On Consistent and Legitimate Multi-issue Group Decisions

Gabriella Pigozzi

LAMSADE - Université Paris-Dauphine - France
(Joint works with Martin Caminada and Umberto Grandi)

Workshop *Judgment Aggregation - Computational Perspectives*
Luxembourg, February 1, 2012

Aggregation problems

- I guess I can skip the judgment aggregation part. ☺
- (One of) my obsession(s): collective decisions don't have only to be logically consistent but also legitimate.
- This obsession started when I first read about the *multiple election paradox* [Brams, Kilgour and Zwicker, 1998].

The multiple elections paradox

Voter 1	yes	yes	no
Voter 2	yes	yes	no
Voter 3	yes	no	yes
Voter 4	yes	no	yes
Voter 5	no	yes	yes
Voter 6	no	yes	yes
Voter 7	no	yes	yes
Voter 8	no	yes	yes
Voter 9	yes	no	no
Voter 10	yes	no	no
Majority	yes	yes	yes

There are also instances with abstention.

Motivation

- The MEP produces arbitrary election outcomes.
- JA outcomes can be logically inconsistent but *also* arbitrary in the sense of the MEP.
- **Research question**: when is a social outcome legitimate? Intuition: when it is compatible with the individual positions.
- Little attention to the notion of legitimacy so far.
- *Inconsistent* outcome ⇒ No group decision can be taken.
- *Illegitimate* outcome ⇒ Agents resist to perform the actions implied by the decision or might not accept to be deemed responsible for the group decision.
- Both problems are of the utmost importance in the design of well-behaved mechanisms for collective decisions.

Argumentation framework [Work with Martin Caminada]

Two parts talk.

- Argumentation framework: a set of arguments and a *defeat* relation among them: $AF = (Ar, def)$.

$$C \rightarrow B \rightarrow A$$

- Argumentation theory identifies the sets of arguments (extensions) that can reasonably survive the conflicts expressed in the argumentation framework.
- Each agent assigns a label to each argument:
  - *in* if he accepts the argument
  - *out* if he rejects it
  - *undec* if he abstains from it

Argument labellings

We use the argument labelling approach to define the argument based semantics. A labelling is a total function

$$\mathcal{L} : Ar \rightarrow \{in, out, undec\}$$

- An argument is *in* iff all its defeaters are *out*.
- An argument is labelled *out* iff it has at least a defeater that is labelled *in*. $\Rightarrow$ *Gunfight rules*
- *A* is *undec* iff it has at least a defeater that is labelled *undec* and has no defeater labelled *in*.

Nixon

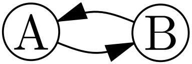

A Nixon is a pacifist because he is a quaker.

B Nixon is not a pacifist because he is republican.

Nixon

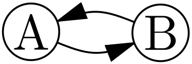

A Nixon is a pacifist because he is a quaker.

B Nixon is not a pacifist because he is republican.

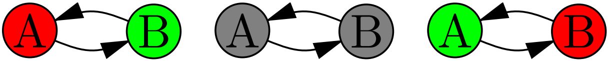

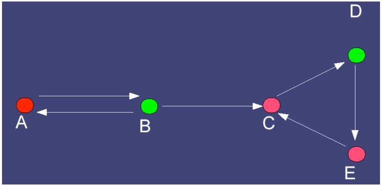
D
A
B
C
E

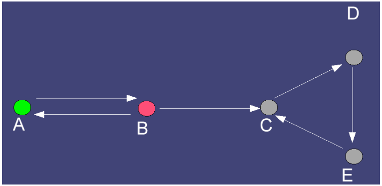
D
A
B
C
E

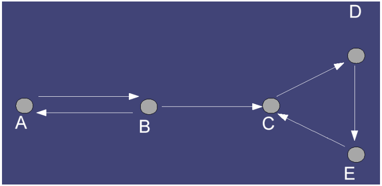
D
A
B
C
E

Labelling based semantics

Definition

Let $\mathcal{L}$ be a labelling of argumentation framework (Ar, def). We say that $\mathcal{L}$ is conflict-free iff for each $A, B \in Ar$, if $\mathcal{L}(A) = in$ and B defeats A, then $\mathcal{L}(B) \neq in$.

Definition

An admissible labelling is a labelling without arguments that are illegally *in* and without arguments that are illegally *out*.

Definition

A complete labelling is a labelling without arguments that are illegally *in*, without arguments that are illegally *out* and without arguments that are illegally *undec*.

Conditions on labelling aggregation

F_{AF} is a labellings aggregation operator that assigns a collective labelling $\mathcal{L}_{Coll}$ to each profile $\{\mathcal{L}_1, \ldots, \mathcal{L}_n\}$.

Conditions (UD, CR, anonymity and independence) for F_{AF}:

- **Universal domain**: The domain of F_{AF} is the set of all profiles of individual labellings belonging to semantics $\mathcal{T}_{conflict-free}$, $\mathcal{T}_{admissible}$ or $\mathcal{T}_{complete}$.
- **Collective rationality**: $F_{AF}(\{\mathcal{L}_1, \ldots, \mathcal{L}_n\})$ is a labelling belonging to semantics $\mathcal{T}_{conflict-free}$, $\mathcal{T}_{admissible}$ or $\mathcal{T}_{complete}$.

The sceptical aggregation (1)

First phase: the sceptical initial labelling ($\mathcal{L}_{sio}$):

- A is labelled *in* if everyone agrees A is *in*.
- A is labelled *out* if everyone agrees A is *out*.
- A is labelled *undec* in all other cases.

Example

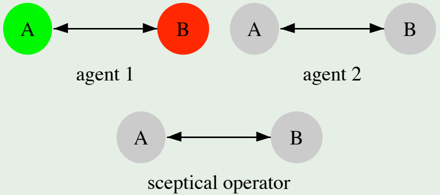

The sceptical aggregation (2)

Definition ($\mathcal{L}_1 \sqsubseteq \mathcal{L}_2$)

$\mathcal{L}_1$ is less or equally committed as $\mathcal{L}_2$ ($\mathcal{L}_1 \sqsubseteq \mathcal{L}_2$) iff $\texttt{in}(\mathcal{L}_1) \subseteq \texttt{in}(\mathcal{L}_2)$ and $\texttt{out}(\mathcal{L}_1) \subseteq \texttt{out}(\mathcal{L}_2)$.

Example

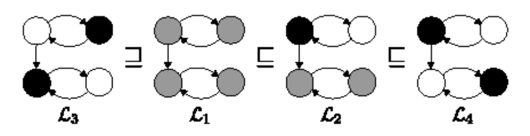

Lemma

$\mathcal{L}_{sio} \sqsubseteq \mathcal{L}_i$

The sceptical aggregation (3)

Problem: $\mathcal{L}_{sio}$ violates collective rationality under any constraint stronger than conflict-freeness.

Example

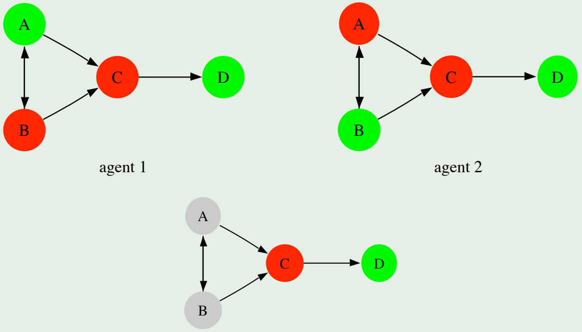

sceptical initial violates collective rationality under admissibility

The sceptical aggregation (4)

Second phase (iteration): at the end the sceptical labelling ($\mathcal{L}_{so}$):

- **Contraction function** relabels an argument from *in* or *out* to *undec* $\Rightarrow$ contraction sequence of labellings until $\mathcal{L}_{so}$.

Theorem

$\mathcal{L}_{so} \sqsubseteq \mathcal{L}_i$.

The group outcome is self-justifying.

$\mathcal{L}_{so}$ satisfies collective rationality under conflict-freeness, admissibility and completeness.

Unanimity

Problem (?): sometimes $\mathcal{L}_{so}$ ignores unanimity.

Example

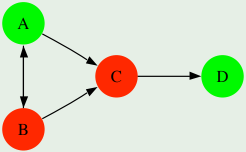

agent 1

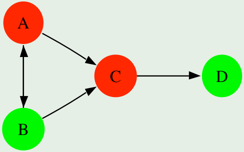

agent 2

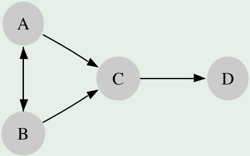

sceptical outcome violates unanimity

The credulous aggregation (1)

First phase: the credulous initial labelling ($\mathcal{L}_{cio}$):

- A is labelled *in* if someone thinks A is *in* and nobody thinks A is *out*.
- A is labelled *out* if someone thinks A is *out* and nobody thinks is *in*.
- A is labelled *undec* in all other cases.

Example

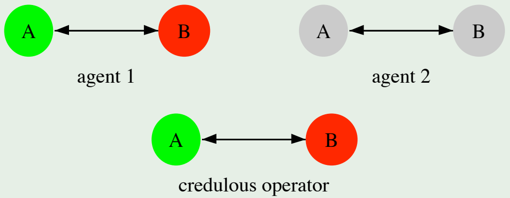

The credulous aggregation (2)

Definition ($\mathcal{L}_1 \approx \mathcal{L}_2$)

$\mathcal{L}_1$ is compatible with $\mathcal{L}_2$ ($\mathcal{L}_1 \approx \mathcal{L}_2$) iff $\texttt{in}(\mathcal{L}_1) \cap \texttt{out}(\mathcal{L}_2) = \emptyset$ and $\texttt{out}(\mathcal{L}_1) \cap \texttt{in}(\mathcal{L}_2) = \emptyset$.

Theorem

$\mathcal{L}_{cio}$ is **compatible** *with each input-labelling.*

$\sqsubseteq$ is stronger than $\approx$: if $\mathcal{L}_1 \sqsubseteq \mathcal{L}_2$, then $\mathcal{L}_1 \approx \mathcal{L}_2$.

Problem: $\mathcal{L}_{cio}$ violates collective rationality even under conflict-freeness (let alone under admissibility and completeness)!

The credulous aggregation (3)

Second phase (iteration): at the end the credulous labelling ($\mathcal{L}_{co}$):

> *Set to abstention all arguments for which the group has no justification to accept or reject.*

$\mathcal{L}_{co} \sqsubseteq \mathcal{L}_{cio}$.

Theorem

$\mathcal{L}_{co}$ *is* **compatible** *with each input-labelling* $\mathcal{L}_i$.

The credulous aggregation (4)

$\mathcal{L}_{co}$ satisfies collective rationality under conflict-freeness and admissibility (but **not** under completeness).

Example

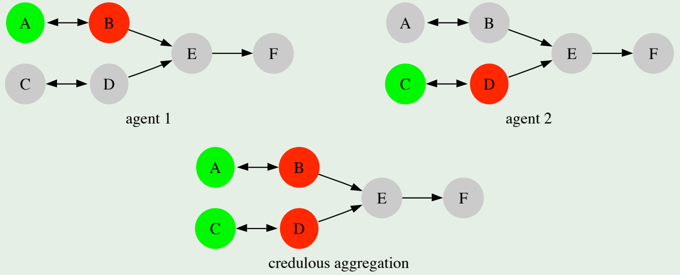

$\mathcal{L}_{co}$ can ignore unanimity.

Sceptical and Credulous Operator

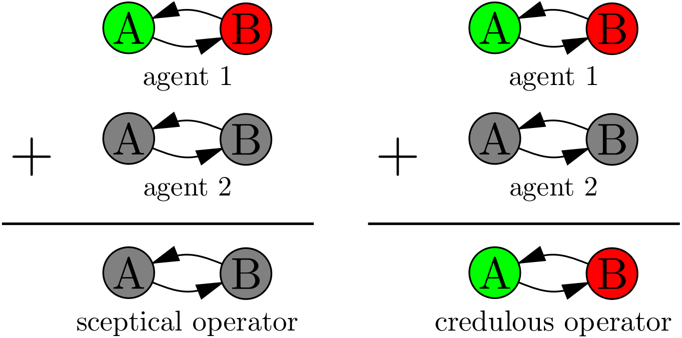

The super credulous aggregation (1)

The super credulous aggregation takes the credulous outcome and expands and make it bigger it, by relabelling illegal *undec*s to *in*s and *out*s.

Example

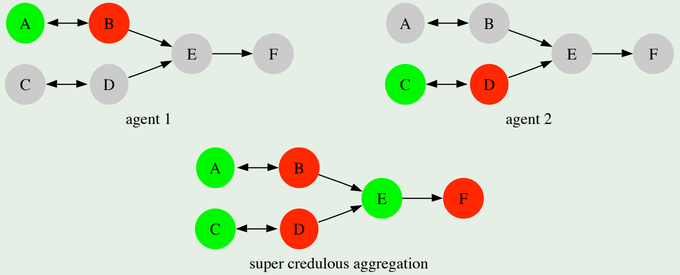

The super credulous aggregation (2)

For any admissible labelling $\mathcal{L}_a$, there exists *at least one* super credulous expansion sequence. The resulting labelling $\mathcal{L}_m$ is a **complete** labelling with $\mathcal{L}_a \sqsubseteq \mathcal{L}_m$.

Theorem

The set of complete labellings that are bigger or equal to $\mathcal{L}_a$ has a unique smallest element, and this is $\mathcal{L}_{sco}$.

Theorem

$\mathcal{L}_{sco} \approx \mathcal{L}_i$

Theorem

$$\mathcal{L}_{so} \sqsubseteq \mathcal{L}_{co} \sqsubseteq \mathcal{L}_{sco}$$

5 definitions of legitimacy (1) [Work with Umberto Grandi]

Definition (Minimal legitimacy)

An agg. procedure satisfies minimal legitimacy if the outcome on every profile is supported by at least one individual.

- Basic intuition of representativeness $\Rightarrow$ No instances of MEP.

Definition (k-legitimacy)

An aggregation procedure satisfies k-legitimacy ($k \leqslant |\mathcal{N}|$) if for every profile the outcome is supported by at least k individuals.

- Avoids instances of MEP where selected outcome has been submitted by the fewest individuals.

Five definitions of legitimacy (2)

We now turn to **incomplete** ballots.

Two ballots B and B' are compatible if they do not disagree (acceptance/rejection) on any issue. (*cfr.* Caminada and Pigozzi)

Definition (Compatibility)

An incomplete agg. procedure satisfies compatibility if for every profile the outcome is compatible with all the individual ballots.

For larger groups we can generalised:

Definition (k-compatibility)

An incomplete agg. procedure satisfies k-compatibility if the outcome is compatible with at least k individual ballots.

Five definitions of legitimacy (3)

The last definition of legitimacy focuses on *single issues* rather than on whole ballots.

Definition (k-legitimacy over issues)

An aggregation procedure (complete or incomplete) satisfies k-legitimacy over issues ($k \leqslant n$) if, for every profile and every issue, the outcome is compatible with at least k individuals.

These definitions don't exhaust the space of possible intuitions of legitimacy. But they are the most natural options.

Four aggregation procedures (1)

Four aggregation procedures that guarantee a legitimate outcome when the issues are independent.

Definition (Average voter rule)

The average voter rule (AVR) chooses the individual ballot that minimises the sum of the Hamming distance to all other individual ballots.

Different from proposition-wise majority voting.

Four aggregation procedures (2)

Example

	p_1	p_2	p_3	p_4	p_5
B_1	1	1	0	1	1
B_2	0	1	1	0	1
B_3	1	0	1	1	0
Maj	1	1	1	1	1

B_1 is selected as AVR.

Proposition

AVR *satisfies minimal legitimacy.*

Proposition

AVR *satisfies U*, A* and M*.* AVR *does not satisfy I*.*

Four aggregation procedures (3)

Three aggregation procedures for *incomplete* ballots.

Definition (Conflict-free rule (CFR))

For every $j \in \mathcal{I}$, let b_j^c be the jth element of the collective outcome $\mathrm{CFR}(\mathbf{B})$:

$$b_j^c = \begin{cases} 1 & \exists i \in \mathcal{N} \text{ s.t. } b_{i,j} = 1 \text{ and } \nexists l \in \mathcal{N} \text{ s.t. } b_{l,j} = 0 \\ 0 & \exists i \in \mathcal{N} \text{ s.t. } b_{i,j} = 0 \text{ and } \nexists l \in \mathcal{N} \text{ s.t. } b_{l,j} = 1 \\ A & \text{otherwise} \end{cases}$$

- Similar to the credulous procedure of Caminada and Pigozzi.
- CFR is a resolute procedure.

Four aggregation procedures (4)

Example

	p	q	r
B_1	1	0	A
B_2	A	A	1
B_3	A	0	1
CFR $(\mathbf{B})$	1	0	1

	p	q	r
B'_1	1	1	A
B'_2	1	0	0
B'_3	A	0	1
CFR $(\mathbf{B}')$	1	A	A

Proposition

CFR *satisfies compatibility.*

Proposition

CFR *satisfies U, I, A and M.*

Four aggregation procedures (5)

The third rule is the *k*-conflict-free rule (*k*-CFR).

Definition (Subprofile of B)

Given a profile $\mathbf{B} = (B_1, \ldots, B_n)$ and a subset of agents $K \subseteq N$, the restriction of $\mathbf{B}$ to K is $\mathbf{B}_K = (B_k, k \in K)$ and is called a subprofile of $\mathbf{B}$.

Definition (*k*-conflict-free rule (*k*-CR))

The *k*-conflict-free rule maps $\mathbf{B}$ to $k\text{-CFR}(\mathbf{B})=\{CR(\mathbf{B}_K) \mid K \subseteq \mathcal{N}\}$, where $CR(\mathbf{B}_K)$ is the CFR over the subprofile $\mathbf{B}_K$.

Four aggregation procedures (6)

Example

Unlike CFR, k-CFR is **not** a resolute procedure. For k=2:

	p	q	r
B_1	1	0	A
B_2	A	A	1
B_3	A	0	1
$F(\mathbf{B}_{2,3})$	A	0	1
$F(\mathbf{B}_{1,2})$	1	0	1
$F(\mathbf{B}_{1,3})$	1	0	1

	p	q	r
B'_1	1	1	A
B'_2	1	0	0
B'_3	A	0	1
$F(\mathbf{B}'_{2,3})$	1	0	A
$F(\mathbf{B}'_{1,2})$	1	A	0
$F(\mathbf{B}'_{1,3})$	1	A	1

Proposition

k-CFR satisfies k-compatibility.

k-CFR satisfies U^*, I^*, A^* and M^*.

Four aggregation procedures (7)

The last rule:

Definition (k-quota rule (k-QR))

Let $1 \leqslant k \leqslant n$ and b_j^c be the jth element of the collective outcome $k-\mathrm{QR}(\mathbf{B})$:

$$b_j^c = \begin{cases} 1 & \text{iff } \exists M \subseteq \mathcal{N},\ |M| \geqslant k \text{ s.t. } \forall i \in M, b_{i,j} = 1 \\ 0 & \text{iff } \exists M' \subseteq \mathcal{N},\ |M'| \geqslant k \text{ s.t. } \forall i \in M', b_{i,j} = 0 \\ A & \text{iff } \exists M'' \subseteq \mathcal{N},\ |M''| \geqslant k \text{ s.t. } \forall i \in M'', b_{i,j} = A \\ A & \text{otherwise} \end{cases}$$

k-QR guarantees a unique result only when $k \geqslant \frac{|\mathcal{N}|}{2}$.

Four aggregation procedures (8)

Example (k-quota rule (k-QR))

	p	q	r
B_1	1	0	A
B_2	A	A	1
B_3	A	0	1
k-QR($\mathbf{B}$)	A	0	1

	p	q	r
B'_1	1	1	A
B'_2	1	0	0
B'_3	A	0	1
k-QR($\mathbf{B}'$)	1	0	A

Proposition

k-QR satisfies k-legitimacy over issues.

Proposition

k-QR satisfies U^*, I^*, M^*, A^* unless $k = |\mathcal{N}|$ or $k = 0$.

A procedure for logically connected issues (1)

- A new consistent and legitimate rule for settings in which individuals do not share the same consistency rule.
- The agenda $\mathcal{I}$ is divided into two sets: $\mathcal{I}_p$ and $\mathcal{I}_c$ of premises and conclusions respectively.

Example

Three automatic trading agents have to decide on whether to buy certain stocks. The first agent thinks they should buy the stock (B) because the revenue is increasing (R) and people are selling a considerable amount of stocks (S): so $B \leftrightarrow (R \wedge S)$. The second agent submits the rule $R \leftrightarrow \neg B$, and the third agent's *conclusion function* is $S \leftrightarrow \neg B$. The profile is as in the doctrinal paradox.

A procedure for logically connected issues (2)

- A procedure that is legitimate and consistent in aggregating the premises and that also outputs a group conclusion function which can be employed to draw collective conclusions.
- To aggregate the premises, we take AVR.

Proposition

The AVR rule is collectively rational on every agenda.

The average voter rule is therefore minimally legitimate and consistent on every agenda $\mathcal{I}$.

A procedure for logically connected issues (3)

To merge the individual conclusion functions, we resort to formula-based belief merging [Konieczny and Perez].

Definition

Let M and K be consistent subsets of $\mathcal{L}_{PS}$:

$$d_D(M,K) = \begin{cases} 0 & M \cup K \text{ is consistent} \\ 1 & \text{otherwise} \end{cases}$$

Given a set of formulas M, let $\mathrm{MAXCONS}\,(M)$ be the set of maximal (w.r.t. inclusion) consistent subsets of M.

Definition (Drastic majority operator)

Let $K_1, \ldots, K_n$ be subsets of propositional formulas, the drastic majority operator is defined as follows:

$$\Delta_D(K_1, \ldots, K_n) = \underset{M \in \mathrm{MAXCONS}(\cup_i K_i)}{\arg\min} \sum_{i=1}^{n} d_D(M, K_i)$$

A procedure for logically connected issues (4)

Definition (Reason-based rule)

$(\varphi_{1,j}, \ldots, \varphi_{n,j})$ = conclusion functions

- $B_p^c = AVR(B_1^p, \ldots, B_n^p)$;
- $\varphi_j^c = \Delta_D(\varphi_{1,j}, \ldots \varphi_{n,j})$;
- $B_c^c = \{B \text{ over } \mathcal{I}_c \text{ s.t. } (B_p^c, B) \models \varphi_j^c \text{ for all } j \in \mathcal{I}_c\}$

Example (Continued)

Two maximal consistent sets of formulas:
$K_1 = \{S \leftrightarrow \neg B, R \leftrightarrow \neg B\}$ and $K_2 = \{(S \wedge R) \leftrightarrow B\}$. The first is consistent with the last two individuals, so it is chosen as collective conclusion function. Using the AVR over premises we get that S is accepted together with R, and therefore the result on B is 0.

Conclusion

- Aggregation paradoxes may not arise only in presence of an inconsistency.
- We urge aggregation procedures able to ensure legitimate **and** consistent social decisions.
- We explored the notion of legitimacy, capturing the intuition of compatibility of the collective outcome with the individual inputs.
- We defined aggregation operators that guarantee compatible outcomes $\Rightarrow$ 'consensus' aggregation operators.
- We considered both logically un/connected agendas.
- *Future work*:
  - Enlarge the space of legitimacy definitions and aggregation procedures
  - Further properties need to be introduced to assess the new procedures.