Judgment Aggregation: June 2018 Project

Sirin Botan

Institute for Logic, Language and Computation
University of Amsterdam

June 4, 2018

Project Info

Project Website: https://staff.science.uva.nl/s.botan/ja2018.html

- This week: 11-13, Mon-Thu in this room
- "Homeworks": One exercise per lecture
- Presentations (schedule coming):
  - Pick a paper from the list
  - Read it
  - Present it to the rest of us (1 hour w/ questions)
- Papers: Too early to think about it!

At the end of today, you can "reserve" a paper for your presentation.

Plan for Today

- We will look at a few other areas in Computational Social Choice so you get a feeling for how JA fits into the bigger picture.
- We'll see the formal framework of judgment aggregation, and prove a simple impossibility result.

As always, consult the Handbook for more details on everything!

Stable Matching

n men and n women, each with preferences over the other gender (full ranking). We want to find a stable matching: no man and woman should both want to leave their current partners and run off together.

Gale-Shapley Algorithm guarantees a stable matching.

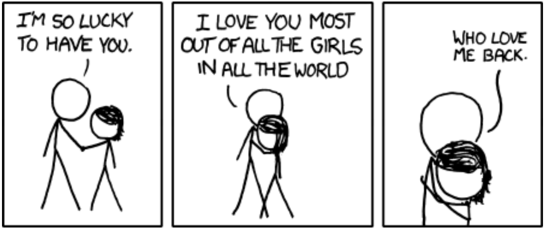

D. Gale and L.S. Shapley. College Admissions and the Stability of Marriage, American Mathematical Monthly, 69, 9-15, 1962.

Stable Matching

n leads and n follows, each with preferences over the other dancers. We want to find a stable matching: no lead and follow should both want to leave their current partners and dance together.

Gale-Shapley Algorithm guarantees a stable matching.

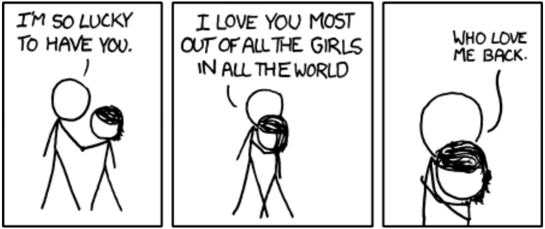

D. Gale and L.S. Shapley. College Admissions and the Stability of Marriage, American Mathematical Monthly, 69, 9-15, 1962.

Matching Applications

- Matching resident doctors and hospitals
- Matching schools and students (Amsterdam)
- Kidney Exchanges: matching kidney donors to patients!

Suppose you have a friend who needs a new kidney and you are willing to donate. What if your kidney is not a match to hers?

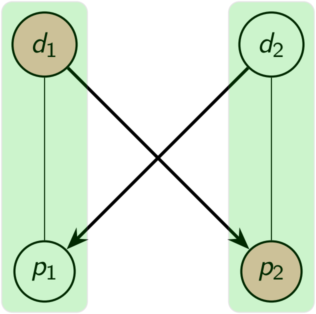

Fair Allocation: Cake Cutting

Cake cutting is about fair allocation of divisible goods. For indivisible goods, see Ch. 12 of the Handbook.

A divisible item

Three indivisible items

How would you fairly divide a cake between two people so that each person feels they got $\geqslant \frac{1}{2}$ when the agents have different preferences about which part of the cake they prefer?

Banach-Knaster Last-Diminisher Procedure

We generalise the "cut-and-choose" procedure to n agents.

- Agent 1 cuts off a piece she considers worth $1/n$ of the cake.
- The piece is passed around. Each agent passes it along (if they consider it $< 1/n$), or trims it down (to what they consider $1/n$).
- When the piece has made the full round, the last person to cut it—the last diminisher—has to take it.
- The procedure is repeated for the remaining agents.

Guarantees proportionality: $1/n$ to each agent according to their own valuation.

How many cuts do we need in the worst case?

H. Steinhaus. The Problem of Fair Division. *Econometrica*, 16, 101-104, 1948

Preferences for Cake Cutting

Note that in the cake cutting scenario, each agent's preferences are modelled as a cardinal utility function.

This is in contrast to representing agent's preferences as an ordinal preference relation, as we saw in the stable matching problem. We'll also see preference relations in voting and judgment aggregation.

Note: On spliddit.org you can play around with different types of fair allocation (for both divisible and indivisible goods).

Voting

n agents each submit their preferences (linear orders) over a set of m alternatives. A voting rule takes as input these preferences and outputs the winning candidate.

Agent 1:	$a \succ b \succ c$
Agent 2:	$b \succ c \succ a$
Agent 3:	$c \succ a \succ b$

Condorcet Principle

If there is an alternative that is preferred to every other alternative by a majority of voters, then it should win.

Such an alternative is unique if it exists, and is called a Condorcet winner. Voting rules which elect the Condorcet winner are called Condorcet extentions.

Agent 1:	$a \succ b \succ c$
Agent 2:	$b \succ c \succ a$
Agent 3:	$c \succ a \succ b$
Majority:	$a \succ b \succ c \succ a$

The collection of preferences above demonstrate what is called the Condorcet Paradox. No Condorcet winner exists.

Borda

Each agent gives $m - 1$ points to their top ranked alternative, $m - 2$ to the alternative she ranks second, and so on. The alternative with the most points (subject to some tie-breaking) is the winner.

3 agents:	$a \succ b \succ c$
2 agents:	$b \succ c \succ a$

- Who is the Borda winner?
- Who is the Condorcet winner?

Positional Scoring Rules

A positional scoring rule (PSR) is given by a scoring vector $\langle s_1, \ldots, s_m \rangle$.

- $s_1 \geqslant s_2 \geqslant \ldots \geqslant s_m$
- $s_1 > s_m$

Borda is the PSR with vector $\langle m-1, m-2, \ldots, 0 \rangle$.

Other examples:

- $\langle 1, 0, \ldots, 0 \rangle$ (Plurality)
- $\langle 1, \ldots, 1, 0 \rangle$ (Antiplurality)
- $\langle 12, 10, 8, 7, 6, 5, 4, 3, 2, 1, 0, \ldots, 0 \rangle$ (Eurovision)

PSRs and Condorcet's Principle

No positional scoring rule satisfies Condorcet's principle.

3 agents:	$a \succ b \succ c$
2 agents:	$b \succ c \succ a$
1 agent:	$b \succ a \succ c$
1 agent:	$c \succ a \succ b$

The Condorcet winner is a. But any PSR (with the right tie-breaking) will give b as the winner.

- (A) $3s_1 + 2s_2 + 2s_3$
- (B) $3s_1 + 3s_2 + s_3$
- (C) $s_1 + 2s_2 + 4s_3$

Condorcet Extensions vs. Borda

Borda does not satisfy Independence of Irrelevant Alternatives (IIA), which states that whether a is preferred to b in the outcome should not depend on c.

But, takes into account all—some might say irrelevant—info. Proponents of Borda argue that the rule avoids "tyranny of the majority" (49 agents with ignored preferences).

51 agents:	$a \succ b \succ c$
49 agents:	$b \succ c \succ a$

Ongoing debate...

And Finally: Judgment Aggregation!

Judgment Aggregation deals with aggregating Yes/No opinions provided by agents into a collective decision that reflects the views of the group.

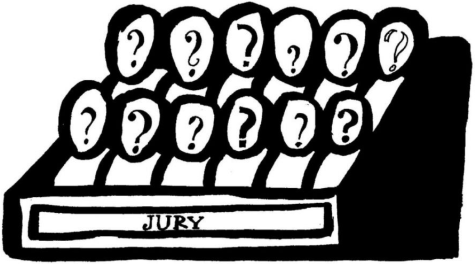

Doctrinal Paradox

$p :=$ 'document is a binding contract'
$q :=$ 'the promise in the document was breached'
$r :=$ 'the defendant is liable'

According to legal doctrine: $p \wedge q \leftrightarrow r$.

	p	q	r
Judge 1	✓	✓	✓
Judge 2	×	✓	×
Judge 3	✓	×	×

How do we aggregate? Is the defendant liable?

Kornhauser, L.A. and Sager, L.G. The One and the Many: Adjudication in Collegial Courts. *California Law Review*, 81(1), 1-59, 1993

Discursive Dilemma

...a related example, without the "constraint" on opinions.

	p	q	$p \rightarrow q$
Judge 1	✓	✓	✓
Judge 2	×	×	✓
Judge 3	✓	×	×
Majority	✓	×	✓

Majority opinion is logically inconsistent. By the end of this lecture we will have a general result about when this can happen.

C. List and P. Pettit. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy*, 18(1), 89-110, 2002.

Formal Framework

Note: Let $\sim\varphi := \varphi'$ if $\varphi = \neg\varphi'$ and $\sim\varphi := \neg\varphi$ otherwise.

- An agenda Φ is a finite, nonempty set of propositional formulas closed under complementation ($\varphi \in \Phi \Rightarrow \sim\varphi \in \Phi$).
- A judgment set J is a subset of Φ. J is:
  - complete if $\varphi \in J$ or $\sim\varphi \in J$ for all $\varphi \in \Phi$
  - complement-free if $\varphi \notin J$ or $\sim\varphi \notin J$ for all $\varphi \in \Phi$
  - consistent if there is an assignment making all $\varphi \in J$ true

  $\mathcal{J}(\Phi)$ is the set of all complete and consistent subsets of Φ.

A set of agents $\mathcal{N} = \{1, \ldots, n\}$ report their judgment sets, giving us a profile $\boldsymbol{J} = (J_1, \ldots J_n)$.

A (resolute) aggregation rule F for an agenda Φ and a set of agents $\mathcal{N}$ is a function mapping a profile of complete and consistent judgment sets to a single (collective) judgment set:

$$F : \mathcal{J}(\Phi)^n \to 2^\Phi$$

.

Back to our example

	p	q	$p \to q$
Judge 1	✓	✓	✓
Judge 2	×	✓	×
Judge 3	✓	×	×
Majority	✓	✓	×

- What is Φ?
- What is J_2?
- What is the profile?
- Is the collective judgment set...
  - ...complete?
  - ...complement-free?
  - ...consistent?

Collective Rationality Requirements

An aggregation rule can “lift” collective rationality requirements.

F is:

- complete if $F(\boldsymbol{J})$ is complete for all profiles $\boldsymbol{J}$
- complement-free if $F(\boldsymbol{J})$ is complement-free for all profiles $\boldsymbol{J}$
- consistent if $F(\boldsymbol{J})$ is consistent for all profiles $\boldsymbol{J}$

Note: We already saw the majority rule is not always consistent.

Axioms in Social Choice

Axioms in social choice identify desirable properties of aggregation methods. Note that desirability of a property is subjective.

We can see the collective rationality requirements as a type of axiom.

The axiomatic method allows us to characterise rules by which properties they satisfy or prove impossibility results.

We will see some of these characterisation results tomorrow. Today we will see a central impossibility result in JA.

Three Basic Axioms

- Anonymity: Treating all agents symmetrically.
  - for any profile $\boldsymbol{J}$ and any permutation $\pi : \mathcal{N} \to \mathcal{N}$, we have that $F(J_1, \ldots J_n) = F(J_{\pi(1)}, \ldots J_{\pi(n)})$.
- Neutrality: Treating all formulas the same.
  - for any $\varphi, \psi \in \Phi$ and any profile $\boldsymbol{J}$, if $\varphi \in J_i \Leftrightarrow \psi \in J_i$ for all $i \in \mathcal{N}$, then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \psi \in F(\boldsymbol{J})$.
- Independence: Outcome on φ depends only on agents' judgment on φ.
  - for any $\varphi \in \Phi$ and any two profiles $\boldsymbol{J}$ and $\boldsymbol{J'}$, if $\varphi \in J_i \Leftrightarrow \varphi \in J'_i$ for all $i \in \mathcal{N}$, then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \varphi \in F(\boldsymbol{J'})$.

Which of these axioms does the majority rule satisfy?

An Impossibility Result

Theorem 1 (List and Pettit, 2002) For $n \geqslant 2$, No judgment aggregation rule for an agenda Φ with $\{p, q, p \wedge q\} \subseteq \Phi$ satisfies anonymity, neutrality, independence, completeness and consistency.

Note: Here p and q are mutually independent propositions and $(p \wedge q)$ can be replaced by $(p \vee q)$ or $(p \rightarrow q)$.

C. List and P. Pettit. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy*, 18(1), 89-110, 2002.

Proof. . .

Note: $N_\varphi^{\boldsymbol{J}}$ is the set of agents who accept φ in profile $\boldsymbol{J}$

Let F be some anonymous, neutral and independent aggregation rule.

- F is independent: whether $\varphi \in F(\boldsymbol{J})$ depends only on $N_\varphi^{\boldsymbol{J}}$.
- F is anonymous: we only need to look at $|N_\varphi^{\boldsymbol{J}}|$.
- F is neutral: the way in which the status of $\varphi \in F(\boldsymbol{J})$ depends on $|N_\varphi^{\boldsymbol{J}}|$, cannot depend on φ.

Then, if φ and ψ are accepted by the same number of individuals, F must either accept both or reject both.

. . . *Proof.*

Let $\{p, q, p \wedge q\} \subseteq \Phi$.

For odd n, consider a profile $\boldsymbol{J}$ where $\frac{n-1}{2}$ accept both p and q, 1 agents accepts p but not q, one agent accepts q but not p, and the remaining $\frac{n-3}{2}$ agents accept neither p nor q.

Then $|N_p^{\boldsymbol{J}}| = |N_q^{\boldsymbol{J}}| = |N_{\neg(p \wedge q)}^{\boldsymbol{J}}|$, and by the previous slide, we have to accept either all or none of them.

- Accept all: Not consistent. 👎
- Accept none: Not complete. 👎

For even n, take any profile $\boldsymbol{J}$ where $|N_p^{\boldsymbol{J}}| = |N_{\neg p}^{\boldsymbol{J}}|$.

- Accept both: Not consistent. 👎
- Accept neither: Not complete. 👎

Do we just give up on JA?

Some of the axioms presented might not be that desirable in all cases. . .

- Does it make sense to require Independence when there are logical dependencies in the agenda?
- Should we use an Anonymous rule if some agents are experts on a subject and some are not?
- What about Neutrality; for a murder trial we might want a jury to unanimously find the defendant guilty to convict her, while a simple majority is enough to convict someone for a traffic violation.

Summary

- We saw several areas of COMSOC:
  - Matching
  - Fair Allocation
  - Voting
  - Judgment Aggregation
- We've gotten a taste for the axiomatic method (more tomorrow)
- We proved the impossibility result of List & Pettit
- We discussed a bit the normative appeal of the axioms we've seen

Homework & Next Lecture

Tomorrow's Lecture: We'll go deeper into the axiomatic method.

- Characterisation results.
- Concrete rules other than the majority rule.

The homework for tomorrow is on the project website. Deadline is tomorrow at 11am.

- We will not be grading the homework, but we'll give you some feedback on it tomorrow at the beginning of lecture.

Wednesday and Thursday:

- Can we restrict the agenda to avoid inconsistent or incomplete collective judgments?
- Strategic behaviour: can agents manipulate the outcome?