

Belief Merging and the Discursive Dilemma: An Argument-Based Account to Paradoxes of Judgment Aggregation

Gabriella Pigozzi *
King's College London
Department of Computer Science, Strand
WC2R 2LS London, UK
email: gabriella@pigozzi.org

Forthcoming in *Synthese*

Abstract

The aggregation of individual judgments on logically interconnected propositions into a collective decision on the same propositions is called judgment aggregation. Literature in social choice and political theory has claimed that judgment aggregation raises serious concerns. For example, consider a set of premises and a conclusion where the latter is logically equivalent to the former. When majority voting is applied to some propositions (the premises) it may give a different outcome than majority voting applied to another set of propositions (the conclusion). This problem is known as the *discursive dilemma* (or paradox). The discursive dilemma is a serious problem since it is not clear whether a collective outcome exists in these cases, and if it does, what it is like. Moreover, the two suggested escape-routes from the paradox — the so-called premise-based procedure and the conclusion-based procedure — are not, as I will show, satisfactory methods for group decision-making. In this paper I introduce a new aggregation procedure inspired by an operator defined in artificial intelligence in order to merge belief bases. The result is that we do not need to worry about paradoxical outcomes, since these arise only when inconsistent collective judgments are not ruled out from the set of possible solutions.

*The title of this paper in an earlier version was “Collective decision-making without paradoxes: a fusion approach”. This research was supported by the Economic and Social Research Council (ESRC), the Alexander von Humboldt Foundation, the German Federal Ministry of Education and Research, and by the German Program for the Investment in the Future (ZIP).

1 Introduction

A department of a prestigious university offers one career development fellowship to a candidate (proposition R) if and only if the candidate proposed a good research project (P) and if she has an excellent track record of publication (Q). The complex decision rule can be expressed as $(P \wedge Q) \leftrightarrow R$. Suppose that there are three members in the departmental committee. Each of them *consistently* casts her vote on R (the conclusion) depending on her judgments on P and Q (the premises). The three members vote as shown in the table below.

| | P=Good project? | Q=Excellent publication? | R=Fellowship? |
|----------|-----------------|--------------------------|---------------|
| Member 1 | Yes | Yes | Yes |
| Member 2 | Yes | No | No |
| Member 3 | No | Yes | No |
| Majority | Yes | Yes | No |

A majority voting on the propositions (*propositionwise majority voting*) produces an inconsistent result, as a majority accepts P , a majority accepts Q (therefore, a majority accepts $P \wedge Q$), but a majority rejects R . Although each member of the committee made a judgment on the propositions P , Q and R consistently with the decision rule $(P \wedge Q) \leftrightarrow R$, the collective outcome violates the same decision rule. The above is an example of the so-called *discursive dilemma* (Kornhauser and Sager 1986; Kornhauser and Sager 1993; Chapman 1998; Brennan 2001; Pettit 2001; List 2006a).¹

In order to escape the paradox and obtain a final decision, two procedures have been suggested (Pettit 2001; Chapman 2002; Bovens and Rabinowicz 2006; List 2006b). One procedure is to let each member publicly vote on each premise and

¹The fellowship example given here is logically equivalent to the one originally used in Kornhauser and Sager 1986 to illustrate the discursive dilemma.

to assign the fellowship to a candidate only if a majority of the committee members believes that the candidate submitted a good research project *and* that she has an excellent publication record track (this is called the *premise-based procedure*). The second procedure is that each member decides about P and Q and then publicly says yes to the conclusion R only if she believes that both P and Q are true (this is called the *conclusion-based procedure*). If the latter procedure is followed, the candidate will obtain the fellowship if and only if a majority in the committee voted for R .

The moral of the story is that a consistent collective outcome can only be defined if either the premise-based procedure or the conclusion-based procedure is adopted. But the two procedures may lead to contradictory results (one saying yes and the other saying no to a certain candidate), depending on whether the majority is taken on the individual judgments of P and Q , or whether the majority is calculated on the individual votes of R . This is obviously a serious problem, not only because it is unclear which of these two (to let the individuals vote on P and Q , or to let them vote on R) is the correct method, but also because it makes the collective outcome open to manipulation through the stipulation of a specific procedure.

Furthermore, the discursive paradox illustrates that collective inconsistencies can be obtained when consistent sets of propositions are aggregated. The idea is that individuals vote (in the form of yes/no) on *logically connected* propositions, and that different (and equally sensible) judgment aggregation procedures give contradictory collective outcomes, though each individual behaves rationally. As in the fellowship example, each member says yes to R only if she says yes both to P and to Q (like member 1). Voting against P (member 3) or against Q (member 2) forces the voter to reject R . All three members conform their judgment to the logical constraints between the premises (P and Q) and the conclusion R expressed by the rule $(P \wedge Q) \leftrightarrow R$. Since the conjunction of the premises is equivalent to the conclusion, one would expect that an aggregation procedure applied on the individual judgments on the premises would give the same result as the aggregation on the individual conclusions. But

this is not the case, and it is precisely where the paradox arises. Moreover, List and Pettit (2002; 2004) have shown that the contradiction does not depend on the specific choice of aggregation procedure. Rather, they prove a general impossibility theorem such that there exists no aggregation function that satisfies a minimal set of desirable conditions.² More impossibility results followed that of List and Pettit (Dokow and Holzman 2005; Gärdenfors 2005; Nehring and Puppe 2005; Pauly and van Hees 2006; Dietrich 2006).

In this paper I will argue that we need not to worry about the discursive dilemma, provided that we fully recognize the logical relations between premises and conclusion, and that we understand the logic of aggregating sets of judgments. The framework I use to define and analyze the aggregation of individual judgments profits from the collaboration of two (so far separate) research areas: belief merging and collective decision. The aggregation procedures used in belief merging and in group decision-making are similar. Belief merging defines a class of operators that produce collective belief from individual (and possibly conflicting) belief bases, and group decision-making studies operators that produce a collective decision from individual preferences. Some of the merging operators proposed in the literature were inspired by some of the voting procedures studied in social choice theory. I believe that also methods from belief merging can be fruitfully imported into group decision-making where, like in judgment aggregation, the individual's decision does not inevitably come as a preference ordering, but has a propositional form.

The main goal of this paper is to reveal that the discursive paradox disappears as soon as we recognize that aggregating logically consistent individual judgments does not guarantee a consistent collective outcome. Additional constraints need to be imposed in order to rule out infeasible group judgments. Moreover, the collective outcome does not always coincide with a unique assignment of yes/no to each proposition under consideration. It may indeed happen that a disjunction of consistent

²We will go back to these conditions in section 5, to show that the new aggregation procedure proposed here relaxes one of the conditions of List and Pettit.

judgments is chosen as the collective outcome. Hence, to make complex collective decisions without running into a paradox we may have to allow to be indecisive among multiple judgments. An outcome in the new aggregation procedure is a set of arguments. An *argument* is a consistent truth-value assignment to a conclusion *and* to the premises supporting that conclusion. The premise-based procedure and the conclusion-based procedure are therefore included in a unitary approach, which I will call *argument-based procedure*.

The structure of the paper is as follows. In section 2 it will be argued that the premise-based procedure and the conclusion-based procedure cannot be rationally justified and therefore do not solve the discursive dilemma. In section 3 the general framework of the new model will be introduced, while the formal framework will be discussed in section 4. In section 5 it will be shown that the discursive dilemma can be avoided when the argument-based procedure is used. This procedure focuses on the argument structure that connects (and logically constraints) the premises with the conclusion. Compared to the premise-based and the conclusion-based procedures, the new approach displays some additional interesting properties which allow, for example, the relaxation of various typically made conditions in judgment aggregation, like the completeness of the sets of judgments, leading to a more realistic account of group decision-making.

2 Neither Premise nor Conclusion-Based Procedure

As I sketched in the previous section, the premise-based and conclusion-based procedures have been adopted as escape-routes from the discursive dilemma. In this section I will criticize these two approaches and claim that we need an aggregation method that provides a conclusion *together* with proper reasons to support it.

The premise-based and the conclusion-based procedures were both discussed by

Pettit (Pettit 2001) within the context of *deliberative democracy* (for an introduction see Elster 1998).³ The ideal of deliberative democracy requires that the whole community be involved in the decision-making process: people have the right to ask for the reasons that supported a decision and to question them. In this respect, the premise-based procedure seems to be more appropriate. Yet there is a serious worry about this method, concerning whether it can guarantee the citizen control of the decision.

As Bovens and Rabinowicz (2006) first noticed, the problem is that it is not clear how one identifies the premises in a complex compound. Suppose that 7 voters express their judgments on the propositions A , B and C according to the decision rule $((A \wedge B) \vee (\neg A \wedge \neg B)) \leftrightarrow C$. Their judgments are shown in the following table:

| | A | B | C | $A \wedge B$ | $\neg A \wedge \neg B$ |
|----------|---|---|---|--------------|------------------------|
| Voter 1 | 1 | 1 | 1 | 1 | 0 |
| Voter 2 | 1 | 1 | 1 | 1 | 0 |
| Voter 3 | 1 | 1 | 1 | 1 | 0 |
| Voter 4 | 1 | 0 | 0 | 0 | 0 |
| Voter 5 | 1 | 0 | 0 | 0 | 0 |
| Voter 6 | 0 | 1 | 0 | 0 | 0 |
| Voter 7 | 0 | 1 | 0 | 0 | 0 |
| Majority | 1 | 1 | 0 | 0 | 0 |

Clearly this is a case of discursive paradox, since a majority says yes to the propositions A and B , yet a majority rejects the conclusion C . But, in addition, majority voting gives two divergent results depending on what we take to be the premises. If the premise-based procedure is applied to the atomic propositions A , B , then the conclusion C should be accepted. If the majority voting is applied to the disjuncts

³For a discussion of these two procedures from an epistemic perspective, see Bovens and Rabinowicz 2006 and List 2005.

$A \wedge B$ and $\neg A \wedge \neg B$, then the conclusion C is rejected.⁴ Needless to say, in the framework of judgment aggregation it is problematic to rationally justify whether we should take a premise to be an atomic proposition or not, and this renders the whole premise-based procedure open to manipulation.

Let us now turn to the conclusion-based procedure. Kornhauser and Sager (Kornhauser 1992; Kornhauser and Sager 1986, 1993) were the first to discover the discursive dilemma and in their juridical example a three-member court has to decide whether a defendant is liable under a charge of breach of a contract. According to the legal doctrine, the defendant is liable (R) if and only if the contract was valid (P) and there was a breach (Q). This case is logically equivalent to our committee example, and thus a possible discursive paradox.

A verdict in a court is a public act. Not only, if convicted, has a defendant the right to know the reasons for which she has been convicted, but also these reasons will guide future decisions — they are patterns for future verdicts. In other words, the final decision must be supported and justified by reasons. For cases similar to this, a method like the conclusion-based procedure that defines the collective outcome without providing support for it, is not satisfactory. We need an approach that provides a collective decision *and* the reasons supporting that decision.

In section 5, I will show that the paradox disappears as soon as exclusively consistent sets of judgments are accepted as candidates for group decision outcomes and, at the same time, we dispose of an aggregation procedure that does not require that a unique assignment of yes/no to the propositions under consideration is selected as the collective outcome. The new method will select the most popular argument (and eventually more than one). By argument I mean a consistent judgment on a conclusion together with the reasons supporting that conclusion. The premise-based procedure and the conclusion-based procedure will be reconciled in a unique approach with no fear of paradoxical results.

⁴Bovens and Rabinowicz referred to this problem as the *instability* of the premise-based procedure.

3 The Framework

The aggregation procedure that will be used in this paper is inspired by a merging operator defined in artificial intelligence. A merging operator combines⁵ various and possibly conflicting belief bases. A *belief base* K_i is a finite set of propositional formulas representing the explicit beliefs of the individual i . The merged base is a set which *consistently* integrates parts of the knowledge from all the initial bases, and satisfies some additional conditions, such as *integrity constraints*.⁶ The items of the resulting collective knowledge base are identified by rules like the majority rule. It should therefore be clear that belief merging and general aggregation problems share a similar objective, i.e. the definition of operators that produce collective belief from individual belief bases, and operators that produce a collective decision from individual judgments (or preferences).⁷ Some researchers have found inspiration for belief merging operators by examining voting procedures from social choice theory (Konieczny 1999; Lin and Mendelzon 1999). Here I want to show how methods from belief fusion can be fruitfully imported into judgment aggregation to tackle specific problems.

First, some terminology is in order. Following the literature on belief merging, the words ‘knowledge’ and ‘belief’ will be used interchangeably in this paper.⁸ This

⁵Konieczny (2000) refers the term ‘combination’ to the syntax-based fusion operators. We will instead use the verbs ‘to combine’, ‘to merge’ and ‘to fuse’ as synonyms.

⁶An integrity constraint (Kowalski 1978; Reiter 1988) is a sentence that has to be satisfied by the merged base. It is usually not required that the individual bases satisfy the integrity constraints, though I will assume it. By doing so, I will maintain the model as close as possible to the original formulation of the discursive paradox, where each voter consistently casts her vote on a set of propositions. We will see that integrity constraints play a crucial role in avoiding irrational sets of collective judgments.

⁷Social choice theory studies how individual preferences are combined in order to select a collectively preferred alternative. Relationships between judgment aggregation and preference aggregation have been investigated (List and Pettit 2004; Dietrich and List 2006). Also the links between belief merging and social choice theory have been explored (Konieczny and Pino-Pérez 2005; Eckert and Pigozzi 2005; Pigozzi 2005).

⁸Also in the belief revision literature (a separate but related area to belief fusion, see Konieczny 1999) the use of the word ‘knowledge’ in a broader sense than in the epistemological literature — such that ‘knowledge’ also covers what is traditionally meant by ‘belief’ — is commonly accepted (Gärdenfors 1988).

is because a fusion operator can be interpreted as an operator merging knowledge or belief bases. Furthermore, the elements of a knowledge base are not required to be true. This would simply be too big an idealization for practical applications in which knowledge is taken to be defeasible — for example when an expert system is defined by merging the knowledge of a group of human experts.

When merging individual belief bases, two cases can occur. If all individual bases are mutually consistent, then the collective outcome can easily be constructed: it is the union of all the individual bases. More interesting, however, is the case when the individual belief bases are in conflict with each other. There are two main approaches to this problem depending on whether all the individuals are treated equally or not.⁹ For the purpose of the paper, fusion operators for individuals who have the same power to influence the final decision will be considered. I will, in fact, restrict myself to the classical discursive dilemma where all voters have equal power.¹⁰

Some merging operators for knowledge bases of equally reliable sources have been proposed by Borgida and Imielinski (1984), Konieczny and Pino-Pérez (1998), Lin and Mendelzon (1999). These operators come in two types. On the one hand, there are majority operators that minimize the level of total dissatisfaction (Konieczny 1999; Lin and Mendelzon 1996). On the other hand, there are arbitration operators, which define rules to equally distribute the level of individual dissatisfaction among the group members (Konieczny 1999). The model presented here will make use of a majority operator studied by Konieczny and Pino-Pérez (1998, 1999, 2002a). This merging operator identifies the collective outcome with the set of models that are closest to the individual bases. The result is the same as the outcome we obtain by propositionwise majority voting.¹¹ Yet, the paradox is avoided by adding domain-specific restrictions (integrity constraints) to the collective base. One of the key points in the literature of knowledge fusion is that merging consistent knowledge bases does

⁹In belief fusion, the latter corresponds to belief bases provided with different reliability values (see, for example, Benferhat *et al.* 1999; Cholvy 1994; Lin 1996; Maynard-Reid and Shoham 1998).

¹⁰For a study of expert rights in judgment aggregation see Dietrich and List 2004.

¹¹This will be clarified in the following of the paper, see section 5.

not guarantee a consistent collective outcome. To overcome this problem, the integrity constraints are imposed on the final bases. This ensures that inconsistent models for the collective are ruled out from the set of possible group decisions. Finally, unlike the existing models for judgment aggregation, *all* possible consistent results for the collective outcome are explored and a ranking of them is defined.

4 The Formal Model

Let $N = \{1, 2, \dots, n\}$ ($n \geq 2$) be a set of individuals. Let \mathcal{L} be a finitary propositional language, built up from a finite set \mathcal{P} of propositional letters and the usual logical connectives (\neg , \wedge , \vee , \rightarrow , and \leftrightarrow). The belief base K_i of an agent i is a consistent and complete¹² finite set of atomic propositions like A and B , and compound propositions like $\neg A$, $A \wedge B$, $(A \vee B) \leftrightarrow C$, and so on. K_i can also be represented as the conjunction of its propositional formulas.

An *interpretation* is a function $\mathcal{P} \rightarrow \{0, 1\}$ and it is represented as the list of the binary evaluations. For example, given three propositional variables P , Q and R , the vector $(0, 1, 0)$ stands for the interpretation in which P and R are false and Q 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. For instance, the interpretation $(1, 0, 0)$ is a model of $(P \wedge Q) \leftrightarrow R$, but it is not a model of $P \wedge Q \wedge R$. Finally, the set of models of a formula is the set of all the interpretations that make that formula true.

IC is the belief base whose elements are the integrity constraints. Given a multi-set¹³ $\underline{K} = \{K_1, K_2, \dots, K_n\}$ and IC , a merging operator \mathcal{F} is a function that assigns a belief base to \underline{K} and IC . By borrowing the term from judgment aggregation, we call \underline{K} a *profile*. Let $\mathcal{F}_{IC}(\underline{K})$ denote the collective belief base resulting from the IC

¹²In the new argument-based procedure the completeness requirement can be relaxed. However, here K_i is assumed to be complete to keep the model as close as possible to the original formulation of the discursive dilemma.

¹³In a multi-set an element can appear several times. Here, it means that two (or more) individuals can hold the same judgments on the propositions about which they have to express their view.

merging on \underline{K} . In this paper I will use the model-based merging operator introduced by Konieczny and Pino-Pérez (1999).¹⁴ The idea of a model-based fusion operator is that models of $\mathcal{F}_{IC}(\underline{K})$ are models of IC , which are preferred according to some criterion depending on \underline{K} . Usually the preference information takes the form of a total pre-order (\leq) on the interpretations induced by a notion of distance $d(\omega, \underline{K})$ between an interpretation ω and the profile \underline{K} .

In a distance-based approach $\leq_{\underline{K}}$ is defined in two steps. The *first* step determines a distance between each possible interpretation (that is, those interpretations that are models of IC) and the models of each K_i in the profile \underline{K} . A 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 $\omega = \omega'$.

Since a belief base can have more than one interpretation that makes it true,¹⁵ the distance between an interpretation ω and a belief base K is the minimal distance d between ω and any model of K . Formally:

$$d(\omega, K) = \min_{\omega' \models K} d(\omega, \omega')$$

Intuitively, the distance is a number that measures the level of dissatisfaction of each individual in the group with respect to a possible collective outcome.

The *second* step is to assign a distance value to an interpretation ω and a profile \underline{K} . This is defined with the help of an aggregation function $D : \mathbb{R}_+^n \rightarrow \mathbb{R}_+$ as $D^d(\omega, \underline{K}) = D(d(\omega, K_1), d(\omega, K_2), \dots, d(\omega, K_n))$ (Konieczny *et al.* 2004). Any such aggregation function induces a total pre-order $\leq_{\underline{K}}$ on the set \mathcal{W} with respect to the distances to a given profile \underline{K} . Thus, an IC majority merging operator \mathcal{F} for a

¹⁴This type of operators are also discussed in (Revesz 1997; Lin and Mendelzon 1999; Liberatore and Schaerf 2000). For syntax-based fusion operators, see Baral *et al.* 1992 and Konieczny 2000.

¹⁵This is true in general, but not for complete bases, like the ones we are taking into consideration in this paper, which have a unique model.

profile \underline{K} can be defined as $\mathcal{F}_{IC}(\underline{K}) = \min([IC], \leq_{\underline{K}})$, i.e., the set of all models of IC (denoted by $[IC]$) with minimal distance D^d to the profile \underline{K} . The minimal distance identifies the final collective outcome, i.e. the belief base with the lowest total level of dissatisfaction among all possible models satisfying IC .

The most widely used majority merging operator in the literature is the operator $\Delta^{d,\Sigma}$ 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. $D^d(\omega, \underline{K}) = \sum_i d(\omega, K_i)$ is the sum of componentwise distances d defined before.

For example, the Hamming distance between $\omega = (1, 0, 0, 1)$ and $\omega' = (0, 1, 0, 1)$ is $d(\omega, \omega') = 2$. We use the Hamming distance because it is a well known, intuitive and easy to define distance. But the Hamming distance is only one among many possible distance functions that we may use. In fact, as shown in (Konieczny and Pino-Pérez, 2002b), the logical properties of a merging operator do not depend on the chosen distance.¹⁶ For the purpose of the present paper, it is interesting to mention that, when there is no paradox, merging individual bases using the Hamming distance gives the same result as propositionwise majority voting. The difference between the two procedures is that — unlike propositionwise majority voting — the merging operator excludes the inconsistent collective judgments (*via IC*), and defines a ranking on all the allowed collective judgments. The observation that, were the inconsistent judgments not excluded by IC , the merging operator and propositionwise majority voting would give the same outcome, shows that the fusion method used to tackle the discursive paradox is not an *ad hoc* method, but really a more fine-grained procedure than propositionwise majority. We are now ready to apply the method of belief merging to judgment aggregation.

¹⁶The logical properties of the merging operators are studied by Konieczny and Pino-Pérez in a number of papers, for example in (1999).

5 The Argument-Based Procedure

The lesson of the discursive dilemma is that propositionwise majority voting may fail to assign a consistent collective judgment to certain individual judgments. List and Pettit (2002) have proven a general impossibility theorem to show that the occurrence of an inconsistency does not depend on a specific aggregation procedure. Their impossibility result holds for all aggregation functions that satisfy some minimal conditions, such as universal domain, anonymity and systematicity. *Universal domain* states that any logically possible individual judgment is accepted as an input by the aggregation procedure. Moreover, *anonymity* ensures that all the voters are equally treated. *Systematicity* is the condition requiring that “the collective judgment on each proposition should depend exclusively on the pattern of individual judgments on that proposition. In particular, the collective judgment on no propositions should be given special weight in determining the collective judgments on others” (List and Pettit 2002, p.98). List and Pettit recognize that the condition of systematicity is controversial and that, when coupled with anonymity, it implies that if two propositions receive the same degree of support (not necessarily from the same individuals) the collective view on the two propositions should be the same.¹⁷ As in the career development fellowship example, where a majority of the members in the committee supports P and *another* majority supports Q , systematicity requires that the collective outcome endorses $P \wedge Q$ in contradiction to the majority voting on the conclusion against R .

Systematicity is relaxed in the argument-based procedure that I propose here. As the next example will make clear, the argument-based procedure gives priority to the integrity constraints in IC over the other propositions. The view of the group is selected from the models of IC according to a distance measure that takes into account each K_i . A set of possible interpretations is assigned to each belief base. Suppose that $K_i = \{(A \vee B) \wedge C\}$. The set ω of interpretations for K_i is $Mod(K_i) =$

¹⁷For a critique on systematicity see Chapman 2002.

$\{(1, 0, 1), (0, 1, 1), (1, 1, 1)\}$. If, for instance, an integrity constraint requires A to be true, the model $(0,1,1)$ will not be considered as a possible outcome.

In the fellowship example all three members consistently make their judgments on the propositions P , Q and R given the decision rule $(P \wedge Q) \leftrightarrow R$. Therefore $\underline{K} = \{K_1, K_2, K_3\}$ and $IC = \{(P \wedge Q) \leftrightarrow R\}$. We can then write:

$$K_1 = \{A, B, C\}$$

$$K_2 = \{A, \neg B, \neg C\}$$

$$K_3 = \{\neg A, B, \neg C\}$$

The models for each belief base are the following:

$$Mod(K_1) = \{(1, 1, 1)\}$$

$$Mod(K_2) = \{(1, 0, 0)\}$$

$$Mod(K_3) = \{(0, 1, 0)\}$$

The table below shows the result of the IC majority fusion operator on $\underline{K} = \{K_1, K_2, K_3\}$. In the first column are all the possible interpretations for the propositional variables P , Q and R . The rows with a shaded background correspond to the interpretations excluded by IC . The numbers in the columns of K_1 , K_2 and K_3 are the Hamming distances of each K_i from the correspondent interpretation. Finally, in the last column is $\sum_i d(\omega, K_i)$.

| | K_1 | K_2 | K_3 | $\Sigma_i d(\omega, K_i)$ |
|---------|-------|-------|-------|---------------------------|
| (1,1,1) | 0 | 2 | 2 | 4 |
| (1,1,0) | 1 | 1 | 1 | 3 |
| (1,0,1) | 1 | 3 | 1 | 5 |
| (1,0,0) | 2 | 2 | 0 | 4 |
| (0,1,1) | 1 | 1 | 3 | 5 |
| (0,1,0) | 2 | 0 | 2 | 4 |
| (0,0,1) | 2 | 2 | 2 | 6 |
| (0,0,0) | 3 | 1 | 1 | 5 |

Because $\Delta_{IC}^{d,\Sigma}(\underline{K})$ is an *IC* merging operator, the collective outcomes are chosen among the interpretations that are models of *IC*. Only (1, 1, 1), (1, 0, 0), (0, 1, 0) and (0, 0, 0) are the available candidates for the collective judgments. Moreover, $\Delta_{IC}^{d,\Sigma}(\underline{K})$ selects the interpretations associated with the minimum distance value.¹⁸ Thus, no paradox arises using the fusion operator: $Mod(\Delta_{IC}^{d,\Sigma}(\underline{K})) = \{(1, 1, 1), (1, 0, 0), (0, 1, 0)\}$. These are the models of the disjunction $K_1 \vee K_2 \vee K_3$. This result can be interpreted as a tie between belief bases. In *this* case, the paradox is avoided at the price of indecision.

The tie means that there is not enough information for the collectivity to choose a unique truth-assignment to P , Q and R . Yet, there are situations in which, despite of the lack of information, the group needs to opt for one of the elements in the tie. Social choice theory, for example, appeals to tie-breaking rules. These are, by definition, arbitrary. However, randomly choosing one of the belief bases in a judgment aggregation context appears difficult to justify. Instead, in cases like our fellowship example, some conditions can be relaxed, namely the assumptions that each voter and each alternatives are equally treated. For instance, degrees of expertise can

¹⁸Minimizing the total distance is a way to capture the intuition that in a majority merging operator the collective outcome endorses a set of opinions if this is supported by a large part of the group.

be assessed for every individual in the group with respect to the propositions under judgment. It is finally worth noticing that an aggregation procedure on n bases K_i that would *always* identify the social outcome with the disjunction $K_1 \vee K_2 \vee \dots \vee K_n$ would admittedly not be satisfactory. But, as the next example will show, this is not the case in the merging approach that we are considering here.

Before moving to the next example, where the merging operator picks one of the K_i s as the group outcome, there are three points that we should note. The *first* is that had it not been discarded by IC , $\Delta_{IC}^{d,\Sigma}(\underline{K})$ would have selected $(1, 1, 0)$ as the model for the collective outcome. This is precisely the result associated with the discursive paradox. The fact that, in the case of aggregation of binary evaluations, the majority voting minimizes the sum of Hamming distances, has already been noticed elsewhere (see Brams *et al.* 2004, Proposition 4). Unfortunately, it may happen that the closest model violates the integrity constraints (this is when the discursive dilemma arises). This, however, is not a danger with $\Delta_{IC}^{d,\Sigma}(\underline{K})$ because the inconsistent options are rejected by IC . The *second* point is that while majority rule on propositions can find only one model, the fusion operator looks at all possible outcomes and orders them according to some criteria (for the majority merging operator the criterion is the minimal distance). This is important especially in those cases where the fusion operator is used to aggregate several bases, since it guarantees that no relevant information is discarded. *Third*, note that normally a fusion operator does not require the belief bases to satisfy IC . Here I assume that each K_i satisfies IC for ease of comparison with the classical discursive paradox. When this assumption is relaxed, the interpretations of the belief bases that violate the integrity constraints are ruled out by the IC imposed on the final outcome. The collectivity is ensured to be rational, despite the fact that some of its members (or all) made irrational judgments.

Let us now show that not only does the merging approach avoid the inconsistent results of the propositionwise majority voting, but it is also not sensitive to the instability problem. As we have seen, the premise-based procedure may lead to

contradictory results. In section 2, the example of a group of people expressing their judgments on the propositions A , B and C consistently with the rule $((A \wedge B) \vee (\neg A \wedge \neg B)) \leftrightarrow C$ illustrated that different outcomes can be obtained by the premise-based procedure, depending on whether the majority is calculated on the atomic propositions A , B , or on the disjuncts $A \wedge B$ and $\neg A \wedge \neg B$. Transforming the judgments of the example in section 2 into interpretations for each belief base, we obtain:

$$\begin{aligned} Mod(K_1) &= Mod(K_2) = Mod(K_3) = \{(1, 1, 1)\} \\ Mod(K_4) &= Mod(K_5) = \{(1, 0, 0)\} \\ Mod(K_6) &= Mod(K_7) = \{(0, 1, 0)\} \end{aligned}$$

The table below displays the results of IC majority fusion operator on $\underline{K} = \{K_1, K_2, \dots, K_7\}$. Interpretations that violate IC do not appear.

| | K_1 | K_2 | K_3 | K_4 | K_5 | K_6 | K_7 | $\Sigma_i d(\omega, K_i)$ |
|---------|-------|-------|-------|-------|-------|-------|-------|---------------------------|
| (1,1,1) | 0 | 0 | 0 | 2 | 2 | 2 | 2 | 8 |
| (1,0,0) | 2 | 2 | 2 | 0 | 0 | 2 | 2 | 10 |
| (0,1,0) | 2 | 2 | 2 | 2 | 2 | 0 | 0 | 10 |
| (0,0,1) | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 14 |

The shaded row is the *unique* model of $((A \wedge B) \vee (\neg A \wedge \neg B)) \leftrightarrow C$ selected as the group outcome. Because $\Delta_{IC}^{d, \Sigma}(\underline{K})$ is a model-based operator, it is not sensitive to the instability problem that affects the premise-based procedure. However, this example also shows that the argument-based procedure always prevents an inconsistent outcome to be chosen as the group's decision and, in some cases, also selects a unique judgment as the social result.

6 Conclusion

Two lessons can be drawn from the discursive paradox. The first one is that the set of propositions on which most group members agree is not guaranteed to be a candidate for the collective decision because the set can fail to satisfy consistency even though each individual consistently expressed her judgments. The second lesson is that, when this happens, we need to look for the second (or third...) best outcome. Complex collective decisions are therefore paradox-free when the logical relations between propositions are treated as integrity constraints on the collective judgment.

The premise-based procedure and the conclusion-based procedure have been presented in the literature on judgment aggregation as the escape-routes from the discursive dilemma. In section 2, I criticized both these procedures and urged an approach that reconciles the premises and the conclusion as part of the same argument.

The approach introduced in this paper profits from the work done on belief merging in artificial intelligence. I have claimed that judgment aggregation and belief fusion are similar processes and that a collaboration between these two areas is definitely fruitful. A majority *IC* merging operator for belief bases has been introduced. The integrity constraints ensure that the merging operator looks for solutions only among the consistent models. Each belief base is confronted with a group decision candidate, and a number is assigned as a result of the comparison. Intuitively, this number captures how much the two bases disagree with each other. Finally, the fusion function selects the model (or the models) that keeps the disagreement of most belief bases at the lowest level.

The value of the argument-based procedure rests upon the exclusion of inconsistent sets of judgments from the set of the candidates apt to become collective judgments, and the definition of a ranking on the remaining candidates.

From these results, implications obtain for all those groups of people — like expert panels, boards, councils, societies, etc. — that have to make decisions on the basis of logically connected propositions, and moreover, want or need to be able to justify

their decisions. In order to define a more realistic model of group decision-making, I plan to investigate extensions of the argument-based procedure. For example, as I already mentioned, belief merging allows us to relax the completeness requirement of the belief bases, on the ground that we want to conceive of the possibility of an individual to be indifferent (with respect to a preference) or ignorant (of a certain matter). Let us suppose that we have a finite number of belief bases expressing the decisions of some individuals about three propositions A , B and C . $K_i = \{A, B\}$ is a belief base even if it is not complete. The agent is indifferent to, or unable to make a decision on, C . Therefore K_i is satisfied by $\{(1, 1, 1), (1, 1, 0)\}$.

Also, an *IC* fusion operator can assign an outcome to belief bases that violate *IC*. Clearly, the models for the bases that violate *IC* will not be taken into consideration as candidates for the final decision. But this is an important feature because it gives us the possibility to merge bases even if some of them (or all) do not conform to *IC*. For instance, an individual who does not obey *IC* has no chance that her own view will be adopted by the group.

Majority fusion operators are not the only operators that can be introduced to merge individuals' judgments where each member has equal power to influence the final decision. An alternative fusion aggregation is the arbitration. The best outcome for an arbitration operation is a belief base that aims at equally distributing the level of dissatisfaction among the individuals. In some cases this allows us to distinguish between two incomparable judgments (see also Konieczny and Pino-Pérez 2002b). For instance, a majority merging operator over $\underline{K} = \{K_1, K_2, K_3\}$ finds that two models have $\sum_i d(\omega, K_i)$ equal to 6. Still, the arbitration operator will prefer the one that equally distributes the dissatisfaction among members (like (2,2,2) to (5,1,0)). A question open to investigation is whether an arbitration operator can be used in any kind of judgment aggregation process. Is the arbitration operator likely to capture only the aggregation of individual desires or can it also be used for the aggregation of beliefs?

Finally, it would be interesting to study the aggregation judgment where the individuals do not all have the same power to influence the final decision, following the work done by Dietrich and List (2004). Setting different levels of expertise over the individuals could be especially useful in those scenarios in which a tie between belief bases is obtained as the social outcome, nevertheless a decision must be made, without invoking an arbitrary tie-breaking rule.

Acknowledgments

This paper was first presented at the workshop *Judgment Aggregation and the Discursive Dilemma* at the University of Konstanz, Germany, June 2004. I wish to thank the organizers, Franz Dietrich and Christian List, and the participants of that workshop for stimulating discussions. Furthermore, I thank Luc Bovens, Fabrizio Cariani, Daniel Eckert, Stephan Hartmann, Sébastien Konieczny, Jérôme Lang, David Makinson, and two anonymous referees for valuable comments on a previous version of this paper.

References

- Baral, C., S. Kraus, J. Minker and V. S. Subrahmanian, 1992. “Combining knowledge bases consisting of first-order theories”, *Computational Intelligence*, 8(1), 45-71.
- Benferhat, S., D. Dubois, H. Prade and M. A. Williams, 1999. “A practical approach to fusing prioritized knowledge bases”, *EPIA 1999*, 223-236.
- Borgida, A., and T. Imielinski, 1984. “Decision making in committees: A framework for dealing with inconsistency and non-monotonicity”, in *Proceedings Workshop on Nonmonotonic Reasoning*, 21-32.
- Bovens, L., and W. Rabinowicz, 2006. “Democratic answers to complex questions.

- An epistemic perspective”, *Synthese*, 150: 131–153.
- Brams, S. J., D. M. Kilgour, and M. R. Sanver, 2004. “A minimax procedure for negotiating multilateral treaties”, preprint, Department of Politics, New York University.
- Brennan, G., 2001. “Collective coherence?”, *International Review of Law and Economics*, 21(2), 197-211.
- Chapman, B., 1998. “More easily done than said: Rules, reason and rational choice”. *Oxford Journal of Legal Studies*, 18: 293-330.
- Chapman, B., 2002. “Rational Aggregation”, *Politics, Philosophy and Economics*, 1(3), 337-354.
- Cholvy, L., 1994. “A logical approach to multi-sources reasoning”, in *Knowledge Representation and Reasoning under Uncertainty: Logic at Work*, ed. by M. Masuch and L. Polos, Springer, LNAI 808, 183-196.
- Dietrich, F., 2006. “Judgment aggregation: (Im)possibility theorems”, *Journal of Economic Theory*, 126(1): 286–298.
- Dietrich, F., and C. List, 2004. “A liberal paradox for judgment aggregation”, *Economics Working Paper Archive at WUSTL*, <http://ideas.repec.org/p/wpa/wuwppe/0405003.html>
- Dietrich, F., and C. List, 2006. “Arrow’s theorem in judgment aggregation”, *Social Choice and Welfare*, forthcoming.
- Dokow, E., and R. Holzman, 2005. “Aggregation of binary evaluations”, Working paper, Technion, Israel. <http://personal.lse.ac.uk/LIST/doctrinalparadox.htm>

- Eckert, D., and G. Pigozzi, 2005. “Belief merging, judgment aggregation, and some links with social choice theory”, in J. Delgrande, J. Lang, H. Rott, and J. M. Tallon (eds.), *Belief Change in Rational Agents: Perspectives from Artificial Intelligence, Philosophy, and Economics*, Dagstuhl Seminar Proceedings 05321, Internationales Begegnungs- und Forschungszentrum (IBFI)
<http://drops.dagstuhl.de/opus/volltexte/2005/333/pdf/05321.PigozziGabriella1.Paper.333.pdf>
- Elster, J. (ed.), 1998. *Deliberative Democracy*, Cambridge University Press.
- Gärdenfors, P., 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, The MIT Press.
- Gärdenfors, P., 2005. “An Arrow-like theorem for voting with logical consequences”, *Economics and Philosophy*, forthcoming.
- Konieczny, S., 1999. *Sur la Logique du Changement: Révision et Fusion de Bases de Connaissance*, Ph.D. dissertation, University of Lille I, France.
- Konieczny, S., 2000. “On the difference between merging knowledge bases and combinig them”, in *Proceedings of KR’00*, Morgan Kaufmann, Breckenridge, Colorado, USA, 135-144.
- Konieczny, S., J. Lang, and P. Marquis, 2004. “DA² merging operators, *Artificial Intelligence*, 157: 49–79.
- Konieczny, S., and R. Pino-Pérez, 1998. “On the logic of merging”, in *Proceedings of KR’98*, Morgan Kaufmann, 488-498.
- Konieczny, S., and R. Pino-Pérez, 1999. “Merging with integrity constraints”, in *Proceedings of ECSQARU’99*, LNAI 1638, 233-244.
- Konieczny, S., and R. Pino-Pérez, 2002a. “Merging information under constraints: a logical framework”, *Journal of Logic and Computation*, 12(5), 773–808.

- Konieczny, S., and R. Pino-Pérez, 2002b. “On the frontier between arbitration and majority”, in *Proceedings of the Eight International Conference on Principles of Knowledge Representation and Reasoning (KR '02)*, pages 109–118.
- Konieczny, S., and R. Pino-Pérez, 2005. “Propositional belief base merging or how to merge beliefs/goals coming from several sources and some links with social choice theory”, *European Journal of Operational Research*, 160(3): 785–802
- Kornhauser, L. A., 1992. “Modeling collegial courts II. Legal doctrine”, *Journal of Law, Economics and Organization*, 8, 441-470.
- Kornhauser, L. A., and L. G. Sager, 1986. “Unpacking the court”, *Yale Law Journal*, 96, 82-117.
- Kornhauser, L. A., and L. G. Sager, 1993. “The one and the many: Adjudication in collegial courts”, *California Law Review*, 81, 1-51.
- Kowalski, R., 1978. “Logic for Data Description”, in *Logic and Data Bases*, ed. by H. G. J. Minker, New York: Plenum, 77-102.
- Liberatore, P., and M. Schaerf, 2000. “Brels: A system for the integration of knowledge bases”, *Proceedings of KR 2000*, 145-152.
- Lin, J., 1996. “Integration of weighted knowledge bases”, *Artificial Intelligence*, 83, 363-378.
- Lin, J., and A. Mendelzon, 1996. “Merging databases under constraints“, *International Journal of Cooperative Information Systems*, 7, 55-76.
- Lin, J., and A. Mendelzon, 1999. “Knowledge base merging by majority”, in *Dynamic Worlds: From the Frame Problem to Knowledge Management*, ed. by R. Pareschi and B. Fronhofer, Kluwer.

- List, C., 2005. “The probability of inconsistencies in complex collective decisions”, *Social Choice and Welfare*, 24(1): 3–32.
- List, C., 2006a. Judgment aggregation - A bibliography on the discursive dilemma, the doctrinal paradox and decisions on multiple propositions.
<http://personal.lse.ac.uk/LIST/doctrinalparadox.htm>
- List, C., 2006b. “The discursive dilemma and public reason”, *Ethics*, 116(2): 362–402.
- List, C., and P. Pettit, 2002. “Aggregating sets of judgments. An impossibility result”, *Economics and Philosophy*, 18, 89-110.
- List, C., and P. Pettit, 2004. “Aggregating sets of judgments. Two impossibility results compared”, *Synthese*, 140, 207-235.
- Maynard-Reid, P., and Y. Shoham, 1998. “From belief revision to belief fusion”, in *Proceedings of the Third Conference on Logic and the Foundations of Game and Decision Theory (LOFT3)*, ICER, Torino, Italy.
- Nehring, K., and C. Puppe, 2005. “Consistent judgment aggregation: A characterization”, working paper, University of Karlsruhe.
- Pauly, M., and M. van Hees, 2006. “Logical constraints on judgment aggregation”, *Journal of Philosophical Logic*, 35.
- Pettit, P., 2001. “Deliberative democracy and the discursive dilemma”, *Philosophical Issues*, 11: 268–299.
- Pigozzi, G., 2005. “Two aggregation paradoxes in social decision making: The Ostrogorski paradox and the discursive dilemma”. *Episteme: A Journal of Social Epistemology*, 2(2): 33–42.

Reiter, R., 1988, “On integrity constraints”, in *Proceedings of the Second Conference on the Theoretical Aspects of Reasoning about Knowledge*, ed. by M. Y. Vardi, San Francisco, Calif.: Morgan Kaufmann, 97-111.

Revesz, P., 1997, “On the semantics of arbitration”, *International Journal of Algebra and Computation*, 7(2), 133-160.