

Belief revision, belief merging and voting

Dov Gabbay and Gabriella Pigozzi and Odinaldo Rodrigues¹

Abstract. In belief revision, an agent is faced with the problem of choosing between several alternatives when trying to restore consistency to theory. Ideally, the choice process is conducted in a way that verifies a number of fairness principles. On the other hand, belief merging concerns with the problem of determining a group's beliefs from individual members' beliefs that are not always compatible with each other. Similarly, in voting systems, a social welfare function takes individual preferences into account in order to produce a collective preference. Here again certain fairness principles are desirable. In this paper, we investigate the relationship between revision, merging and voting.

1 Introduction

This paper makes a connection between voting, belief revision and merging by applying general the principles/ideas to a particular problem and seeing how the principles interact. We recognize the following three scenarios:

1. *Voting.* There are several conflicting demands/preferences and we are looking for a collective compromise.
2. *Belief revision.* We are facing an inconsistent or unacceptable logical theory, and we are looking for a way out.
3. *Belief merging.* We are trying to aggregate knowledge bases which together are possibly inconsistent.

Our best bet in bringing the above three areas together is to look at a problem which can be equally considered by each one of them and compare how they would deal with it. In this way we can learn from each point of view and export ideas to the other two.

Our strategy for investigation can be outlined as follows.

Step 1. Start with an example of theory revision of a theory T by an input formula τ .

Step 2. See this as a voting problem, i.e., consider τ as the formula expressing the voting rules, and the voters as a classical logic theory T . If the voters' preferences are incompatible, the revision of T by τ will seek a compromise consistent with the voting rules. This, in the case of a maxichoice revision function, will be obtained from a maximal subset of T consistent with τ . We compare the a priori philosophical demands of revision theory and voting and find that while revision theory may pick a dictator voting does not want one.

Step 3. Now that we have a connection between voters and the voting rules (by seeing them as logical structures) we can apply the same reasoning for the case of belief merging. In this case, we see voters as individual belief bases and the voting rules as integrity constraints. We also find that this approach is not entirely satisfactory.

Step 4. We continue by analysing how the way of thinking in one view can enrich the way of thinking in the other views.

Essentially this problem involves reasoning about orders and there are two natural moves one can make with logic theories involving orders. One has to do with the way we interpret the order itself whereas the other has to do with the kind of logic we use to represent and manipulate it. We note that

a. Logical theories need not be complete or associated with any particular *strict* linear order.

b. Theories can be made numerically valued of fuzzy values (i.e., values need not be restricted to just *true* or *false*, but can be taken from a range of points in-between).

From the voting point of view, (a.) means the voters may be uncertain about how all outcomes compare with each other, but may be clear about how some of them do, i.e., a voter may not have a particular preference between outcomes x and y , but may prefer z to either. In addition, voters may give conditional preferences (for example: if a voter prefers a to b , then he/she might also expect to prefer c to d).

As for (b.), numerical values do appear in *range voting* in a natural way. In this paper, we will concentrate on (a.) leaving (b.) to be explored in future work.

A number of authors have investigated ways of combining preference relations [2, 17]. The difference in our approach is that we bring the mechanism for representing preferences to the object level itself and hence can analyse it from an entirely logical perspective.

2 Background Review

Since our approach combines principles from voting theory, belief revision and belief merging, we start by introducing some concepts that will be used in the remaining of this paper.

2.1 Voting

Voting is concerned with the aggregation of individual preferences in order to select a collectively preferred alternative. This problem is extensively studied by social choice theory [3, 4, 18]. Probably the most famous method for the aggregation of preferences is the one proposed in the 18th century by the Marquis de Condorcet. Given a set of individual preferences, we compare each of the alternatives in pairs. For each pair, we determine the winner by majority voting, and the final collective ordering is obtained by a combination of all partial results. Unfortunately, this method led to the first aggregation problem, known as *the Condorcet paradox*: the pairwise majority rule can lead to cycles in the collective ordering. In other words, this ordering cannot be used to select an overall preferred candidate.

Preferences over some set X of alternatives can be formalised as follows. Let $<$ be a binary relation on $X \times X$, where $x < y$ denotes that alternative x is preferred to y . The desired properties of preferences corresponding to strict linear orders are given below, where $\{x, y, z\}$ range over elements of X .

¹ Department of Computer Science, King's College London, email: {dov.gabbay, gabriella.pigozzi, odinaldo.rodrigues}@kcl.ac.uk

- (P1) $\forall x, y, z((x < y \wedge y < z) \rightarrow x < z)$ (transitivity)
(P2) $\forall x, y(x \neq y \rightarrow (x < y \vee y < x))$ (totality)
(P3) $\forall x, y((x < y) \rightarrow \neg(y < x))$ (asymmetry)

With the above formalisation, the Condorcet paradox can be expressed as follows. Suppose that there are three possible candidates a, b and c (that is, $\{a, b, c\} \in X$) and three voters, who express their total preferences in the following way:

$$\begin{aligned} V_1 &= \{a < b, b < c\} \\ V_2 &= \{b < c, c < a\} \\ V_3 &= \{c < a, a < b\} \end{aligned}$$

According to Condorcet's method, $a < b$ has the majority of the voters (V_1 and V_3), so does $b < c$ (V_1 and V_2) and, so does, $c < a$ (V_2 and V_3). This leads us to the collective outcome $a < b, b < c$ and $c < a$, which together with transitivity (P1) violates (P3) (asymmetry).

Unfortunately, this is not a particular problem of Condorcet's method. More recently, the aggregation of preferences was investigated by K. Arrow, who proved an important result which became known as "Arrow's impossibility theorem", stated below.

Let X be a non-empty set of mutually exclusive social states and \leq_i be a total, reflexive and transitive preference relation for an individual i over the states in X (and $<_i$ its strict counterpart).

Suppose there are n individuals in society. A social welfare function (SWF) is a function that produces a total, reflexive and transitive social preference relation \preceq from a given n -tuple of individual orderings $\{\leq_1, \dots, \leq_n\}$ (again we use \prec to denote \preceq 's strict counterpart).

Arrow's impossibility theorem states that whenever $|X| > 2$, the following conditions become incompatible:

(Universal domain) The social preference function should be able to cover all admissible individual preference relations.

(Independence of irrelevant alternatives - IIA) The social preference on any pair of alternatives depends exclusively on the individual preferences over that pair.

(Non-dictatorship) There is no individual i such that for each $\{x, y\} \in X$, $x <_i y$ implies $x \prec y$.

(Weak Pareto principle) if for all i , $x <_i y$, then $x \prec y$

The above condition is also called *unanimity*.

2.2 Belief revision

The main object of study of theory of belief revision is the dynamics of the process of belief change: when an agent is faced with new information which contradicts his/her current beliefs, he/she will have to retract some of the old beliefs in order to accommodate the new belief consistently. This can generally be done in several ways. Rationality postulates were initially proposed by Alchourrón, Gärdenfors and Makinson and became known as the *AGM postulates for belief revision*, presented below. There is extensive literature on the subject, including [1, 9, 10].

Let K be a set of sentences closed under logical consequence (a *belief set*) and A and B well formed formulae (*beliefs*). The revision of K by A is denoted by $K \circ A$ and should satisfy the following requirements:

- (K°1) $K \circ A$ is a belief set
(K°2) $A \in K \circ A$
(K°3) $K \circ A \subseteq \text{Cn}(K \cup \{A\})$
(K°4) If $\neg A \notin K$, then $\text{Cn}(K \cup \{A\}) \subseteq K \circ A$
(K°5) $K \circ A = K_{\perp}$ only if A is contradictory

- (K°6) If $A \equiv B$, then $K \circ A \equiv K \circ B$
(K°7) $K \circ (A \wedge B) \subseteq \text{Cn}((K \circ A) \cup \{B\})$
(K°8) If $\neg B \notin K \circ A$, then $\text{Cn}(K \circ A \cup \{B\}) \subseteq K \circ (A \wedge B)$

These postulates are well known. In [8], we have argued that the postulates can be somewhat simplified. In particular, (K°3)–(K°5) have something to say only when $K \cup \{A\}$ is consistent, or when it is inconsistent even though A is non-contradictory. The particular way of writing the postulates given above makes use of technical properties of classical logic (the way inconsistent theories prove everything). (K°3)–(K°4) effectively mean the following:

(K_{3,4}°) If A is consistent with K , then $K \circ A = \text{Cn}(K \cup \{A\})$.

If K is finitely representable, it can be taken as a formula and the postulate above corresponds to (R2) in Katsuno and Mendelzon's rephrasing of the AGM postulates for belief sets represented by finite bases [11, p. 187].

Similarly, postulates (K°7)–(K°8) do not tell us anything new (beyond what we can deduce from earlier postulates), except in the case where B is consistent with $K \circ A$, when (K°7) and (K°8) together are equivalent to the postulate below:

(K_{7,8}°) If B is consistent with $K \circ A$, then $\text{Cn}((K \circ A) \cup \{B\}) = K \circ (A \wedge B)$

which again corresponds to Katsuno and Mendelzon's (R6).

One of the foundations of the AGM formalism is the idea of informational economy, that states that old beliefs should not be given up unless strictly necessary in order to consistently accommodate the new belief. One possibility of defining revision functions is then to think in terms of the minimal mutilation required to accommodate a new belief — the contraction of the old belief set — followed by an expansion of the contracted set by the new belief.² This is formalised as follows.

Definition 2.1 (Maximal subsets that fail to imply a sentence)

Let K be a belief set and $\neg A$ a belief. A set K' is a maximal subset of K that fails to imply $\neg A$ iff the following conditions are met:

- $K' \subseteq K$
- $\neg A \notin \text{Cn}(K')$
- $\forall K'', K' \subset K'' \subseteq K$ implies $\neg A \in \text{Cn}(K'')$

In other words, the closure of any subset of K larger than such a K' would result in a theory that entails $\neg A$. It should always be possible to find such subsets unless $\neg A$ is a tautology, in which case there is no K' meeting the above conditions. The set of all subsets of a belief set K that do not imply a sentence $\neg A$ is usually denoted $K_{\perp} \neg A$. This set is used to define a number of contraction functions as explained next. In all cases, a selection function s is used to pick an appropriate collection of elements of $K_{\perp} \neg A$ if it is not empty or K itself otherwise. The behaviour of s defines three classes of contraction operations as follows:

- a *maxichoice contraction* is obtained when s picks one element of $K_{\perp} \neg A$
- a *full meet contraction* is obtained when s returns the intersection of all elements of $K_{\perp} \neg A$
- and finally, a *partial meet contraction* is obtained when s returns the intersection of *some* appropriately selected elements of $K_{\perp} \neg A$

² This is revision defined in terms of contraction — the well known *Levi identity*.

Based on these contraction operations, maxichoice, partial and full meet revision operations can be defined via the subsequent expansion of the contracted belief set by the new belief.

2.3 Belief merging

The aggregation of finite sets of information into a collective one is studied by a recent discipline called belief merging [12, 14, 15]. The aggregation procedure in belief merging faces problems similar to those addressed in voting theory. Links between these two disciplines have been investigated in [12, 16, 5].

A particular type of aggregation is called *model-based merging*. Its formal framework consists of a propositional language \mathcal{L} , built up from a finite set \mathcal{P} of propositional letters and the usual logical connectives. Given a finite set of individuals, each person i states his/her own beliefs as a consistent finite set of propositional formulae denoted by K_i (his/her *belief base*).³

An interpretation is a function $\mathcal{P} \rightarrow \{0, 1\}$ which is extended to complex formulae in the usual way. Let $\mathcal{W} = \{0, 1\}^{\mathcal{P}}$ be the set of all interpretations and $2^{\mathcal{W}}$ denote the power set of \mathcal{W} . For any formula $\varphi \in \mathcal{L}$, $[\varphi] = \{\omega \in \mathcal{W} \mid \omega \models \varphi\}$ denotes the set of models of φ , i.e., the set of interpretations w such that $w(\varphi) = 1$. Conversely, for any set of models $M \subseteq \mathcal{W}$, let $form(M)$ denote a propositional formula (up to logical equivalence) such that $[form(M)] = M$, i.e., a formula whose set of models are precisely M . \mathcal{K} will be used to denote the set of all consistent belief bases.

In a model-based framework a merging operator $\Delta : \mathcal{K}^n \rightarrow \mathcal{K}$ is defined by a function $m : \mathcal{K}^n \rightarrow 2^{\mathcal{W}}$ from the set of all possible collections of consistent belief bases (called *profiles*) to the power set of \mathcal{W} such that for all profiles $\underline{K} \in \mathcal{K}^n$, $\Delta(\underline{K}) = form(m(\underline{K}))$. For reasons of convenience, a merging operator $\Delta : \mathcal{K}^n \rightarrow 2^{\mathcal{W}}$ and its defining function $m : \mathcal{K}^n \rightarrow 2^{\mathcal{W}}$ will be used interchangeably when the context is clear.

Most model-based majoritarian⁴ merging operators $m : \mathcal{K}^n \rightarrow 2^{\mathcal{W}}$ are defined in terms of the minimisation of an appropriate notion of *distance*. The idea is to select the interpretations whose distance to a model in the collection of sets of models $\{[K_1], [K_2], \dots, [K_n]\}$ is minimal. These models correspond to a profile $\underline{K} = (K_1, K_2, \dots, K_n) \in \mathcal{K}^n$ of belief bases.

A distance between interpretations is a function $\underline{d} : \mathcal{W} \times \mathcal{W} \rightarrow \mathbb{R}_+$ such that for all $\omega, \omega' \in \mathcal{W}$:

1. $\underline{d}(\omega, \omega') = \underline{d}(\omega', \omega)$
2. $\underline{d}(\omega, \omega') = 0$ iff $\omega = \omega'$.

The distance d between an interpretation ω and a belief base K is the minimal distance \underline{d} between ω and any model of K , i.e., $d(\omega, K) = \min_{\omega' \in [K]} \underline{d}(\omega, \omega')$. The distance between an interpretation ω and a profile \underline{K} is defined with the help of an aggregation function $D : \mathbb{R}_+^n \rightarrow \mathbb{R}_+$ as $D^d(\omega, \underline{K}) = D(d(\omega, K_1), d(\omega, K_2), \dots, d(\omega, K_n))$ [13]. Any such aggregation function induces a total pre-order $\leq_{\underline{K}}$ on the set \mathcal{W} with respect to the distances to a given profile \underline{K} . Thus, a majoritarian merging operator m for a profile \underline{K} can be defined as $m(\underline{K}) = \min(\mathcal{W}, \leq_{\underline{K}})$, i.e., the set of all interpretations with minimal distance D^d to the profile \underline{K} .

Obviously, the properties of the merging operator essentially depend on the functions \underline{d} and D . The most widely used merging operator used in the literature is the operator $\Delta^{\underline{d}, \Sigma}$ defined as follows:

³ Since each K_i is finite, it is identified with the conjunction of its elements.

⁴ Intuitively, a majoritarian merging operator aims at satisfying the view of the majority.

1. \underline{d} is the Hamming distance — the number of propositional letters on which two interpretations differ, i.e., $\underline{d}(\omega, \omega') = |\{\pi \in \mathcal{P} \mid \omega(\pi) \neq \omega'(\pi)\}|$ and
2. $D^d(\omega, \underline{K}) = \sum_i d(\omega, K_i)$ is the sum of componentwise distances d defined in terms of \underline{d} as given before.

It is also possible to impose some restrictions on the result of the merging operator. A set of integrity constraints IC is a satisfiable set of formulae. Given a set of IC , $\Delta_{IC}(\underline{K})$ restricts the image of the merging operator Δ to $[IC]$, i.e., $\Delta_{IC}(\underline{K})$ gives the models of IC that differ minimally from models of each $K_i \in \underline{K}$.

3 Interactions between belief revision, merging and voting

One way of analysing the interaction between belief revision, merging and voting is to express voting principles in a logical framework and then consider what belief revision and belief merging would do in specific voting scenarios. We start by considering a logic theory of order and its relation with belief revision.

Consider the language of predicate logic with binary relation $<$; the constants a, b, c and the equality symbol $=$. Assume the axioms $\forall x(x = a \vee x = b \vee x = c)$ and $a \neq b \neq c$ (this means $\neg(a = b) \wedge \neg(b = c) \wedge \neg(a = c)$). Let T be $Cn(\{a < b, b < c, c < a\})$ and consider an input τ to T saying that $<$ is the strictly linear order of the three elements a, b, c , i.e., $\tau = P1 \wedge P2 \wedge P3$.

It can be clearly seen that both $a < c$ and $c < a$ follow from $T + \tau$ and this contradicts P3, hence $T + \tau$ is not consistent. If we want to analyse what aspects of T are compatible with a strict linear order of a, b, c , we can consider the revision of T by τ . This would replace $T + \tau$ with a new consistent theory $T \circ \tau$ containing τ , by making minimal changes in T . As we saw in Section 2.2, the AGM postulates constrain how the new theory $T \circ \tau$ is related to τ and to T . The new theory $T \circ \tau$ is closed under logical consequence, i.e., if $T \circ \tau \vdash A$ then $A \in T \circ \tau$, but the AGM framework does not give an algorithm for how to find any such $T \circ \tau$. One algorithm which can do the job is given below.

Starting with $T_{\perp} \neg \tau = \{T_1, T_2, T_3, \dots\}$, $T \circ \tau$ can be constructed from any $T_i \in T_{\perp} \neg \tau$, say $Cn(T_1 \cup \{\tau\})$ (this would give a maxichoice revision of T by τ). We can find such $T_i \cup \{\tau\}$ by listing all sentences which T proves as the list A_1, \dots, A_n, \dots and defining a sequence $S_0 \subseteq S_1 \subseteq S_2, \dots$ as follows. Let $S_0 = \{\tau\}$ and for $n \geq 0$, define S_{n+1} in the following way

$$S_{n+1} = \begin{cases} S_n \cup \{A_{n+1}\}, & \text{if this set is consistent} \\ S_n, & \text{otherwise} \end{cases}$$

Finally, let $S = \bigcup_{i \in \mathbb{N}} S_i$.

If we want to look at what we retain from our original T , we see that $S - \{\tau\} \subseteq T$ is a maximal subtheory of T consistent with τ , i.e. $S - \{\tau\} = T_i$ for some i . Which T_i we get depends on the way we present T as a sequence.

Let us now see what happens if we apply these procedures to our concrete example. τ says that $\{a, b, c\}$ is strictly linearly ordered. T says that $a < b$ and $b < c$ and $c < a$. T is not consistent with τ . The maximal subtheories of T consistent with τ include:

$$\begin{aligned} T_1 &= Cn(\{a < b, b < c\}) \\ T_2 &= Cn(\{b < c, c < a\}) \\ T_3 &= Cn(\{a < b, c < a\}) \end{aligned}$$

When τ is added to these, we get the three options for revision below:

$$\begin{aligned}
V_1 &= \text{Cn}(\{a < b, b < c, \tau\}) \\
V_2 &= \text{Cn}(\{b < c, c < a, \tau\}) \\
V_3 &= \text{Cn}(\{c < a, a < b, \tau\})
\end{aligned}$$

Note that this logical revision philosophy/approach is entirely compatible with AGM revision and hence uses three basic assumptions:

1. We must replace the inconsistent $T + \tau$ by a single consistent theory $T \circ \tau$.
2. This replacement contains τ and as much of T as possible. Certainly we do not want anything not in T to be admitted to $T \circ \tau$, even if consistent with it.
3. We are dealing with two valued logic. In other words, preferences have to be represented as yes/no statements (as opposed to numerical, probabilistic or fuzzy values).

From the revision point of view, our voting example consists of three candidate options a, b and c and several voters who express their total preferences regarding these options. When put together these preferences result in the theory T . So, for example, we could have had the following preferences:

$$\begin{aligned}
\text{Voter 1} &— a < b, b < c \\
\text{Voter 2} &— b < c, c < a \\
\text{Voter 3} &— c < a, a < b
\end{aligned}$$

Since we need to make a group decision here, we require a compromise functional H based on the preferences of Voter 1, Voter 2 and Voter 3, motivated by some general principles, such that:

$$H(\text{Voter 1, Voter 2, Voter 3}) = \text{Some compromise preference, (i.e., technically some new voter).}$$

We have some reasonable conditions on H , for instance, those given in the latter part of Section 2.1. One such condition is that it does not choose as compromise one of the voters — the non-dictatorship requirement. In practice, this means that we do not want H to be a projection. Another condition is the principle of independence of irrelevant alternatives (IIA), i.e., the group decision on how two distinct elements x and y relate ($x < y$ or $y < x$) depends only on how the different voters voted on their relationship. Note that whereas the principle of non-dictatorship is a purely meta-level one on the function H and does not make use of the contents of the theories T_i , (IIA) relates to the properties of the order predicate of T_i .

Let us now look at our revision example from the voting point of view. The consistent theories T_1, T_2, T_3 can stand for voters. The sentence τ is a statement of the layout of the voting system. It specifies the alternatives $\{a, b, c\}$ and says that the combination of the voters preferences is strictly linearly ordered. We immediately observe that the theory T can be obtained back from the voters as the result of majority vote.

$$\begin{aligned}
a < b &\text{ is voted by } V_1, V_3 \\
b < c &\text{ is voted by } V_1, V_2 \\
c < a &\text{ is voted by } V_2, V_3
\end{aligned}$$

We now have a voting interpretation of our revision theory situation. What does maxichoice logical revision do in this situation? It simply chooses a dictator. This is not always the case. We can construct a consistent theory T from a number of voters V_1, V_2, \dots that is

incompatible with the voting rules τ , but whose subsequent revision by τ will not necessarily pick a dictator even if the revision turns out to be maxichoice. This is illustrated below.

Let $V_1 = \{a < b, b < c, a < c\}$ and $V_2 = \{c < b, b < a, c < a\}$ and τ be the voting rules as before. Now take $T = V_1 \cup V_2 = \{a < b, b < c, a < c, c < b, b < a, c < a\}$. T is consistent, since it does not know about the properties of linear orders. If we now enforce these, i.e., revise T by τ , a maxichoice revision would look at $T_{\perp} \neg \tau$. One of the sets in $T_{\perp} \neg \tau$ is, for instance, $\{a < c, c < b, a < b\}$ which together with τ would result in a strict linear order which does not correspond to either V_1 or V_2 . In the voting example, this is may be a desirable outcome.

Let us return to the expectations of voting theory from the point of view of revision theory. Voting theory expects some compromise vote satisfying certain conditions. Belief revision tries to find some compromise between all the $T_i \subseteq T$ that are consistent with τ . It will seek some compromise theory S_{comp} which will be acceptable to all. This is left mostly for the selection function s presented in Section 2.2.

Maxichoice revision operations look at all $T_i \uparrow s$ and sets $T \circ \tau$ as $T_i + \tau$ for some T_i . In general, this T_i is not constrained at all on containing consequences of all of the voters. Full meet revisions will be too restrictive and comprise only the consequences of the voting system τ (since $s(T_{\perp} \neg \tau) = \emptyset$). On the other hand, partial meet revisions would be based on the particular subsets T_i picked by s which again could leave the wishes of some voters out — an unfair prospect. Therefore, an acceptable S_{comp} from the voting point of view would have to rely on some meta-level principle in the case of AGM revision functions. What can we then expect of the relationship between S_{comp} and AGM? In summary,

1. If we stick with AGM, certain conditions of the voting system (τ) can be enforced, but we cannot ensure a fair outcome unless we also adopt some meta-level principles.
2. However, the AGM postulates may hold for a desirable S_{comp} even though they do not incorporate themselves any fairness principles from the voting point of view.

Let us now show what result we obtain when we apply a majoritarian model-based merging operator to the same example. As before, we want to restrict the set of all possible outcomes to the set of linear preference orderings satisfying transitivity, totality and asymmetry. That is, we make $IC = \tau$. This results in the six possible preference orderings K_1 — K_6 illustrated in the following table. In order to simplify the presentation, we consider only three propositions $a < b, b < c$ and $a < c$. These are sufficient to represent all possible linear orders with the three elements a, b and c (note that $(a < b) = 0$ iff $(b < a) = 1$). We use $x < y < z$ as a shorthand for $\text{Cn}(\{x < y, y < z, \tau\})$.

	$a < b$	$b < c$	$a < c$	order
K_1	0	0	0	$c < b < a$
K_2	0	1	0	$b < c < a$
K_3	0	1	1	$b < a < c$
K_4	1	0	0	$c < a < b$
K_5	1	0	1	$a < c < b$
K_6	1	1	1	$a < b < c$

Each voter V_i of our example is satisfied exactly by one of these models:

$$\begin{aligned} V_1 &= \{(1, 1, 1)\} (K_6) \\ V_2 &= \{(0, 1, 0)\} (K_2) \\ V_3 &= \{(1, 0, 0)\} (K_4) \end{aligned}$$

When we calculate the distances between each V_i and the possible social outcomes, we obtain the following result:

	$d(\cdot, V_1)$	$d(\cdot, V_2)$	$d(\cdot, V_3)$	$D^d(\cdot, \underline{V})$
K_1	3	1	1	5
K_2	2	0	2	4
K_3	1	1	3	5
K_4	2	2	0	4
K_5	1	3	1	5
K_6	0	2	2	4

There are three social orderings with minimal distance to the profile $\underline{V} = \{V_1, V_2, V_3\}$. These are K_2, K_4 and K_6 , which coincide respectively with V_2, V_3 and V_1 , respectively. The result of the belief merging operator is a tie: $\Delta_{IC}(\underline{V}) = \{V_1 \vee V_2 \vee V_3\}$. This means that, although belief merging (with the help of the IC) avoids the paradoxical result, this is done at the price of indecision, i.e., there is no procedure to decide which V_i should be taken to represent the collective preference — this effectively means no election.

4 Specifying preferences for multiple outcomes

The connection with revision theory where voters turn out to be logical subtheories gives us the idea that maybe voters need not vote on all the preferential options, but be allowed to give constraints on the resulting compromise social vote. So instead of saying $a < b$ and $b < c$, the voter might give a conditional constraint “If you make $a < b$ then you must also make $b < c$ ”. We call this a *conditional preference*. From the point of view of revision theory, where voters are theories, the condition $\{a < b \rightarrow b < c\}$ is a perfectly legitimate vote-theory.

Let us give an example. Consider W_1, W_2, W_3 below:

$$\begin{aligned} W_1 &= \{a < b \rightarrow b < c\} \\ W_2 &= \{b < c \rightarrow c < a\} \\ W_3 &= \{c < a \rightarrow a < b\} \end{aligned}$$

So the voters do not give clear preferences but only certain constraints. We know that many people do that. From the logical point of view, each voter is giving several alternative orderings, those which satisfy their individual preference conditions. Here is a list of the options for each voter W_1, W_2 and W_3 :

$$\begin{array}{lll} W_1 = a < b < c & W_2 = b < c < a & W_3 = c < a < b \\ & b < a < c & c < b < a & a < c < b \\ & c < b < a & a < c < b & b < a < c \\ & b < c < a & c < a < b & a < b < c \end{array}$$

In semantical terms, the set of options for each W_i will correspond to classes of models (as seen in the previous section) — those satisfying his/her constraints. Let us refer to the models of W_1 as follows

$$\begin{aligned} \mathbf{m}_1^1 &= a < b < c \\ \mathbf{m}_2^1 &= b < a < c \\ \mathbf{m}_3^1 &= c < b < a \\ \mathbf{m}_4^1 &= b < c < a \end{aligned}$$

We can now view our problem in the following way: given a number of consistent theories, which are not consistent together, can we find a compromise combination of them? As we have seen, we do not want this combination to be simply one of the theories itself (this would be a strict dictator). We will address this point of view later in the paper.

Our problem is to consider whether there is a compromise ordering acceptable to all. Let each voter give value according to how many options support his/her preference. Let us consider each pairwise constraint in turn.

One of the options compatible with W_1 supports $a < b$ and three are against it (\mathbf{m}_1^1 is in favour and $\mathbf{m}_2^1, \mathbf{m}_3^1$ and \mathbf{m}_4^1 are against it). That is, -2 in total. From the options compatible with W_2 the overall support is 0 and from the options compatible with W_3 the overall support is $+2$. The collective support for $a < b$ is hence 0.

For $a < c$, we have that W_1 gives 0; W_2 gives -2 ; W_3 gives $+2$. The collective support is hence 0.

For $b < c$, W_1 gives $+2$; W_2 gives -2 and W_3 votes 0. The collective support is hence also 0.

It is easy to check that the collective support for $b < a, c < a$ and $c < b$ will also be 0. Since all pairwise constraints will have the same level of collective support, the resulting theory will contain all of them, i.e., $\{a < b, b < a, a < c, c < a, b < c, c < b\}$ which is again not consistent with τ . Its maximal subsets consistent with τ are:

$$\begin{aligned} T_1^* &= \{b < a, a < c\} \\ T_2^* &= \{a < c, c < b\} \\ T_3^* &= \{c < b, b < a\} \\ T_4^* &= \{a < b, b < c\} \\ T_5^* &= \{b < c, c < a\} \\ T_6^* &= \{c < a, a < b\} \end{aligned}$$

Together with τ we get:

$$\begin{aligned} V_1 &= \{b < a < c\} \\ V_2 &= \{a < c < b\} \\ V_3 &= \{c < b < a\} \\ V_4 &= \{a < b < c\} \\ V_5 &= \{b < c < a\} \\ V_6 &= \{c < a < b\} \end{aligned}$$

V_1 is acceptable to W_1, W_3 . V_2 is acceptable to W_2, W_3 . V_3 is acceptable to W_1, W_2 . V_4 is acceptable to W_1, W_3 . V_5 is acceptable to W_1, W_2 . V_6 is acceptable to W_2, W_3 . That is, each outcome is acceptable by two of the voters. We are getting another circular compromise.

It is easy to verify that belief merging would give the same result, that is, a tie over all of the possible preference orderings V_i .

If instead of voting for each individual preference $x < y$, we restrict the voting to a complete ordering $x < y < z$, we need to say how the voters are going to vote given that they have only a list of alternatives and not a clear cut opinion. Logically, we can say let them vote for possible total linear orderings according to whether a total linear order satisfies their constraints. The possible orders are the following:

$$\begin{aligned} L_1 : a < b < c & \quad L_2 : b < c < a & \quad L_3 : b < a < c \\ L_4 : a < c < b & \quad L_5 : c < a < b & \quad L_6 : c < b < a \end{aligned}$$

and they correspond exactly to V_1 – V_6 given before. Again there is no option preferred by a majority of voters. Nevertheless, this result is not discouraging. It simply reflects in the logical paradigm the underlying difficulties highlighted by Condorcet and Arrow. But can the logical approach provide us with anything new?

5 Sympathetic dictatorship

As we have just seen, allowing the representation of conditional preferences by the voters does not necessarily avoid the problem with cycles. In this section, we propose a more general procedure that makes use of the logical machinery of a particular domain to conciliate possibly conflicting requirements. In the case of voting, the idea is to let the machinery itself dictate what preferences can be consistently combined.

In our example, we have a world of three elements $\{a, b, c\}$ and voters put forward a complete ordering on the world, for instance, $a < b < c$; or give some constraints, for example, $a < b \rightarrow b < c$. Since from the logical point of view a voter is a theory T , we are free to represent voters in any way allowed by the logic. When we have several voters we get several theories T_1, \dots, T_n . We are looking for a reasonable result for the vote. Note that the voting mechanism itself can be one of the theories, say T_v . If $\bigcup_i T_i$ is consistent, this means the wishes of all the voters can be accommodated by $T = \bigcup_i T_i$ (possibly within the legality of the voting process if $T_v = T_i$ for some i). If the union is not consistent, then we seek some compromise solution. In the revision and merging approaches the problem is the following. Given the theories T_1, \dots, T_n such that $\bigcup_i T_i$ is inconsistent, find a compromise theory T_{comp} . Revision theory would give priority to one of the T_i 's (the input formula) and so would belief merging (the integrity constraints). The only difference between the two is in the way the original T_i 's are considered — collectively by belief revision and individually by belief merging.

We now present an *ad hoc* algorithm for finding one such T_{comp} . Later in the section, we explain how the idea can be further developed. At a first approximation we can list the set of formulas which each T_i proves as A_1^i, A_2^i, \dots and try to build a consistent theory by choosing elements from each sequence. For example, we can make one big joint sequence of these n sequences as traditionally done in set theory:

$$A_1^1, \dots, A_1^n, A_2^1, \dots, A_2^n, A_3^1, \dots, A_3^n, \dots$$

Let B_1, B_2, B_3, \dots be a renaming of this sequence. Let $S_1 = \{B_1\}$ and let $S_{n+1} = S_n \cup \{B_{n+1}\}$ if consistent and $S_{n+1} = S_n$ otherwise. Then $T_{comp} = \bigcup_n S_n$ yields a maximal compromise subset of $\bigcup_i T_i$. The actual content of T_{comp} depends on the sequencing of A_k^i . Obviously, we will get more of T_i participating in T_{comp} if we sequence its elements earlier than the elements of other T_j . In fact, this can be used to enforce the mechanism of the voting process by having T_v listed first or to model a dictator T_d by listing him/her first instead (thus overruling T_v , if he/she so wishes). The difference between the approach above and the one given in Section 3 is that it allows for consequences of the other theories (voters) to be incorporated as long as they do not violate the wishes of the theories appearing earlier in the sequence. It is hence *sympathetic* in this sense.

If we compare the above procedure with what is done in voting theory, we see there are two differences in principle:

1. Voting compromises on a T which may not necessarily have been chosen if we were to look at the problem from the point of view of one or of a group of voters.
2. The choice of T is motivated not only just by sequencing the elements of T_i , which is external to the logical content of T_i , but also by what the theories T_i actually say (i.e., in the case of T_i talking about order then the choice must satisfy principles relating to the ordering, etc.).

Thus, if our theories talk about electrical circuits, the construction of T_{comp} would involve circuits considerations (compromise on a cheaper circuit?). This revision process is therefore *context sensitive*. Our problem is how to develop a general theory of revision which looks into the content of the theories involved (like voting does) without sinking into the level of case by case analysis (voting, circuits, agricultural theories, etc.). This cannot be done without having a general methodology. The answer can be found within the methodology of Labelled Deductive Systems [6, 7], in which formulae have labels. The declarative units have the form $t : A$, where t is a label for a formula A . The label can convey information about the formula as well as whatever “specific” content sensitive conditions we want to impose on the revision process. In general, the labels come from an abstract algebra of labels (which can be the algebraic form of another logic). This kind of labelled revision can be developed in the abstract but is too general for our needs. A more specific way of labelling is to give fuzzy values. We will consider these ideas in a future paper.

6 Conclusions and future work

In this paper, we looked at the problem of voting from a number of different perspectives. In particular, by expressing voters preferences and the layout of the voting system as logical theories. In trying to conciliate these different preferences, we found that this problem has a corresponding counterpart in different formalisms.

We started by considering the problem related to the well known Condorcet paradox. We showed that the theory consisting of the voters' preferences and the voting principles is inconsistent and we considered different ways of restoring consistency. In all cases, we sought to ensure that the voting principles remained satisfied.

For the case of belief revision, the principles can be enforced by revising the preferences of the individual voters by a sentence representing the layout of the system. We claimed that we can only obtain satisfactory revisions (from the voting point of view), by imposing fairness voting conditions at the meta-level. In addition, we showed that full meet revisions will simply result in the voting model itself, without keeping the preferences of any of the individual voters and that partial meet revisions will in general fail to consider the preferences of all individuals. As a consequence, the direct application of an AGM compliant operator is not adequate to this particular kind of problem. However, we believe that an investigation of the applicability of the principles of one of the areas to the other is indeed promising.

We then considered a more general setting in which voters can express *conditional* preferences — they state a number of possible *desirable* outcomes. We found that, in general, the collective outcome of the voting process would still encounter the same cyclical problems related to the Condorcet paradox. On the other hand, model-based belief merging circumvents the paradox by providing a result that is, in a sense, too cautious — it simply does not make a choice between the alternatives. To overcome these difficulties, we proposed a conciliation process, that even though biased towards one of the theories representing the voting process or the voters themselves, compromises on the wishes of the remaining ones. We called this a *sympathetic dictator*.

We found from our analysis that it is possible to revise according to the rules specifying the underlying layout of the system. In other words, how we revise depends on the subject matter, which suggests that we should develop postulates for *context sensitive* revision theory.

In the specific scenario of the voting problem, the development of

a logic-based framework where one can express not only the rules of the voting system but also how individuals formulate and change their preferences would allow us to model the dynamics of the election process as a whole. In particular, the elaboration, clarification and modification of candidates' manifestoes. This is important since it is well known that voters rank candidates according to how well their manifestoes match the voter's preferences. In fact, voters may even apply some reliability parameters in the process. As we saw in this paper, voters can be viewed as logical theories and hence a natural extension would be to consider theories with fuzzy values. We find that the fuzziness should be of the Dempster-Shafer type and that there is a connection with the geometrical cross ratio of projective geometry. We can then offer voting theory a way to compile the wishes of several voters by geometrical means taking the preferences of each voter as a total package.

Finally, techniques from belief merging can help candidates optimise their number of supporters through changes in their manifestoes based on results obtained in opinion polls. It may well be the case that there is an optimal manifesto that can guarantee the election of the candidate who chooses it.

REFERENCES

- [1] C. A. Alchourrón, P. Gärdenfors, and D. Makinson, 'On the logic of theory change: partial meet contraction and revision functions', *The Journal of Symbolic Logic*, **50**, 510–530, (1985).
- [2] H. Andreka, M. D. Ryan, and P.-Y. Schobbens, 'Operators and laws for combining preference relations', *Journal of Logic and Computation*, **12**, 13–53, (2002).
- [3] K. Arrow, *Social choice and individual values*, Cowles Foundation Monograph Series, second edn., 1963.
- [4] K. Arrow, A. K. Sen, and K. Suzumura, *Handbook of Social Choice and Welfare*, volume 1, Elsevier, 2002.
- [5] D. Eckert and G. Pigozzi, 'Belief merging, judgment aggregation, and some links with social choice theory', in *Belief Change in Rational Agents: Perspectives from Artificial Intelligence, Philosophy, and Economics*, eds., J. Delgrande, J. Lang, H. Rott, and J.M. Tallon, Dagstuhl Seminar Proceedings 05321, Internationales Begegnungs- und Forschungszentrum (IBFI), (2005). Available at <http://drops.dagstuhl.de/opus/volltexte/2005/333/pdf/05321.PigozziGabriella1.Paper.333.pdf>.
- [6] D. Gabbay, *Labelled Deductive Systems: Principles and Applications. Basic Principles*, volume 1, Oxford University Press, 1996.
- [7] D. Gabbay, *Labelled Deductive Systems: Principles and Applications. Further Developments*, volume 2, Oxford University Press, In preparation.
- [8] D. Gabbay, O. Rodrigues, and A. Russo, *Information, Uncertainty, Fusion*, chapter Revision by Translation, 3–31, Kluwer Academic Publishers, 1999. ISBN: 0-7923-8590-X.
- [9] P. Gärdenfors, *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, A Bradford Book - The MIT Press, Cambridge, Massachusetts - London, England, 1988.
- [10] P. Gärdenfors, ed. *Belief Revision*, number 29 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1992.
- [11] H. Katsuno and A. O. Mendelzon, 'On the difference between updating a knowledge base and revising it', in *Belief Revision*, ed., P. Gärdenfors, 183–203, Cambridge University Press, (1992).
- [12] S. Konieczny, *Sur la Logique du Changement: Révision et Fusion de Bases de Connaissance*, Ph.D. dissertation, University of Lille, France, 1999.
- [13] S. Konieczny, J. Lang, and P. Marquis, 'Da² merging operators', *Artificial Intelligence*, **157**, 49–79, (2004).
- [14] S. Konieczny and R. Pino-Pérez, *Proceedings of KR '98*, chapter On the logic of merging, 488–498, Morgan Kaufmann, 1998.
- [15] S. Konieczny and R. Pino-Pérez, 'Merging information under constraints: a logical framework', *Journal of Logic and Computation*, **12**(5), 773–808, (2002).
- [16] S. Konieczny and R. Pino-Pérez, 'Propositional belief base merging or how to merge beliefs/goals coming from several sources and some links with social choice theory', *European Journal of Operational Research*, **160**(3), 785–802, (2005).
- [17] F. Rossi, M. S. Pini, K. B. Venable, and T. Walsh, 'Aggregating preferences cannot be fair', in *Proceedings of the 10th Conference on Theoretical Aspects of Rationality and Knowledge*. ACM Digital Library, (June 2005).
- [18] A. K. Sen, *Collective Choice and Social Welfare*, Holden Day, 1970.