Computational Social Choice: Spring 2015

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

Plan for Today

We will have a closer look at the concept of *collective rationality:* the preservation of rationality requirements during aggregation.

This will provide yet another opportunity for investigating how *axioms* interact with *structural properties* of the domain of aggregation.

We will work with *binary aggregation with integrity constraints* and start by introducing this framework in some more detail than we had done in the first lecture on the topic.

Preference Aggregation

Expert 1: $\triangle \succ \bigcirc \succ \square$

Expert 2: $\bigcirc \succ \square \succ \triangle$

Expert 3: $\square \succ \triangle \succ \bigcirc$

Expert 4: $\square \succ \triangle \succ \bigcirc$

Expert 5: $\bigcirc \succ \square \succ \triangle$

?

Judgment Aggregation

	p	$p \to q$	q
Judge 1:	True	True	True
Judge 2:	True	False	False
Judge 3:	False	True	False
		?	

Multiple Referenda

	fund museum?	*fund school?*	*fund metro?*
Voter 1:	Yes	Yes	No
Voter 2:	Yes	No	Yes
Voter 3:	No	Yes	Yes

?

[Constraint: we have money for *at most two projects*]

General Perspective

The last example is actually pretty general. We can rephrase many aggregation problems as problems of *binary aggregation:*

Do you rank option △ above option ○? Yes/No

Do you believe formula "$p \rightarrow q$" is true? Yes/No

Do you want the new school to get funded? Yes/No

Each problem domain comes with its own *rationality constraints:*

Rankings should be transitive and not have any cycles.

The accepted set of formulas should be logically consistent.

We should fund at most two projects.

The *paradoxes* we have seen show that the *majority rule* does not *lift* our rationality constraints from the *individual* to the *collective* level.

Binary Aggregation with Integrity Constraints

Let $\mathcal{I} = \{1, \ldots, m\}$ be a finite set of *issues* on which to take a decision and let $\mathcal{D} := \{0, 1\}^m$ be the binary combinatorial domain defined by $\mathcal{I}$.

A *ballot* is a vector $B = (b_1, \ldots, b_m) \in \mathcal{D}$, indicating for each issue $j \in \mathcal{I}$ whether it is *accepted* $(b_j = 1)$ or *rejected* $(b_j = 0)$.

Associate a set $PS = \{p_1, \ldots, p_m\}$ of *propositional symbols* with $\mathcal{I}$ and let $\mathcal{L}_{PS}$ be the *language* of propositional logic over PS.
Note that *models* for formulas in $\mathcal{L}_{PS}$ are isomorphic to ballots.

An *integrity constraint* (IC) is a formula $\Gamma \in \mathcal{L}_{PS}$. For a given Γ, we say that ballot B is *rational* if $B \in \mathrm{Mod}(\Gamma)$ (that is, if $B \models \Gamma$).

Now a finite set $\mathcal{N} = \{1, \ldots, n\}$ of *individuals*, with $n \geqslant 2$, each report a rational ballot, producing a *profile* $\boldsymbol{B} = (B_1, \ldots, B_n)$.

A (resolute) *aggregation rule* for n individuals, issues $\mathcal{I}$, and IC Γ is a function $F : \mathrm{Mod}(\Gamma)^n \to \mathcal{D}$, mapping profiles to single ballots/models.

Example

Our multiple-referenda example is formalised as follows:

- Three individuals: $\mathcal{N} = \{1, 2, 3\}$
- Three issues/prop. symbols: $\mathcal{I} = \{\texttt{museum}, \texttt{school}, \texttt{metro}\}$.
- Integrity constraint: $\Gamma = \neg(\texttt{museum} \wedge \texttt{school} \wedge \texttt{metro})$
- Profile: $\boldsymbol{B} = (B_1, B_2, B_3)$ with

$$\begin{array}{rcl} B_1 & = & (1,1,0) \\ B_2 & = & (1,0,1) \\ B_3 & = & (0,1,1) \end{array}$$

  Note that $B_i \models \Gamma$ for all $i \in \{1, 2, 3\}$

- However, $F_{\textsf{maj}}(\boldsymbol{B}) = (1,1,1)$ and $(1,1,1) \not\models \Gamma$.

Paradoxes

We are now able to give a general definition of "paradox" that captures many of the paradoxes in the literature on social choice theory.

> A *paradox* is a triple $\langle F, \Gamma, \boldsymbol{B} \rangle$, consisting of an aggregation rule F, a profile $\boldsymbol{B}$, and an integrity constraint Γ, such that $B_i \models \Gamma$ for all individuals $i \in \mathcal{N}$ but $F(\boldsymbol{B}) \not\models \Gamma$.

Collective Rationality

An aggregation rule F is *collectively rational* for integrity constraint $\Gamma \in \mathcal{L}_{PS}$ if $B_i \models \Gamma$ for all $i \in \mathcal{N}$ implies $F(\boldsymbol{B}) \models \Gamma$.

That is, F is collectively rational for Γ, if there exists not profile $\boldsymbol{B}$ such that $\langle F, \Gamma, \boldsymbol{B} \rangle$ is a paradox.

We also say: F can *lift* Γ from the individual to the collective level.

Remark: As we have defined F only on rational profiles in $\mathrm{Mod}(\Gamma)^n$, technically the condition $B_i \models \Gamma$ always holds. But for many rules (e.g., majority) it is natural to think of them as being defined also on irrational profiles, and the "lifting" metaphor also makes sense then.

Axioms for Binary Aggregation

Some (mostly) familiar axioms, adapted to this framework:

- *Unanimity:* For any profile of rational ballots $\boldsymbol{B} = (B_1, \ldots, B_n)$ and any $v \in \{0, 1\}$, if $b_{ij} = v$ for all $i \in \mathcal{N}$, then $F(\boldsymbol{B})_j = v$.
- *Anonymity:* For any rational profile $\boldsymbol{B} = (B_1, \ldots, B_n)$ and any permutation $\pi : \mathcal{N} \to \mathcal{N}$, we get $F(\boldsymbol{B}) = F(B_{\pi(1)}, \ldots, B_{\pi(n)})$.
- *Independence:* For any issue $j \in \mathcal{I}$ and any two rational profiles $\boldsymbol{B}, \boldsymbol{B'}$, if $b_{ij} = b'_{ij}$ for all $i \in \mathcal{N}$, then $F(\boldsymbol{B})_j = F(\boldsymbol{B'})_j$.
- *Issue-Neutrality:* For any two issues $j, j' \in \mathcal{I}$ and any rational profile $\boldsymbol{B}$, if $b_{ij} = b_{ij'}$ for all $i \in \mathcal{N}$, then $F(\boldsymbol{B})_j = F(\boldsymbol{B})_{j'}$.
- *Domain-Neutrality:* For any two issues $j, j' \in \mathcal{I}$ and any rational profile $\boldsymbol{B}$, if $b_{ij} = 1 - b_{ij'}$ for all $i \in \mathcal{N}$, then $F(\boldsymbol{B})_j = 1 - F(\boldsymbol{B})_{j'}$.

Note that we had not considered domain-neutrality before.

Template for Results

Let $\mathcal{L} \subseteq \mathcal{L}_{PS}$ be a *language of integrity constraints*. By fixing $\mathcal{L}$ we fix a range of possible domains of aggregation (one for each $\Gamma \in \mathcal{L}$).

Two ways of defining classes of aggregation rules:

- The class of rules defined by a given list of *axioms* AX:

$$\mathcal{F}_{\mathcal{L}}[\text{AX}] \quad := \quad \{F : \text{Mod}(\Gamma)^n \to \mathcal{D} \mid \Gamma \in \mathcal{L} \text{ and } F \text{ satisfies AX}\}$$

- The class of rules that *lift* all integrity constraints in $\mathcal{L}$:

$$\mathcal{CR}[\mathcal{L}] \quad := \quad \{F : \text{Mod}(\Gamma)^n \to \mathcal{D} \mid F \text{ is collect. rat. for all } \Gamma \in \mathcal{L}\}$$

What we want:

$$\mathcal{CR}[\mathcal{L}] \;=\; \mathcal{F}_{\mathcal{L}}[\text{AX}]$$

U. Grandi and U. Endriss. Lifting Integrity Constraints in Binary Aggregation. *Artificial Intelligence*, 199–200:45–66, 2013.

Lifting Conjunctions of Literals

Proposition 1 *F will lift all integrity constraints that can be expressed as a conjunction of literals ("cube") iff F is unanimous:*

$$\mathcal{CR}[\text{cubes}] \quad = \quad \mathcal{F}_{\text{cubes}}[\text{Unanimity}]$$

Proof: Immediate from the definitions. ✓

Discussion: While technically almost trivial, conceptually this is a very nice link between two completely separate worlds: syntactic structure of formulas and a very fundamental economic principle.

More Results

Characterisation results:

- F lifts all constraints $p_j \leftrightarrow p_k$ *iff* F is *issue-neutral*
- F lifts all constraints $p_j \leftrightarrow \neg p_k$ *iff* F is *domain-neutral*

Negative results:

- there exists *no language* that characterises *anonymous* rules
- there exists *no language* that characterises *independent* rules

Discussion

It is an open research question of how the results on *lifting IC's* discussed today and the *agenda characterisation* results discussed earlier (for formula-based JA) relate to each other precisely.

What they have in common:

- Both kinds of results relate *axioms* (on the mechanics of the rule) with *structural properties* of the domain of aggregation.

Some observations regarding differences:

- Lifting results tend to be about structural properties that can be expressed in terms of *syntactic* features, while agenda properties used for characterisation results are clearly *semantic*.

- Known results on agenda characterisation tend to apply to natural *combinations of axioms*. Known lifting results instead tend to focus on a *single axioms* at the time.

Lifting Arbitrary Formulas

Are there rules that will lift *every* integrity constraint? *Yes!*

Define a *generalised dictatorship* as any rule F for which there exists a function $g : \mathcal{D}^n \to \mathcal{N}$ such that $F(\boldsymbol{B}) = B_{g(\boldsymbol{B})}$ for all profiles $\boldsymbol{B}$.
Thus: g picks a local dictator for any given profile.

Example: For $g \equiv i^\star$, we get a proper (Arrovian) dictatorship.

Proposition 2 *F will lift all IC's iff F is a generalised dictatorship:*

$$\mathcal{CR}[\mathcal{L}_{PS}] \quad = \quad \mathrm{GDIC}$$

Proof: Immediate. ✓

Discussion: Is this a positive or a negative result?

Representative-Voter Rules

The class of generalised dictatorships is large and includes many obviously terrible rules (e.g., proper dictatorships). But some look ok:

- *Average-voter rule:* from the ballots supplied, pick the one minimising the sum of the Hamming distances to all others.

  Thus: max-sum (Kemeny) applied only to the support of the profile

- *Majority-voter rule:* from the ballots supplied, pick the one minimising the Hamming distance to the majority outcome.

  Thus: max-number (Slater) applied only to the support of the profile

Call a generalised dictatorship with an intuitively appealing choice of the agent-picking function g a *most-representative voter rule*.

So the "lifting perspective" motivated a new family of rules.

U. Endriss and U. Grandi. Binary Aggregation by Selection of the Most Representative Voter. *Proc. AAAI-2014*.

Example

The average-voter rule (AVR) and the majority-voter rule (MVR) really can sometimes give different outcomes:

Issue:	**1**	**2**	**3**	**4**	**5**	**6**
1 voter:	1	0	0	0	0	0
10 voters:	0	1	1	0	0	0
10 voters:	0	0	0	1	1	1
Majority:	0	0	0	0	0	0
MVR:	1	0	0	0	0	0
AVR:	0	1	1	0	0	0

Summary

We have focused on *collective rationality* in binary aggregation by investigating what *integrity constraints* are *lifted* from the individual to the collective level by what kinds of rules (characterised by axioms):

- Results linking *axioms* to *syntactic properties* of constraints (impossible for certain axioms)
- Design of aggregation rules ("*representative-voter rules*") inspired by requirement to lift all constraints

What next?

We will discuss *strategic behaviour* in judgment aggregation:

- How can we define the incentives of agents to strategise in JA?
- How can we analyse strategic behaviour with the axiomatic method?
- How complex is it to strategise?