Common Foundations for belief revision, belief merging and voting (extended abstract)

D. Gabbay[1], G. Pigozzi[2], O. Rodrigues[1]

[1] King's College London, Department of Computer Science
The Strand, London, WC2R 2LS, UK
{dov.gabbay,odinaldo.rodrigues}@kcl.ac.uk
[2] University of Luxembourg
Computer Science and Communications (CSC)
6 rue Richard Coudenhove Kalergi
L-1359 Luxembourg-Kirchberg
Luxembourg
gabriella.pigozzi@uni.lu

Abstract. In this paper, we consider a number of different ways of reasoning about voting as a problem of conciliating contradictory interests. The mechanisms that do the reconciliation are belief revision and belief merging. By investigating the relationship between different voting strategies and their associated counterparts in revision theory, we find that whereas the counting mechanism of the voting process is more easily done at the meta-level in belief merging, it can be brought to the object level in base revision. In the former case, the counting can be tweaked according to the aggregation procedure used, whereas in base revision, we can only rely on the notion of minimal change and hence the syntactical representation of the voters' preferences plays a crucial part in the process. This highlights the similarities between the revision approaches on the one hand and voting on the other, but also opens up a number of interesting questions.

Keywords. belief revision, belief merging, voting, social choice theory

1 Introduction

In [GPR06], we pointed out that there are a number of similarities between the areas of belief revision, belief merging and voting. These include

- **resolution of conflicts,** which in the case of revision and merging has to do with avoiding/resolving inconsistency and in the case of voting it has to do with conciliation of contradictory voting interests
- **integrity constraints,** which are implicitly given in belief revision in the form of postulates and explictly given in the case of belief merging. In voting these are conditions on the aggregation procedures, e.g., fairness, non-dictatorship, etc.

Dagstuhl Seminar Proceedings 07351
Formal Models of Belief Change in Rational Agents
http://drops.dagstuhl.de/opus/volltexte/2007/1217

– **representation and manipulation of preferences** which gives priority to the new information in belief revision and to the integrity constraints in belief merging. In the case of voting this has to do with the individual voters preferences for candidate options.

The similarities are worth exploring for a number of reasons. In some cases there is a very clear analogy (e.g., computation of a Kemeny consensus and merging of knowledge bases in model-based belief merging). More importantly, by investigating the similarities between the areas, we can bring solutions to problems of one area to the others.

In this paper we seek to explore the underlying principles which are common to these three processes by considering the voting problem under the perspective of belief revision and belief merging and analysing the effects of the operations on the voting scenario. We start by introducing the voting problem used in the remaining of the paper.

2 Voting: aggregating preferences

The problem of 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 [Arr63,ASS02,Sen70]. 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.

Let $C = \{c_1, \ldots, c_k\}$ be the set of candidates. There are exactly $|C| \times (|C| - 1)$ distinct ordered pairs of candidates $\langle c_i, c_j \rangle$. In general, we speak of a binary relation $<$ on C^2, where $c_i < c_j$ denotes that candidate c_i is (strictly) preferred to candidate c_j. The desired properties of preference relations associated to strict linear orders are given below, where the variables $\{x, y, z\}$ range over elements of C.

(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)

In Section 4, conditions (P1)–(P3) will be expressed in our propositional language, but for now we leave them in the meta-level.

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, $C = \{a, b, c\}$) and three voters V_1, V_2 and V_3, who express their total preferences in the following way:

$V_1 = \{a < b, b < c\}$
$V_2 = \{b < c, c < a\}$
$V_3 = \{c < a, a < b\}$

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 $V = \{V_1, \ldots, V_n\}$ in the 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 $\{\leq_1, \ldots, \leq_n\}$ (again we use $\prec$ to denote $\preceq$'s strict counterpart). A tuple of n rankings one for each individual over the set of alternatives is called a *profile*.

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 $\{x, y\} \in X$ and every profile $\langle <_1, \ldots, <_i, \ldots, <_n \rangle$, $x <_i y$ implies $x \prec y$.

(Weak Pareto principle) if for all i, $x <_i y$, then $x \prec y$ (this principle is also called *unanimity*).

3 Aggregating preferences via belief revision

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 $\mathrm{Cn}(\{a < b, b < c, c < a\})$ and consider an input τ to T saying that $<$ is the strictly linear order of the three elements a, b, c, i.e., $\tau = \mathrm{P1} \wedge \mathrm{P2} \wedge \mathrm{P3}$.

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 **??**, 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 $T_\perp \neg\tau = \{T_1, T_2, T_3, ...\}$, $T \circ \tau$ can be constructed from any $T_i \in T_\perp \neg\tau$, say $\mathrm{Cn}(T_1 \cup \{\tau\})$ (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 $A_1, \ldots, A_n, \ldots$ and defining a sequence $S_0 \subseteq S_1 \subseteq S_2, \ldots$ as follows. Let $S_0 = \{\tau\}$ and for $n \geq 0$, define 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:

$T_1 = \mathrm{Cn}(\{a < b, b < c\})$
$T_2 = \mathrm{Cn}(\{b < c, c < a\})$
$T_3 = \mathrm{Cn}(\{a < b, c < a\})$

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

$V_1 = \mathrm{Cn}(\{a < b, b < c, \tau\}$
$V_2 = \mathrm{Cn}(\{b < c, c < a, \tau\}$
$V_3 = \mathrm{Cn}(\{c < a, a < b, \tau\}$

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:

Voter 1 — $a < b$, $b < c$
Voter 2 — $b < c$, $c < a$
Voter 3 — $c < a$, $a < b$

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(Voter 1, Voter 2, Voter 3) = 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. 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.

$a < b$ is voted by V_1, V_3
$b < c$ is voted by V_1, V_2
$c < a$ is voted by V_2, V_3

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, ... 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 a selection function.

Maxichoice revision operations look at all $T_{i'}$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.

4 Voting as belief merging

We now want to express the ideas presented in Section 2 in the context of belief merging. For this we use a propositional language $\mathcal{L}$ and associate to each pair of candidates c_i, c_j taken from C, a propositional variable "$c_i < c_j$" implicitly meaning that candidate c_i is preferred to candidate c_j.

We use $\mathcal{P}$ to denote the set of all propositional variables of $\mathcal{L}$ constructed in this way. Complex formulae of $\mathcal{L}$ are defined as usual. We assume the usual semantics for L and use the symbol $\boldsymbol{W}$ to denote the set of all of its valuations.

A set of propositional variables Δ is a faithful representation of a strict total order $<$ on C if the following conditions are met

(P1): $c_i < c_j \in \Delta$ and $c_j < c_m \in \Delta$ implies $c_i < c_m \in \Delta$ (P1)
(P2) and (P3): for every pair $\{c_i, c_j\}$ of distinct elements $c_i, c_j \in C$, either $c_i < c_j \in \Delta$ or $c_j < c_i \in \Delta$, but not both
(MIN) no other propositional variables appear in Δ

In general, we would like to consider sets containing any complex formulae, so we need to impose some conditions on these sets as to what constitutes a faithful representation of a strict total order on C. We do this by defining the integrity constraint τ in the language $\mathcal{L}$ as follows.

$$\tau = \tau_1 \wedge \tau_2$$

where

$$\tau_1 = \bigwedge_{i \neq j \neq m} [(c_i < c_j \wedge c_j < c_m) \rightarrow c_i < c_m]$$

and

$$\tau_2 = \bigwedge_{i \neq j} [c_i < c_j \vee c_j < c_i] \wedge [c_i < c_j \rightarrow \neg c_j < c_i]$$

τ_1 and τ_2 must be constructed for every distinct pair of candidates taken from C. The first conjunct of τ_2 can be rewritten as $\neg c_i < c_j \rightarrow c_j < c_i$, and hence τ_2 can be rewritten as $c_i < c_j \leftrightarrow \neg c_j < c_i$.

Proposition 1. *For every strict linear order $<'$ on C, there is a valuation w in $\mathrm{mod}(\tau)$ such that $x <' y$ iff $w \Vdash x < y$.*

Proof. This is easy to see. We can simply construct the valuation w, by considering every pair of candidates x, y and making $w(x < y) = 1$, $w(y < x) = 0$ if $x <' y$ and $w(x < y) = 0$, $w(y < x) = 1$, otherwise. We then need to show that $w \in \mathrm{mod}(\tau)$, but this is straighforward. By construction, τ_2 is satisfied by w. τ_1 is also satisfied by the fact that $<'$ is transitive.

Proposition 2. *For every valuation w in $\mathrm{mod}(\tau)$, there exists a strict linear order $<_w$ such that $w \Vdash x < y$ iff $x <_w y$.*

Proof. Pick a valuation $w \in \mathrm{mod}(\tau)$. By construction, w will assign values to every variable $x < y$ associated to every ordered pair $\langle x, y \rangle$ in C. We can then contruct a strict linear order $<_w$, by making $x <_w y$ iff $w \Vdash x < y$. We then need to show that $<_w$ is not reflexive, that it is antisymmetric, transitive and total. Construction guarantees that $<_w$ is not reflexive and that it is total. Since $w \in \mathrm{mod}(\tau)$, then $w \Vdash \tau_1 \wedge \tau_2$. τ_1 guarantees transitivity and τ_2 guarantees antisymmetry.

If there are no extra propositional variables in $\mathcal{L}$ apart from the ones necessary to represent $<$, then there is a correspondence between $\mathrm{mod}(\tau)$ and the set of all strict linear orders. We will assume this is the case here and use $w_<$ to denote the valuation associated with a particular strict linear order $<$ and $<_w$ to denote the strict linear order with a particular valuation w.[1]

Example 1. Suppose our candidates are represented by the set $C = \{a, b, c\}$. For this configuration, τ as defined above would be $\tau = \tau_1 \wedge \tau_2$, where

$$\begin{aligned} \tau_1 = [&((a < b \wedge b < c) \rightarrow a < c) \wedge \\ &((a < c \wedge c < b) \rightarrow a < b) \wedge \\ &((b < a \wedge a < c) \rightarrow b < c) \wedge \\ &((b < c \wedge c < a) \rightarrow b < a) \wedge \\ &((c < a \wedge a < b) \rightarrow c < b) \wedge \\ &((c < b \wedge b < a) \rightarrow c < a)] \end{aligned}$$

$$\begin{aligned} \tau_2 = [&(a < b \leftrightarrow \neg b < a) \wedge \\ &(a < c \leftrightarrow \neg c < a) \wedge \\ &(b < c \leftrightarrow \neg c < b) \wedge \end{aligned}$$

Note that $\mathrm{mod}(\tau)$ has exactly the following six valuations:

[1] There is only significance in the uniqueness when we come to consider distances between valuations for pairs of orders. If there is more than one valuation associated with each order $<$, for a pair of orders $<_1$ and $<_2$, we will be interested in the *minimum* distance between any two valuations associated with them.

prop/val	w_1	w_2	w_3	w_4	w_5	w_6
$a < b$	0	0	0	1	1	1
$a < c$	0	0	1	0	1	1
$b < c$	0	1	1	0	0	1
$b < a$	1	1	1	0	0	0
$c < a$	1	1	0	1	0	0
$c < b$	1	0	0	1	1	0
order	$c < b < a$	$b < c < a$	$b < a < c$	$c < a < b$	$a < c < b$	$a < b < c$

This is easy to see, since τ_2 forces valuations to have symmetrical values for $c_i < c_j$ and $c_j < c_i$. We then only need to consider the eight possible combinations of values for $a < b$, $a < c$ and $b < c$. From these, $w_7(c < b) = 1$, $w_7(b < a) = 1$ and $w_7(a < c) = 1$, and $w_8(a < b) = 1$, $w_8(b < c) = 1$ and $w_8(c < a) = 1$, do not satisfy τ_1. $w_1, \ldots, w_6$ correspond exactly to all possible strict total orders on a set of three elements.

Definition 1. *A* voter V_i *is a set of propositional formulae.*

Voters in general need not be consistent, nor complete in their preferences. However, some voters have special characteristics.

Definition 2. *An* opinionated voter *is a conjunction of propositional formulae.*

In other words, an opinionated voter wants all of his preferences to be considered *atomically*. It will prove useful to relate an opinionated voter with his more flexible counterpart V_i and so we will denote the former using the symbol $\bar{V}_i$. Similarly, a voter may have more or less information about his preferences over candidates:

Definition 3. *A* well-informed voter *is a voter whose propositional formulae admit exactly one chain in C.*

Definition 4. *Let Δ be a set of formulae.*
$\mathrm{mod}(\Delta) = \{w \in W \mid w \Vdash \Delta\}$

In order to use a specific merging operator introduced in [Kon99], we recall the definition of the Hamming distance [Dal88]. However, the Hamming distance is only one among many possible distance functions that we may use. In fact, as shown in [KPP02], the logical properties of a merging operator do not depend on the chosen distance.

Definition 5. *Let w_1 and w_2 be two valuations in $\boldsymbol{W}$. The distance d between w_1 and w_2 is the number of propositional variables p, for which $w_1(p) \neq w_2(p)$.*

We now extend this notion to cater for distances between a single valuation and a set of valuations:

Definition 6. *Let $w \in \boldsymbol{W}$ and $W \subseteq \boldsymbol{W}$. The* pointwise distance *between w and W, in symbols $\mathring{d}(w, W)$ is defined as*

$$\min\{\mathring{d}(w, w') \mid w' \in W\}$$

We now define the model-based majority merging operator, as introduced by Konieczny and Pino-Pérez [KPP98].

Definition 7. *Let $V_1,\ldots,V_n$ be n voters and IC a set of integrity constraints. The majoritarian merging of $V_1,\ldots,V_2$ with integrity constraints IC, in symbols $M_{IC}(V_1,\ldots,V_n)$, is defined as*

$$M_{IC}(V_1,\ldots,V_n) = \{w \in \mathrm{mod}(IC) \mid \sum_{1\leq i\leq n} \mathring{d}(w, \mathrm{mod}(V_i)) \textit{ is minimum}\}$$

Definition 8. *The* Kemeny distance between two profiles P_1 and P_2, *in symbols $d_K(P_1,P_2)$, is defined as*

$$d_K(P_1,P_2) = |\{(c_1,c_2) \in C^2 \mid (c_1,c_2) \in (P_1 \cup P_2) - (P_1 \cap P_2)\}|$$

Proposition 3. *For any two profiles P_1 and P_2,*

$$d_K(P_1,P_2) = d(w_{P_1}, w_{P_2})$$

Proof. Straightforward, since d simply counts the propositional variables with different truth-values in two valuations.

The result above was also proved in [EM05].

Definition 9. *The* Kemeny distance between a profile P and a set of profiles $P_1,\ldots,P_n$, *in symbols $d_K(P,\langle P_1,\ldots,P_n\rangle)$, is defined as*

$$d_K(P,\langle P_1,\ldots,P_n\rangle) = \sum_{1\leq i\leq n} d_K(P,P_i)$$

A profile P is called a Kemeny consensus *of a set of profiles $P_1,\ldots,P_n$, if $d_K(P,\langle P_1,\ldots,P_n\rangle)$ is minimum.*

Definition 10. *A candidate c is called a* Kemeny winner *if there exists a Kemeny consensus P in which c is minimum.*

Proposition 4. *Let $V_1,\ldots,V_n$ be n well-informed voters. $w_p \in M_\tau(V_1,\ldots,V_n)$ iff w_p represents a Kemeny consensus P_w of $\langle V_1,\ldots,V_n\rangle$.*

Proof. Since each voter V_i is well-informed, we can associate him/her to a single profile P_{V_i}. Whether or not $\mathrm{mod}\,(V_i)$ is a singleton is irrelevant, since $\mathring{d}(w, \mathrm{mod}$ (V_i) will pick the element in $\mathrm{mod}(V_i)$ with minimum distance to w. Now, it is easy to see that for each $w \in \mathrm{mod}(\tau)$, there is a profile P_w which represents w and such that

$$\mathring{d}(w, mod(V_i)) = d_K(P_w, P_{V_i})$$

and hence $\sum_{1\leq i\leq n} \mathring{d}(w, mod(V_i)) = \sum_{1\leq i\leq n} d_K(P_w, P_{V_i})$. Since the former is minimum, so is the latter and therefore, P_w is a Kemeny consensus. The converse is proved in a similar way.

Proposition 5. *A candidate x is a Kemeny winner for $\langle V_1, \ldots, V_n \rangle$ iff $\exists w \in M_\tau(V_1, \ldots, V_n)$, such that $w \in \mathrm{mod}(\wedge_{x \neq y} x < y)$.*

Proof. *($\Leftarrow$) Since $w \in M_\tau(V_1, \ldots, V_n)$, by Proposition 4, w represents a Kemeny consensus. Since $w \in \mathrm{mod}(\wedge_{x \neq y} x < y)$, then x is minimum, therefore it is a Kemeny winner.*

($\Rightarrow$) Suppose x is a Kemeny winner for $\langle V_1, \ldots, V_n \rangle$. Then there is a Kemeny consensus P in which x is minimum. Pick w_P, such that w_P represents P. By Proposition 4, $w_p \in M_\tau(V_1, \ldots, V_n)$. Since x is minimum in P, then $x < y$ for every $y \neq x$, and hence $w_p \Vdash \wedge_{x \neq y} x < y$.

Note that in the above characterization, the actual election result was mostly calculated in the meta-level. It was mostly the distance values (calculated outside the logical theory) that determined how the votes were counted. To be more general, we need to include the mechanism of counting the votes in the object level itself.

5 Voting as belief revision

In the previous section, we had one preference relation $<$ and the voters expressed their opinion as to how they *wished* $<$ to be like. τ imposed further conditions on such relation. Belief merging allowed us to do this directly, because the preferences of the different voters could be distinguished implicitly by placing them in different belief bases. This is not the case with belief revision. In order to distinguish the preferences of each voter we will need to do represent them explicitly in the language.

In so doing, we will then also need to code the machinery that computes the overall $<$ from each of the individual voter's preference relations. In order to distinguish the overall social preference relation from that of the individual voters, we will re-formulate the problem by introducing new symbols as follows. We replace $<_M$ for $<$ in τ ($\tau_1 \wedge \tau_2$) and introduce a collection of propositional symbols $\boldsymbol{V_i} = \{c_i <_i c_j\}$, for each distinct pairs of candidates $c_i, c_j \in \mathcal{C}$ and voter v_i.

The set of the candidates and the set of the voters do not necessarily coincide. When they do the above scenario represents a different aggregation problem [AT05].

Now, in general, we will want to represent a voter's V_i preferences through one or more formulae taken from $\boldsymbol{V_i}$. This could be done either by a set of formulae or by a conjunction of formulae. We normally will not want the representation of a voter's preference to be a closed theory.[2] In fact, a number of issues need to be considered when choosing the most appropriate representation. Base revision, for instance is *syntax-dependent* and this has a direct effect on the concept of *minimality*. Revision is done by comparing consistent subsets of the original belief base. In particular,

[2] To be expanded...

a) if we choose the representation of a voter's preferences to be a single formulae, then by doing base revision, the result will incorporate all or none of that voter's preferences. This is the *opinionated* version.
b) if we choose the representation of a voter's preferences to be the set of propositional variables associated with each of his/her preferences for a candidate c_i over c_j, it may be possible to retain at least some of these preferences in the revision process. This approach is more flexible.

As we mentioned at the end of the previous section, we now want to simplify the revision/merging mechanism by placing some of the machinery in the object level. So we will now turn to this problem by considering *base revision* (instead of belief merging), and hence we need to include in the *object-level* some information about what it means to win the election according to some criteria.

Arguably, the simplest principle to incorporate is that of *majority*. We would like to express under what circumstances the majority of voters prefer one candidate to another. This of course depends on the number of voters. In our example with three voters, a candidate c_i is preferred over candidate c_j by the majority of voters if and only if any two voters prefer c_i over c_j:

$$\begin{aligned} c_i <_M c_j \leftrightarrow [&((c_i <_1 c_j) \wedge (c_i <_2 c_j))\vee \\ &((c_i <_1 c_j) \wedge (c_i <_3 c_j))\vee \\ &((c_i <_2 c_j) \wedge (c_i <_3 c_j))] \end{aligned}$$

We need a formula such as the above for each combination c_i, c_j such that $i \neq j$. In other words, majority can be represented as the formula θ_M

$$\begin{aligned} \theta_M = \wedge_{c_i, c_j \in \mathcal{C}, i \neq j}\{c_i <_M c_j \leftrightarrow [&((c_i <_1 c_j) \wedge (c_i <_2 c_j)) \vee \\ &((c_i <_1 c_j) \wedge (c_i <_3 c_j)) \vee \\ &((c_i <_2 c_j) \wedge (c_i <_3 c_j))]\} \end{aligned}$$

This will only tell us that a particular candidate c_i is preferred (by the majority) over another candidate c_j, when one of the disjuncts on the right of the formula is true, i.e., two out of three of the voters. For larger number of voters, one needs to calculate what the majority number m of voters is and write formulae accordingly for each combination of m voters – a tedious and repetitive process, but nevertheless easy to do. In order to be a *Condorcet winner*, a candidate c_i needs to be preferred by the majority over *every* other candidate c_j. Hence, for each $c_i \in \mathcal{C}$,

$$CW(c_i) = \wedge_{i \neq j} c_i <_M c_j$$

And there will be such a winner in a voting scenario, if one of the candidates is a Condorcet winner, i.e.,

$$ECW = \vee_{c_i \in \mathcal{C}} CW(c_i)$$

Now that the machinery (i.e., the voting mechanism) is expressed in the logic's language itself, we can consider different combinations of revision mechanisms and constraints and compare what we get as the result of the revision.

Note that we can at this point, choose to revise the machinery itself, but let us leave this point for later and as such, require that the machinery is preserved by the revision process. In order to accomplish this, all we need is to revise *by* the machinery, because of the success postulate. We then need to decide what kind of revision we want.

AGM belief revision itself is independent of the syntax form of the formulae. This is due to the fact that what is revised is in fact a *theory*. As a result, voters' preferences contain a lot more information than can be at first realised and one needs to be careful to constrain exactly what is allowed to remain from that original theory. In our case, each propositional symbol in a voter's preference representation is associated with a vote of that voter for one candidate over another. In doing the revision we want to minimise the loss of these propositions. Therefore, one reasonable way of seeing the process is to consider *base revision* instead. Proper subsets of a voter's representation will be associated with the failure to satisfy all of that voter's original preferences.

There are two ways of calculating the change to the original base. The first one is to consider all maximal subsets of a belief base K that fail to imply a sentence α, in symbols, $K_\perp\alpha$:

$$K_\perp\alpha = \{K' \subseteq K \mid K' \not\vdash \alpha \text{ and for every } \; K'' \supset K',\; K'' \vdash \alpha\}$$

Using the above notion, it is possible to define the revision of a belief base K by a formula α, in symbols $K \star \alpha$ by picking one element of $K_\perp\neg\alpha$ if any exists and then expand it by α. If $K_\perp\neg\alpha$ is empty, we can simply take α itself as the result of the revision:

$$K \star \alpha = \begin{cases} K' \cup \alpha \text{ for some } K' \in K_\perp\neg\alpha \\ \{\alpha\} \qquad \text{if } K_\perp\neg\alpha = \emptyset \end{cases}$$

This is the *traditional* way of evaluating minimal change to K. One could also consider the subsets of K that are consistent with α and have *maximum* cardinality (i.e., failure of set inclusion between sets is not sufficient to qualify a set as maximally consistent). This can be formalised as follows:

$$K_{\perp_c}\alpha = \{K' \subseteq K \mid K' \not\vdash \alpha \text{ and for every } \; K'' \subseteq K,\; |K''| > |K'|,\; K'' \vdash \alpha\}$$

In [Kon99, Chapter 9], Konieczny's investigated syntactic fusion operators. One of the operators defined was a syntactic merging operator that selects a consistent set with the maximum cardinality among all maximally consistent subsets.

The base revision according to this policy can be defined in a similar way.

$$K \star_c \alpha = \begin{cases} K' \cup \alpha \text{ for some } K' \in K_{\perp_c}\neg\alpha \\ \{\alpha\} \qquad \text{if } K_{\perp_c}\neg\alpha = \emptyset \end{cases}$$

Definition 11 (Young winner). *A Young winner is a candidate that can be made a Condorcet winner by the least removal of voters.*

Proposition 6. *Let* $B = \cup_i V_i$ *and* $\alpha = \tau \wedge ECW \wedge \theta_M$. $B \star_c \alpha \vdash CW(c_i)$ *iff* c_i *is a Young winner for* $\langle V_1, \ldots, V_n \rangle$.

Proof. ($\Rightarrow$) Suppose that $B \star_c \alpha \vdash CW(c_i)$. We show that c_i is a Young winner for $\langle V_1, \ldots, V_n \rangle$.

We first look at the sentence revising the belief base:

- τ says that $<$ is a linear order
- θ_M defines what majority of one candidate over another means
- ECW requires that either a, b or c has a majority over every other candidate

These three formulae together basically require the revision to pick one total linear order over the three candidates (this will be determined by the selection function for $\star_c$). The complications now have to do with the way we are allowed to "count" the propositions in B (i.e., the votes).

Let us assume that the voters are *opinionated*. We know from the properties of belief revision that $B \star_c \alpha$ is consistent (since α is consistent), and that any eventual inconsistency with B will be resolved by removing sentences from B. Each sentence removed means "a voter" (since voters are opinionated). Suppose that $B \star_c \alpha \vdash CW(c_i)$, but c_i is not a Young winner. Further to this, take B' to be the set in $B_{\perp_c} \neg\alpha$ chosen by the selection function of $\star_c$

Since c_i is a Condorcet winner, there must be another candidate $c_j \neq c_i$ who can be made a Condorcet winner by the removal of less voters from B than the number of voters removed to make c_i a Condorcet winner. In other words, if B'' is set with the remaining voters, then it must be the case that $|B''| > |B'|$, but this is a contradiction, since $B' \in B_{\perp_c} \neg\alpha$.

($\Leftarrow$) This is easier to prove. Suppose that c_i is a Young winner for $\langle V_1, \ldots, V_n \rangle$, but that $B \star_c \alpha \not\vdash CW(c_i)$. Since c_i is a Young winner for $\langle V_1, \ldots, V_n \rangle$, then it can be made a Condorcet winner for $\langle V_1, \ldots, V_n \rangle$ by the least removal of voters from B. Since each voter is represented by a single sentence, then we must have that

1. there is some $B' \in B_{\perp_c} \neg\alpha$, such that the cardinality of B' is maximum and
2. c_i is a Condorcet winner for B'

Therefore, $B \star_c \alpha = B' \cup \alpha \vdash CW(c_i)$, a contradiction.

However, notice that picking a voter as a set of propositional variables as opposed to a conjunction of those symbols may potentially change the result (needs to be proven!). Also, maybe a majoritarian number of voters have more similar preferences and hence a dictator may not be elected.

Example 2. Consider the scenario described earlier on on this paper. $\tau = \tau_1 \wedge \tau_2$ is exactly as in Example 1, except that we replace $<_M$ for $<$:

$$\begin{aligned}
\tau_1 = [&((a <_M b \wedge b <_M c) \rightarrow a <_M c) \wedge \\
&((a <_M c \wedge c <_M b) \rightarrow a <_M b) \wedge \\
&((b <_M a \wedge a <_M c) \rightarrow b <_M c) \wedge \\
&((b <_M c \wedge c <_M a) \rightarrow b <_M a) \wedge \\
&((c <_M a \wedge a <_M b) \rightarrow c <_M b) \wedge \\
&((c <_M b \wedge b <_M a) \rightarrow c <_M a)]
\end{aligned}$$

$$\begin{aligned}
\tau_2 = [&(a <_M b \leftrightarrow \neg b <_M a) \wedge \\
&(a <_M c \leftrightarrow \neg c <_M a) \wedge \\
&(b <_M c \leftrightarrow \neg c <_M b) \wedge
\end{aligned}$$

The three voters are represented as the sets $V_1 = a < b < c = \{a <_1 b, b <_1 c, a <_1 c\}$; $V_2 = b < c < a = \{b <_2 c, c <_2 a, b <_2 a\}$ and $V_3 = c < a < b = \{c <_3 a, a <_3 b, c <_3 b\}$. θ_M will be the conjunction of formulae such as the one below

$$a <_M b \leftrightarrow [(a <_1 b \wedge a <_2 b) \vee (a <_1 b \wedge a <_3 b) \vee (a <_2 b \wedge a <_3 b)]$$

and we will have one conjunct for each distinct pair $c_i, c_j \in \mathcal{C}$. A candidate c_i is a Condorcet winner if the formula $CW(c_i)$ is true. For, say candidate a, this is

$$CW(a) = a <_M b \wedge a <_M c$$

and there will be a Condorcet winner if the formula above is true for one of the candidates, i.e.,

$$(a <_M b \wedge a <_M c) \vee (b <_M a \wedge b <_M c) \vee (c <_M a \wedge c <_M b)$$

Now what happens when we revise $\boldsymbol{V} = V_1 \cup V_2 \cup V_3 = \{a <_1 b, b <_1 c, a <_1 c, b <_2 c, c <_2 a, b <_2 a, c <_3 a, a <_3 b, c <_3 b\}$ by $\tau \wedge ECW \wedge \theta_M$? Since $\boldsymbol{V} \wedge \theta_M \vdash a <_M b \wedge b <_M c \wedge c <_M a$ and the latter formula is inconsistent with τ, we need to retract some formulae and because of the way the revision operation is defined, these can only be removed from $\boldsymbol{V}$.

The inconsistency here can be solved by breaking the cycle $a <_M b \wedge b <_M c \wedge c <_M a$, i.e., by removing from $\boldsymbol{V}$ formulae that support the derivation of any of $a <_M b$ or $b <_M c$ or $c <_M a$. These were obtained by the formulae in θ_M, which defines majority. Therefore, the options are the ones enclosed in a square in the tree below.

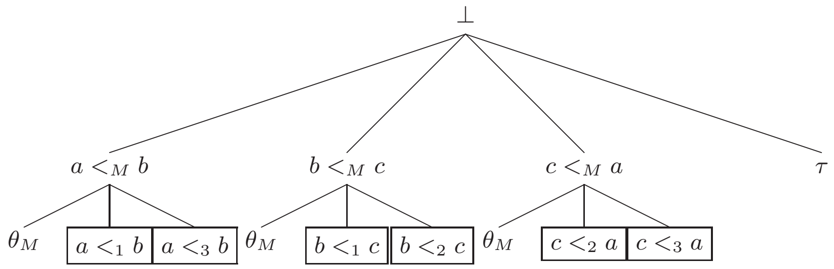

There are hence six subsets of $\boldsymbol{V}$ that are maximally consistent with $\tau \wedge ECW \wedge \theta_M$ (with respect to set inclusion). One just needs to remove any of the formulae in the boxes above from $\boldsymbol{V}$. Each removal will be associated with a linearization of the candidates a, b and c in $<_M$. For instance, $B' = B \backslash \{a <_1 b\} \wedge \alpha \vdash b <_M c \wedge b <_M a$, and hence we also get that $CW(b)$. In this example, any difference between $\star$ and $\star_c$ is only determined by the choice made by the selection function. In other words, the maximally consistent subsets will all have maximal cardinality.

Note that in this particular example, any of the three candidates can be made a Young winner by removing exactly two of the voters (the other two). As we said, the choice of winner will depend on the selection function used in $\star_c$.

6 Conclusions

In this paper, we have considered a number of different ways of reasoning about voting as a problem of conciliating contradictory interests. The machineries that do the reconciliation are belief revision and belief merging. The basic idea is to code a particular voting scenario in a logical theory and analyse what belief revision and belief merging do there.

Both belief merging and belief revision are based on an underlying notion of minimal change which supports informational economy. However, there are two major differences: 1) in the case of model-based belief merging, the principle of minimal change is applied on top of an extra aggregation step performed to compute the overal "distance" of a model of the integrity contraints to the collection of belief bases 2) belief merging in general makes an explicit distinction of the bases to be merged whereas in belief revision there is just one object of revision. These differences need to be taken into account when we see the problem of voting under the perspective of information change.

Since there is no structural distinction between bases in belief revision, when we represent the voting problem in this scenario the only way we can differentiate between voters's preferences is to enrich the language, such that their preferences can be kept apart, e.g., $a <_1 b$ and $b <_2 a$ for voters 1 and 2. In addition, we also need to "code" the counting mechanism in the logic, for instance by saying what it means for the majority of voters to prefer a to b.

Now to mimick the behaviour of the voting procedure in belief revision (resp., belief merging) we need to devise a logical theory of voting based on the particular voting mechanism in such a way that the elementary unit of change in belief revision (belief merging) matches that of the voting scenario. We showed that this is possible for the Kemeny and Young procedures.

By bringing the voting mechanims to the object level, we gain the ability to reason about them in the language and more importantly the ability to use the belief revision (resp., belief merging) m echanism itself to *modify* it.

So far we have limited ourselves to mimicking voting procedures in the information change context. We would like to investigate in more detail the relationship between other voting procedures and revision and merging to see what kind of voting procedure we would obtain for a particular revision/merging strategy.

Acknowledgements

We would like to thank Jérôme Lang for invaluable comments made on an initial version of this paper presented at the Multidisciplinary Workshop on Advances in Preference Handling at ECAI 2006.

References

Arr63. K. Arrow. *Social choice and individual values.* Cowles Foundation Monograph Series, second edition, 1963.

ASS02. K. Arrow, A. K. Sen, and K. Suzumura. *Handbook of Social Choice and Welfare*, volume 1. Elsevier, 2002.

AT05. A. Altman and M. Tennenholtz. The axiomatic foundations of ranking systems, 2005.

Dal88. M. Dalal. *Proceedings of AAAI-88*, chapter Investigations into a theory of knowledge base revision: preliminary report, pages 475–479. 1988.

EM05. D. Eckert and J. Mitlöhner. *Multidisciplinary IJCAI-05 Workshop on Advances in Preference Handling*, chapter Logical representation and merging of preference information, pages 85–87. 2005.

GPR06. D. M. Gabbay, G. Pigozzi, and O. Rodrigues. Belief revision, belief merging and voting. In *Proceedings of the Seventh Conference on Logic and the Foundations of Games and Decision Theory (LOFT06)*, pages 71–78. University of Liverpool, 2006.

Kon99. S. Konieczny. *Sur la Logique du Changement: Révision et Fusion de Bases de Connaissance*. PhD thesis, University of Lille, France, 1999.

KPP98. S. Konieczny and R. Pino-Pérez. *Proceedings of KR'98*, chapter On the logic of merging, pages 488–498. Morgan Kaufmann, 1998.

KPP02. S. Konieczny and R. Pino-Pérez. Merging information under constraints: a logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.

Sen70. A. K. Sen. *Collective Choice and Social Welfare*. Holden Day, 1970.