Formal ex-post rationalization – A complete conclusion-based procedure for judgment aggregation

Gabriella Pigozzi gabriella.pigozzi@uni.lu
Marija Slavkovik marija.slavkovik@uni.lu
Leendert van der Torre leendert@vandertorre.com

Individual and Collective Reasoning
Computer Science and Communication
University of Luxembourg
6, Rue Richard Coudenhove Kalergi, L-1359 – LUXEMBOURG

Abstract:
Judgment aggregation is a formal theory reasoning about how a group of agents can aggregate their individual judgments on connected propositions into a collective set of judgments on these propositions. There are three procedures for aggregating judgments sets : premise-based procedure, conclusion-based procedure and distance-based merging. The conclusion-based procedure has been little investigated because it provides a way to aggregate the conclusions, but not the premises, thus it outputs an incomplete judgment set. The goal of this paper is to present a conclusion-based procedure outputting complete judgment sets. Our procedure is composed of two steps. First, majority voting is used to aggregate the conclusion as it has been previously done by the incomplete conclusion-based procedure. Once a collective conclusion is reached, distance-based merging is used to aggregate the premises using the chosen conclusion as an integrity constraint to select only those premise sets which support it. An additional integrity constraint can be added in the merging phase to ensure that unanimity is heeded when selecting the premises. Additionally we discuss the issue of manipulability of the proposed procedure by drawing parallels with existing work on manipulability of belief merging operators.

Keywords: judgment aggregation, conclusion-based procedure, distance-based merging

1 Introduction

Judgment aggregation [13] is an emerging research area in economics. It is a formal theory reasoning about the aggregation of judgments of agents in small groups such as expert panels, legal courts, boards and councils. Judgment aggregation has recently attracted attention in multiagent systems and computer science, which often deal with the combination of information coming from equally reliable sources.

We talk about judgment aggregation whenever a group of individuals needs to make a collective decision on a finite set of issues, and these propositions are logically connected. The propositions are of two kinds : *premises* and a *conclusion*. The first serve as supporting reasons to derive a certain judgment on the conclusion. If, for example [2], your department has to hire a new lecturer and the decision rule is such that a candidate X will be hired only if the candidate is good at teaching and good at research, we will say that “hiring X" is the conclusion while “good at teaching" and “good at research" are the premises.

How shall we derive a group decision given the individuals’ opinions on premises and conclusion ? It is assumed that each individual expresses yes/no opinions on the propositions while respecting the logical relations. If we now define the group opinion as the majority view on the issues, it turns out that the collectivity may have to endorse an inconsistent position. This means that your department may have to face a situation in which the majority does not deem X a good candidate. However, it will not be possible to provide reasons for this as a majority of people agrees that X is actually good at teaching and (another) majority deems X to be good at research. An example of such situation is presented in Table 1.

	prof. A	prof. B	prof. C	Majority
X is good at teaching	*yes*	*yes*	*no*	*yes*
X is good at research	*no*	*yes*	*yes*	*yes*
hire X	*no*	*yes*	*no*	*no*

TAB. 1 – Hiring committee example. The candidate X is hired only if X is good at teaching and X is good at research.

The problem is avoided if we decide to let the majority vote on the premises to dictate the final decision on the hiring process or if we decide just on the conclusion and ignore the reasons given by the group members. Unlike the aggregation procedure on the premises only [6], the aggregation on the individual judgments on the conclusion has been too quickly dismissed.

We claim that in many decision problems, the final conclusion is arguably more relevant than the reasons for it. For example, when deciding which candidate to hire in your department, you may be more concerned of which new colleague you will have in your department than of the reasons for choosing her. Considering only the individual judgments on the conclusions has also the advantage that it is a strategy-proof procedure : if you disagree with the conclusion, you have no incentive to be insincere. The best you can do to increase the possibility that the group will agree with you is to truthfully state your view about the conclusion. The same does not hold when you aggregate on the premises.

The problem this paper addresses is how a group can make decisions on the conclusion while providing reasons in support of the collective conclusion. Our procedure prioritizes the individual judgments on the conclusion and outputs sets of premises that support the collective decision.

The paper is structured as follows : in Section 2 we present the problem of judgment aggregation. Section 3 is devoted to our formal framework, and in Section 4 we prove some results about our procedure. Section 5 relates our approach with existing work and, finally, Section 6 concludes the paper and outlines directions for future work.

2 Judgment aggregation

In this section we first present the original aggregation paradox from which judgment aggregation originated. Agents are required to express judgments (in the form of yes/no or, equivalently, 1/0) over propositions that have different status. More specifically, some propositions (called *premises*) provide the reasons for some other propositions (the *conclusions*). As in [21], to represent the distinction between premise and conclusion in our language, and in contrast to the existing literature on judgment aggregation, we distinguish between premise variables $a, b, c, p, q \ldots$, and conclusion variable x.

The problem of judgment aggregation was discussed by Kornhauser and Sager [10, 11]. In their example, a court has to make a decision on whether a person is liable of breaching a contract (proposition x, or *conclusion*). The judges have to reach a verdict following the legal doctrine. This states that a person is liable if and only if there was a contract (first *premise* a) and there was a conduct constituting breach of such a contract (second *premise* b). The legal doctrine can be formally expressed by the rule $(a \wedge b) \leftrightarrow x$. Each member of the court expresses her judgment on the propositions a, b and x such that the rule $(a \wedge b) \leftrightarrow x$ is satisfied.

Suppose now that the three members of the court make their judgments according to Table 2.

Agenda	A	B	C	Majority
a	1	0	1	1
b	0	1	1	1
$x = (a \wedge b)$	0	0	1	0

TAB. 2 – Doctrinal paradox. Premises : a = There was a contract, b = There was conduct constituting breach of such a contract. Conclusion : $x = (a \wedge b)$ = There was a breach of contract.

Each judge expresses a consistent opinion, i.e. she says yes to x if and only if she says yes to both a and b. However, propositionwise majority voting (consisting in the separate aggregation of the votes for each proposition a, b and x via majority rule) results in a majority for a and b and yet a majority for $\neg x$. This is an inconsistent collective result, in the sense that $\{a, b, \neg x, (a \wedge b) \leftrightarrow x\}$ is inconsistent in propositional logic. The paradox lies in the fact that majority voting can lead a group of rational agents to endorse an irrational collective judgment. The literature on judgment aggregation refers to such problems as the *doctrinal paradox*.

The relevance of such aggregation problems goes beyond the specific court example, because it applies to all situations in which individual binary evaluations need to be combined into a group decision. Furthermore, the problem of aggregating individual judgments is not restricted to majority voting, but it applies to all aggregation procedures satisfying some seemingly desirable conditions. For an overview, the reader is referred to [15].

The paradox originates from the fact that it is assumed that the aggregation on the premises should be logically equivalent to the aggregation on the conclusion, i.e. the group of agents should say yes to x if and only if the group says yes both to a and b. However, by applying the majority rule on each proposition separately, the logical relations between premises and conclusion are disregarded. The aggregation of logically related propositions into a consistent outcome cannot be achieved by imposing that all propositions should be treated independently of each other. The independence condition that is imposed on the aggregation rules in the literature is a reminiscence of the independence of irrelevant alternatives in social choice theory. In the aggregation of judgments, where the propositions are connected, the independence condition is the source of the inconsistent group outcomes.

2.1 Premise vs conclusion-based procedure

Two ways to avoid the inconsistency are the *premise-based procedure* and the *conclusion-based procedure* [19, 4]. According to the premise-based procedure, each agent votes on each premise. The conclusion is then inferred from the rule $(a \wedge b) \leftrightarrow x$ and from the judgment of the majority of the group on a and b. If the judges of the example followed the premise-based procedure, the defendant would be declared liable of breaching the contract.

In the premise-based procedure, the aggregation of a premise proposition is independent from the other premise propositions and the aggregation of the conclusion depends only on the aggregation results of all the premises. For example, if judge A and judge B both would argue that there was no contract and there was no conduct constituting breach of such a contract, then the conclusions of the individual judges would be like in Table 2. Yet, unlike Table 2, there would be no inconsistent collective judgment on a, b and x, and the defendant would be declared not liable.

Because the aggregation of the conclusion does not depend on the individual votes on the conclusion, it is possible for the premise-based procedure to violate the majority, or even the unanimity, of the individual votes on the conclusion.

In [18] Nehring presents a variation on the discursive dilemma, which he calls the Paretian dilemma for short. In his example, a three-judges court has

> to decide whether a defendant has to pay damages to the plaintiff. Legal doctrine requires that damages are due if and only if the following three premises are established : 1) the defendant had a duty to take care, 2) the defendant behaved negligently, 3) his negligence caused damage to the plaintiff. [p.1]

Suppose that the judges vote as in Table 3.

Agenda	A	B	C	Majority
a	1	0	1	1
b	1	1	0	1
c	0	1	1	1
$x = (a \wedge b \wedge c)$	0	0	0	0

TAB. 3 – Paretian dilemma. Premises : a = duty, b = negligence., c = causation. Conclusion : $x = (a \wedge b \wedge c)$ = damages.

The Paretian dilemma is disturbing because, if the judges would follow premise-based procedure, they would condemn the defendant to pay damages contradicting the *unanimous* belief of the court that the defendant is *not* liable.

A conclusion-based procedure would not lead to such a unanimity violation. According to the conclusion-based procedure, the judges decide privately on a and b and only express their opinions on x publicly. The judgement of the group is then inferred from applying the majority rule to the agent judgments on x. The defendant will be declared liable if and only if a majority of the judges actually believes that she is liable. In the example, contrary to the premise-based procedure, the application of the conclusion-based procedure would free the defendant. However, no reasons for the court decision could be supplied.

Unlike the premise-based procedure [16, 6], conclusion-based procedure did not receive much attention in the literature. In this paper we aim at filling this gap. We propose a procedure that attempts to overcome the major limit of conclusion-based procedure, that is the lack of reasons supporting the decision.

3 Framework

In this section we introduce our formal framework to represent judgment aggregation problems. A set of agents $N = \{1, 2, \ldots, n\}$, with $n \geq 3$ and odd, makes judgments on logically interconnected propositions. The set $\mathcal{P}$ of atomic propositions is defined as the union of two disjoint sets : $\mathcal{P}_p$ containing variables $a, b, c, \ldots, p, q$ for the premises, and $\mathcal{P}_c$ being a singleton $\{x\}$, where x is the variable for the conclusion. We assume that the conclusion is an atomic formula. $\mathcal{L}$ is a language built from $\mathcal{P}$, including complex formulas as $\neg a, (a \wedge b), (a \vee b), (p \rightarrow q), (a \leftrightarrow p)$.

The set of issues on which the judgments have to be made is called *agenda* and is denoted by $\Phi \subseteq \mathcal{L}$. The agenda is closed under negation : if $a \in \Phi$, then $\neg a \in \Phi$.[1] We split the agenda in two parts : one containing the premises (Φ_p), and one containing the conclusion (Φ_c). We exclude agenda items such as $a \rightarrow x$, i.e. formulas containing premises and conclusion. Our procedure consists of two different aggregations : one on the individual judgments on Φ_p and one on the individual judgments on Φ_c.

A subset $J \subseteq \Phi$ is the *collective judgment set* and contains the set of propositions believed by the group. Similarly, we define individual i's judgment set $J_i \subseteq \Phi$. A collective judgment set is *consistent* if it is a consistent set in $\mathcal{L}$, and is *complete* if, for any $a \in \mathcal{L}$, $a \in J$ or $\neg a \in J$ (consistent and complete individual judgment sets are defined in the same way). We only consider consistent complete judgment sets.

A *decision rule* $\mathcal{R}$ is a formula of $\mathcal{L}$ that represents the logical connections between premises and conclusion. More precisely, $\mathcal{R}$ has the form $\Psi \leftrightarrow x$, where $\Psi \in \mathcal{L}/\{x\}$. The decision rule is not an item of the agenda. This means that the group members do not vote on $\mathcal{R}$, but each individual is required to give judgments that satisfy the given rule (see [1] for an approach in which agents can disagree with the rule).

Like the agenda, each judgment set is split in two disjoint subsets : $J_{i,p}$ and $J_{i,c}$. The first is the individual i's judgment set on the premises, and $J_{i,c}$ is the individual i's judgment set on the conclusion. The collective judgment sets on premises and conclusion will be denoted respectively by J_p and J_c.

We say that a premise a (resp. a conclusion x) is *unanimously supported* if $a \in J_{i,p}$ for all $J_{i,p} \subseteq \Phi$ (resp. $x \in J_{i,c}$ for all $J_{i,c} \subseteq \Phi$).

A *profile* $\underline{J}$ is an n-tuple $(J_1, J_2, \ldots, J_n)$ of agents' judgment sets. An *aggregation rule* F assigns a set of collective judgment sets J to each profile $\underline{J}$. To define our procedure, we define two aggregation rules : one for the aggregation of the individual premises and one for the conclusion. To relate the two aggregation rules, we have set of integrity constraints IC to govern the premise-aggregation rule. Also, to allow for situations in which the aggregated judgment set is not unique, i.e. there are ties, we aggregate the profiles into *sets* of aggregated judgment sets.

A *premise profile* $\underline{J_p}$ is an n-tuple $(J_{1,p}, J_{2,p}, \ldots, J_{n,p})$ of agent judgment sets on premises. A *premise aggregation rule* F_{IC} assigns a set of collective judgment sets J_p to each premise profile $(J_{1,p}, J_{2,p}, \ldots, J_{n,p})$ and set of integrity constraints IC. Conclusion profiles $(J_{1,c}, J_{2,c}, \ldots, J_{n,c})$ and conclusion aggregation rules F_c are defined similarly.

3.1 Complete conclusion-based procedure

Each individual provides, simultaneously, the set of premises and conclusion that she believes. Our two-step procedure first performs a standard conclusion-based procedure, i.e. it aggregates the individual judgments on the conclusion by majority rule. This means that x (resp. $\neg x$) is the collective conclusion iff there are at least $\frac{n+1}{2}$ agents voting for x (resp. $\neg x$). The second step consists in determining the set of reasons which support the collective conclusion. This is done by applying a *distance-based merging operator* to $J_{i,p}$.

Distance minimization merging procedures have been already applied to judgment aggregation problems [20]. In this section we briefly present a majority merging operator with integrity constraints following [9, 8]. Whereas merging operators in general merge sets of propositional sentences called knowledge bases, we are interested in the particular case of merging maxi-consistent sets of literals called complete knowledge bases. When there are ties, a belief merging operator returns the disjunction of the alternatives.

An *interpretation* is a function $v : \mathcal{P} \rightarrow \{0, 1\}$ and it is represented as the list of the binary evaluations. For example, given three propositional variables a, b and c, the vector (0,1,0) stands for the interpretation in which a and c are false and b is true. Let $\mathcal{W} = \{0, 1\}^{\mathcal{P}}$ be the set of all interpretations. An interpretation is a *model* of a propositional formula if and only if it makes the formula true in the usual truth functional way.

Let us suppose that $\Phi_p = \{a, \neg a, b, \neg b, c, \neg c\}$, and that agent 1 believes that $a, \neg b$ and $\neg c$, i.e. $J_{1,p} = \{a, \neg b, \neg c\}$. We represent $J_{1,p}$ as a 0-1 vector of length equal to the number of propositions in $J_{1,p}$, i.e. $(1, 0, 0)$. Suppose also that $\mathcal{R} = ((a \vee b) \wedge c) \leftrightarrow x$ and that, unlike agent

[1] To increase readability, in the tables we list only the positive issues, and assume that, for any issue in the agenda, an individual deems that issue to be true if and only if she deems its negation to be false.

1, the majority of the individuals voted in favor of x. Hence, the first step of our procedure sets $v(x) = 1$. We now want to define an aggregation on $J_{i,p}$ such that the collective judgment set on the premises is one of the models of $((a \vee b) \wedge c) \leftrightarrow x$ where $v(x) = 1$. This means that J_p must be one of the following interpretations : (1,1,1), (0,1,1), (1,0,1). The set of premises supporting the collective conclusion will constrain the aggregation procedure on $J_{i,p}$. We indicate the set of admissible interpretations as IC (*integrity constraints*).

Given a premise profile $\underline{J_p}$ and IC, $F_{IC}(\underline{J_p})$ denotes a set of collective judgment sets on the premises resulting from the IC merging on $\underline{J_p}$. The idea of a distance minimization merging operator is that $F_{IC}(\underline{J_p})$ will select those interpretations in IC, which are at minimal distance from $\underline{J_p}$. A distance $d(\omega, \underline{J_p})$ between an interpretation ω and the premise profile $\underline{J_p}$ induces a total pre-order ($\leq$) on the interpretations.

In order to obtain the total pre-order on the interpretations, we first need to determine a pseudo-distance between each admissible interpretation and each $J_{i,p}$. Then, we need to aggregate all these values in order to obtain a pseudo-distance value between an interpretation and $\underline{J_p}$. Let us see this in detail (we follow [9, 8]).

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

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

A pseudo-distance between an interpretation ω and $\underline{J_p}$ is defined with the help of an aggregation function $D : \mathbb{R}^n_+ \rightarrow \mathbb{R}_+$ as $D^d(\omega, \underline{J_p}) = D\left(d(\omega, J_{1,p}), d(\omega, J_{2,p}), \ldots, d(\omega, J_{n,p})\right)$ [8]. Any such aggregation function induces a total pre-order $\leqslant_{\underline{J_p}}$ on the set $\mathcal{W}$ with respect to the pseudo-distances to a given $\underline{J_p}$. Thus, an *IC majority merging operator* for a profile $\underline{J_p}$ can be defined as $\Delta_{IC}(\underline{J_p}) = \min([IC], \leqslant_{\underline{J_p}})$, i.e., the set of all models of IC (denoted by $[IC]$) with minimal pseudo-distance D^d to $\underline{J_p}$. The minimal pseudo-distance identifies the final collective outcome on the premises, i.e. the set of premises that support the conclusion voted by the majority of the agents and with the minimal distance among all possible models satisfying IC.

A majority merging operator, often mentioned in the literature, is the operator $\Delta^{d,\Sigma}_{IC}$ defined as follows :

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

For example, the Hamming distance between $\omega = (1, 0, 0)$ and $\omega' = (0, 1, 0)$ is $d(\omega, \omega') = 2$. In the following we use the Hamming distance because it is a well known and intuitive distance. But the Hamming distance is only one among many possible distance functions that we may use.

The premise aggregation rule F_{IC} outputs the disjunction of the sets selected by $\Delta^{d,\Sigma}_{IC}$ as the reasons in support of the conclusion voted by the majority of the agents. Given a premise profile $\underline{J_p}$, F_{IC} is defined as :

$$F_{IC}(\underline{J_p}) = \bigvee \Delta^{d,\Sigma}_{IC}(\underline{J_p})$$

The constraint IC is defined as $IC = \mathcal{R} \wedge \hat{x}$, where $\hat{x}$ is the conclusion chosen by the majority.

The best way to illustrate our procedure is with an example. For simplicity, we take the original doctrinal paradox.

Example 1 *Individuals vote as in Table 2. Their judgment sets are then :* $J_1 = \{(1, 0, 0)\}$*,* $J_2 = \{(0, 1, 0)\}$ *and* $J_3 = \{(1, 1, 1)\}$*. A majority is in favor of* $\neg x$*, so the first step of our procedure sets that the group decided* $v(x) = 0$*.*

We now want to provide a set of reasons in support of $v(x) = 0$*. In order to select them, we calculate the minimal distance between* $IC = \{(1, 0), (0, 1), (0, 0)\}$ *(all the admissible sets of reasons for* $v(x) = 0$*) and each* $J_{i,p}$*,* $1 \leq i \leq 3$*. This is calculated in the table below.*

The final group decision will support $v(x) = 0$ *and the reasons will be* $(0, 1) \vee (1, 0)$*, representing a tie between* $(0, 1)$ *and* $(1, 0)$*.*

ω	$J_{1,p}$	$J_{2,p}$	$J_{3,p}$	$\Sigma_i\ d(\omega, J_{i,p})$
(0,1)	2	0	1	3
(1,0)	0	2	1	3
(0,0)	1	1	2	4

TAB. 4 – Selection of premises for the doctrinal paradox with $v(x) = 0$.

The example illustrates that, when aggregating the premises, we do not only take into account the judgment sets of the agents that support the aggregated conclusion, but also the judgment sets of agents that do not support the conclusion. Consider for example the selection of premises for the doctrinal paradox with $v(x) = 0$ in Table 2. We take also judgment set $J_{3,p}$ into account, although the third judge voted for $v(x) = 1$.

The justification for taking all individual judgments on the premises into account is twofolded. On the one hand, from the perspective of probability theory, if all judgments are independent, then more judgment sets mean a higher chance to get a better judgment. On the other hand, from the perspective of democracy, involving agents whose conclusion is not supported will give broader basis for the decision. However, we do not exclude the possibility that there are situations in which only the individuals' judgments that actually supported the aggregated conclusion should be taken into account when determining the reasons for that conclusion.

We now want to show that our procedure does not always provide a set of reasons, but that can also select a unique justification for the collective conclusion.

Example 2 *Consider a* collegium medicum *that wishes to eliminate the possibility of a patient suffering from condition X before administering a treatment. We take $v(x) = 0$ if the patient is free of X. The doctors consider the three relevant alternative medical conditions a, b and c the patient may suffer from. The patient is free of X if medical conditions a, b and c are present ($v(a) = 1$, $v(b) = 1$ and $v(c) = 1$), if all three medical conditions are absent ($v(a) = 0$, $v(b) = 0$ and $v(c) = 0$) or if the last condition is present while the previous two are absent ($v(a) = 0$, $v(b) = 0$ and $v(c) = 1$). In all other cases the patient is likely to suffer from X. Table 5 gives the truth table of $\mathcal{R}$. Three equally qualified members of the* collegium medicum *give their opinions shown in Table 6. As Table 6 shows, the group is facing a dilemma. The majority of the conclusions from the doctors opinions indicates that the patient does not suffer from X though the majority on the premises supports the opposite conclusion.*

Our procedure (see Table 7) selects the reasons that are most compatible with the doctors' different opinions, i.e. the judgment set (1,1,1).

a	b	c	x
0	0	0	0
0	0	1	0
0	1	0	1
0	1	1	1
1	0	0	1
1	0	1	1
1	1	0	1
1	1	1	0

TAB. 5 – The truth table of $\mathcal{R}$ for the doctor example.

Condition	Dr. A	Dr. B	Dr. C	Majority
a	1	0	1	1
b	1	0	1	1
c	1	0	0	0
x	0	0	1	0

TAB. 6 – The dilemma faced by the doctors.

	$J_{1,p}$	$J_{2,p}$	$J_{3,p}$	$\Sigma_i\ d(\omega, J_{i,p})$
(1,1,1)	0	3	1	4
(0,0,0)	3	0	2	5
(0,0,1)	2	1	3	6

TAB. 7 – Selection of the premise set from the lab results under the constraint $v(x) = 0$.

4 Results

We now show some properties which hold for the premise aggregation rule F_{IC} we had defined in the previous section.

We start by noticing that, in the case of the aggregation of binary evaluations, there is an obvious correspondence between proposition-wise majority voting and distance minimization. This has been already observed in several contexts (see, e.g., [3]), and can be generalized to the following folk theorem.

Theorem 4.1 *Let $\underline{J} = (J_1, \ldots, J_n)$ be a profile over the agenda Φ. Let $J^{maj} \subset \Phi$ be a complete and consistent set. Let it hold that for every premise $a \in J^{maj}$, $a \in J_{i,p}$ for at least $\frac{n+1}{2}$ premise sets in the profile J_p. Also, for the conclusion $x \in J^{maj}$, let it holds that $x \in J_{i,c}$ for at least $\frac{n+1}{2}$ conclusion sets in the profile J_c. The sum of Hamming distances from J^{maj} to the judgment sets in $\underline{J}$ is minimal.*

This means that, in the absence of a Paretian dilemma (i.e. when J^{maj} satisfies the decision rule $\mathcal{R}$), proposition-wise majority voting, distance-based merging and our procedure coincide.

4.1 Unanimity preservation

One of the desirable properties for a judgment aggregation procedure is the heeding of unanimity. If all the agents unanimously support an agenda item, then it is natural to expect the unanimously supported item will be adopted as the collective judgment. However, the premise-based procedure does not necessarily preserve unanimity on the conclusion (as it was the case with the Paretian dilemma shown in Table 3).

The premise-based procedure aggregates each premise independently from the other premises, but the aggregation on the conclusion depends on the collective judgments on the premises. Therefore the unanimity on the premises will be preserved, but the unanimity on the conclusion may be violated.

When aggregating according to the conclusion-based procedure, unanimity on the conclusion will always be maintained, but unanimity on the premises may be violated. However, our procedure offers the option to heed unanimity on the premises as well, by constraining the models which do not support unanimity.

We begin by giving a formal definition on when a premise aggregation rule F_{IC} preserves unanimity. Whether or not the unanimity on the premises is preserved by our F_{IC} depends on the rule $\mathcal{R}$ as well as the agenda Φ. We show two decision rules for which the unanimity is preserved and than we use an example to show that in the case of an arbitrary rule and agenda, the unanimity of the premises is not guaranteed.

Definition 1 *Let $\underline{J_p} = (J_{1,p}, \ldots, J_{1,n})$ be a premise profile on the agenda Φ and p a premise from the agenda. A premise aggregation rule F_{IC} heeds unanimity on the premises if and only if the following holds :*
If $p \in J_{i,p}$ for all $i = \{1, \ldots, n\}$ then $p \in F_{IC}(\underline{J_p})$.

Note that, since F_{IC} can select more than one premise judgment set, p needs to be in all of them for unanimity to be preserved.

The following theorem indicates two decision rules $\mathcal{R}$, and an agenda, in the presence of which unanimity is preserved on the premises by F_{IC}.

Theorem 4.2 *Let Φ be an agenda in which all the elements are atoms or negations of atoms. Let $\mathcal{R}$ be a decision rule of the form $(a_1 \wedge \ldots \wedge a_n) \leftrightarrow x$ or of the form $(a_1 \vee \ldots \vee a_n) \leftrightarrow x$. $\{a_1, \ldots, a_n\} \subseteq \Phi$ are premises and $x \subseteq \Phi$ is a conclusion. F_{IC} preserves unanimity on the premises for any profile $\underline{J}$ over Φ and $\mathcal{R}$.*

Proof *We give the proof for the decision rule $(a_1 \wedge \ldots \wedge a_n) \leftrightarrow x$. The proof for the decision rule $(a_1 \vee \ldots \vee a_n) \leftrightarrow x$ can be constructed symmetrically.*

The aggregated conclusion for x can be either $v(x) = 1$ or $v(x) = 0$. We construct a proof by cases :

`Case` $v(x) = 1$

There is exactly one set of evaluations on premises that can be selected which is consistent with $v(x) = 1$, i.e. the premise set J_p^+ in which every premise is evaluated to be true. Any unanimously supported premise in this case, can only be unanimously supported to have the value true and is necessarily included in J_p^+. Since $x = 1$, the majority of the premise judgment sets of the profile $\underline{J}$ are necessarily the premise set J_p^+. Consequently, the premise judgment set selected by the merging operator will select precisely the set J_p^+.

`Case` $v(x) = 0$

We construct a proof by contradiction.

`Assumption` *: there is profile $\underline{J}$ and a premise p such that $p \in J_{i,p}$ for every judgment set i in the profile, and $\neg p \in F_{IC}(\underline{J_p})$. $F_{IC}(\underline{J_p})$ has the minimal distance to the judgment sets of the agents.*

As a consequence of the assumption, J_p must be the model setting $v(x) = 0$,

$v(a_j) = 1$ *for all* $a_j \neq p$.

`Contradiction` **:** *Since there is a majority for* $v(x) = 0$*, there is an agent whose judgment set contains* $v(x) = 0$*, and thus, in this agents judgment set, there is a premise* q *such that* $v(q) = 0$*. Now consider the judgment set* J' *with* $v(a_k) = 1$ *for all* $a_k \neq q$*, and* $v(q) = 0$ *while* $v(x) = 0$ *.* J' *is necessarily closer to the profile* $\underline{J}$ *than* J*. This is a contradiction with the assumption.*

Given an arbitrary agenda, a decision rule $\mathcal{R}$ corresponding to that agenda and an arbitrary profile $\underline{J}$, the merging operator does not necessary preserves unanimity. We show this through an example.

Consider the profile presented in Table 8. The rule $\mathcal{R}$ is such that the value of x is 1 if and only if the premises are assigned as in the first column of Table 9. For all other evaluations of premises, x is evaluated to 0.

	Agent A	Agent B	Agent C	Majority
p_1	1	0	0	0
p_2	0	1	0	0
p_3	0	0	1	0
p_4	1	0	0	0
p_5	0	1	0	0
p_6	0	0	1	0
p_7	1	0	0	0
p_8	0	1	0	0
p_9	0	0	1	0
p_{10}	1	0	0	0
p_{11}	0	1	0	0
p_{12}	0	0	1	0
p_{13}	1	1	1	1
x	1	1	1	1

TAB. 8 – A case in which unanimity on premises will be violated by the complete CBP.

Our procedure preserves the unanimity on the conclusion and selects $v(x) = 1$, but gives an aggregation for the premises which violates the unanimity on premise p_{13} (Table 9).

The preservation of unanimously held premises can be imposed by IC. This is done by making $IC = \mathcal{R} \wedge \hat{x} \wedge a^*$ (see Section 3.1). Admissible outcomes for J_p then are those supporting the conclusion voted by the majority and containing the premise(s) unanimously chosen.

	$J_{1,p}$	$J_{2,p}$	$J_{3,p}$	$\Sigma_i d()$
(1,0,0,1,0,0,1,0,0,1,0,0,1)	0	8	8	16
(0,1,0,0,1,0,0,1,0,0,1,0,1)	8	0	8	16
(0,0,1,0,0,1,0,0,1,0,0,1,1)	8	8	0	16
(0,0,0,0,0,0,0,0,0,0,0,0,0)	5	5	5	15

TAB. 9 – Selection of premises for the counterexample.

4.2 Manipulability

Another property which is of interest when dealing with aggregation procedures is that of manipulability. A judgment aggregation procedure is called manipulable if an agent, who would not obtain a desired outcome by submitting her "honest" premise set, can obtain a desired outcome by choosing to submit a set of premises different than her "honest" premise set. Under the context of complete-conclusion based procedures, we will distinguish between full and preferred manipulability.

Full manipulability means that we distinguish only whether the aggregated premise set entirely corresponds to an agent's judgments on premises, or not. A procedure is fully manipulable if an agent can obtain her complete "honest" premise set as an output from the procedure by submitting another (insincere) premise set that supports the same conclusion. Formally, let $\underline{J_p} = (J_{1,p}, J_{2,p}, \ldots, J_{i,p} \ldots, J_{n,p})$ be a premise profile. Let $F_{IC}(\underline{J_p}) = \{J^\circ_{1,p}, \ldots, J^\circ_{m,p}\}$, i.e. the merging operator selects the premise sets $J^\circ_{1,p}, \ldots, J^\circ_{m,p}$. Let $J_{i,p}$ be the "honest" premise set of an agent i.

Definition 2 *Assume that a premise set* $J^*_{i,p} \neq J_{i,p}$ *exists, such that* $J^*_{i,p}$ *supports the same conclusion as the premise set* $J_{i,p}$*. The operator* F_{IC} *is fully manipulable if* $J_{i,p} \in F_{IC}(J_{1,p}, J_{2,p}, \ldots, J^*_{i,p} \ldots, J_{n,p})$ *but* $J_{i,p} \notin F_{IC}(J_{1,p}, J_{2,p}, \ldots, J_{i,p} \ldots, J_{n,p})$.

Theorem 4.3 F_{IC} *is not fully manipulable.*

Proof *The proof of this theorem is given in [7], under strategy proofness for complete bases when the merging operator is* $\Delta^{d,\Sigma}_{\mu}$.

Let us now assume that an agent has a premise p which she holds most important (has a strong preference on the evaluation of this premise). We say that a procedure is *preferred manipulable* if an agent can ensure that the preferred projection $w(p)$ is included in the output by submitting another premise set that supports the same conclusion. Since we do not represent the preferred premise explicitly in our framework, any premise can be the preferred one, and preferred manipulability therefore means that the agent is able to change her premise set in a way such that one premise which is not a member of the aggregated set becomes member of it.

Definition 3 *Assume that a premise set $J^*_{i,p} \neq J_{i,p}$ exists, such that $J^*_{i,p}$ supports the same conclusion as the premise set $J_{i,p}$ and premise p_{pref} is in both of the premise sets. The operator F_{IC} is preferred manipulable if p_{pref} is in at least one premise set $J^\circ_{j,p} \in F_{IC}(J_{1,p}, J_{2,p}, \ldots, J^*_{i,p} \ldots, J_{n,p})$, but $\neg p_{pref}$ is in all of the premise sets selected by $F_{IC}(J_{1,p}, J_{2,p}, \ldots, J_{i,p} \ldots, J_{n,p})$.*

Theorem 4.4 *F_{IC} is preferred manipulable.*

Proof *To show that F_{IC} is preferred manipulable it is sufficient to show that there exists a case in which an agent can ensure that the preferred premise appears in the aggregated result by misrepresenting her premise judgment set.*

Assume the rule $\mathcal{R}$ is $x \leftrightarrow (p_1 \wedge p_2 \wedge p_3 \wedge p_4) \vee (p_1 \wedge p_2 \wedge \neg p_3 \wedge \neg p_4) \vee (\neg p_1 \wedge p_2 \wedge p_3 \wedge \neg p_4)$. The profile in which every agent submits the "honest" judgment sets is given in Table 10.

	Agent A	Agent B	Agent C	Majority
p_1	0	1	1	1
p_2	0	1	1	1
p_3	0	1	0	0
p_4	0	1	0	0
x	0	1	1	1

TAB. 10 – Preferred manipulability on premises when all agents vote "honestly".

The full conclusion-based procedure will select $x = 1$ and the premise set $\{p_1 \wedge p_2 \wedge \neg p_3 \wedge \neg p_4\}$ (see Table 11).

Assume that Agent B holds premise p_3 as specially important and has the incentive to see it in the selected premise judgment set. If this agent submits (0,1,1,0) instead of (1,1,1,1) as in Table 12, then the selected premise judgment sets is

	$J_{1,p}$	$J_{2,p}$	$J_{3,p}$	$\Sigma_i\ d(\omega, J_{i,p})$
(0,0,0,0)	0	4	2	6
(1,1,1,1)	4	0	2	6
(1,1,0,0)	2	2	0	4
(0,1,1,0)	2	2	2	6

TAB. 11 – Selection of premises for the "honest" profile.

$\{0,0,0,0\}$, $\{0,1,1,0\}$ $\{1,1,0,0\}$ *(Table 13). Premise p_3 will be included in one of the selected premise judgment sets.*

	Agent A	Agent B	Agent C	Majority
p_1	0	0	1	1
p_2	0	1	1	1
p_3	0	1	0	0
p_4	0	0	0	0
x	0	1	1	1

TAB. 12 – Preferred manipulability on premises when Agent B manipulates.

	$J_{1,p}$	$J_{2,p}$	$J_{3,p}$	$\Sigma_i\ d(\omega, J_{i,p})$
(0,0,0,0)	0	2	2	4
(0,1,1,0)	2	0	2	4
(1,1,0,0)	2	2	0	4
(1,1,1,1)	4	2	2	8

TAB. 13 – The result from the merging of the "manipulated" profile.

Full manipulability is a relatively weak condition, in the sense that it is relatively easy to satisfy. Everaere *et al.* [7] mention that their satisfaction relations corresponding to complete manipulability are the "most meaningful" ones, but in our setting, it seems that stronger notions of manipulability may be called for. Moreover, this definition of manipulability seems to conflict with the intuition of the distance measure used to aggregate the premises, which does take such distinctions into account. However, preferred manipulability is a very strong condition, since it means in practice that an agent should not be able to improve any premise (since this premise may happen to be the preferred one). Other notions of manipulability could be studied, such as the improvement of a preferred premise by changing the judgment on this premise only.

5 Related work

One of the noted shortcomings of the conclusion-based procedure is that it is susceptible to path-dependence [17]. Path-dependent decisions are decisions whose outcome depends on the order in which propositions are considered. An order of priority over the propositions is assumed. For any proposition, the collective judgment on it is decided by majority rule (or by any other suitable aggregation rule) unless this conflicts with the collective judgments of previously aggregated propositions. In the latter case, the collective value of that proposition is deduced by logical implication from the previously aggregated propositions. List [12] provided necessary and sufficient conditions for path-dependence. Furthermore, in [5] it has been shown that the absence of path-dependence is equivalent to strategy-proofness.

Path-dependent procedures may be employed to define complete conclusion-based procedures. If a priority order is assumed on the premises, we can ensure a complete and consistent collective set of premises supporting the collective conclusion. However, here we propose a complete conclusion-based procedure without assuming any order over the premises. It is indeed arguable where such order comes from and, since different orders lead to different outcomes, path-dependend aggregation procedure are prone to strategic manipulation. In the present work we aimed at a procedure that treats all premises in an even-handed way. The conclusion-based procedure we presented in Section 3.1 avoids the problem of path-dependance by selecting the premises as a complete set instead of one premise after another. The absence of full manipulability is coherent with the results of [5].

The non-manipulability (of the outcome for the conclusion) of the conclusion-based procedure is one of its advantages over the premise-based procedure. The question of manipulability under operators used for merging of propositions has been treated extensively in [7]. There, Everaere *et al.* explore a broad spectrum of manipulability for various merging operators over complete and incomplete sets of beliefs (propositions, or in our case judgments). Our work uses results from [7] on complete sets of beliefs under model-based merging operators that use the sum of the distances between belief bases i.e. “the aggregation function Σ".

6 Conclusions and future work

By presenting a complete conclusion-based procedure for judgment aggregation, we show that conclusion-based procedures deserve more attention than they have received in the past. Our complete conclusion-based procedure keeps the desirable properties of non-manipulability (over the conclusion) and it can be modified to heed unanimity on the premises.

What can be considered a shortcoming of the procedure is that the procedure may select more than one premise judgment set to support the collective conclusion. Such “ties" in the output from aggregation are known to be resolved with an additional approval vote [3] or by random selection. A random selection is not a desirable tie-breaking solution in cases when the decisions on premises can influence some future decision making process. The approval voting requires more information to be injected in the framework and opens the questions of what incentives an agent may have to prefer one premise judgment set over another. In the case when there is a small number of agents, it may happen that each agent has her own judgment set in the tie and will not cast her vote of approval to any of the other premise judgment sets.

The issue which we plan to explore in future work is the relevance that current group decisions can have on future decisions when their agendas have a nonempty intersection. This “evolutionary" impact over the decision making process has been an important issue in the work that gave rise to the interest in judgment aggregation [10, 11], but it has fallen out of scope in the more formal study of judgment aggregation [16, 14]. It would be of interest, especially from the aspect of manipulability, to reintroduce the issue of “evolutionary" impact. This would require an extension of the current formal framework to allow for the “evolutionary" considerations to be part of the aggregation procedure.

Acknowledgment

Thanks to Franz Dietrich and Christian List for discussions on the issues raised in this paper.

Références

[1] F. Benamara, S. Kaci, and G. Pigozzi. Group decision procedures reactive to individual opinions. *Working paper*, 2009.

[2] L. Bovens and W. Rabinowicz. Democratic answers to complex questions. an epistemic perspective. *Synthese*, 150 :131–153, 2006.

[3] S. Brams, D. Kilgour, and M. Sanver. A minimax procedure for negotiating multilateral treaties. In M. Wiberg, editor, *Reasoned Choices : Essays in Honor of Hannu Nurmi*, pages 265–282. Finnish Political Science Association, 2004.

[4] B. Chapman. Rational aggregation. *Politics, Philosophy and Economics*, 1(3) :337–354, 2002.

[5] F. Dietrich and C. List. Judgment aggregation by quota rules : majority voting generalized. *Journal of Theoretical Politics*, 19(4) :391–424, 2007.

[6] F. Dietrich and P. Mongin. The premiss-based approach to judgment aggregation. Research memoranda, Maastricht : METEOR, Maastricht Research School of Economics of Technology and Organization, 2008.

[7] P. Everaere, S. Konieczny, and P. Marquis. The strategy-proofness landscape of merging. *J. Artif. Intell. Res. (JAIR)*, 28 :49–105, 2007.

[8] S. Konieczny, J. Lang, and P. Marquis. Da^2 merging operators. *Artificial Intelligence*, 157 :49–79, 2004.

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

[10] L. Kornhauser and L. Sager. Unpacking the court. *Yale Law Journal*, 96 :82–117, 1986.

[11] L. Kornhauser and L. Sager. The one and the many : Adjudication in collegial courts. *California Law Review*, 81 :1–51, 1993.

[12] C. List. A model of path-dependence in decisions over multiple propositions. *American Political Science Review*, 98(3) :495–513, 2004.

[13] C. List. Judgment aggregation - a bibliography on the discursive dilemma, the doctrinal paradox and decisions on multiple propositions. *http ://personal.lse.ac.uk/LIST/doctrinalparadox.htm*, 2007.

[14] C. List and P. Pettit. Aggregating sets of judgments : An impossibility result. *Economics and Philosophy*, 18 :89–110, 2002.

[15] C. List and C. Puppe. Judgment aggregation : A survey. In P. Anand, C. Puppe, and P. Pattanaik, editors, *Oxford Handbook of Rational and Social Choice*. Oxford University Press, forthcoming.

[16] P. Mongin. Factoring out the impossibility of logical aggregation. *Journal of Economic Theory*, 141(1) :100–113, July 2008.

[17] J. R. Nash. A context-sensitive voting protocol paradigm for multimember courts. *Stanford Law Review*, 56 :75–159, 2003.

[18] K. Nehring. The (im)possibility of a Paretian rational. Economics working papers, Institute for Advanced Study, School of Social Science, Nov 2005.

[19] P. Pettit. Deliberative democracy and the discursive dilemma. *Philosophical Issues*, 11 :268–299, 2001.

[20] G. Pigozzi. Belief merging and the discursive dilemma : an argument-based account to paradoxes of judgment aggregation. *Synthese*, 152(2) :285–298, 2006.

[21] G. Pigozzi, M. Slavkovik, and L. van der Torre. Conclusion-based procedure for judgment aggregation satisfying premise independence. In W. v. d. H. G. Bonanno, B. Lowe, editor, *Proceedings of the Eighth International Conference on Logic and the Foundations of Game and Decision Theory, LOFT 2008*, page 35, Amsterdam, The Netherlands, 2008.