwise. Since K_{i+1} has some models, the minimal value of $M_V(I,j,i) - V(i+1)$ is zero, proving that the minimal value of $M_V(I,j,i)$ is $V(i+1)$. This proves that $D_C(i+1) = V(i+1)$. □

Lemma 18 *If reinforcement revision generates a revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ from some total preorder, then it also generates the same sequence from a preorder in which the minimal initial class of models of P_1 is 0 or 1.*

Proof. If $K_0 \wedge P_1$ is consistent then P_1 has models in K_0, which is class zero.

Otherwise, let the class indexes of models of P_1 be $k_1 < k_2 < k_3 < \ldots$ Revision by P_1 decreases all these numbers by k_1, making them $k_1 - k_1 = 0$, $k_2 - k_1$, $k_3 - k_1$, etc.

From the given initial preorder, a new one can be generated by reducing each model of P_1 of $k_1 - 1$ classes. This new initial preorder generates the same revision sequence. The models that do not satisfy P_1 are not changed of initial class, and after revising by P_1 are still moved up one class.

The models of P_1 are in classes $k_1 - (k_1 - 1)$, $k_2 - (k_1 - 1)$, $k_3 - (k_1 - 1)$, etc. of the new preorder. Since k_1 is the minimal index of models of P_1, none of these indexes is zero. Therefore, no model of P_1 enters K_0.

The minimal class of models of P_1 in the new initial preorder is 1, proving that $D_C(1) = 1$. This implies that revising by P_1 decreases the class of the models of P_1 by one. As a result, the indexes of these classes are $k_1 - (k_1 - 1) - 1$, $k_2 - (k_1 - 1) - 1$, $k_3 - (k_1 - 1) - 1$, etc. and these coincide with $k_1 - k_1 = 0$, $k_2 - k_1$, $k_3 - k_1$, etc. These are the same classes obtained revising the original preorder by P_1. The models that do not satisfy P_1 are in the same initial class and are still moved up one class. This proves that the ordering after revising by P_1 is the same as before. Therefore, from this point on the revision sequences are identical. □

Lemma 19 *If reinforcement revision generates the sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ from a total preorder, it is also generates the same sequence from a total preorder C such that $D_C(i+1) \leq D_C(1) + \ldots + D_C(i) + i + 1$.*

Proof. The claim is proved in two cases separately: P_{i+1} has some models that satisfy some K_j with $j < i$, or it does not.

Case 1: some models of P_{i+1} satisfy K_j with $j < i$. Such models are in class zero at step j. Therefore, they are at most in class $i - j$ at step i. This proves that P_{i+1} has some models in class $i - j$, so the minimal class of its models is $i - j$ or less: $D_C(i+1) \leq i - j$, which is less that $D_C(1) + \ldots + D_C(i) + i + 1$ because $i - j < i + 1$.

Case 2: no model of P_{i+1} is in any K_j with $j < i$. Let $C(l)$ be the minimal initial class of the models of P_{i+1}. If $l > D_C(1) + \ldots + D_C(i)$, then all models of P_{i+1} can be decreased of $l-(D_C(1)+\ldots+D_C(i))-1$ classes without affecting the generated sequence. Models of class l move to class $l-(l-(D_C(1)+\ldots+D_C(i)-1)) = D_C(1)+\ldots+D_C(i)+1$. The maximal class such a model may reach at step i is $D_C(1)+\ldots+D_C(i)+1+i$, since each step may at most increase the class of a model of one. Therefore, $D_C(i+1) \leq D_C(1)+\ldots+D_C(i)+i+1$.

What remains to be proved is that this change does not affect the revision results. Regarding the steps before $i+1$, since the models of P_{i+1} are in the initial class $D_C(1) + \ldots + D_C(i)+1$ or greater, while the models of K_1 are in the initial class $D_C(1)$, the minimal class of models of P_1 is still $D_C(1)$ and the result of revision is still K_1. Since $D_C(1)$ is the same as before, a similar line of reasoning can be applied to the models of K_2 and to $D_C(2)$, then to K_3 and so on until K_i, proving that neither $K_1, \ldots, K_i$ nor $D_C(1), \ldots, D_C(i)$ are affected by the change. Since $D_C(1), \ldots, D_C(i)$ tell how the models are moved, at step i the models of $\neg P_{i+1}$ are in the same classes as before the change. The models of P_{i+1} are lowered of $D_C(i+1)$ classes. Since the change has not altered their relative initial positions, it does not modify their relative positions at step i. The preorder at step $i+1$ is therefore the same as before. □

Lemma 20 *Reinforcement revision generates the revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ from a total preorder if and only if it generates the same sequence from a total preorder such that each $D_C(i)$ is bounded by $2^i - 1$.*

Proof. It has been proved that if $D_C(1)$ is not 1, the preorder can be modified to make it so.

Since $D_C(i+1)$ is bounded by $D_C(1) + \ldots + D_C(i) + i + 1$, assuming the claim for $D_C(1), \ldots, D_C(i)$ leads to $D_C(i+1) \leq (2^1-1) + \cdots + (2^i-1) + i + 1$, which is the same as $D_C(i+1) \leq (2^0 + 2^1 + \cdots + 2^i) - 2^0 - 1 * i + i + 1 = 2^1 + \cdots + 2^i$, which is $2^{i+1}-1$. □

Theorem 8 *Establishing the existence of a total preorder generating the revision sequence $[K_0, P_1, \ldots, P_n, K_n]$ by reinforcement revision is in Σ_2^p, and is in* coNP *if n is a constant.*

Proof. By the above lemma, a revision sequence is generated by a preorder if and only if it is generated by a preorder for which $D_C(i)$ is between 0 and $2^i - 1$. As a result, the problem can be solved by guessing a reinforcement mover $V = [V(1), \ldots, V(n)]$ for the sequence, where each $V(i)$ is between 0 and $2^i - 1$, since this implies not only the existence of a total preorder C generating the sequence but also that $D_C(i) = V(i)$ for all indexes. If n is a constant the guessing can be replaced by a disjunction. Checking whether V is a reinforcement mover amount to check whether, for all I, i and j, it holds that $I \models K_i$ implies $M_V(I,i,j) = 0$ if $I \models K_j$ and $M_V(I,i,j) > 0$ otherwise. Calculating $M_V(I,i,j)$ can be done in polynomial time since the only checks it requires are in the form $j > i$ and $I \models P_l$. Therefore, this verification is in coNP, and the whole problem in Σ_2^p because of the guessing of V. □

Theorem 9 *Checking the existence of a preorder generating a reinforcement revision sequence $[K_0, P_1, K_1, \ldots, P_n, K_n]$ is* coNP*-complete, if n is a constant.*

Proof. Membership is proved by the previous theorem. Hardness can be proved from the problem of propositional unsatisfiability. Given a formula F, its corresponding revision sequence is $[K_0, P_1, K_1]$ with $K_0 = a \vee F$, $P_1 = b$ and $K_1 = a \wedge b$, where a and b are fresh variables not occurring in F. In this sequence, $K_0 = a \vee F$ is consistent with $P_1 = b$. Therefore, $K_1 = a \wedge b$ should be equivalent to their conjunction $(a \vee F) \wedge b$. If F is unsatisfiable this is the case. Otherwise, the model of F extended by assigning a to false and b to false is a model of $K_0 \wedge P_1$ that is not a model of K_1. This proves that the sequence is generated by some preorder if and only if F is unsatisfiable. □

References

Alchourròn, C., & Makinson, D. (1982). On the logic of theory change: Contraction functions and their associated revision functions. *Theoria*, *48*(1), 14–37.

Areces, C., & Becher, V. (2001). Iterable AGM functions. In Rott, H., & Williams, M.-A. (Eds.), *Frontiers in Belief Revision*, Applied Logic Series, pp. 261–277. Kluwer Academic Publisher.

Baltag, A., Gierasimczuk, N., & Smets, S. (2011). Belief revision as a truth-tracking process. In *Proceedings of the Thirteenth Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2011)*, pp. 187–190.

Benferhat, S., Kaci, S., Le Berre, D., & Williams, M.-A. (2004). Weakening conflicting information for iterated revision and knowledge integration. *Artificial Intelligence*, *153*, 339–371.

Booth, R., & Meyer, T. (2006). Admissible and restrained revision. *Journal of Artificial Intelligence Research*, *26*, 127–151.

Booth, R., & Nittka, A. (2008). Reconstructing an agent's epistemic state from observations about its beliefs and non-beliefs. *Journal of Logic and Computation*, *18*, 755–782.

Boutilier, C. (1996). Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic*, *23*, 263–305.

Carson, R., & Louviere, J. (2011). A common nomenclature for stated preference elicitation approaches. *Environmental and Resource Economics*, *49*(4), 539–559.

Darwiche, A., & Pearl, J. (1997). On the logic of iterated belief revision. *Artificial Intelligence Journal*, *89*(1–2), 1–29.

Delgrande, J., Dubois, D., & Lang, J. (2006). Iterated revision as prioritized merging. In *Proceedings, Tenth International Conference on Principles of Knowledge Representation and Reasoning, KR-2006*, pp. 210–220.

Eiter, T., & Gottlob, G. (1996). The complexity of nested counterfactuals and iterated knowledge base revisions. *Journal of Computer and System Sciences*, *53*(3), 497–512.

Fagin, R., Ullman, J. D., & Vardi, M. Y. (1983). On the semantics of updates in databases. In *Proceedings of the Second ACM SIGACT SIGMOD Symposium on Principles of Database Systems (PODS'83)*, pp. 352–365.

Fermé, E., & Hansson, S. (2011). AGM 25 years - Twenty-five years of research in belief change. *Journal of Philosophical Logic*, *40*(2), 295–331.

Gärdenfors, P. (1988). *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books, MIT Press, Cambridge, MA.

Hild, M., & Spohn, W. (2008). The measurement of ranks and the laws of iterated contraction. *Artificial Intelligence*, *172*(10), 1195–1218.

Jin, Y., & Thielscher, M. (2007). Iterated belief revision, revised. *Artificial Intelligence Journal*, *171*(1), 1–18.

Katsuno, H., & Mendelzon, A. O. (1991). Propositional knowledge base revision and minimal change. *Artificial Intelligence*, *52*, 263–294.

Konieczny, S., & Pino Pérez, R. (2000). A framework for iterated revision. *Journal of Applied Non-Classical Logics*, *10*, 339–367.

Lehmann, D., & Magidor, M. (1992). What does a conditional knowledge base entail? *Artificial Intelligence*, *55*, 1–60.

Liberatore, P. (2014a). Belief revision by examples. *Computing Research Repository (CoRR)*, *abs/1409.5340*.

Liberatore, P. (2014b). Belief revision by reliability assessment. Manuscript.

Nayak, A. (1994). Iterated belief change based on epistemic entrenchment. *Erkenntnis*, *41*, 353–390.

Nebel, B. (1992). *Syntax-Based Approaches to Belief Revision*, pp. 52–88. Cambridge University Press.

Nittka, A., & Booth, R. (2008). A method for reasoning about other agents' beliefs from observations. In *Logic and the foundation of game and decision theory*, Vol. 3 of *Texts in logic and games*, pp. 153–182.

Papini, O. (2001). Iterated revision operations stemming from the history of an agent's observations. In *Frontiers in belief revision*, Vol. 22 of *Applied Logic Series*, pp. 279–301. Springer.

Peppas, P. (2008). *Belief revision*, pp. 317–359. Elsevier.

Rott, H. (2009). Shifting priorities: Simple representations for twenty-seven iterated theory change operators. In *Towards Mathematical Philosophy*, Vol. 28, pp. 269–296. Springer.

Sandholm, T., & Conen, W. (2010). Preference elicitation in combinatorial auctions. US Patent 7,742,971.

See, K., Morrison, W., Rothman, N., & Soll, J. (2011). The detrimental effects of power on confidence, advice taking, and accuracy. *Organizational Behavior and Human Decision Processes*, *116*(2), 272–285.

Spohn, W. (1988). Ordinal conditional functions: A dynamic theory of epistemic states. In *Causation in Decision, Belief Change, and Statistics*, pp. 105–134. Kluwer Academics.

Suzumura, K. (1976). Remarks on the theory of collective choice. *Economica, New Series*, *43*, 381–390.

SyncML (2002). SyncML sync protocol, version 1.1.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: the conjunction fallacy in probability judgment. *Psychological review*, *90*(4), 293–315.

Wang, H., Zhang, J., & Johnson, T. R. (2000). Human belief revision and order effect. In *Proceedings of the 22th Annual Conference of the Cognitive Science Society*.

Williams, M. (1994). Transmutations of knowledge systems. In *Proceedings of the Fourth International Conference on the Principles of Knowledge Representation and Reasoning (KR'94)*, pp. 619–629.

Winslett, M. (1988). Reasoning about actions using a possible models approach. In *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI'88)*, pp. 89–93.

Zhang, D. (2004). Properties of iterated multiple belief revision. In *Proceedings of the Seventh International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2004)*, pp. 314–325.