

Efficient minimal preference change

Natasha Alechina* Fenrong Liu[†] Brian Logan*

Abstract

In this paper, we study a minimal change approach to preference dynamics. We treat a set of preferences as a special kind of theory, and define minimal change preference contraction and revision operations in the spirit of the Alchourrón, Gärdenfors, and Makinson 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 then define minimal contraction by a set of preferences, and show that the problem of a minimal contraction by a set of preferences is NP-hard.

1 Introduction

Preferences are central both to individual decision making, and to strategic interactions between rational agents. They have been studied in many research fields, including philosophy, AI, and social choice theory. Since the 1990s, the dynamics of preference change has become a major focus of research. For example, Hansson [18] has proposed postulates for several preference change operations in the spirit of the Alchourrón, Gärdenfors and Makinson (AGM) theory of belief revision. Van Benthem & Liu [6] have proposed a dynamic epistemic logic (DEL)-based dynamic preference logic to model the changes before and after some informational event takes place. A comprehensive collection of research papers in preference change from different fields can be found in [16]. In this paper, we also focus on the of dynamics of preference, but we take more computationally oriented approach. As in [18], we are inspired by an analogy between preference change and the AGM theory of belief revision, however we seek *efficient* algorithms (at most polynomial in the size of the agent's preference set) for minimal preference change which may be employed by feasible, resource-bounded agents.

*School of Computer Science, University of Nottingham, UK, {nza, bs1}@cs.nott.ac.uk

[†]Department of Philosophy, Tsinghua University, Beijing, China, fenrong@tsinghua.edu.cn

We focus on two cases: the contraction and revision of an agent's set of preferences by a single preference, and by a set of preferences. In *preference contraction*, the agent wishes to remove one or more preferences from its set of preferences, together with any preference(s) from which the target preference(s) can be derived by transitivity. A *minimal preference contraction* is a contraction that removes the smallest number of preferences necessary to make the target preference(s) undervivable. Contraction may be triggered by a change in the agent's information about a domain leading to a change in its preferences, for example if it no longer has any reason to prefer one alternative to another, or as a result of the agent revising its preferences, that is, when the agent acquires a new preference that is inconsistent with its existing preferences and must remove one or more preferences to restore consistency. A minimal revision is a revision that removes the smallest number of preferences necessary to restore consistency.

As an example of preference revision, consider the following scenario:

Example 1 *When choosing between the fruits apple (A), orange (O), pear (P) and banana (B), Alice initially prefers oranges to bananas ($B < O$), apples to oranges ($O < A$), pears to bananas ($B < P$), apples to pears ($P < A$) and pears to oranges ($O < P$). In addition, we assume that Alice is rational and accepts transitivity as a property of her preferences. For instance, she prefers apples to bananas ($B < A$).*

Now Alice receives new information that the health benefits of bananas far outweigh those of apples, so she has a new preference $A < B$ which is inconsistent with her current preferences.

To restore consistency, Alice must contract by her existing preference $B < A$. There are several ways in which this could be done. For example, she could remove $O < A$, $B < P$ and $O < P$. However such a contraction is not minimal. Alice can make $B < A$ undervivable by removing only two preference statements: either $B < O$ and $B < P$ in which B lies on the left side, or alternatively the preference statements $O < A$ and $P < A$ in which A lies on the right side. We will show below that these are minimal contractions. We also give an efficient algorithm for minimal contraction based on this idea.

We follow the tradition of representing preferences as binary relations between two alternatives. In other words, we interpret preference in a qualitative sense rather than in terms of utility. In addition, we assume the preference relation is closed with respect to some intuitive properties, for example, transitivity for strict preference. Most notably, we do not assume completeness or connectedness of preference relation. This means, for two alternatives A and B , there could be no preference relation between them. We interpret minimal change in the most straightforward sense of 'minimal': as the minimal cardinality change to the preference relation.

This paper extends [1] where we defined minimal preference contraction and showed how to define revision in terms of contraction. We gave postulates for rational minimal preference contraction, proved a representation theorem, and gave a linear time minimal preference contraction algorithm. To the best of our knowledge, [1] contained the first representation theorem for minimal change preference contraction in the literature. We also investigated the problem of contracting by a set of preferences, and gave a polynomial time algorithm for a special case of that problem where the set of preferences to be removed are uncoupled (essentially, when the order in which contractions by elements of the set are performed does not matter). In this paper, we prove a new result that the problem of minimal contraction by a set of preferences is NP-hard. This means that for the general case, it is impossible to do minimal contraction by a set of preferences in polynomial time (unless $P = NP$). We also consider the problem of minimal preference aggregation, and discuss the implications of adopting a different primitive preference relation. The latter discussion is based on the results in the informal proceedings of LAMAS 2014 [2].

The paper is organised as follows. In Section 2 we introduce preference relations and preference sets. In Section 3 we consider the problem of minimal preference contraction, giving a set of postulates and an efficient algorithm, and in Section 4 we briefly discuss the notion of minimal preference revision. In Section 5 we turn to the problem of minimal contraction by a set of preferences, and show that it is NP-hard. We can however characterise a special case of contraction by a set of preferences, and we give postulates and a polynomial algorithm for this special case. In Section 6 we discuss preference aggregation, and in Section 7 we discuss an alternative choice for the primitive preference relation (\leq). We briefly survey related work in Section 8, and conclude in Section 9.

2 Formal Preliminaries

We assume that an agent's preferences are given by a set of binary relations 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 involving preference relations. Here we assume that we have a set of preference relations (and corresponding connectives in atomic sentences). The relations are $<$ (where $A < B$ means that B is strictly preferred to A), \equiv (where $A \equiv B$ means that A and B are equally preferred), and $\#$ (where $A \# B$ means that A and B are incomparable). These basic relations were taken as primitive in, for example, [3].

In addition to atomic sentences built using preference relations, a preference set

may contain a special sentence \perp , which is used to indicate a problem (derivability of an inconsistency).

We assume that a preference set may be incomplete (it is possible that no relation holds between some alternatives A and B), but that the agents are rational, i.e., they don't accept $A < B$ and $B < A$ or $A \# B$ at the same time, and that they can complete their preference sets using transitivity of $<$ and \equiv and symmetry of $\#$, etc. (so the sets are deductively closed with respect to the corresponding rules).

We postulate the following natural set of *rational reasoning rules* or integrity constraints in the sense of [16] for preference relations. Rule 1 states that $\#$ is symmetric, rules 2–4 state that \equiv is an equivalence relation, rule 5 states that $<$ is transitive, and the remaining rules state that at most one of $\#, \equiv, <, >$ can hold between two alternatives.¹

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

¹Rules of the following form (i.e., which mix two different connectives):

$$A \equiv B, B < C \Rightarrow A < C$$

$$A \equiv B, B \# C \Rightarrow A \# C$$

are excluded for two reasons. First, we do not have an efficient minimal contraction algorithm for the ‘mixed’ case. Second, one might consider such rules to be too strong as a closure condition. If the agent uses different criteria to form preferences, it may have a preference regarding B and C , and consider A and B indistinguishable, but may not have a preference regarding A and C . For example, an agent may consider ice cream and sorbet to be equally nice as deserts, and prefer fruit salad to ice cream on health grounds, but not have any opinion or information on the relative healthiness of sorbet and fruit salad.

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 $S = Cn(S)$. Note that $Cn(S)$ is a closure under logical consequence, but in a very weak logic (weaker than the classical logic; in particular, $Cn(S)$ is not closed under modus tollens).

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 $n \leq 2$):

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

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

In what follows, we 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 Minimal Contraction

In this section, we introduce the operation of minimal contraction of a preference set by a single preference.

Definition 1 (*Minimal contraction*) Given a preference set S and a preference p , such that $\not\vdash p$, a minimal contraction of S by p is any operation – that returns a set $S - p$ such that:

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

3.1 Minimal Contraction Postulates

Before stating the postulates characterising minimal contraction, we introduce the following abbreviations: by $A_S^<$ we denote $\{C : A < C \in S\}$; by $A_S^>$ we denote $\{C : C < A \in S\}$; and by A_S^{\equiv} we 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:

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

- $c_S(A\#B) = 2$

We briefly outline the reasons why cost is defined in this way (for full details see the proof of Theorem 1). To start with the simplest case, in order to make $A\#B$ underivable, we need to remove $A\#B$ and $B\#A$ from S , hence $c_S(A\#B) = 2$. In order to make $A < B$ underivable, we need to remove one of the premises in each possible application of the transitivity rule $A < C, C < B \Rightarrow A < B$. The minimal number of premises to remove is equal to the number of such transitivity rule applications, which is equal to $|\{C : C \in A_S^< \cap B_S^>\}|$. Similarly, to remove $A \equiv B$ we need to remove $A \equiv C$ and $C \equiv A$ for all C in the equivalence class of A and B .

The following postulates characterise minimal contraction.

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. 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 C-Minimality. 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 $A < B$ removes exactly $|A^< \cap B^>| + 1$ preferences;
- a minimal contraction by $A \equiv B$ removes exactly $2 * |A^{\equiv}|$ 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 using the transitivity rule. 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 all premises of the form $A < C$ or all premises of the form $C < B$. 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 once we

destroyed all one step derivations by removing a premise, there is no possibility of that premise being re-derivable. Consider $A < C$ for $C \in A^{< \cap B >}$. If it is derivable, then one premise in the derivation of $A < C$ is $A < D$ with $D < C$ (because $C < B$ and $D < C$), so $D \in A^{< \cap B >}$. But then $A < D$ has also been 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 the 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

3.2 Minimal Contraction Algorithm

We give an algorithm for the minimal contraction of S by a preference p such that $\not\vdash p$.

The algorithm for computing $S - p$ is given by cases (see Algorithm 1). Note that the first case $p \notin S$ is not strictly necessary since if $p \notin S$ then the set which the algorithm removes from S is empty, hence the result of contracting S by p is S .

Algorithm 1 Minimal preference contraction algorithm

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

Theorem 2 *Algorithm 1 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 set returned by the algorithm 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 implemented as an array of length $|\mathcal{A}|$ mapping from alternatives to sets (lists) of alternatives). Then we can determine in time linear in $|\mathcal{A}|$ 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 $2|\mathcal{A}|$ operations) 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 $|\mathcal{A}|$ operations. The running time of the algorithm is $O(|\mathcal{A}|)$ hence also linear in the size of the preference set. \square

4 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, that is, we wish to remove as few sentences as possible to restore consistency. As in AGM belief revision, we define revision in terms of contraction by a preference sentence. Before we do this, we need to define the notion of an S -complement of p (intuitively, a preference in S which together with p derives \perp):

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

Note that if S does not contain \perp , and $S \cup \{p\}$ is inconsistent, then p^{-S} contains a single preference. We will abuse notation slightly and use p^{-S} to refer to this preference.

Revision of a preference set S by a preference p , $S * p$ is defined as follows:

$$S * p = \begin{cases} Cn(S \cup \{p\}) & \text{if } S \cup \{p\} \not\vdash \perp \\ Cn((S - p^{-S}) \cup \{p\}) & \text{otherwise} \end{cases}$$

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. This is essentially the Levi identity [26, 14]: $S * p = (S - \neg p) + p$.

5 Minimal Contraction by a Set of Preferences

In this section we turn to the problem of contracting by a set of preferences. As in the case of single preferences, we concentrate on 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 any other set S' such that $S' \subseteq S$ and $S' \not\vdash p$ for any $p \in X$ where $\not\vdash p$, it holds that $|S'| \leq |S - X|$.

A minimal revision of a preference set S by a set of preferences S' can be defined analogously to Hansson’s *consolidation* [19]: first compute $Cn(S \cup S')$, then minimally contract by contradictions (remove enough sentences to make \perp underivable). 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:

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

A minimal contraction of S by $A < B$ computed by Algorithm 1 is $S - A < B = \{C < B\} \cup \{A \equiv A \mid A \in \mathcal{A}\}$. 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 $\{A \equiv A \mid A \in \mathcal{A}\}$ which is the result of removing three preferences from S . However, it is possible to make $A < B$ and $C < B$ underivable 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.

Here we show that minimal contraction by a set of preferences which only involve the \equiv relation is already an NP-hard problem. The decision version of the problem is as follows: given a preference set S , a set of preferences X , and a number k , is there a set $S' \subseteq S$ such that $|S \setminus S'| \leq k$ and S' does not derive any preferences in X ?

Theorem 4 *The problem of minimal contraction by a set of preferences is NP-hard, even if all preferences are of the form $A \equiv B$.*

Proof. The proof is by reduction from a known NP-complete problem: minimal edge multicut. The problem is as follows. Given a graph $G = (V, E)$ and a set of pairs of vertices $H = \{(s_1, t_1), \dots, (s_n, t_n)\}$, remove a minimal number of edges from G so that there is no path between s_i and t_i for all pairs (s_i, t_i) in H . (The decision version is: given G , H and k , is there a set of edges of cardinality $\leq k$ such that removing this set from G makes all pairs in H disconnected.) This problem is NP-complete even if G is a clique, that is of the form (V, V^2) [23]. The reduction to the problem of minimal contraction by a set of preferences is as follows. Let $G = (V, V^2)$, and $H = \{(s_1, t_1), \dots, (s_n, t_n)\}$. Let $\mathcal{A} = V$ and $S = \{s \equiv s' \mid s, s' \in V\}$. The set X of preferences that we are going to minimally contract S by is $X = \{s_i \equiv t_i \mid (s_i, t_i) \in H\}$. Let $X' = S \setminus (S - X)$. We need to show that $H' = \{(s, s') \mid s \equiv s' \in X'\}$ is a minimal multicut of G . It is clear that with edges in H' removed from G , there is no path between any of the pairs (s_i, t_i) in H (otherwise there would have been a corresponding derivation by transitivity in $S - X$). It is also minimal. Assume there is a smaller set of edges that can be removed from G so that there is no path between all pairs in H . Then the corresponding set of preferences could have been removed from S and destroy all derivations by transitivity of preferences in X . \square

5.1 Minimal Contraction by an Uncoupled Set of Preferences

However, we can perform efficient minimal contraction by a set of preferences and characterise minimal set contraction in an important special case where 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 , where each X_i is an element in the partition of $X(<)$, $X(\equiv)$, or $X(\#)$ as explained below.

Definition 3 (*Uncoupled set of preferences*) *A set X is an uncoupled set of preferences if, and only if, it satisfies the following properties:*

- the subset $X(<)$ of X is partitioned into subsets $X(A_1, A_n, <)$ of the form

$$\{A_1 < A_2, A_2 < A_3, A_1 < A_3, \dots, A_1 < A_n\}$$

(where all A_i are linearly ordered between A_1 and A_n) 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

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

and no alternative occurs in two different partitions of $X(\equiv)$

As a notational convenience for the postulates below, $X(\#)$ 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 $A \# B, B \# A \in X$.

5.2 Postulates for Minimal Contraction by an Uncoupled Set

We provide a representation theorem and an efficient algorithm for contraction of S by an uncoupled set of preferences X . Essentially, minimal contraction by an uncoupled set X can be reduced to minimal contractions by single sentences corresponding to an element of the partition of X . Since these single sentences represent disjoint subsets of X , and never require the removal of the same preferences from S , the order of these single preference contractions does not matter.

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

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 $|S - X| = |S| - \sum_i c_S(X_i)$ where the costs of contracting by each X_i are defined as follows:

- $c_S(X(A_1, A_n, <)) = c_S(A_1 < A_n)$
- $c_S(X(A, \equiv)) = (|A_{\bar{X}}| + 1) * |A_{\bar{S}}|$, 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 5 *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. We first show 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 $X(A_1, A_n, <)$, it is sufficient and necessary to remove $\{A_1 < C \mid C \in A_1^< \cap A_n^>\}$ to make all sentences in $X(A_1, A_n, <)$ underivable.
- 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 $|A^{\equiv}| = m$ 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

Note that 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 $X_1 = \{A \equiv B, A \equiv C, B \equiv C\}$, $X_2 = \{A \equiv B, A \equiv C\}$, and $S = \{A \equiv B, A \equiv C, B \equiv C, A \equiv A, B \equiv B, C \equiv C\}$. Clearly $Cn(X_1) = Cn(X_2) = S$. However, $S - X_1 = \{A \equiv A, B \equiv B, C \equiv C\}$, and $S - X_2 = \{B \equiv C, A \equiv A, B \equiv B, C \equiv C\}$. This is because contracting by X is not the same as contracting by X and all its (non-trivial) consequences: some consequences of X that already belong to S may

remain after we make every single sentence of X underivable from S . So if the closures of two sets are the same, they may differ in whether they contain any non-trivial consequences of themselves, and the results of contracting by them may be different.

5.3 Minimal Contraction by an Uncoupled Set Algorithm

We also give a concrete polynomial time algorithm for contraction by an uncoupled set of preferences (see Algorithm 2).

Algorithm 2 Minimal preference set contraction algorithm

```

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

6 Preference Aggregation

In this section we briefly consider the problem of preference aggregation. We assume that there are n agents, and each agent i has a deductively closed consistent set of preferences S_i ; however $S_1 \cup \dots \cup S_n$ is not guaranteed to be consistent. The problem of *minimal change aggregation* is how to produce a consistent joint preference set while retaining as many of the original preferences as possible.

Definition 4 (*Minimal change preference aggregation*) Given preference sets S_1, \dots, S_n , a minimal change preference aggregation of these sets is any operation \oplus that satisfies the following properties:

1. $S_1 \oplus \dots \oplus S_n \subseteq Cn(S_1 \cup \dots \cup S_n)$
2. $S_1 \oplus \dots \oplus S_n \not\vdash \perp$
3. for every other S' such that $S' \subseteq Cn(S_1 \cup \dots \cup S_n)$ and $S' \not\vdash \perp$, it holds that $|S'| \leq |S_1 \oplus \dots \oplus S_n|$

Unfortunately, as in the case of iterated preference contraction, the minimal change preference aggregation problem is a combinatorial problem, so there is no efficient algorithm for computing $S_1 \oplus \dots \oplus S_n$. The most obvious (and obviously exponential) algorithm for implementing a minimal change preference aggregation is:

Algorithm 3 Minimal preference aggregation algorithm

```

procedure MINIMAL-AGGREGATION( $S_1, \dots, S_n$ )
   $U := Cn(S_1 \cup \dots \cup S_n)$ 
  for each  $k \in \{0, \dots, |U|\}$  do
    for each  $X \subseteq U, |X| = k$  do
      if  $U \setminus X \not\vdash \perp$  return  $U \setminus X$ 

```

An efficient preference aggregation algorithm that is *not* guaranteed to produce a minimal cardinality result is possible. The lexicographic aggregation proposed in [3] is shown to be the only aggregation operation that satisfies a set of desirable properties and does not suffer from Arrow's paradox. (The reason Arrow's result does not apply is because agents' preferences are not assumed to be total orders.)

For the un-prioritised case that we are considering in this paper (i.e., all agents are assumed to be equally important), lexicographic aggregation is equivalent to taking the intersection of preference sets S_1, \dots, S_n . Although the intersection operation on preference relation may remove too many preferences from the result of the aggregation and is not guaranteed to be minimal, it does have attractive properties.

Let us define the *intersection aggregation* of S_1, \dots, S_n as $S_1 \cap \dots \cap S_n$.

Fact 1 *If S_1, \dots, S_n are consistent and deductively closed, then the intersection aggregation satisfies properties 1-2 of Definition 4.*

It is also easy to implement an efficient algorithm for computing the intersection aggregation:

Fact 2 *There exists an algorithm for computing $S_1 \cap \dots \cap S_n$ which runs in time $O(|S_1 \cup \dots \cup S_n|)$ (if we assume that the elements of S_1, \dots, S_n are sorted).*

It is also straightforward to characterise the intersection aggregation axiomatically by the following, rather obvious, postulate:

Intersection-Aggregation $p \in S_1 \oplus \dots \oplus S_n$ iff $p \in S_i$ for $i \in \{1, \dots, n\}$

This postulate looks similar to unanimity [4] (which says that if a preference is in every agent's set of preferences, then it is also in the result of the aggregation). It is however much stronger because it requires that *only* such preferences are in the aggregated set of preferences.

It was shown in [3] that in addition to unanimity, this operation also satisfies several desirable properties for a preference aggregation mechanism, such as no dictator.

7 Choice of Preference Relations

We have chosen $<$, \equiv and $\#$ as our primitive preference relations. A natural question to ask is whether, for example, \leq would be a good (or better) choice. Clearly, $A < B$ is definable as $(A \leq B) \wedge \neg(B \leq A)$ and $A \equiv B$ as $(A \leq B) \wedge (B \leq A)$. (On the other hand, according to our semantics, $A \# B$ is *not* definable as $\neg(A \leq B) \wedge \neg(B \leq A)$, since there is a difference between the absence of any preference between A and B , and A and B being considered incomparable.) In order to define $<$ in terms of \leq , we need to add explicit negation to the language (i.e., in addition to $A \leq B$, the set of preference sentences may contain sentences of the form $\neg(A \leq B)$). In addition, the closure rules need to include reflexivity and transitivity of \leq :

Refl $A \leq A$

Trans $A \leq B, B \leq C \Rightarrow A \leq C$

and also say that $A \leq B$ and $\neg(A \leq B)$ are incompatible:²

Neg $A \leq B, \neg(A \leq B) \Rightarrow \perp$

Let us denote the derivability relation with respect to **Refl**, **Trans**, and **Neg** \vdash_{\leq} , and the corresponding closure operation Cn_{\leq} .

Definition 5 (Minimal contraction by a \leq -preference) Given a preference set S and a \leq -preference p , such that $\not\vdash_{\leq} p$, a minimal contraction of S by p is any operation – that returns a set $S - p$ such that:

(1) $S - p \subseteq S$

(2) $S - p \not\vdash_{\leq} p$

(3) for any other set S' such that $S' \subseteq S$ and $S' \not\vdash_{\leq} p$, it holds that $|S'| \leq |S - p|$.

Similarly to the case of $<$, let A_S^{\leq} denote $\{C \mid A \leq C \in S\}$ and A_S^{\geq} denote $\{C \mid C \leq A \in S\}$. The cost $c_S(A \leq B)$ for $A \leq B \in S$ is defined as follows:

$$c_S(A \leq B) = |A_S^{\leq} \cap B_S^{\geq}| + 1$$

The cost $c_S(\neg(A \leq B))$ of removing $\neg(A \leq B)$ is always 1 since negative preferences cannot be derived.

The following postulates characterise minimal contraction of preferences where the only preference relation is \leq :

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

C-Inclusion $S - p \subseteq S$

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

²Although the following inference rules also make sense:

Neg-Trans1 $A \leq B, \neg(A \leq C) \Rightarrow \neg(B \leq C)$

Neg-Trans2 $B \leq C, \neg(A \leq C) \Rightarrow \neg(A \leq B)$

we do not adopt them purely for reasons of efficiency (it does not appear to be possible to provide a linear time contraction algorithm for $\neg(A \leq B)$ if it can be derived using several rules which have premises both of the form $A \leq B$ and $\neg(A \leq B)$).

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

C-Equivalence If $Cn_{\leq}(p_1) = Cn_{\leq}(p_2)$, then $S - p_1 = S - p_2$

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

Theorem 6 *The result of any minimal contraction by a \leq -preference satisfies the minimal preference contraction postulates above, and every contraction satisfying these postulates is a minimal contraction by a \leq -preference.*

The proof is very similar to the proof for the multiple preferences relations.

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 5, 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 5 (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 5 (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 5 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 for all atomic non-tautological p_1, p_2 , $Cn_{\leq}(p_1) \neq Cn_{\leq}(p_2)$ if $p_1 \neq p_2$ (because $Cn_{\leq}(p) = \{p\} \cup \{A \leq A \mid A \in \mathcal{A}\}$). For tautological p_1, p_2 contraction is not defined since it is impossible to construct a deductively closed preference set which does not contain them. (Alternatively, we could have defined $S - A \leq A = S$, in which case again C-Equivalence would hold.)

Now let us consider the minimality postulates. We need to prove that any minimal contraction by $A \leq B$ removes exactly $|A^{\leq} \cap B^{\geq}| + 1$ preferences.

In order to contract by $A \leq B$, we need to remove $A \leq B$ itself from S . However $A \leq B$ may still be derivable using the transitivity rule. The number of possible derivations of $A \leq B$ using the rule $A \leq C, C \leq B \Rightarrow A \leq B$ is exactly $|A^{\leq} \cap B^{\geq}|$. 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 all premises of the form $A \leq C$ or all premises of the form $C \leq B$. So any contraction satisfying (1) and (2) needs to remove at least $|A^{\leq} \cap B^{\geq}| + 1$ preferences (1 is for $A \leq B$ itself). Conversely, if one of the preferences for each possible derivation is

removed, then $A \leq B$ is no longer derivable, so the operation already satisfies (1) and (2). (Note that if $A \leq C$ for $C \in A^{\leq} \cap B^{\geq}$ is itself derivable, one premise in the derivation of $A \leq C$ is $A \leq D$ where $D \leq C$ since $C \leq B, D \leq C$, so $D \in A^{\leq} \cap B^{\geq}$, so $A \leq D$ will be removed and hence $A \leq C$ is not re-derivable.) Hence, in order to satisfy (3), the operation should not remove anything else. Hence any minimal contraction removes exactly $|A^{\leq} \cap B^{\geq}| + 1$ preferences.

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 5. 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

The algorithm for computing $S - p$ is given below. It assumes that p is not tautological (in that case contraction is not defined). Note that if $p \notin S$, the set $\{C \mid C \in A^{\leq} \cap B^{\geq}\}$ is empty so $S \setminus \{A \leq C \mid C \in A^{\leq} \cap B^{\geq}\} = S$.

Algorithm 4 Minimal preference contraction algorithm for \leq

```

procedure MINIMAL-CONTRACTION- $\leq(S, p)$ 
  case  $p \notin S$ 
    return
  case  $p = \neg(A \leq B)$ 
     $S := S \setminus \{\neg(A \leq B)\}$ 
  case  $p = A \leq B$ 
     $A^{\leq} := \{C \mid A \leq C\}$ 
     $B^{\geq} := \{C \mid C \leq B\}$ 
    for each  $C \in A^{\leq} \cap B^{\geq}$  do
       $S := S \setminus \{A \leq C\}$ 
    end for
     $S := S \setminus \{A \leq B\}$ 

```

Theorem 7 Algorithm 4 computes a minimal preference contraction.

Proof. We show that the result of applying the algorithm to a preference set S and p which is not of the form $A \leq A$, always satisfies the conditions in Definition 5. 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 6. \square

Theorem 8 The time complexity of Algorithm 4 is in $O(|\mathcal{A}|)$.

Proof. We assume that we can order the alternatives in some order (e.g., lexicographic order) and we can recover the ordered set of alternatives to which an alternative A is related by \leq in constant time (e.g., a hash table mapping from alternatives to sets (lists) of alternatives A^{\leq} and A^{\geq}). The maximum size of A^{\leq} and B^{\geq} is bounded by $|\mathcal{A}|$, since A and B can be related to at most $|\mathcal{A}|$ alternatives by \leq . Computing the set of alternatives $C \in A^{\leq} \cap B^{\geq}$ is also linear in $|\mathcal{A}|$ (to be precise it requires at most $2|\mathcal{A}|$ operations) and the number of such alternatives C is again bounded by $|\mathcal{A}|$. Removing the preferences $A \leq C$ for $C \in A^{\leq} \cap B^{\geq}$ 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. \square

The problem of minimal contraction of a set of \leq preferences S by a set of \leq preferences X can be shown to be NP-hard (just as in the case of multiple preferences relations).

Theorem 9 *The problem of minimal contraction by a set of \leq preferences is NP-hard.*

Proof (sketch). The proof is very similar to the proof of Theorem 4, with encoding of (s, t) edges as pairs of preferences $s \leq t, t \leq s$ instead of $s \equiv t$. \square

8 Related Work

In this section we briefly discuss related work, focusing on the main ideas and similarities to our approach, rather than providing a full-fledged comparison.

Van Benthem and Liu [6] propose a dynamic preference logic. They adopt the framework of dynamic epistemic logic, and introduce dynamic operators to interpret the triggers of preference change. In later work, Liu [28, 29] presents a structured model of preference, called the *two-level model*, which can account for both preference and reasons for preference, and investigates the dynamics of preference at both levels. Lang and van der Torre [25] study the relation between belief revision and preference change in general. They identify four different cases for preference change, and propose AGM postulates for the case in which preference change is triggered by belief revision. Grüne-Yanoff and Hansson [16] also consider this issue, and argue that preference change cannot be reduced to belief change, and must consider priorities. Preference change under social peer pressure is studied by Liang and Seligman [27].

In [18] Hansson describes four types of preference change: contraction and revision of preference relations, and addition and subtraction of alternatives. In our approach, we do not consider changes to alternatives and we focus on contraction

and revision of preference relations. 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. On the technical side, Hansson introduces a measure of similarity between preference relations R_1 and R_2 in order to define a *minimal* preference revision operator. In the case of finite sets of alternatives, the measure is minimising the cardinality of the symmetric difference $\Delta(R_1, R_2)$ between R_1 and R_2 (for any sets X and Y , $X\Delta Y$ is equal to $(X\setminus Y) \cup (Y\setminus X)$). A similar idea is used in the field of social choice theory, for example in [21, 22], where the distance between two preference relations essentially counts the minimal number of (pairwise) “inversions” of alternatives necessary to transform one binary relation into the other. Note that if R_2 is the result of the contraction of R_1 by a preference or a set of preferences, we have $R_2 \subseteq R_1$, so minimising the symmetric difference corresponds to minimising the cardinality of the set of removed preferences. Our work can thus be seen as using the same measure of distance between preference relations as Hansson’s and other approaches that minimise the number of tuples in the symmetric difference, such as [21, 22]. However, Hansson considers a full logical language with negations, disjunctions etc. of preferences, and the computational complexity of the preference change operations he defines is clearly much higher than ours. In essence, our approach is the same as Hansson’s but reasoning only about atomic preferences with simple inference rules and not closing preference sets under full classical reasoning.

Rationality constraints, or inference rules relating specifically to preferences, have been discussed by Hallden [17], von Wright [35] and Grüne-Yanoff and Hansson [16], among others. In our approach, we adopt only the basic rules proposed in [16].

Our approach is also related to work on minimal belief base contraction, for example, Rott [33], where minimal contraction of a finite belief base H by ϕ involves identifying maximal subsets of H which do not entail ϕ . However, the meanings of ‘minimal’ and ‘maximal’ is different from those used in our approach. For us, ‘maximal’ refers to cardinality; for belief base revision, H' is a maximal subset of H not entailing ϕ if for every $\psi \in H \setminus H'$, $H' \cup \{\psi\}$ entails ϕ . One example of when this definition of maximal does not imply maximal cardinality is where $H = \{p, q, p \wedge q \rightarrow r, p \wedge q\}$ and $\phi = r$. Then one maximal subset of H which does not entail r is $\{p, q, p \wedge q\}$ of cardinality 3, and another is $\{p, p \wedge q \rightarrow r\}$ which is of cardinality 2.

There exists considerable work on iterated belief revision, see, for example [8, 10, 11, 20]. There are some similarities between our work and the work of Ma et al. [30], who study belief revision postulates for revision of a partial preorder by a partial preorder. Our preference sentences correspond to their units, and our

revision by a set of preference sentences corresponds to their revision of a preorder by a preorder. However there are also some differences: Ma et al. assume similar closure properties, but in addition they require closure under the rule

$$A \equiv B, B < C \Rightarrow A < C$$

which we do not have. The concrete revision operations they propose are exponential in the size of the original preorder (preference set) since they require considering all possible permutations of that set. Booth and Meyer [7] highlight potential connections between iterated belief revision and preference aggregation, as they study revision of a total preorder in the context of iterated belief revision.

There is a very extensive body of literature in social choice theory on preference aggregation, e.g., [32, 9, 15]. We focus here on Andreka et al. [3], as we share the basic setting of preference relations between two alternatives. However, the difference with our setting is that they consider complete preference sets (for each pair A, B , one of $A < B$, $B < A$, $A \equiv B$ or $A \# B$ holds), and their setting includes priorities over agents and preference criteria. They use lexicographic ordering for aggregating preferences. They do not consider computational complexity, but their preference aggregation (for finite sets) can be implemented in polynomial time.

9 Conclusion and Future Work

In this paper, we consider a simple setting of preference change where it is possible to define minimal preference contraction. We give rationality postulates and an efficient algorithm for that setting. Then we study minimal contraction by a set of preferences and show that this problem is NP-hard. We provide a characterisation and an efficient algorithm for the case where the set of preferences is uncoupled. We briefly comment on the problem of preference aggregation and on our choice of basic preference relations. Finally, we discuss related work, highlighting similarities and differences with our approach.

There are a number of possible directions for future work. In our approach, we adopt only the basic rationality constraints on preferences proposed in [16]. More advanced constraints proposed in [16] involve priorities, and domain specific constraints. In future work, we would like to extend our approach to incorporate these more advanced constraints. The definition of minimality (of change) assumed in our paper is based on the minimal *number* of preferences which have changed. This corresponds to using the Hamming distance between preference relations to measure their dissimilarity. There is a large literature on alternative distance measures between preferences, for example [5, 13, 21, 22, 24, 31, 36]. It has been argued in for example [34, 12], that the Hamming distance as a measure of dissimilarity may

not be appropriate when the elements in the preference set are logically related. For example, $A \equiv B$ and $B \equiv A$ are logically related (logically equivalent in fact). It can be argued that logically equivalent formulas should only count once when the distance between two preference sets is computed. In future work, we plan to consider efficient minimal preference revision for different measures of minimality which take into account logical relations between preferences.

Acknowledgement The authors thank the three anonymous LORI'13 referees and the special issue referees for their useful comments.

Fenrong Liu is supported by the Project of National Social Science Foundation of China (NO.13AZX018), and Tsinghua University Initiative Scientific Research Program (20131089292). Natasha Alechina would like to thank Xueying Liu, an undergraduate student at the University of Nottingham, who implemented the preference contraction algorithms and inspired the idea of the NP-hardness proof by generating test cases for set contraction with \equiv preferences.

References

- [1] Natasha Alechina, Fenrong Liu, and Brian Logan. Minimal preference change. In Davide Grossi, Olivier Roy, and Huaxin Huang, editors, *Proceedings Fourth International Workshop on Logic, Rationality and Interaction (LORI-IV)*, volume 8196 of *LNCS*, pages 15–26, Hangzhou, China, October 2013. Springer.
- [2] Natasha Alechina, Fenrong Liu, and Brian Logan. Postulates and a linear-time algorithm for minimal preference contraction. In *Workshop on Logical Aspects of Multi-Agent Systems (LAMAS 2014), Informal Proceedings*, 2014.
- [3] Hajnal Andréka, Mark Ryan, and Pierre-Yves Schobbens. Operators and laws for combining preference relations. *Journal of Logic and Computation*, 12(1):13–53, 2002.
- [4] K. J. Arrow. A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):328–346, 1950.
- [5] Nick Baigent. Preference proximity and anonymous social choice. *The Quarterly Journal of Economics*, 102(1):161–170, 1987.
- [6] Johan van Benthem and Fenrong Liu. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logic*, 17:157–182, 2007.

- [7] Richard Booth and Thomas Meyer. How to revise a total preorder. *Journal of Philosophical Logic*, 40:193–238, 2011.
- [8] Craig Boutilier. Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic*, 25(3):263–305, 1996.
- [9] Vincent Conitzer, Jérôme Lang, and Lirong Xia. Hypercubewise preference aggregation in multi-issue domains. In *In Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-2011)*, 2011.
- [10] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.
- [11] 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.
- [12] Conal Duddy and Ashley Piggins. A measure of distance between judgment sets. *Social Choice and Welfare*, 39(4):855–867, 2012.
- [13] Peter C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.
- [14] Peter Gärdenfors. *Knowledge In Flux: Modeling the Dynamics of Epistemic States*. MIT Press, 1988.
- [15] Umberto Grandi. *Binary Aggregation with Integrity Constraints*. PhD thesis, ILLC, University of Amsterdam, 2012.
- [16] Till Grüne-Yanoff and Sven Ove Hansson, editors. *Preference Change: Approaches from Philosophy, Economics and Psychology*. Theory and Decision Library. Springer, 2009.
- [17] Sören Hallden. *On the Logic of “Better”*. Lund, 1957.
- [18] Sven Ove Hansson. Changes in preference. *Theory and Decision*, 38:1–28, 1995.
- [19] Sven Ove Hansson. Semi-revision (invited paper). *Journal of Applied Non-Classical Logics*, 7(2), 1997.
- [20] Sven Ove Hansson. Multiple and iterated contraction reduced to single-step single-sentence contraction. *Synthese*, 173(2):153–177, 2010.

- [21] John G. Kemeny. Mathematics without numbers. *Daedalus*, 88(4):577–591, 1959.
- [22] Christian Klamler. A distance measure for choice functions. *Daedalus*, 30:419–425, 2008.
- [23] Martin Lackner, Reinhard Pichler, Stefan Rümmele, and Stefan Woltran. Multicut on graphs of bounded clique-width. In Guohui Lin, editor, *Combinatorial Optimization and Applications - 6th International Conference, COCOA 2012, Banff, AB, Canada, August 5-9, 2012. Proceedings*, volume 7402 of *Lecture Notes in Computer Science*, pages 115–126. Springer, 2012.
- [24] Cline Lafage and Jérôme Lang. Propositional distances and compact preference representation. *European Journal of Operational Research*, 160(3):741–761, 2005.
- [25] Jérôme Lang and Leendert van der Torre. From belief change to preference change. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI-2008)*, pages 351–355, 2008.
- [26] Isaac Levi. Subjunctives, dispositions and chances. *Synthese*, 34:423–455, 1977.
- [27] Zhen Liang and Jeremy Seligman. A logical model of the dynamics of peer pressure. *Electronic Notes in Theoretical Computer Science*, 278:275–288, 2011.
- [28] Fenrong Liu. *Reasoning about Preference Dynamics*, volume 354 of *Synthese Library*. Springer, 2011.
- [29] Fenrong Liu. A two-level perspective on preference. *Journal of Philosophical Logic*, 40:421–439, 2011.
- [30] Jianbing Ma, Salem Benferhat, and Weiru Liu. Revision over partial preorders: A postulational study. In Eyke Hüllermeier, Sebastian Link, Thomas Fober, and Bernhard Seeger, editors, *Scalable Uncertainty Management - 6th International Conference, SUM 2012, Marburg, Germany, September 17-19, 2012. Proceedings*, volume 7520 of *Lecture Notes in Computer Science*, pages 219–232. Springer, 2012.
- [31] Hannu Nurmi. A comparison of some distance-based choice rules in ranking environments. *Theory and Decision*, 57(1):5–24, 2004.

- [32] Maria Silvia Pini, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Incompleteness and incomparability in preference aggregation: Complexity results. *Artificial Intelligence*, 175(7-8):1272–1289, 2011.
- [33] Hans Rott. Modellings for belief change: Base contraction, multiple contraction, and epistemic entrenchment. In David Pearce and Gerd Wagner, editors, *Logics in AI, European Workshop, JELIA '92, Berlin, Germany, September 7-10, 1992, Proceedings*, volume 633 of *Lecture Notes in Computer Science*, pages 139–153. Springer, 1992.
- [34] Marija Slavkovic and Thomas Ågotnes. Measuring dissimilarity between judgment sets. In Eduardo Fermé and João Leite, editors, *Proceedings of the 14th European Conference on Logics in Artificial Intelligence (JELIA 2014)*, volume 8761 of *Lecture Notes in Computer Science*, pages 609–617. Springer, 2014.
- [35] Georg Henrik von Wright. *The Logic of Preference*. Edinburgh: Edinburgh University Press, 1963.
- [36] Peyton Young and Arthur Levenglick. A consistent extension of Condorcet's election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.