Postulates and a linear-time algorithm for minimal preference contraction
EXTENDED ABSTRACT

Natasha Alechina[1], Fenrong Liu[2], and Brian Logan[1]

[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

1 Introduction

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 operation in the spirit of minimal change as advocated by the Alchourrón, Gärdenfors, and Makinson (AGM) theory of belief revision [1]. To be precise, a contraction of a set of preferences S by a preference p is minimal if the cardinality of a set of preferences removed from S in order to make p underivable is minimal. We characterise minimal preference contraction by a set of postulates and prove a representation theorem. We also give a linear time algorithm which implements minimal contraction. This extended abstract is based on the paper published in [2] but considers a simpler preference language (with only the $\leq$ preference relation).

2 Minimal contraction of preferences

We assume that an agent's preferences are given by a binary relation $\leq$ over some finite set of alternatives $\mathcal{A}$. For $A, B \in \mathcal{A}$, we will write $A \leq B$ to mean that B is preferred to A. We call $A \leq B$ a preference sentence. An agent's preference state is represented by a *preference set* consisting of preference sentences (or simply preferences).

We assume that the agents are rational, i.e., they can complete their preference sets using reflexivity and transitivity of $\leq$:

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

We denote by $Cn_{\leq}(S)$ the closure of a set S under the rules above. Formally, $Cn_{\leq}(S)$ is the set of preferences which contains S, $A \leq A$ for every $A \in \mathcal{A}$, and in addition $A \leq B, B \leq C \in Cn_{\leq}(S)$, then $A \leq C \in Cn_{\leq}(S)$. A set of preferences S is *deductively closed* iff $S = Cn_{\leq}(S)$. In what follows, we assume that preference sets are always deductively closed.

Sometimes we will use the notation $S \vdash_{\leq} p$ to say that p can be derived from S and the reasoning rules **Refl** and **Trans**. Clearly, $\vdash_{\leq} A \leq A$ (since $A \leq A$ is derivable from an empty S). We will refer to sentences derivable from an empty set as *tautological*.

Note that we *do not* assume that the agent can also reason about its preferences in propositional logic using a language that has negation, disjunction etc. The problem of preference revision in this simple setting appears, at first sight, trivial. Since there is no way to derive a contradiction, revision of S by p is always the same as expansion of S by p and can be defined as $Cn_{\leq}(S \cup \{p\})$. However, we would like to argue that the problem of *contracting by a preference* still does make sense. An agent may want to give up a preference $A \leq B$ because, for example, it no longer believes in the reason which led it to prefer B to A. For example, if the agent is choosing between alternatives A and B on the grounds of cost, and looked up the prices in a catalogue where B appeared to be significantly cheaper than A, it may have decided that $A \leq B$. However, if it turns out later that the catalogue was out of date, or did not even refer to the A and B the agent is concerned with, but to some similarly named but different alternatives, then the agent has no reason to prefer B to A or to consider them equally preferred. It just needs to change its preference set so that this set no longer contains (and entails) $A \leq B$. This is the problem of contraction of S by $A \leq B$, and we would like to consider the notion of a *minimal* contraction which removes the minimal number of preferences from S. In what follows, we give a formal definition, a set of postulates for the minimal contraction by a preference, and an efficient (linear time) algorithm for performing the minimal contraction.

The definition of minimal contraction of S by a non-tautological p is given below and requires that p should be removed from S and made underivable from S and that the resulting subset of S has the maximal possible cardinality (is obtained by removing as few sentences as possible from S):

Definition 1. *(Minimal contraction by a $\leq$-preference) Given a preference set S and a 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' satisfying (1) and (2), $|S'| \leq |S - p|$.

Note that it follows from the definition that if $S \not\vdash_{\leq} p$ (which given that S is deductively closed, is equivalent to $p \notin S$), then $S - p = S$.

3 Postulates

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

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

(The reason why the cost is like this is because in order to make $A \leq B$ underivable by transitivity, for every pair $A \leq C$ and $C \leq B$ we need to remove at least one of the premises, plus we need to remove $A \leq B$ itself.)

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 $A \leq A$, 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 $|S - p| = |S| - c_S(p)$

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 [5]. 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 $S - p \nvdash p$ (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 $S - p \nvdash p$ and $S \vdash q$, it follows that $(S - p) \cup \{q\} \nvdash p$. 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 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, namely 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 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. □

4 Algorithm for computing a minimal preference contraction

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 $\{C \mid C \in A^{\leq} \cap B^{\geq}\}$

Algorithm 1 Minimal preference contraction algorithm for $\leq$

```
procedure MINIMAL-CONTRACTION-≤(S, p)
    A≤ := {C | A ≤ C}
    B≥ := {C | C ≤ B}
    for each C ∈ A≤ ∩ B≥ do
        S := S \ {A ≤ C}
    end for
    S := S \ {A ≤ B}
```

is empty so $S \setminus \{A \leq C \mid C \in A^{\leq} \cap B^{\geq}\} = S$.

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 which is not of the form $A \leq A$, 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. □

Theorem 3. *The time complexity of Algorithm 1 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 $2|\mathcal{A}|$) 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. □

5 Related work

In [4], Hansson 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. [3] 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.

References

1. Carlos Eduardo Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
2. Natasha Alechina, Fenrong Liu, and Brian Logan. Minimal preference change. In Davide Grossi, Olivier Roy, and Huaxin Huang, editors, *Logic, Rationality, and Interaction - 4th International Workshop, LORI 2013, Hangzhou, China, October 9-12, 2013, Proceedings*, volume 8196 of *Lecture Notes in Computer Science*, pages 15–26. Springer, 2013.
3. 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.
4. Sven Ove Hansson. Changes in preference. *Theory and Decision*, 38:1–28, 1995.
5. David Makinson. On the status of the postulate of recovery in the logic of theory change. *Journal of Philosophical Logic*, 16:383–394, 1987.