Belief merging, judgment aggregation and some links with social choice theory*

Daniel Eckert
Institute of Public Economics, University of Graz
daniel.eckert@uni-graz.at

Gabriella Pigozzi†
Department of Computer Science, King's College London
gabriella.pigozzi@kcl.ac.uk

Abstract

In this paper we explore the relation between three areas: judgment aggregation, belief merging and social choice theory. Judgment aggregation studies how to aggregate individual judgments on logically interconnected propositions into a collective decision on the same propositions. When majority voting is applied to some propositions (the premises) it may however give a different outcome than majority voting applied to another set of propositions (the conclusion). Starting from this so-called doctrinal paradox, the paper surveys the literature on judgment aggregation (and its relation to preference aggregation), and shows that the application of a well known belief merging operator can dissolve the paradox. Finally, the use of distances is shown to establish a link between belief merging and preference aggregation in social choice theory.

*A preliminary version of this paper has been presented at the Dagstuhl seminar *Belief change in rational agents: perspectives from artificial intelligence, philosophy and economics* in August 2005. We thank the organizers of the seminar, James Delgrande, Jérôme Lang and Hans Rott, and the participants for stimulating discussion.

†Gabriella Pigozzi's research was supported by the Economic and Social Research Council (UK), the Alexander von Humboldt Foundation, the German Federal Ministry of Education and Research, and by the German Program for the Investment in the Future.

Dagstuhl Seminar Proceedings 05321
Belief Change in Rational Agents: Perspectives from Artificial Intelligence, Philosophy, and Economics
http://drops.dagstuhl.de/opus/volltexte/2005/333

1 Introduction

Social choice theory (Arrow 1963, Arrow *et al.* 2002, Sen 1970) studies the aggregation of individual preferences in order to select a collectively preferred alternative. Way back in 1770, the Marquis de Condorcet proposed a method for the aggregation of preferences which led to the first aggregation problem: the voting paradox. 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 social ordering is obtained by a combination of all partial results. The paradoxical result is that the pairwise majority rule can lead to cycles.

More recently, scholars in law and social choice theorists have become interested in other aggregation problems like the doctrinal paradox (or discursive dilemma).[1] It has been shown that the discursive dilemma is a generalization of the paradox of voting (List and Pettit 2002). However, unlike the voting paradox, the doctrinal paradox arises when the members of a group have to make a judgment (in the form of yes or no) on specific propositions rather than express preferences among candidates. A possibility result (Pigozzi 2004) is obtained when an operator defined in artificial intelligence to merge knowledge bases is applied to the doctrinal paradox.

This paper surveys the literature on judgment aggregation, summarizes the results obtained when a distance-based merging operator is applied to the discursive dilemma, and establishes the link to distance-based approaches in social choice theory. Section 2 provides an example of the doctrinal paradox. Section 3 summarizes the major results on the relation between judgment aggregation and preference aggregation. Section 4 illustrates how the doctrinal paradox dissolves when a well known majority merging operator is imported into judgment aggregation. Section 5 establishes some links to distance-based approaches in social choice theory.

2 The doctrinal paradox

The doctrinal paradox can emerge when the members of a group have to make a judgment (in the form of yes or no) on several logically interconnected propositions, and the individually logically consistent judgments need to be combined into a collective decision. For example, consider a set of propositions, where some (the ‘premises’) are taken to be equivalent to another proposition (the ‘conclusion’). When majority voting is applied to the

[1] For a comprehensive bibliography of the rapidly growing body of literature on the doctrinal paradox, see List (2005).

premises it may give a different outcome than majority voting applied to the conclusion. This phenomenon did first draw the attention of researchers in law (Kornhauser 1992, Kornhauser and Sager 1986, 1993), who illustrated the paradox with the following example. Suppose that three judges have to decide whether a defendant is liable for breaching a contract. According to the legal doctrine, a person is liable of a certain action X (this is the proposition R) if and only if the defendant performed the action X (P) and had contractual obligation not to do X (Q). Now assume that each judge makes a consistent judgment over the propositions P, Q and R, as the following table shows:

	P	Q	R	$(P \wedge Q) \leftrightarrow R$
Judge 1	Yes	Yes	Yes	Yes
Judge 2	Yes	No	No	Yes
Judge 3	No	Yes	No	Yes
Majority	Yes	Yes	No	Yes

Each individual consistently assigns a truth value to each proposition P, Q and R (saying yes to R if and only if both P and Q are believed to be true). However, if majority voting is applied only to the premises (P, Q) of the argument (this procedure is called *premise-based procedure*), the result is that there is a majority that believes both P and Q to be true (and, therefore, because of $(P \wedge Q) \leftrightarrow R$, that majority is held to believe that R is also true). At the same time, if majority voting is applied only to the conclusion R (*conclusion-based procedure*), the majority of the group believes that R is false, which conflicts with the aggregation of the premises.[2] The paradox lies precisely in the fact that the two procedures may lead to contradictory results (one accepting and the other rejecting R), depending on whether the majority is taken on the individual judgments of P and Q, or whether the majority is calculated on the individual votes of R. The question is then whether a collective outcome exists in these cases, and if it does, what it is like.[3]

[2]See Pigozzi (2004) for a criticism of the premise-based procedure and the conclusion-based procedure.

[3]Condorcet Jury Theorem provides an epistemic justification of majority voting. Suppose that any two pairs of alternatives are given, each member of a group has a probability greater than 0.5 to vote for the right choice and the group votes through a majority rule. Under specific additional (independence) assumptions, the theorem states that the group's probability of choosing the right alternative increases with the size of the group itself and approaches 1 in the limit. List and Goodin (2001) have generalized the Condorcet Jury Theorem to multiple propositions.

A possibility result (see Section 4) is obtained when a merging operator, originally introduced in computer science to merge several finite sets of information, is imported and applied to the problem of judgment aggregation. The justification for this move is that the theory of information merging and group decision-making share a similar difficulty, viz. the definition of operators that produce collective knowledge from individual knowledge bases, and operators that produce a collective decision from individual decisions.

In the next session we will mention some of the results concerning the relation between judgment aggregation and preference aggregation.

3 Judgment and preference aggregation

Let us first introduce some terminology. Given a finite set of n individuals, the finite set X of propositions on which the individuals have to make their judgments is called an *agenda*. A (individual or collective) *judgment set* is a subset $A \subseteq X$, where $p \in A$ means that the proposition p is accepted in A. Individual judgment sets are usually assumed to be consistent. Following Dokow and Holzmann (2005), a judgment set can be represented by a *binary evaluation* $a : X \rightarrow \{0, 1\}$ where, for all $p \in X$, $a(p) = 1$ iff $p \in A$.

A *profile* is a n-tuple $(A_1, \ldots, A_n)$ of individual judgment sets. Finally, a (judgment) *aggregation rule* is a function f that assigns to each profile $(A_1, \ldots, A_n)$ a collective judgment set $f(A_1, \ldots, A_n) \subseteq X$.

An impossibility theorem for judgment aggregation has been proved by List and Pettit (2002). It says that there exists no aggregation procedure (generating complete, consistent and deductively closed collective sets of judgments) which satisfies the following conditions: universal domain, anonymity and systematicity. *Universal domain* states that an aggregation procedure accepts as admissible input any logically possible profile of individual sets of judgments. *Anonymity* ensures that all individuals have equal weight in determining the collective sets of judgments, while *systematicity* is the condition requiring that the aggregation procedure treats all propositions in an evenhanded way.

Social choice theory studies how to define a collective preference relation from individual preference relations. Arrow's theorem is the famous impossibility theorem in social choice. Even though judgment aggregation and preference aggregation are distinct research areas, they show some similarities. As the problem of preference aggregation is illustrated by Condorcet's paradox, so the difficulties of judgment aggregation are illustrated by the doctrinal paradox. A natural question is then whether there is a relation and, if so, of what kind, between the two problems. List and Pettit (2002) show that judgment aggregation is a generalization of the preference aggregation framework.

In particular, Dietrich and List (2005) show that Arrow's theorem (for strict preferences) can be stated in the judgment aggregation framework, and prove Arrow's theorem to be a corollary of a general impossibility theorem on judgment aggregation.

It is straight-forward to see that Condorcet's paradox and most preference aggregation problems can be embedded into a judgment aggregation framework by representing preference orderings as sets of binary ranking judgments. While preferences over some set S of alternatives are usually represented as binary relations $R \subset S \times S$, they can as well be given a logical representation by defining the agenda such that $X = \{(x > y) | x \neq y\}$ is a set of atomic preference judgments. Then, for any individual preference judgment set $A \subseteq X$, any preference relation R such that $\{(x, y) | (x > y) \in A\} \subseteq R$ can be interpreted as a model of A.

The usual properties of preferences can be imposed with the help of preference postulates. The postulates characterizing preference judgment sets corresponding to linear orders are given, for all distinct $x, y, z \in S$ and all $A \subseteq X$ by:

P1: $(x > y) \in A \vee (y > x) \in A$ (completeness)
P2: $(x > y) \in A \rightarrow (y > x) \notin A$ (asymmetry)
P3: $(x > y) \in A \wedge (y > z) \in A \rightarrow (x > z) \in A$ (transitivity) (see Hansson 1995).

The following table lists the preference judgment sets corresponding to the set of all linear orders on a set of three alternatives:

	$x > y$	$y > x$	$y > z$	$z > y$	$x > z$	$z > x$
a_1	1	0	1	0	1	0
a_2	0	1	1	0	0	1
a_3	1	0	0	1	0	1
a_4	1	0	0	1	1	0
a_5	0	1	0	1	0	1
a_6	0	1	1	0	1	0

The mere logical representation of preferences and the reformulation of the problem of preference aggregation in terms of the aggregation of atomic preference judgments and preference postulates does of course not dissolve the paradoxes of preference aggregation, as the example of the aggregation of three sets of preference judgments leading to the doctrinal paradox shows.

	$x > y$	$y > x$	$y > z$	$z > y$	$x > z$	$z > x$	P
a_1	1	0	1	0	1	0	1
a_2	0	1	1	0	0	1	1
a_3	1	0	0	1	0	1	1
Majority	1	0	1	0	0	1	1

(P stands for the preference postulates for linear orders)

This example of judgment aggregation corresponds to the problem of the aggregation of three transitive preferences where the application of the majority rule leads to the familiar Condorcet paradox. Majority voting on these propositions produces an inconsistent result, as the preference postulates for linear orders are unanimously accepted, while the outcome under proposition wise majority voting violates transitivity.

Fortunately the fact that preference relations are the natural models of preference states facilitates the use of model-based merging operators, which can be shown to dissolve judgment aggregation paradoxes (Pigozzi 2004). This is addressed in the next section.

4 The doctrinal paradox and belief merging

Belief merging formally investigates how to aggregate a finite number of belief bases into a collective one. Its formal framework consists of a propositional language $\mathcal{L}$ which is built up from a finite set $\mathcal{P}$ of propositional letters standing for atomic propositions and the usual connectives ($\neg$, $\vee$, $\wedge$, $\rightarrow$, $\leftrightarrow$). A *belief base* K_i is a consistent finite set of propositional formulas from the agent i (and corresponds to the individual judgment set in judgment aggregation). A *belief set* is a set $E = \{K_1, K_2, \ldots, K_n\}$ (what is called a profile in judgment aggregation). Given a set of integrity constraints IC (i.e., some extra conditions imposed on the result of the merging operator), Δ (similarly as the aggregation rule F in the previous section) maps E and IC into a new (collective) belief base $\Delta_{IC}(E)$. The outcome should keep maximal information from each K_i. This is achieved using a distance-based approach such that the total distance between the individual bases and the collective outcome is minimized. In a model-based framework the models of $\Delta_{IC}(E)$ are models of IC that differ minimally from the models of each K_i.

An interpretation is a function $\mathcal{P} \rightarrow \{0, 1\}$. Let $\mathcal{W} = \{0, 1\}^{\mathcal{P}}$ denote the set of all interpretations. For any formula $\varphi \in \mathcal{L}$, $[\varphi] = \{\omega \in \mathcal{W} | \omega \models \varphi\}$ denotes the set of models of φ, i.e. the set of interpretations which validate φ in the usual classical truth functional way.[4] Conversely, for any set of

[4]If a formula φ is validated by exactly one model $\omega \in \mathcal{W}$ the latter will be written ω_φ.

models $M \subset \mathcal{W}$ let $form(M)$ denote the propositional formula (up to logical equivalence) such that $[form(M)] = M$, i.e. the formula the models of which are precisely M. A belief base K_i is a finite set of propositional formulae which will be identified with the conjunction of its elements. Denote by $\mathcal{K}$ the set of all consistent belief bases. (It is usually assumed that belief bases are consistent.)

In a model-based framework a merging operator (with integrity constraint $IC \in \mathcal{K}$) $\Delta_{IC} : \mathcal{K}^n \to \mathcal{K}$ is defined by a correspondence $m_{IC} : \mathcal{K}^n \to [IC] \subset \mathcal{W}$ from the set of all possible profiles of belief bases to the set of all models of IC such that for all $\underline{K} \in \mathcal{K}^n$, $\Delta_{IC}(\underline{K}) = form(m_{IC}(\underline{K}))$.[5]

Most model-based merging operators $m_{IC} : \mathcal{K}^n \to [IC] \subset \mathcal{W}$ are based on the selection of the interpretation(s) that minimize(s) distance to the collection of sets of models $\{[K_1], [K_2], ..., [K_n]\}$ corresponding to a profile $\underline{K} = (K_1, K_2, ..., K_n) \in \mathcal{K}^n$ of belief bases.

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

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

For any interpretation $\omega \in \mathcal{W}$ and any belief base $K \in \mathcal{K}$, let $d(\omega, K) = \min_{\omega' \in [K]} d(\omega, \omega')$, i.e. the distance between an interpretation and a belief base is the minimal distance over the models of the latter.

For any interpretation $\omega \in \mathcal{W}$ and any profile of belief bases $\underline{K} \in \mathcal{K}^n$ the distance between an interpretation and a profile can now be defined with the help of an aggregation function $D : \mathbb{R}^n_+ \to \mathbb{R}_+$ as $D^d(\omega, \underline{K}) = D\left(d(\omega, K_1), d(\omega, K_2), ..., d(\omega, K_n)\right)$ (Konieczny, Lang and Marquis 2004). Any such aggregation function gives rise to a total pre-order $\leqslant_{\underline{K}}$ of the set $\mathcal{W}$ of all interpretations with respect to their distance to a given profile of belief bases. Thus the merging operator defined for every $\underline{K} \in \mathcal{K}^n$ by $m_{IC}(\underline{K}) = \min([IC], \leqslant_{\underline{K}})$ selects the set of all interpretations with minimal distance to the profile of belief bases.

Obviously, the properties of the merging operator essentially depend on the distance functions d and D. However, no specific characterization results are available to evaluate the appropriateness of a given merging operator for a given aggregation problem.

The most widely used merging operator in the belief merging literature is the operator denoted by $\Delta^{d,\Sigma}$, where:

[5] For reasons of convenience, a merging operator $\Delta : \mathcal{K}^n \to \mathcal{K}$ will be identified in the following with the correspondence $m : \mathcal{K}^n \to \mathcal{W}$ that defines it.

1. d is the Hamming distance, which is defined by the number of propositional letters on which two interpretations differ, or, formally, $d(\omega, \omega') = |\{\pi \in \mathcal{P} | \omega(\pi) \neq \omega'(\pi)\}|$, and where

2. $D^d(\omega, \underline{K}) = \sum_i d(\omega, K_i)$ is the sum of componentwise Hamming distances.[6]

In the original example of the doctrinal paradox, all three judges agree that a person is liable (R) only if he did a certain action X (P) and had contractual obligation not to do X (Q); that is they accept $(P \wedge Q) \leftrightarrow R$. Therefore $E = \{K_1, K_2, K_3\}$ and $IC = \{(P \wedge Q) \leftrightarrow R\}$. Each judge makes a judgment on P, Q and R that satisfies the integrity constraint. The three belief bases are:

$K_1 = \{P, Q, R\}$
$K_2 = \{P, \neg Q, \neg R\}$
$K_3 = \{\neg P, Q, \neg R\}$

with corresponding interpretations $\{(1, 1, 1)\}, \{(1, 0, 0)\}$ and $\{(0, 1, 0)\}$.

The table below shows the result of the IC majority merging operator on $E = \{K_1, K_2, K_3\}$. The first column lists all the consistent evaluations for the propositional variables P, Q and R. The numbers in the columns of $d(., K_1)$, $d(., K_2)$ and $d(., K_3)$ are the Hamming distances of each K_i from the correspondent interpretation. Finally, in the last column is the sum of the distances over all belief bases.

	$d(., K_1)$	$d(., K_2)$	$d(., K_3)$	$D^d(., \underline{K})$
$(1, 1, 1)$	0	2	2	4
$(1, 0, 0)$	2	2	0	4
$(0, 1, 0)$	2	0	2	4
$(0, 0, 0)$	3	1	1	5

Thus, $\Delta^{d,\Sigma}_{IC}(\underline{K}) = form\,(\min([IC], \leqslant_{\underline{K}})) = form(\{(1, 1, 1), (1, 0, 0), (0, 1, 0)\}) = K_1 \vee K_2 \vee K_3$. Any of the individual judgment set is of minimal distance to the profile of judgments. The result is that the application of the merging operator $\Delta^{d,\Sigma}_{IC}$ allows to avoid the paradox and to obtain a tie instead.

[6] Early references to this operator include (Lin and Mendelzon 1999) and (Revesz 1997).

5 Some links with distance-based approaches in social choice theory

The use of distances establishes a close link between belief merging and social choice theory.[7] As the original Condorcet paradox can be expressed in a similar way as the doctrinal paradox (see Section 3), it comes as no wonder that it can also be dealt with the help of the same distance-based merging operator $\Delta_{IC}^{d,\Sigma}$. Computation of the distance between the evaluations for preference judgments (corresponding to linear orders) and the profile of preference judgment sets in our example yields the following:

	$d(.,K_1)$	$d(.,K_2)$	$d(.,K_3)$	$D^d(.,\underline{K})$
a_1	0	4	4	8
a_2	4	0	4	8
a_3	4	4	0	8
a_4	2	6	2	10
a_5	6	2	2	10
a_6	2	2	6	10

Thus, $\Delta_P^{d,\Sigma}(\underline{K}) = form\,(\min([P], \leqslant_{\underline{K}})) = form(\{a_1, a_2, a_3\}) = K_1 \vee K_2 \vee K_3$. Any of the individual preferences in our example of the Condorcet paradox is of minimal distance to the profile of preference judgment sets such that the application of the merging operator $\Delta^{d,\Sigma}$ dissolves the paradox at the price of indecision.

On the one hand typical social choice problems can be dealt with a distance-based merging operator like the merging operator $\Delta^{d,\Sigma}$. On the other hand $\Delta^{d,\Sigma}$ can be shown to be equivalent to a widely known distance-based social aggregation rule, Kemeny's rule (for the following, see Eckert and Mitlöhner 2005). This rule operates with the minimization of a distance function which is based on the symmetric difference between two binary relations considered as sets of ordered pairs of alternatives.

Definition 1 *Consider the set $L(X)$ of all linear orders on some finite set X of alternatives. Define for any $L, L' \in L(X)$ the distance $d_K(L, L') = |L\backslash L' \cup L'\backslash L|$ as the cardinality of the symmetric difference between the two linear orders. For every profile $\underline{L} = (L_1, L_2, ..., L_n) \in L(X)^n$ the Kemeny ranking $\leqslant_{\underline{L}}^K$ on $L(X)$ is defined by $L' \leqslant_{\underline{L}}^K L''$ if and only if $\sum_i d_K(L', L_i) \leq$*

[7] For a survey on distance-based approaches in social choice theory and recent applications see Baigent (2005) and Klamler (2005).

$\sum_i d_K(L'', L_i)$. *Then the (set valued) aggregation rule* $f : L(X)^n \to L(X)$ *is Kemeny's rule if for every* $\underline{L} \in L(X)^n$ $f(\underline{L}) = \min(L(X), \leqslant^K_{\underline{L}})$.

Proposition 2 *For the above preference aggregation problem, the belief merging operator* $\Delta^{d,\Sigma}$ *is equivalent to Kemeny's rule.*

Proof. *Consider the mapping* $c : L(X) \to \mathcal{K}$*, which associates with every linear order* $L \in L(X)$ *the complete preference judgment set* $K \in \mathcal{K}$ *such that, for all* $x, y \in X$*,* $x > y \in K$ *if and only if* $(x, y) \in L$*. It is easily verified that for any liner orders* $L, L' \in L(X)$ *with corresponding preference judgment sets* K *and* K' *and any atomic preference judgment* $x > y$*,* $\omega_K(x > y) \neq \omega_{K'}(x > y)$ *if and only if* $(x, y) \in L \backslash L' \cup L' \backslash L$*. Thus, the mapping* $c : L(X) \to \mathcal{K}$ *induces, for every profile of linear preference orders* $\underline{L} \in L(X)^n$ *and every profile of preference states* $\underline{K} = (c(L_1), c(L_2), ..., c(L_n)) \in \mathcal{K}^n$*, an order-embedding from* $\left(L(X), \leqslant^K_{\underline{L}}\right)$ *into* $(\mathcal{K}, \leqslant_{\underline{K}})$*.* ■

This equivalence of the belief merging operator $\Delta^{d,\Sigma}$ with Kemeny's rule allows to apply characterization results of the latter to the evaluation of the former. The well-known fact (Young and Levenglick 1978) that Kemeny's rule is the only preference aggregation rule that is neutral, consistent and satisfies the Condorcet property[8], might in particular be adduced as a justification for the use of the belief merging operator $\Delta^{d,\Sigma}$ when order information is involved.

Distance-based approaches do not only allow to construct merging operators and aggregation rules with certain properties, they can also be used to justify these rules as approximate implementations of some universally acknowledged principles (like the unanimity principle). The merging operator $\Delta^{d,\Sigma}$ is a good example.

In the case of the aggregation of binary evaluations there is an obvious consistency between majority voting and distance minimization which has been observed in several contexts (see e.g. Brams et al. 2004) and can be generalized to the following folk theorem.

Theorem 3 *Majority voting minimizes the sum of Hamming distances for the aggregation of binary evaluations.*

Two corollaries immediately follow from this theorem:

[8]A preference aggregation rule satisfies the Condorcet property if, whenever an alternative x defeats another alternative y in pairwise majority voting, it can never be the case in the social preference that y is ranked immediately above x.

Corollary: Propositionwise majority voting is the only judgment aggregation rule which satisfies anonymity, neutrality and strict monotonicity.

Corollary: Propositionwise majority voting is strategyproof.

The first corollary is an application of May's (1952) classical theorem on majority voting in the case of two alternatives to the similar case of the binary evaluation of propositions. The second corollary establishes the absence of any incentive for the agents to misreport their true evaluations and is a direct implication of the strict monotonicity of propositionswise majority voting.

The positive significance of this theorem is however limited as the discursive dilemma and the discursive version of the Condorcet paradox show that majority voting can lead to inconsistent judgment sets if the judgments are not logically independent, which preference judgments cannot be if conditions like transitivity are imposed.

When an aggregation rule does not satisfy a desirable property on a subset of the domain, an obvious escape route is the relaxation of the universal domain condition (see Gaertner 2001). Distance-based approaches can now help to compensate for the relaxation of the universal domain condition, by extending a restricted aggregation rule to a mapping defined on the whole domain.

Let $f : \mathcal{K}^n \supset \mathcal{K}^* \rightarrow \mathcal{K}_{IC}$ denote an aggregation rule defined on a subset of the profiles, which assigns to every element of the restricted domain an outcome satisfying an integrity constraint. Then this aggregation rule can be extended to the whole domain by constructing a new aggregation rule $F : \mathcal{K}^n \rightarrow \mathcal{K}_{IC}$, based on a distance $\underline{D} : \mathcal{K}^n \times \mathcal{K}^n \rightarrow R_+$ such that, for every $\underline{K} \in \mathcal{K}^n$ and all $\underline{K}^* \in \mathcal{K}^*$, $F(\underline{K}) = f(\arg\min \underline{D}(\underline{K}^*, \underline{K}))$.

A candidate for such a restricted aggregation rule with metric extension would be any mapping $f : \mathcal{K}^n \supset \mathcal{K}^* \rightarrow \mathcal{K}_{IC}$ which incorporates a universally accepted principle of social choice. An example of such a principle is the principle of unanimity, which requires that each unanimous profile of belief bases $\underline{K} = (K', K', ..., K')$ be associated the corresponding belief base K'.

In the case of preference aggregation it is easily seen that Kemeny's rule and hence the belief merging operator $\Delta^{d,\Sigma}$ can be reconstructed as a metric extension of the principle of unanimity: Kemeny's rule assigns to each profile of preference judgment sets the social preference that would be assigned to the closest unanimous profile according to the unanimity principle.

The search for such metric extensions of principles of social choice might appear as an attractive escape route as it seems to extend in an "approximative" way the appeal of such a principle to an aggregation rule defined on the whole domain. In the literature on distance-based social choice there exists however a result that may be considered as a caveat for this type of approximation strategy: Baigent (1987) (see also Baigent and Eckert, 2004) shows that there does not exist a social aggregation rule that satisfies typical properties in the spirit of the Arrovian conditions of non-dictatorship and Pareto together with a condition of proximity preservation, essentially requiring that smaller changes in profiles should not lead to larger changes in the social outcome (than larger changes). This result immediately implies that the proximity between profiles on which the metric extension is based cannot be preserved by any "reasonable" aggregation rule.

Thus, further research is required to weigh the benefits and shortcomings of distance-based approaches both in belief merging and in social choice.

References

Arrow, K.J. (1963) *Social Choice and Individual Values*, Wiley, New York, second edition.

Arrow, K.J., A.K. Sen, and K. Suzumura, eds. (2002) *Handbook of Social Choice and Welfare*, Vol.1, Elsevier.

Baigent, N. (1987) "Preference proximity and anonymous social choice", *Quarterly Journal of Economics* 102, 161-169.

Baigent, N. (2005) "Topological theories of social choice", to appear in Handbook of Social Choice and Welfare, Arrow, K.J., A.K. Sen, and K. Suzumura, eds., Vol. 2, Elsevier.

Baigent, N., and D. Eckert (2004) "Abstract aggregations and proximity preservation", *Theory and Decision*, 56, 359-366.

Brams, S.J., D.M. Kilgour, and M.R. Sanver (2004) "A Minimax Procedure for Negotiating Multilateral Treaties", preprint, Department of Politics, New York University.

Dietrich, F., and C. List (2005) "Arrow's theorem in judgment aggregation". Working paper. http://personal.lse.ac.uk/LIST/doctrinalparadox.htm

Dokow, E., and R. Holzman (2005) "Aggregation of binary evaluations". Working paper. http://personal.lse.ac.uk/LIST/doctrinalparadox.htm

Eckert, D., and J. Mitlöhner (2005) "Logical representation and merging of preference information", *Multidisciplinary IJCAI-05 Workshop on Advances in Preference Handling*, 85-87.

Gaertner, W. (2001) *Domain conditions in social choice theory*, Cambridge University Press.

Hansson, S.O. (1995) "Changes in preferences", *Theory and Decision*, 28, 1-28.

Klamler, C. (2005), "Borda and Condorcet: Some distance results", to appear in *Theory and Decision.*

Konieczny, S. (1999) *Sur la Logique du Changement: Révision et Fusion de Bases de Connaissance*, Ph.D. dissertation, University of Lille I, France.

Konieczny, S., J. Lang, and P. Marquis (2004) "DA^2 merging operators", *Artificial Intelligence*, 157, 49–79.

Konieczny, S., and R. Pino-Pérez (1998) "On the logic of merging", in *Proceedings of KR'98*, Morgan Kaufmann, 488-498.

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

Kornhauser, L. A. (1992) "Modelling collegial courts. II. Legal doctrine", *Journal of Law, Economics and Organization*, 8, 441-470.

Kornhauser, L.A., and L.G. Sager (1986) "Unpacking the court", *Yale Law Journal* 82.

Kornhauser, L.A., and L.G. Sager (1993) "The one and the many: Adjudication in collegial courts", *California Law Review* 81(1), 1-59.

Liberatore, P, and M. Schaerf (1998) "Arbitration (or how to merge knowledge bases", in *ProIEEE Transactions on Knowledge and Data Engineering*, 10.

Lin, J., and A. Mendelzon (1999) "Knowledge base merging by majority", in *Dynamic Worlds: From the Frame Problem to Knowledge Management*, R. Pareschi and B. Fronhoefer, eds., Kluwer.

List, C. (2005) Judgment Aggregation: a Bibliography on the Discursive Dilemma, the Doctrinal Paradox and Decisions on Multiple Propositions. http://personal.lse.ac.uk/LIST/doctrinalparadox.htm

List, C., and R.E. Goodin (2001) "Epistemic democracy: Generalizing the Condorcet Jury Theorem", *Journal of Political Philosophy*, 9(3), 277-306.

List, C., and P. Pettit (2002) "Aggregating sets of judgments: An impossibility result", *Economics and Philosophy*, 18, 89-110.

May, K. (1952) "A set of independent necessary and sufficient conditions for simple majority decision", *Econometrica*, 20, 680-684.

Pigozzi, G. (2004) "Collective decision-making without paradoxes: A fusion approach". Working paper. King's College London.
http://www.dcs.kcl.ac.uk/staff/pigozzi/publications.html

Revesz, P. (1997) "On the semantics of arbitration", *International Journal of Algebra and Computation*, 7, 133–160.

Sen, A. (1970) *Collective Choice and Social Welfare*, Holden Day, San Francisco.

Young, H.P. and Levenglick, A. (1978) "A consistent extension of Condorcet's election principle", *SIAM Journal of Applied Mathematics*, 35, 285–300.