

Minimal preference change

Natasha Alechina¹, Fenrong Liu², and Brian Logan¹

¹ School of Computer Science, University of Nottingham

² Department of Philosophy, Tsinghua University, Beijing, China

nza@cs.nott.ac.uk, fenrong@tsinghua.edu.cn, bsl@cs.nott.ac.uk

Abstract. We propose a novel approach to preference change. We treat a set of preferences as a special kind of theory, and define minimal change contraction and revision operations in the spirit of minimal change as advocated by the Alchourron, Gärdenfors, and Makinson (AGM) theory of belief revision. We characterise minimal contraction of preference sets by a set of postulates and prove a representation theorem. We also give a linear time algorithm which implements minimal contraction by a single preference. We also define minimal contraction by a set of preferences, and for a significant special case state postulates, prove a representation theorem, and provide an efficient algorithm implementing minimal contraction by a set of preferences.

Key words: Preferences, preference change, belief revision

1 Introduction

Preference plays a crucial role in agents' reasoning and their intelligent interaction with other agents. In this paper we consider the problem of *preference change*: the contraction and revision of an agent's set of preferences by a single preference and by a set of preferences. We are motivated by an analogy between preference change and the Alchourrón, Gärdenfors, and Makinson (AGM) theory of belief revision. Both contraction and revision require maintaining consistency of the agent's set of preferences. When a new preference is inconsistent with an agent's existing preferences, a rational agent should remove as few preferences from its set of preferences as possible to restore consistency. We are interested in *efficient* algorithms (at most polynomial in the size of the agent's preference set) for minimal contraction and revision that may be employed by feasible, resource-bounded reasoners.

The contribution of this paper is as follows. We define minimal preference contraction and show how to define revision in terms of contraction. Next, we give postulates for rational minimal preference contraction and prove a representation theorem. As far as we know, this is the first representation theorem for minimal change preference contraction in the literature. We then give a linear time minimal preference contraction algorithm. We also investigate the problem of contracting by a *set* of preferences rather than by a single preference, or minimal iterative preference contraction. We define minimal iterative preference set contraction and state postulates characterising an important special case in which contraction is by an uncoupled set of preferences. We also give a polynomial time algorithm to compute minimal contraction by a set for that case.

2 Formal Preliminaries

We assume an agent's preferences are given by a binary relation over some finite set of alternatives \mathcal{A} . An agent's preference state is represented by a *preference set* consisting of preference sentences (or simply preferences) which are atomic statements of the form $A < B$ (B is preferred to A), $A \equiv B$ (A and B are equally preferred) or $A \# B$ (A and B are incomparable). In using these three basic relations, we essentially follow [1].³ In addition, the agent's preference set may contain a special sentence \perp , which is used to indicate a problem (derivability of an inconsistency).

We assume that preference sets are not necessarily complete, in the sense that they may include no sentences expressing a relation between A and B , for $A, B \in \mathcal{A}$. We take this to be a natural feature of resource-bounded agents. We do assume that the agents are rational, i.e., they don't accept $A < B$ and $B < A$ or $A \# B$ at the same time, they can complete their preference sets using transitivity of $<$ and \equiv and symmetry of $\#$, etc.

The agent's *rational reasoning rules* or integrity constraints in the sense of [7] are given below. Rule 1 states that $\#$ is symmetric, rules 2-4 state that \equiv is an equivalence relation, rule 5 states that $<$ is transitive, and the rest of the rules state that at most one of $\#, \equiv, <, >$ can hold between two alternatives⁴.

1. $A \# B \Rightarrow B \# A$
2. $A \equiv A$
3. $A \equiv B \Rightarrow B \equiv A$
4. $A \equiv B, B \equiv C \Rightarrow A \equiv C$
5. $A < B, B < C \Rightarrow A < C$
6. $A < B, B < A \Rightarrow \perp$
7. $A \equiv B, A < B \Rightarrow \perp$
8. $A \equiv B, A \# B \Rightarrow \perp$
9. $A \# B, A < B \Rightarrow \perp$

We denote by $Cn(S)$ the closure of a set S under the rules above. Formally, $Cn(S)$ is the set of preferences which contains S , $A \equiv A$ for every $A \in \mathcal{A}$, and in addition for every rule $p_1, \dots, p_n \Rightarrow p$ above, if $p_1, \dots, p_n \in Cn(S)$, then $p \in Cn(S)$. A set of preferences S is *deductively closed* iff $S = Cn(S)$.

Sometimes we will use the notation $S \vdash p$ to say that p can be derived from S and the reasoning rules above by application of the following inference rule (where $n \leq 2$):

³ One may start with a different initial setting, for instance, taking as $A \leq B$ (B is at least as good as A , cf. [8]) as the primitive relation, then define other relations and explore the similar questions. We leave this possibility for another occasion.

⁴ Note that we do not have the following rules:

$$\begin{aligned} A \equiv B, B < C &\Rightarrow A < C \\ A \equiv B, B \# C &\Rightarrow A \# C \end{aligned}$$

the agent may have a preference regarding B and C , and consider A and B indistinguishable, but may not have a preference regarding A and C .

$$\frac{p_1, \dots, p_n \quad p_1, \dots, p_n \Rightarrow p}{p}$$

Clearly for any p , $\vdash p$ (p is derivable from an empty set) if, and only if, p is of the form $A \equiv A$. Note that we do not assume any logical connectives or any other inference rules.

In what follows, we always assume that the agent's set of preferences S is deductively closed. The set of preferences is consistent if and only if it does not contain \perp .

3 Preference Revision

Clearly, if an agent acquires a new preference, its preference set may become inconsistent. For example, if the agent used to prefer B to A ($A < B$) and C to B ($B < C$) and has decided that it prefers A to C , its preference set is inconsistent since it contains both $A < C$ by transitivity from the old preferences and $C < A$ (the new preference). In order to incorporate the new preference and have a consistent preference set, the agent needs to remove some of the old preferences. We are interested in *minimal* preference revision, namely removing as few sentences as possible to restore consistency. As in AGM belief revision, we define revision in terms of contraction by a preference sentence.

Revision of a preference set S by a preference p is defined as adding p to S if the result is consistent and deductively closing the resulting set, otherwise first contracting S by p^{-S} , denoted $S - p^{-S}$, where p^{-S} is the S -complement of p . For a consistent set of preferences S and a preference p such that $S \cup \{p\}$ is inconsistent, the S -complement of p is defined as follows:

- $A \equiv B^{-S} = S \cap \{A < B, B < A, A \# B\}$
- $A < B^{-S} = S \cap \{A \equiv B, B < A, A \# B\}$
- $A \# B^{-S} = S \cap \{A \equiv B, A < B, B < A\}$

Contracting S by the S -complement of p makes p consistent with the result, and we can add p to the resulting set and close it under consequence. Revision $S * p$ of S by p is thus defined as $Cn(S - p^{-S} \cup \{p\})$. This is essentially the Levi identity [12, 6] $S * p = (S - \neg p) + p$.

3.1 Minimal Contraction

Definition 1. (*Minimal contraction*) Given a preference set S and a preference p , such that $\not\vdash p$, the result of a minimal contraction of S by p is a set $S - p$ such that:

- (1) $S - p \subseteq S$
- (2) $S - p \not\vdash p$
- (3) for any other set S' satisfying (1) and (2), $|S'| \leq |S - p|$.

The removal of minimal number of preferences is similar to Hansson's definition of revising to the most similar preference relation (where the distance between preference relations is defined as the symmetrical difference between the sets of pairs of alternatives in the two relations) [9].

3.2 Minimal Contraction Postulates

Before we can state the postulates characterising minimal contraction, we need to introduce the following abbreviations. By $A_S^<$ we will denote $\{C : A < C \in S\}$. By $A_S^>$ we will denote $\{C : C < A \in S\}$. By A_S^{\equiv} we will denote $\{C : A \equiv C \in S\} \setminus \{A\}$. The *cost* $c_S(p)$ of $p \in S$ (intuitively, the number of preferences a contraction by p has to remove from S) is defined as follows:

$$\begin{aligned} - c_S(A < B) &= |A_S^< \cap B_S^>| + 1 \\ - c_S(A \equiv B) &= 2 * |A_S^{\equiv}| \\ - c_S(A \# B) &= 2 \end{aligned}$$

The following postulates characterise minimal contraction. For readability, we will omit subscript S when it is unambiguous.

C-Closure $S - p = Cn(S - p)$

C-Inclusion $S - p \subseteq S$

C-Vacuity If $p \notin S$, $S - p = S$

C-Success If p is not of the form $A \equiv A$, then $p \notin S - p$

C-Equivalence If $Cn(p_1) = Cn(p_2)$, then $S - p_1 = S - p_2$

C-Minimality If $p \in S$, then $|S - p| = |S| - c_S(p)$

The postulates of C-Closure, C-Inclusion, C-Vacuity, C-Success and C-Equivalence are standard postulates for contraction of beliefs. Recovery ($S \subseteq Cn((S - p) \cup p)$) does not hold, but this postulate has always been considered controversial [13]. The C-Minimality postulates characterise specifically minimal contraction of preferences, because for preferences it is possible to predict the cardinality of the resulting set.

Theorem 1. *The result of any minimal contraction satisfies the minimal preference contraction postulates above, and every contraction satisfying these postulates is a minimal preference contraction.*

Proof. For the case when $p \notin S$, clearly the minimal contraction is S itself, and all the postulates hold for $S - p = S$ trivially.

Let us consider the case when $p \in S$. We show first that every minimal contraction satisfies the postulates. C-Inclusion holds by Definition 1, and C-Vacuity trivially since $p \in S$. To show that C-Closure holds, assume by contradiction that $S - p$ is a minimal contraction and it is not deductively closed. Since $S - p \not\vdash p$ (by Definition 1 (2)) and $S - p$ is not deductively closed, then there must be a consequence q of $S - p$ such that $q \notin S - p$. Since $S - p \not\vdash p$ and $S \vdash q$, it follows that $(S - p) \cup \{q\} \not\vdash p$. Since $S - p \subseteq S$ (by Definition 1 (1)), $S \vdash q$, and since S is deductively closed, $q \in S$. Hence there is a set $S' = (S - p) \cup \{q\}$ such that conditions (1) and (2) of Definition 1 hold for S' , and its cardinality is greater than that of $S - p$. Hence $S - p$ is not a minimal contraction because it violates condition (3): a contradiction. C-Success holds for all p which are not derivable from an empty preference set because there is always a subset of S which does not derive p (in the worst case, \emptyset). C-Equivalence holds rather trivially because the only cases when two syntactically different preferences have the same set of consequences are: $Cn(A \equiv B) = Cn(B \equiv A)$ and $Cn(A \# B) = Cn(B \# A)$;

due to symmetry rules, any successful contraction by one of $A \equiv B$, $B \equiv A$ has to get rid of both of them, similarly for $A \# B$, $B \# A$. Now let us consider the minimality postulates. We need to prove that any minimal contraction removes exactly $|S| - |S - p|$ sentences for each of the cases. In particular, we need to prove that:

- a minimal contraction by $A \equiv B$ removes exactly $2 * |A^{\equiv}|$ preferences;
- a minimal contraction by $A < B$ removes exactly $|A^{<} \cap B^{>}| + 1$ preferences;
- a minimal contraction by $A \# B$ removes exactly 2 preferences.

Let us consider the easiest case first. If $A \# B \in S$ and we want to remove it and make sure that $S \not\vdash A \# B$, we need to remove $A \# B$ itself, and $B \# A$ (note that since $A \# B \in S$ and S is deductively closed, $B \# A \in S$). Clearly if one of those preference is left in S then it would be possible to derive $A \# B$. So both $A \# B$ and $B \# A$ *have to* be removed. On the other hand, from the inspection of the reasoning rules, there is no other way to derive $A \# B$. So these two preferences are the *only* ones which have to be removed. Hence any contraction satisfying (2) will remove these 2 sentences, and any contraction satisfying (3) will only remove these 2 sentences.

Now consider the case of $A < B \in S$. In order to contract by $A < B$, we need to remove $A < B$ itself from S . However $A < B$ may still be derivable, namely using the transitivity rules. The number of possible derivations of $A < B$ using the rule $A < C, C < B \Rightarrow A < B$ is exactly $|A^{<} \cap B^{>}|$. We need to ‘destroy’ each such derivation, and in order to do this we need to remove *at least one* of the premises in each derivation, namely either $A < C$ or $C < B$. So any contraction satisfying (1) and (2) needs to remove at least $|A^{<} \cap B^{>}| + 1$ preferences (1 is for $A < B$ itself). Conversely, if one of the preferences for each possible derivation is removed, then $A < B$ is no longer derivable, so the operation already satisfies (1) and (2). (Note that if $A < C$ for $C \in A^{<} \cap B^{>}$ is itself derivable, one premise in the derivation of $A < C$ is $A < D$ where $D < C$ since $C < B, D < C$, so $D \in A^{<} \cap B^{>}$, so $A < D$ will be removed and hence $A < C$ is not re-derivable.) Hence, in order to satisfy (3), the operation should not remove anything else. Hence any minimal contraction removes exactly $|A^{<} \cap B^{>}| + 1$ preferences.

In the case when $A \equiv B \in S$, any contraction operation needs to remove $A \equiv B$ and $B \equiv A$. However after this $A \equiv B$ may still be derivable by transitivity, using $A \equiv C, C \equiv B \Rightarrow A \equiv B$. The number of such derivations is the number of elements in $A^{\equiv} \setminus \{B\}$ (we are only considering uses of transitivity rule where C is different from both A and B). If for some of those derivations, both premises are left in S , then $A \equiv B$ can be re-derived. So any contraction satisfying (1) and (2) needs to remove at least one of the premises, either $A \equiv C$ or $C \equiv B$. Note that in order to properly remove $A \equiv C$, we also need to remove $C \equiv A$, otherwise $A \equiv C$ will be rederivable by symmetry. This means that any contraction needs to remove at least $2 * |A^{\equiv}|$ preferences: $A \equiv B, B \equiv A$, and $2 * (|A^{\equiv} \setminus \{B\}|)$. To show that this number of removed preferences is sufficient, and hence that no minimal contraction needs to remove more, we exhibit a concrete contraction which satisfies (1) and (2) and removes only $2 * |A^{\equiv}|$ preferences. Namely, consider a contraction which removes A from its equivalence class in S : it removes all $A \equiv C, C \equiv A$ for $C \in A^{\equiv}$. In the resulting set, A is not connected by \equiv to any other alternative, hence $A \equiv B$ is not derivable.

The other direction: if an operation satisfies the postulates, it is a minimal contraction. Clearly, since the operation satisfies C-Closure, C-Inclusion and C-Success, it satisfies conditions (1)-(2) of Definition 1. To show that it satisfies (3), we need to prove that there is no set of strictly larger cardinality than $S - p$ which still satisfies (1)-(2), in other words that every successful contraction has to remove at least as many preferences as is stated in C-Minimality postulates. The argument is exactly as above. \square

3.3 Minimal Contraction Algorithm

We give an algorithm for the case when p is derivable from S , and $\not\vdash p$.

The algorithm for computing $S - p$ is given by cases (see Algorithm 1).

Algorithm 1 Minimal preference contraction algorithm

```

procedure MINIMAL-CONTRACTION( $S, p$ )
  case  $p \notin S$ 
    return
  case  $p == A < B$ 
     $A^< := \{C \mid A < C\}$ 
     $B^> := \{C \mid C < B\}$ 
    for each  $C \in A^< \cap B^>$  do
       $S := S \setminus \{A < C\}$ 
    end for
     $S := S \setminus \{A < B\}$ 
  case  $p == A \equiv B$ 
     $A^\equiv := \{C \mid A \equiv C, C \neq A\}$ 
    for each  $C \in A^\equiv$  do
       $S := S \setminus \{A \equiv C, C \equiv A\}$ 
    end for
  case  $p == A \# B$ 
     $S := S \setminus \{A \# B, B \# A\}$ 

```

Theorem 2. *The minimal preference contraction algorithm computes a minimal preference contraction.*

Proof. We show that the result of applying the algorithm to a preference set S and $p \in S$, p not of the form $A \equiv A$, always satisfies the conditions in Definition 1. Condition (1) holds because the algorithm only removes sentences from S . Condition (2) holds because the algorithm removes a premise from every possible derivation of p . Condition (3) holds because the algorithm result satisfies the minimal contraction postulates hence it is a minimal contraction by Theorem 1. \square

Theorem 3. *The time complexity of the algorithm for minimal contraction is in $O(|\mathcal{A}|)$.*

Proof. We assume that we can order the alternatives in some order (e.g., lexicographic order) and for each relation ($<$, \equiv , $\#$) we can recover the ordered set of alternatives to which an alternative A is related in constant time (e.g., a hash table for each relation/position mapping from alternatives to sets (lists) of alternatives).

Then we can determine in constant time whether $p \notin S$ (recall that S is deductively closed).

For the $A < B$ case, the maximum size of $A^<$ and $B^>$ is bounded by $|\mathcal{A}|$, since A and B can be related to at most $|\mathcal{A}| - 1$ alternatives by $<$. Computing the set of alternatives $C \in A^< \cap B^>$ is also linear in $|\mathcal{A}|$ (to be precise it requires at most $2|\mathcal{A}|$) and the number of such alternatives C is again bounded by $|\mathcal{A}|$. Removing the preferences $A < C$ for $C \in A^< \cap B^>$ requires at most $|\mathcal{A}|$ operations (if the set of preferences is implemented as, e.g., a linked list) and replacing the new set in the map is constant time. For the $A \equiv B$ case, replacing the entry for A in the \equiv map is a constant time operation. For the $A \# B$ case, we need to remove a single entry from the set of preferences for A in the $\#$ map. This requires at most $|\mathcal{A}|$ steps. \square

4 Minimal Set Contraction

In this section we turn to the problem of contracting by a *set* of preferences, which is similar to the problem of iterated belief revision [3–5]. As in the case of single preferences, we concentrate on the contraction rather than revision by a set of preferences, since ‘minimal change’ has a more intuitive and straightforward interpretation in the case of contraction.

We define a minimal contraction of a preference set S by a set of preference sentences X as follows:

Definition 2. (*Minimal contraction by a set*) An operation $-$ is a minimal contraction of S by a set X if it satisfies the following properties:

1. $S - X \subseteq S$
2. if $p \in X$ and $\not\vdash p$, then $S - X \not\vdash p$
3. for every other set S' which satisfies properties (1)-(2) above, $|S'| \leq |S - X|$.

A minimal revision of a preference set S by a set of preferences S' can be defined analogously to Hansson’s *consolidation* [10]: first compute $Cn(S \cup S')$, then minimally contract by contradictions. Note that contracting S' by all sentences X inconsistent with S may not be enough to make \perp underivable from $S \cup (S' \setminus X)$.

A natural question to ask is whether a minimal contraction of S by p_1 followed by a minimal contraction of $S - p_1$ by p_2 is a minimal contraction of S by $\{p_1, p_2\}$. The answer is negative. Consider the following example:

- $S = \{A < B, A < C, C < B\} \cup \{A \equiv A : A \in \mathcal{A}\}$
- $p_1 = A < B$
- $p_2 = C < B$

A minimal contraction of S by $A < B$ computed by Algorithm 1 is $S - A < B = \{C < B\} \cup \{A \equiv A : A \in \mathcal{A}\}$. It removes two preferences, $A < B$ itself and $A < C$.

A minimal contraction of this set by $C < B$ removes $C < B$. The set $(S - p_1) - p_2$ is $\{A \equiv A : A \in \mathcal{A}\}$ which is the result of removing three preferences from S . However, it is possible to make $A < B$ and $C < B$ undervivable from S by removing just two preferences: $A < B$ and $C < B$. Recall that Algorithm 1 makes a particular choice in contraction by $A < B$: it removes sentences of the form $A < C$ where $C \in A^< \cap B^>$. It could have just as well removed sentences of the form $C < B$; for a single step contraction it does not matter which choice is made, since the number of removed sentences would be the same in each case. However for the iterated case, we need to look ahead to decide which choice to make. The problem of computing a minimal contraction by a set is, of course, decidable, but may require considering exponentially many (in $|X|$) choices. The problem only arises when several preferences in X share alternatives, as $A < B$ and $C < B$ above.

We can however characterise minimal set contraction in an important special case, when X has a specific form (which we call uncoupled) defined below. This special case covers, for example, contraction by a set of (some other) agent's preferences when that agent has a linearly ordered preference set. Given a set X , we will denote by $X(<)$ all elements of X which are of the form $A < B$, by $X(\equiv)$ all elements of X which are of the form $A \equiv B$ (we assume X does not contain tautologies $A \equiv A$), and by $X(\#)$ all elements of X of the form $A \# B$. The set X can therefore be represented as a set of disjoint sets X_1, \dots, X_k .

Definition 3. (*Uncoupled set of preferences*)

- the subset $X(<)$ of X is partitioned into subsets $X(A, <)$ of the form

$$\{A < A_1, A_1 < A_2, A < A_2, \dots, A < A_n\}$$

(where all A_i are linearly ordered between A and A_n) and $X(B, >)$ of the form

$$\{B_1 < B, B_1 < B_2, B_2 < B, \dots, B_n < B\}$$

(where all B_i are linearly ordered between B_1 and B) and no alternative occurs in two different partitions of $X(<)$

- the subset $X(\equiv)$ of X is partitioned into subsets $X(A, \equiv)$ of the form

$$\{A \equiv A_1, A_1 \equiv A_2, \dots, A_{n-1} \equiv A_n\}$$

(or an equivalent way of stating that A, A_1, \dots, A_n form an equivalence class) and no alternative occurs in two different partitions of $X(\equiv)$

- $X(\#)$ is partitioned into two parts, $X(\#)_1$ which contains $A \# B$ such that $B \# A \notin X$, and $X(\#)_2$ which contains $A \# B, B \# A$ such that $A \# B, B \# A \in X$.

Note that the last condition is not a restriction on X , just a notational convenience for the postulates below.

4.1 Minimal Set Contraction Postulates

It is possible to provide a representation theorem and an efficient algorithm for the case of contraction of S by an uncoupled set of preferences X . Essentially, minimal

contraction by an uncoupled set X can be reduced to an unordered set of minimal contractions by single sentences, where those single sentences correspond to partitions of X .

The following postulates characterise a minimal contraction of S by an uncoupled $X \subseteq S$.

CX-Closure $S - X = Cn(S - X)$

CX-Inclusion $S - X \subseteq S$

CX-Vacuity If $X \cap S = \emptyset$, $S - X = S$

CX-Success If $p \in X$ is not of the form $A \equiv A$, then $p \notin S - X$

CX-Minimality If $X \subseteq S$, and X is uncoupled, then $|S - X| = |S| - \sum_i c_s(X_i)$ where the costs of contracting by each X_i are defined as follows:

- $c_S(X(A, <)) = c_S(A < A_n)$
- $c_S(X(B, >)) = c_S(B_1 < B)$
- $c_S(X(A, \equiv)) = (|A_{\bar{X}}| + 1) * |A_{\bar{S}}|$, where by $A_{\bar{X}}$ we denote the set of alternatives occurring in $X(A, \equiv)$
- $c_S(X(\#)) = 2 * |X(\#)_1| + |X(\#)_2|$

Theorem 4. *The result of any minimal set contraction by an uncoupled set of preferences satisfies the minimal set contraction postulates above, and every contraction by an uncoupled set satisfying these postulates is a minimal set contraction.*

Proof. Let us prove that every minimal set contraction satisfies the postulates. The proof for CX-Closure, CX-Inclusion, CX-Vacuity, CX-Success is very similar to Theorem 1. For CX-Minimality, observe that since the partitions X_i do not share alternatives, the sets of sentences which have to be removed to contract by each X_i are disjoint. Note that

- for each $X(A, <)$, it is sufficient and necessary to remove $\{A < C : C \in A^< \cap A_n^>\}$ to make all sentences in $X(A, <)$ undervivable
- for each $X(B, >)$, it is sufficient necessary to remove $\{B < C : C \in B_1^< \cap B^>\}$ to make all sentences in $X(B, >)$ undervivable
- for each $X(A, \equiv)$, it is sufficient and necessary to remove connections between alternatives occurring in $X(A, \equiv)$ and other members of the equivalence class of A in S , so assuming that $|A_{\bar{S}}| = m$ and $X(A, \equiv)$ contains occurrences of alternatives A, A_1, \dots, A_n , then we need to remove $(n + 1) * m$ sentences ($2 * m$ for removing sentences connecting A to the equivalence set, $2 * (m - 1)$ for removing sentences connecting $A_1, \dots, 2 * (m - n)$ for removing sentences connecting A_n).
- for the whole of $X(\#)$, we need to remove $2 * |X(\#)_1|$ and $|X(\#)_2|$.

For the other direction, assume an operation satisfies the postulates for minimal set contraction. Then it clearly satisfies (1) and (2) of Definition 2. It also satisfies (3), since any other contraction by X has to remove at least as many preferences. \square

A postulate corresponding to C-Equivalence: if $Cn(X_1) = Cn(X_2)$, then $S - X_1 = S - X_2$ does not hold. For example, let $X_1 = \{A < B, B < C, A < C\}$ $X_2 = \{A < B, B < C\}$, and $S = \{A < C, A \equiv A, B \equiv B, C \equiv C\}$. Clearly $Cn(X_1) = Cn(X_2)$. However, $S - X_1 = S \setminus \{A < C\}$ and $S - X_2 = S$.

4.2 Minimal Set Contraction Algorithm

We can also give a concrete polynomial time algorithm for contraction by an uncoupled set of preferences.

Algorithm 2 Minimal preference set contraction algorithm

```

procedure MINIMAL-SET-CONTRACTION( $S, X$ )
  for each  $X_i \subseteq X$  do
    case  $X_i == X(A, <)$ 
       $A^< := \{C \mid A < C\}$ 
       $A_n^> := \{C \mid C < A_n\}$ 
      for each  $C \in A^< \cap A_n^>$  do
         $S := S \setminus \{A < C\}$ 
      end for
       $S := S \setminus \{A < A_n\}$ 
    case  $X_i == X(B, >)$ 
       $B_1^< := \{C \mid B_1 < C\}$ 
       $B^> := \{C \mid C < B\}$ 
      for each  $C \in B_1^< \cap B^>$  do
         $S := S \setminus \{C < B\}$ 
      end for
       $S := S \setminus \{B_1 < B\}$ 
    case  $X_i == X(A, \equiv)$ 
       $A^\equiv := \{C \mid A \equiv C\}$ 
       $A^{X,\equiv} := \{D \mid D \text{ occurs in } X(A, \equiv)\}$ 
      for each  $D \in A^{X,\equiv} \cup \{A\}$  do
        for each  $C \in A^\equiv \setminus \{D\}$  do
           $S := S \setminus \{D \equiv C, C \equiv D\}$ 
        end for
      end for
    case  $X_i == X(\#)$ 
      for each  $A\#B \in X(\#)$  do
         $S := S \setminus \{A\#B, B\#A\}$ 
      end for
  end for

```

The algorithm contracts by each $X_i \subseteq X$ in turn; since X_i are disjoint, in the worst case there are $|X|$ members of the partition. Each contraction by X_i is linear in $|\mathcal{A}|$, by an argument similar to the proof of Theorem 3. This means that the time complexity of Algorithm 2 is $O(|X| \times |\mathcal{A}|)$.

5 Related work

In this section we compare our results to those works which are most closely related to our own, focusing on the main ideas rather than providing a full-fledged comparison, and highlighting some future research directions.

[9] describes four types of preference change: contraction and revision of preference relations, and addition and subtraction of alternatives. We do not consider changes in alternatives in our framework, thus we compare with the first two kinds. Hansson defines contraction in terms of revision with the intuition that “to contract your state of preference by α means to open it up for the possibility that $\neg\alpha$ ” and gives postulates for this operation. To define a *minimal* preference revision operator, Hansson introduces a measure of similarity between preference relations. This involves a calculation of the symmetric difference between two sets X and Y ($X\Delta Y$), which is equal to $(X\setminus Y)\cup(Y\setminus X)$. The result of the preference change is a preference relation that has as small a distance from the original relation as possible. This idea inspired our notion of minimal contraction. Since Hansson considers a full logical language with negations, disjunctions etc. of preferences, the complexity of his operations is clearly much higher than ours. [7] discuss logical constraints on preference — formal requirements that a preference state has to satisfy. These are called *reasoning rules* in our framework. A further distinction between logical constraints, input constraints that come with a specific input, and priorities has been made in the same discussion, and various ways of formalizing those aspects in logical models are proposed. In our work, we consider reasoning involving merely logical constraints. It would be interesting to study how to modify our algorithms to incorporate other kinds of constraints.

There exists considerable work on iterated belief revision, see for example [3–5, 11]. We consider set contraction in this paper. However our focus is a special case of contraction by an uncoupled set of preferences. [2] point out potential connections between this area and preference aggregation as they study revision of a total preorder in the context of iterated belief revision. We would like to extend our result to preference aggregation in the future.

6 Conclusion

In this paper, we introduce a simple setting of preference change where it is possible to define minimal preference contraction (revision). We propose rationality postulates and an efficient algorithm for that setting. Then we study contraction by a set of preferences and provide a characterisation and an efficient algorithm for the case where the set of preferences is uncoupled. Finally, we compare our work with some related work and highlight some directions for future research.

References

1. Hajnal Andr eka, Mark Ryan, and Pierre-Yves Schobbens. Operators and laws for combining preference relations. *Journal of Logic and Computation*, 12(1):13–53, 2002.
2. Richard Booth and Thomas Meyer. How to revise a total preorder. *Journal of Philosophical Logic*, 40:193–238, 2011.
3. Craig Boutilier. Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic*, 25(3):263–305, 1996.
4. Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.

5. James P. Delgrande, Didier Dubois, and Jérôme Lang. Iterated revision as prioritized merging. In Patrick Doherty, John Mylopoulos, and Christopher A. Welty, editors, *Proceedings, Tenth International Conference on Principles of Knowledge Representation and Reasoning, Lake District of the United Kingdom, June 2-5, 2006*, pages 210–220. AAAI Press, 2006.
6. Peter Gärdenfors. *Knowledge In Flux: Modeling the Dynamics of Epistemic States*. MIT Press, 1988.
7. Till Grüne-Yanoff and Sven Ove Hansson. From belief revision to preference change. In *Preference Change: Approaches from Philosophy, Economics and Psychology*, pages 159–184. Springer, 2009.
8. Meiyun Guo and Jeremy Seligman. Making choices in social situations. In *Logic and Interactive Rationality: Yearbook 2011*, pages 176–202. ILLC, University of Amsterdam, 2012.
9. Sven Ove Hansson. Changes in preference. *Theory and Decision*, 38:1–28, 1995.
10. Sven Ove Hansson. Semi-revision (invited paper). *Journal of Applied Non-Classical Logics*, 7(2), 1997.
11. Sven Ove Hansson. Multiple and iterated contraction reduced to single-step single-sentence contraction. *Synthese*, 173(2):153–177, 2010.
12. Isaac Levi. Subjunctives, dispositions and chances. *Synthese*, 34:423–455, 1977.
13. David Makinson. On the status of the postulate of recovery in the logic of theory change. *Journal of Philosophical Logic*, 16:383–394, 1987.

Acknowledgement The authors thank the three anonymous referees for their useful comments. Fenrong Liu is supported by the Project of National Social Science Foundation of China (NO.13AZX018), and Tsinghua University Project (NO. 2012WHYX003).