Premise Independence in Judgment Aggregation

Gabriella Pigozzi and Leendert van der Torre
University of Luxembourg, Computer Science and Communications
{gabriella.pigozzi,leon.vandertorre}@uni.lu

ABSTRACT

Judgment aggregation studies how agent opinions on logically interconnected propositions can be mapped into a collective judgment on the same propositions, and is plagued by impossibility results. In this paper we study the central notion of *independence* in these impossibility results. First, we argue that the distinction between the premises and conclusions play an important role in the benchmark examples of judgment aggregation. Second, we consider the notion of independence in judgment aggregation frameworks, and we observe that the distinction between premises and conclusion is not taken into account. Third, based on our analysis, we introduce independence assumptions that distinguish premises from conclusion. We show that, by introducing new operators that satisfy our independence assumptions, the problematic impossibility results no longer hold.

1. INTRODUCTION

Judgment aggregation [12] is an emerging research area in economics. It is a formal theory reasoning about the aggregation of judgments of agents in expert panels, legal courts, boards, and councils. Judgment aggregation has recently attracted attention in multiagent systems and artificial intelligence, in particular due to the relations with belief merging [18], for example for the combination of opinions of equally reliable agents.

Independence plays a central role in judgment aggregation, since the aggregation problems arising in judgment aggregation are a generalization of the aggregation issues discussed in Arrow's social choice theory [13, 6]. Whereas social choice theory studies the aggregation of individual preferences in order to select a collectively preferred alternative, judgment aggregation explores how to combine agents' judgments on specific propositions. Just like Arrow's property of "Independence of Irrelevant Alternatives" plays a central role in his famous impossibility results, independence leads to impossibility results in judgment aggregation.

However, these formal frameworks ignore an important distinction between aggregating judgments and preferences, which becomes apparent in our study of the benchmark examples of these areas. The benchmark example of social choice theory is the Condorcet paradox. Given a set of agent 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. Similarly, when aggregating agents' judgments on propositions, a seemingly reasonable procedure, such as propositionwise majority voting, cannot ensure a consistent collective decision. This observation leads to the following research questions of this paper:

1. What is the role of independence in the benchmark examples of judgment aggregation? We argue that premises and conclusions must be distinguished.

2. What is the role of independence in judgment aggregation frameworks? We illustrate that the distinction between premises and conclusions does not play a role in the frameworks.

3. How to bridge the gap between the individual examples and the frameworks? We propose weaker independence assumptions and we show their consistency by defining operators for them.

Since independence is the central notion in judgment aggregation, and the literature on judgment aggregation (like social choice theory) concentrate on formal impossibility results rather than discussing the examples in depth, a large part of this paper is concerned with analyzing and criticizing the benchmark examples and existing frameworks. We use the framework of Mongin [15] to formulate our new independence assumptions, because his notion of independence of irrelevant propositional alternatives enables us to define our assumption and operators based on a partition between premise, intermediate and conclusion propositions.

As an appetizer of the results in this paper, consider our notion of conclusion dictator. In this judgment aggregation operator, one of the agents is a dictator for the conclusion, in the sense that whatever his judgment on the conclusion, it will be the collective judgment. However, for the premises, he may follow the majority in most of the cases. In an analogous way, a conclusion majority operator is defined, which follows the majority for the conclusion, but not always for the premises. We show that these operators satisfy our weakened notion of independence, without satisfying the other independence assumptions discussed in the literature.

The layout of this paper is as follows. In Section 2 we discuss independence in judgment aggregation examples, and we argue for the importance of distinguishing premises from conclusions. In Section 3 we discuss and criticize the independence assumptions in judgment aggregation frameworks of List and Pettit [13], and of Mongin [15]. In Section 4 we introduce our weakened independence assumptions, and we show by two operators that they are compatible with the other assumptions of the framework.

2. INDEPENDENCE

In this section we first present the original aggregation paradox from which judgment aggregation originated. We then discuss various other examples from the literature. The purpose is to examine the independence (and, especially, dependence!) relations among propositions in different instances of judgment aggregation.

Dagstuhl Seminar Proceedings 07351
Formal Models of Belief Change in Rational Agents
http://drops.dagstuhl.de/opus/volltexte/2007/1216

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 to some other propositions (the *conclusions*). A typical condition imposed on the aggregation procedure is that it should treat all propositions in an even-handed way. However, we notice that premises are often independent from each other, but they are never independent from the conclusion (and vice versa). This observation is the basis for the properties of premise and conclusion independence introduced in Section 4.

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 $a, b, c, \ldots$, intermediate variables $e, f, g, \ldots$ and conclusion variables $x, y, \ldots$. Logical constraints relate premises to intermediate variables, intermediate variables to conclusion variables, or directly premises variables to conclusion variables.

2.1 Doctrinal paradox

The problem of judgment aggregation was first identified by Kornhauser and Sager [9, 10]. 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 1.

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

Table 1: 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, proposition-wise 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 problem as the *doctrinal paradox*.

As illustrated in Section 2.3.1, 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.

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. However, in the aggregation of judgments, where the propositions are connected, the independence condition is the source of the inconsistent group outcomes.

2.1.1 Premise vs conclusion-based procedure

The first two ways to avoid the inconsistency that have been suggested are the *premise-based procedure* and the *conclusion-based procedure* [17, 3]. 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. In case 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 still independent from the other premise propositions, but the aggregation of the conclusion is not. 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 1. Yet, unlike Table 1, there would be no inconsistent collective judgment on a, b and x, and the defendant would be declared not liable. The paradox in Table 1 is due to the fact that judges A and B agree on the fact that there was no breach of contract, but for different reasons. Therefore, the premises and conclusions are not independent and cannot be treated as such.

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. Moreover, no reasons for the court decision could be supplied. In the conclusion-based procedure, if there are multiple conclusion propositions, then the aggregation of a conclusion proposition is independent of the other conclusions. Since there is no aggregation of the premises, the procedure says nothing about the dependence or independence between premises and conclusions.

2.2 Discursive dilemma

The discursive dilemma [17] is a variant of the doctrinal paradox, in which the judges also vote on the logical connection rule $x \leftrightarrow (a \wedge b)$. In Table 2, we model the rule using an intermediate variable e, which represents that the rule is neither a premise nor a conclusion, but a relation between them. The discursive dilemma in Table 2 is a generalization of the doctrinal paradox in Table 1, in the sense that in the latter all agents are assumed to vote for the rule. By contrast, in the case of the doctrinal paradox it is possible that some of them vote against it. Consequently, a general framework to model examples like the discursive dilemma in Table 2 is more restricted than a general framework to model examples like the doctrinal paradox in Table 1. This is because, unlike in the doctrinal paradox, the discursive dilemma explicitly considers the notion of a rule.

We are able to discuss the independence relations in the discursive dilemma, due to the distinction between premise and conclusion propositions, and extending it with the intermediate proposition e. For instance, a premise-conclusion based procedure can be

	a	b	$x \leftrightarrow (a \wedge b)$	x
Judge A	1	0	1	0
Judge B	0	1	1	0
Judge C	1	1	1	1
Majority	1	1	1	0

Table 2: Discursive dilemma. Premises: a = There was a contract, b = There was conduct constituting breach of such a contract. Intermediate variable: $e = (x \leftrightarrow (a \wedge b))$ = legal doctrine. Conclusion: x = There was a breach of contract.

defined. This, at the best of our knowledge, has not been discussed explicitly in the literature. In this procedure, each agent votes on each premise and conclusion, and the intermediate variable is then inferred from the rule $e \leftrightarrow ((a \wedge b) \leftrightarrow x)$ and from the judgment of the majority of the group on a, b and x. In the example given in Table 2, the aggregated judgment would reject the rule. The premises and conclusions are independent, but they both depend on the intermediate propositions.

Finally, in the literature on judgment aggregation there is no discussion of what voting for or against a rule means. Consider, for example, the rule "If carbon dioxide emissions are above threshold k, then there will be global warming" (see Section 2.3.4). The conditional $a \rightarrow x$ is false when it is the case that "carbon dioxide emissions are above threshold k, and there will be no global warming". However, it is questionable if this is what an agent means when stating that $a \rightarrow x$ is false. The fact that, in many contexts, the material interpretation of the implication operator is not natural, has been observed in [5], and a characterization of admissible quota rules on a large class of agendas with subjunctive implications is given.

2.3 Variants

The importance of the judgment aggregation literature is based on various generalizations of the doctrinal paradox and the discursive dilemma. It has been used outside of the legal domain, other propositions have been voted on, it has been related to other paradoxes in the social science literature, and other rules than the majority rule have been considered.

2.3.1 *Beyond the legal domain*

Some of the other examples given in the literature with precisely the same logical structure as the original contract paradox have been listed in Table 3. Some of them come from other domains than the legal one. For instance, Bovens and Rabinowicz [2] illustrate the case of a hiring committee that agrees that someone is worthy of tenure if he or she is worthy of tenure on teaching and worthy of tenure on research.

In the examples, it seems relevant whether the agents are more concerned with the outcome of the vote or with the need of providing justifications for their final decision. If the agents of the hiring committee are external to the institution where the candidate will be hired, they may be more inclined to follow a procedure that makes explicit the reasons for the final decision, like the premise-based procedure. Instead, when the decision-makers work for the same institution, they may be view their conclusion more relevant and, therefore, let the aggregation of the premises depend on that conclusion.

In practice, it seems very hard to construct examples where such dependency considerations do not play a role. Consider, for instance, a situation from knowledge discovery where expert opinions from scientists must be aggregated. One may expect that scientists would give their opinion without any strategic considerations in mind. However, collective decisions may influence their prestige and their abilities to attract future funding, and therefore also in these cases they may be biased towards a more conclusion-based procedure.

2.3.2 *Three-premise examples*

Nehring [16] presents a variant with three premises from the legal domain, and List and Pettit [13] discuss an example from group decision making. These are both summarized in Table 5. On the one hand, as argued by Chapman for the original doctrinal paradox, in a legal domain example as in [16], a premise-based approach seems more appropriate. Indeed, the legal code requires the judges to provide arguments for their decision. On the other hand, in an example like the one in [13], a conclusion-based approach seems defendable since the agents are more affected by the final decision than by the reasons that supported that decision.

Table 4 illustrates an instance of the *Paretian Dilemma* [16] with three premises. In a Paretian Dilemma the premise-based procedure contradicts a unanimously supported conclusion.

	a	b	c	x
Agent 1	1	0	1	0
Agent 2	0	1	1	0
Agent 3	1	1	0	0
Majority	1	1	1	0

Table 4: Premises: a = Serious danger?, b = effective measure?, c = Bearable loss? Conclusion: x = Pay sacrifice?

2.3.3 *Disjunctive variants*

Chapman [3] considers the disjunctive variant of a doctrinal paradox, i.e. a case in which a majority would reject both the atomic propositions a and b as false, and yet accept the compound disjunction $(a \vee b)$ as true. Suppose that a panel of three judges have to decide whether a given tribunal had jurisdiction to hear some legal dispute. Suppose as well that there are only two ways for the tribunal to take jurisdiction. Let us call these two ways J_1 and J_2, each of which, if available, is sufficient to settle the jurisdictional dispute. Thus, there will be jurisdiction J if and only if $(J_1 \vee J_2)$. Table 6 shows the votes of the three judges and the resulting majority voting on each of the three propositions that are in dispute.

While Judges A and B form a majority in favor of finding jurisdiction J, they find this for conflicting reasons. Moreover, the problem is that the court reaches the majority view that neither J_1 nor J_2 are adequate reasons for that tribunal to take jurisdiction.

	a	b	$x = (a \vee b)$
Judge A	1	0	1
Judge B	0	1	1
Judge C	0	0	0
Majority	0	0	1

Table 6: Premises: a = There is J_1, b = There is J_2. Conclusion: x = There is J (i.e., $(J_1 \vee J_2)$).

This example shows once more that if the premise variables a and b are independent from each other, they are not independent on the conclusion x (and vice versa). To demand an aggregation rule that satisfies independence in an aggregation problem like this, inevitably opens the way to inconsistent collective judgments.

	a	b	x
	There was a contract	There was conduct constituting breach of such a contract	There was a breach of contract.
[7]	The defendant killed the victim	The defendant was sane at the time	The defendant is guilty
[4]	Iraq hides weapons of mass destruction	The war can be won with acceptable military losses	Invasion
[2]	Candidate is worthy of tenure on teaching	Candidate is worthy of tenure on research	The candidate is worthy of tenure tout court

Table 3: Doctrinal paradox: variants

	a	b	c	x
[13]	Serious danger?	Effective measure?	Bearable loss?	Pay sacrifice?
[16]	Defendant had duty to take care	Defendant behaved negligently	His negligence caused damage	Damages are due

Table 5: Three premise examples: variants

2.3.4 *Implication based examples*

An aggregation paradox variant with an implication has been given in [6]. The example considers a three-member committee that has to express agent and collective judgments on three connected propositions:
a: Carbon dioxide emissions are above the threshold k.
$a \rightarrow x$: If carbon dioxide emissions are above the threshold k, then there will be global warming.
x: There will be global warming.

	a	$e = (a \rightarrow x)$	x
Agent 1	1	1	1
Agent 2	1	0	0
Agent 3	0	1	0
Majority	1	1	0

Table 8: Premise: a, Intermediate $e = a \rightarrow x$, Conclusion: x

As shown in Table 8, the first agent accepts all three propositions; the second accepts a but rejects $a \rightarrow x$ and x; the third accepts $a \rightarrow x$ but rejects a and x. Then, the judgments of each agent are individually consistent, and yet the majority judgments on the propositions are inconsistent: a majority accepts a, a majority accepts $a \rightarrow x$, but a majority rejects the conclusion x.

If we analyze this example in terms of premise and conclusion, we see that there is a premise a (or cause), a rule (a, x) (or norm) and a conclusion x (or evidence). a and x are clearly not independent. Requiring that these three issues should be treated in the same way and such that the group outcome on each of the three issues is independent on the judgments on the two others is untenable. The premise variable is not independent from the norm $a \rightarrow x$ or from the conclusion (and vice versa). As also the discussion on the Jørgensen's dilemma in deontic logic [8] shows, to look at $a \rightarrow x$ as a proposition is a sure source of trouble.

2.4 Independence

In the contract example, it is arguable that $(a \wedge b) \leftrightarrow x$ should be in the agenda of the issues on which the judges have to vote. Judges must obey to the legal doctrine when they deliberate. To accept or reject the legal doctrine is not an option for the judges. Likewise, in most of the examples found in the literature, the rules are fixed before voting and the agents are committed to express opinions that do not violate such rules.

Despite the apparent generality of the judgment aggregation problem, surprisingly the examples are often rather biased. For instance: the original contract paradox is based on a constitutive norm defining a legal concept, and thus the issue of how to formalize legal reasoning becomes part of the aggregation problem. In order to fully understand the example, it is necessary to know that the court needs to justify its decision. On the other hand, many other examples are based on general decision problems, where other consideration may play a role.

Generally speaking, it seems that in some situations, we can make a choice between the premise and conclusion-based procedures. In legal reasoning a premise-based procedure may be called for, whereas in decision-theoretic problems a conclusion-based procedure may be favored. However, this cannot represent a solution to the discursive dilemma for the following reasons: 1. If in legal reasoning a premise-based procedure is used, what is the scope of introducing legal ontology defining institutional facts? Likewise, if in decision theoretic problems a conclusion-based procedure is opted for, why should people be concerned about the arguments? 2. The dilemma is not which one of these two extremes we should opt for, but the fact that we have only the choice between these two extremes. 3. For some examples, it is less clear whether a premise or conclusion-based procedure should be preferable.

In order to aggregate premises and conclusions, we need to better understand the *relation* between them. To study how premises and conclusion are related amounts to investigate the justification for the independence conditions found in the literature on judgment aggregation. This is what we address in the next section.

3. JUDGMENT AGGREGATION

3.1 The first impossibility theorem

In this section we introduce the formal framework of judgment aggregation. A set of agents $N = \{1, 2, \ldots, n\}$, with $n \geq 3$, has to make judgments on logically interconnected propositions. $\mathcal{L}$ is a language with atomic propositions $a, b, c, \ldots$, including the complex formulas $\neg a, (a \wedge b), (a \vee b), (a \rightarrow b), (a \leftrightarrow b)$.

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 $\neg a \in \Phi$. A subset $J \subseteq \Phi$ is called (agent or collective) *judgment set* and it is the set of propositions believed by the agents or the group. A 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 $\neg a \in J$. An n-tuple $(J_1, J_2, \ldots, J_n)$ of agent judgment sets is called *profile*.

A *judgment aggregation rule* F assigns a collective judgment set J to each profile $(J_1, J_2, \ldots, J_n)$ of agent judgment sets.

As in social choice theory, a set of rational and desirable conditions are imposed on the aggregation rules and then, typically, an impossibility result is derived. The first impossibility theorem of judgment aggregation appeared in [13]. It states that there exists no

a	b	$x = a \vee b$
(1) There is J_1	(2) There is J_2	(3) There is J (i.e., $(J_1 \vee J_2)$)
The product was defectively manufactured	The product was sold with an inadequate warning	Recovery of damages from the defendant manufacturer

Table 7: Disjunctive examples: variants

	a	$e = (a \rightarrow x)$	x
[6]	CO_2 emissions above threshold k	If CO_2 emissions above threshold k, then global warming	There will be global warming
[14]	Current CO_2 emissions lead to global warming	If current CO_2 emissions lead to global warming, then we should reduce CO_2 emissions	We should reduce CO_2 emissions.

Table 9: Implication examples: variants

aggregation rule F satisfying the following conditions:

Universal Domain (UD): the domain of F is the set of all profiles of consistent and complete judgment sets.

Anonymity: For any profiles $(J_1, \ldots, J_n), (J'_1, \ldots, J'_n)$ in the domain that are permutations of each other, $F(J_1, \ldots, J_n) = F(J'_1, \ldots, J'_n)$. Intuitively, this means that all agents have equal weight.

Systematicity: For any $a, b \in \Phi$ and any profiles $(J_1, \ldots, J_n)$, $(J'_1, \ldots, J'_n)$ in the domain, if [for all $i \in N, a \in J_i \iff b \in J'_i$], then [$a \in F(J_1, \ldots, J_n) \iff b \in F(J'_1, \ldots, J'_n)$]. This condition ensures that the collective judgment on each proposition depends only on the agent judgments on that proposition, and that the aggregation rule is the same across all propositions.

Let us illustrate systematicity with an example. The agent evaluations of a and b in the two situations below are swapped. Systematicity ensures that the collective judgment on a in the first situation is the same as the group outcome on b in the second situation (they are both mapped into 1).

a	b	$a \wedge b$
0	1	0
1	0	0
1	1	1
1	0	0

a	b	$a \wedge b$
1	0	0
0	1	0
1	1	1
0	1	0

Table 10: Systematicity

Systematicity is a very strong condition. Not only systematicity requires that the collective judgment of each proposition should depend exclusively on the agent judgments on that propositions (and not on other — assumed to be independent — propositions), but also that the way in which the collective judgment is determined should be the same across all the propositions.

3.2 Independence condition relaxed

Among the conditions imposed on F, systematicity is the most controversial. In subsequent impossibility results, systematicity has been weakened to the independence of irrelevant alternatives, which captures the independence condition imposed on preference aggregation functions in social choice theory:

Independence of Irrelevant Alternatives (IIA): For any $a \in \Phi$ and any profiles $(J_1, \ldots, J_n), (J'_1, \ldots, J'_n)$ in the domain, if [for all $i \in N, a \in J_i \iff a \in J'_i$], then [$p \in F(J_1, \ldots, J_n) \iff a \in F(J'_1, \ldots, J'_n)$].

IIA is systematicity without the neutrality condition, requiring that all propositions are equally treated. As in Table 10, the agent evaluations for a and b in the two situations in Table 11 have been swapped. The collective judgment on a (resp. b) still depends only on the agent opinions on a (resp. b). However, because the neutrality part of systematicity has been dropped, it is possible to have different aggregation rules on different propositions. For instance, in the table below, we apply majority voting on a and minority voting on b:

a	b	$a \wedge b$
0	1	0
1	0	0
1	1	1
1	0	0

a	b	$a \wedge b$
1	0	0
0	1	0
1	1	1
1	0	0

Table 11: Independence of Irrelevant Alternatives

Yet, IIA remains a contentious condition, and its main justification is that it prevents possible manipulations. The urge for a theory of judgment aggregation on normatively defensible conditions has been claimed by Mongin [15]:

> [IIA] remains open to a charge of irrationality. One would expect society to pay attention not only to the agents' judgments on ϕ, but also to their reasons for accepting or rejecting this formula, and these reasons may be represented by other formulas than ϕ in the agent sets. Before deciding that two profiles call for the same acceptance or rejection, society should in general take into account more information than is supposed in the condition. ([15] p.7)

Thus, Mongin recognizes that propositional formulas are not independent when they share propositional variables. This leads Mongin to weaken the IIA. Nevertheless, as we shall see soon, his new independence condition is not weak enough to ensure possibility results.

Mongin introduces a new IIA condition that is restricted to the atomic propositions of the language. It is then shown that this new condition together with unanimity and universal domain give a dictatorial aggregation rule.

The formal framework of [15] considers a set $\mathcal{P} = \{a_1, \ldots, a_m, \ldots\}$ of proposition variables (p.v.) of $\mathcal{L}$. As usual, *literals* are either p.v. or negations of p.v.. A *literal value* for any $a \in \mathcal{P}$ is a choice between a or $\neg a$ and is denoted by $\tilde{a}$. If Φ is an agenda, $\Phi_0 = \Phi \cap \mathcal{P}$.

DEFINITION 1. *For any $a, b \in \Phi_0$, we say that $\tilde{a}, \tilde{b}$ are connected in terms of a k-disjunction of literals if Φ contains some disjunction of k disjuncts, among which $\tilde{a}$ and $\tilde{b}$, the other disjuncts*

— if there are any — being also literals. To illustrate, $\tilde{a}$ and $\tilde{b}$ are connected in terms of a 2-disjunction iff Φ contains $\tilde{a} \vee \tilde{b}$, and in terms of a 3-disjunction if Φ contains $\tilde{b} \vee \tilde{c} \vee \tilde{a}$ for some literal $\tilde{c}$. ([15] p.4)

The following *Closure Condition on* Φ is defined:

DEFINITION 2. *(i) Closure Under Propositional Variable: if $\phi \in \Phi$, and $p \in \mathcal{P}$ appears in ϕ, then $a \in \Phi_0$. (ii) Limited Disjunctive Closure: in every 3-element subset of Φ_0, there is an element a such that each literal value of a is connected in terms of 2- or 3-disjunctions with each literal value of the other two elements, b and c. (iii) The previous condition holds with at least one 3-disjunction, i.e., if $\{a, b, c\} \subseteq \Phi_0$, there is at least one choice of literals $\tilde{a}, \tilde{b}, \tilde{c}$ for which $\tilde{a} \vee \tilde{b} \vee \tilde{c} \in \Phi$. ([15] p.4)*

Hence, for example, the agenda of the discursive dilemma as in Table 2 would contain the literals: $a, b, x, \neg a, \neg b, \neg x$, the rule $x \leftrightarrow (a \wedge b)$ would be expanded as $\neg a \vee \neg b \vee x, a \vee \neg x, b \vee \neg x$. Finally, the agenda would contain more constraints: $a \vee x, b \vee x, \neg a \vee \neg b \vee \neg x$.

Independence of Irrelevant Propositional Alternatives (IIPA): $\forall a \in \Phi_0, \forall (J_1, \ldots, J_n), (J'_1, \ldots, J'_n)$ in the domain [for all $i \in N$, $a \in J_i \iff a \in J'_i] \Rightarrow [a \in J \iff a \in J']$. Intuitively, IIPA amounts to reserving IIA only to the p.v.

Table 12 illustrates IIPA, where the independence condition holds only for the propositional variables. Hence, the complex formula $(a \wedge b) \leftrightarrow x$ is mapped (by a hypothetical aggregation rule) into 1 in the first case and into 0 in the second situation.

a	b	$(a \wedge b) \leftrightarrow x$	x
1	1	0	0
		1	
		1	
		1	

a	b	$(a \wedge b) \leftrightarrow x$	x
1	1	0	0
		1	
		1	
		0	

Table 12: Independence of Irrelevant Propositional Alternatives

The last condition imposed by Mongin is the following, intuitively desirable, unanimity preservation:

Unanimity Preservation (UP): For all $\phi \in \Phi$ and all $(J_1, \ldots, J_n)$ in the domain, for all $i \in N, \phi \in J_i \Rightarrow \phi \in J$.

It is important to observe that, unlike IIPA that is imposed only on the p.v. of the agenda, UP is imposed on *all* formulas in the agenda Φ. The best way to explain UP is with a counterexample. In Table 13 propositionwise majority voting on $(a \wedge b \wedge c)$ violates UP. Though the agents unanimously voted against $(a \wedge b \wedge c)$, the majority outcome on the atomic propositions will force the group to accept $(a \wedge b \wedge c)$ (Paretian Dilemma).

a	b	c	$(a \wedge b \wedge c)$
1	1	0	0
1	0	1	0
0	1	1	0
1	1	1	1

Table 13: Unanimity Preservation

We can now finally state Mongin's impossibility theorem:

THEOREM 1. *Assume that $|\Phi| \geq 3$ and that the Closure Condition (i), (ii), and (iii) hold. If F satisfies IIPA and UP, F is dictatorial.*

An aggregation function F is *dictatorial* if there is an agent i (the dictator) among the voters, such that, for every profile $(J_1 \ldots J_n)$, $J_j = F(J_1 \ldots J_n)$.

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

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

Table 14: Dictatorship

In Table 14, no matter what the other agents in the group say, the group judgment coincides with the agent judgment $(1, 0, 0)$. That agent is the dictator.

Anti-dictatorship does not forbid a collective outcome to be one of the agents' judgment sets. Rather, an aggregation rule is dictatorial if the *same* agent judgment is selected to be the group judgment for *any* profile in the domain.

4. INDEPENDENCE RECONSIDERED

As we have seen, the impossibility results in judgment aggregation make use of independence conditions, either in the strong form of systematicity or in the weaken form of the independence of irrelevant alternatives or even in the further weakened independence of irrelevant propositional alternatives. This suggests that relaxing such conditions is also a way to achieve more possibility results. This is precisely what we explore in this section. More specifically, we introduce further weakened independence conditions: a *premise independence of irrelevant propositional alternatives* and a *conclusion independence of irrelevant propositional alternatives*. The motivation is to define independence conditions that capture the intuition that premises are often independent from each other, but are not independent from the conclusion (and vice versa).

4.1 Strong and weak premise independence

The formal framework we use to formalize our new independence assumptions is Mongin's framework, extended with the partitioning of the p.v. into three classes. Let us decompose Φ_0 as the union of the p.v. that are premises, p.v. that are intermediate, and of the p.v. that are conclusions of the aggregation problem at hand: $\Phi_0 = \Phi_0^P \cup \Phi_0^I \cup \Phi_0^C$. In this section we discuss only the case that there are no intermediate propositions, $\Phi_0^I = \emptyset$.

Strong Premise Independence (SPI): $\forall a \in \Phi_0^P, \forall (J_1, \ldots, J_n)$, $(J'_1, \ldots, J'_n)$ in the domain [for all $i \in N$, $a \in J_i \iff a \in J'_i] \wedge [\forall x \in \Phi_0^C, x \in J \iff x \in J] \Rightarrow [a \in J \iff a \in J']$.

Table 15 shows two profiles with the same agent judgments on the premise b. However, since the collective judgment on the conclusion x is different in the two profiles (it is 0 in the first profile and 1 in the second), also the collective judgments on b in the two profiles will not coincide.

This is a first proposal to relax the independence condition in such a way to capture the logical dependency between premises and conclusion. IIA fails in this example since the collective evaluation of b in the first profile is different from the collective judgment of b in the second profile.

a	b	$x = a \wedge b$
1	0	0
0	1	0
1	1	1
1	0	0

a	b	$x = a \wedge b$
1	0	0
1	1	1
1	1	1
1	1	1

Table 15: Strong Premise Independence

We can further relax SPI and impose that, given any two profiles, the collective evaluations of a premise in the two profiles coincide if and only if the agent judgments on that premise *and* on the conclusion are the same in the two profiles. Formally, this can be expressed as follows:

Weak Premise Independence (WPI): $\forall a \in \Phi_0^P, \forall (J_1, \ldots, J_n)$, $(J'_1, \ldots, J'_n)$ in the domain [for all $i \in N, a \in J_i \iff a \in J'_i] \wedge$ $[\forall x \in \Phi_0^C, \forall i \in N, x \in J_i \iff x \in J'_i] \Rightarrow [a \in J \iff a \in J']$.

Table 16 shows two profiles with the same agent judgments on the premise b. However, since the agent judgments on the conclusion x is different in the two profiles (it is 0 for the second agent in the first profile and 1 in the second), also the collective judgments on b in the two profiles will not coincide. Strong premise independence fails in this example since the collective evaluation of b in the first profile is different from the collective judgment of b in the second profile, whereas the collective judgment of x is the same.

a	b	$x = (a \wedge b)$
1	0	0
0	1	0
1	1	1
1	0	0

a	b	$x = (a \wedge b)$
0	0	0
1	1	1
1	1	1
0	1	0

Table 16: Weak premise independence

4.2 Strong and weak conclusion independence

The strong and weak conclusion independence assumptions are defined analogously to the premise independence assumptions. Since this assumption trivially holds for any example in which there is only one conclusion, and this holds for all the benchmark examples in the literature, we do not present any examples.

Strong Conclusion Independence (SCI): $\forall x \in \Phi_0^C, \forall (J_1, \ldots, J_n)$, $(J'_1, \ldots, J'_n)$ in the domain [for all $i \in N, x \in J_i \iff x \in J'_i] \wedge$ $[\forall a \in \Phi_0^P, a \in J \iff a \in J] \Rightarrow [x \in J \iff x \in J']$.

Weak Conclusion Independence (WCI): $\forall x \in \Phi_0^C, \forall (J_1, \ldots, J_n)$, $(J'_1, \ldots, J'_n)$ in the domain [for all $i \in N, x \in J_i \iff x \in J'_i] \wedge$ $[\forall a \in \Phi_0^P, \forall i \in N, a \in J_i \iff a \in J'_i] \Rightarrow [x \in J \iff x \in J']$.

4.3 Conclusion-based procedure revised

In the revised conclusion-based procedure we consider in this section, the judgment on a takes the same evaluation as the conclusion x unless there is unanimity on a. The collective judgments on b and x are obtained by propositionwise majority voting. In contrast to the conclusion-based procedure discussed in the literature and in Section 2.1.1, we not only decide by majority but we also give an explanation for it. The explanation in this case is that b is considered in some sense more important and therefore follows the majority, whereas a is in some sense less important and therefore may be decided against the majority. For example, b may be considered as a more sensitive issue where an explanation against the majority will not be accepted by the agents.

a	b	$x = (a \wedge b)$
1	0	0
1	1	1
0	0	0
0	0	0

a	b	$x = (a \wedge b)$
1	1	1
1	1	1
0	0	0
1	1	1

Table 17: Revised Conclusion-Based Procedure

We do have strong and weak premise independence, which follows directly from the definition. However, we do not have independence of irrelevant propositional alternatives, as illustrated in Table 17. The agent judgments on proposition a are the same, but the collective judgments on the proposition are not. Since also the collective judgments on the conclusion x are distinct, the example does not violate the premise independence assumptions.

The example illustrates that we do not have a dictator, because in the first table the collective judgment coincides with the judgment of the third agent, whereas in the second table it coincides with the judgments of the first two agents. Finally it follows directly from the definition that anonymity is satisfied.

4.4 Conclusion dictator

A conclusion dictator is an agent in the group such that, whatever her opinion on the conclusion, it will be adopted as the group judgment on the conclusion. As for the collective judgments on the premises, a takes the same evaluation as the collective conclusion unless there is a unanimous judgment for a.

a	b	$x = (a \wedge b)$
1	0	0
0	0	0
0	0	0
0	0	0

a	b	$x = (a \wedge b)$
1	1	1
0	0	0
0	0	0
1	1	1

Table 18: Conclusion Dictator

We do have strong and weak premise independence, which follows directly from the definition. However, we do not have independence of irrelevant propositional alternatives, as illustrated in Table 18. The agent judgments on proposition a are the same, but the collective judgments on the proposition are not. Since also the collective judgments on the conclusion x are distinct, the example does not violate the premise independence assumptions.

The example illustrates that we do not have a dictator, because in the first table the collective judgment coincides with the judgment of the last two agents, whereas in the second table it coincides with the judgments of the first agent. Moreover, it follows directly from the definitions that anonymity is satisfied.

5. CONCLUSIONS

We study independence in judgment aggregation by examining a large set of benchmark examples from the literature, where we observe that to model these examples we have to distinguish between premises and conclusion. This is surprising, since judgment aggregation has been developed as a generalization of preference aggregation in Arrow's social choice, and there does not seem to be a need to distinguish premises from conclusion in the benchmark examples of preference aggregation such as Condorcet's paradox. This suggests that there is a fundamental distinction between judgment aggregation on the one hand, and preference aggregation on the other hand. The distinction is also absent in related probabilistic examples such as the lottery paradox [11].

We consider the role of independence in the impossibility results in judgment aggregation. The systematicity property of the first judgment aggregation framework of List and Pettit has already been weakened to independence of irrelevant alternatives and, by Mongin, to the independence of irrelevant propositional alternatives but, together with a unanimity assumption, it still leads to an impossibility result. We extend Mongin's framework with a distinction between premise, intermediate and conclusion variables, and we use the extended framework to define notions of independence that take the distinction between premise and conclusion into account. In particular, we define:

1. A strong notion of premise independence, where the aggregated premise propositions depend on the agent judgments of this proposition, as well as on the aggregated conclusion propositions.

2. A weak notion of premise independence, where the aggregated premise propositions depend on the agent judgments of this proposition, as well as on the agent and aggregated conclusion propositions.

3. A strong notion of conclusion independence, where the aggregated conclusion propositions depend on the agent judgments of this proposition, as well as on the aggregated premise propositions.

4. A weak notion of conclusion independence, where the aggregated conclusion propositions depend on the agent judgments of this proposition, as well as on the agent and aggregated premise propositions.

Since in most cases there will be several premise propositions and only one conclusion proposition, as in the running example of the doctrinal paradox or the discursive dilemma, the first two assumptions will be more often used than the latter ones. Of the first two, the strong notion already enables several interesting operators, which we illustrate using the revised conclusion procedure and the output dictator. In both cases the aggregation may in some cases depend on the aggregated conclusion. For example, the output dictator ensures that the joint decision is always determined by himself, but he allows the other agents to influence the explanation why his decision is justified.

More independence assumptions can be defined if we also take the intermediate variables into account. In the same way as we defined premise and conclusion independence, we can define independence notions that make a premise proposition depend on the intermediate propositions only, on both the intermediate propositions and the conclusion propositions, and so on.

A topic of further research is to apply our framework to actual judgment aggregation problems of agents, and to develop a logic to specify judgment aggregation problems. In particular, the logic of judgment aggregation developed by Ågotnes *et al.* [1] can be extended to cover Mongin's framework, and our extension with premise, intermediate and output propositions. For example, such a logic could be used to give a formalization of Chapman's informal analysis of judgment aggregation problems using modes ponens and modes tollens [3]. Moreover, graphical representations of independence between variables can be used, as in Bayesian, causal and utility networks, to have more detailed representations of dependencies between variables, and formal logics for independence can be used to reason about such detailed independence models.

6. REFERENCES

[1] T. Ågotnes, W. van der Hoek, and M. Wooldridge. Reasoning about judgment and preference aggregation. In *Proc. of the AAMAS-2007 Conference*, pages 554–561, 2007.

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

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

[4] C. Claussen and O. Røisland. Collective economic decisions and the discursive paradox. *Working paper - Norges Bank*, 2005.

[5] F. Dietrich. The possibility of judgment aggregation on agendas with subjunctive implications. *Journal of Economic Theory*, 2006.

[6] F. Dietrich and C. List. Arrow's theorem in judgment aggregation. *Social Choice and Welfare*, 2006.

[7] E. Dokow and R. Holzman. Aggregation of binary evaluations. *Working paper - Technion Israel University of Technology*, 2005.

[8] J. Jørgensen. Imperatives and logic. *Erkenntnis*, 7:288–296, 1938.

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

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

[11] I. Levi. List and Pettit. *Synthese*, 140:237–242, 2004.

[12] 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.

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

[14] C. List and C. Puppe. Judgment aggregation: A survey. *Working paper*, 2007.

[15] P. Mongin. Factoring out the impossibility of logical aggregation. *Working paper, Princeton University*, 2006.

[16] C. Nehring. The impossibility of a paretian ratonal. *Working paper*, 2005.

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

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