Computational Social Choice 2020

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

[http://www.illc.uva.nl/~ulle/teaching/comsoc/2020/]

Plan for Today

We often need to model and reason about the preferences of voters over *sets of alternatives*, not just individual alternatives. Examples:

- A voter considering to manipulate an irresolute voting rule has to compare *sets of tied winners* ($\hookrightarrow$ *optimistic*, *pessimistic*, *cautious*).
- In multiwinner voting and participatory budgeting, voters have preferences over *committees* and *sets of projects*, respectively.

Today we approach this topic of *set preferences* more systematically and review ideas from two different strands of the literature:

- *Preference Extensions:* Given a voter's preferences over alternatives, what can we say about her preferences over sets of alternatives?
- *Voting in Combinatorial Domains:* What are good approaches for modelling a voter's preferences over sets of alternatives?

Preference Extensions: Ranking Sets of Objects

Given: A set X of *alternatives* and an agent ranking its elements.

Question: *How would the same agent rank nonempty subsets of X?*

Let us model the agent's preferences on X as a total order $\succcurlyeq$ on X.
We want to *extend* $\succcurlyeq$ to a weak order $\succcurlyeq^E$ on $2^X \setminus \{\emptyset\}$.

Think of sets $S \in 2^X \setminus \{\emptyset\}$ as representing a situation where the agent will eventually get one element $x \in S$ but cannot control which one.
(Other options: you choose / you get everything / freedom of choice)

Exercise: *Test your intuitions! Suppose we know that $a \succ b \succ c \succ d$.*

(1) *Can we infer $\{a\} \succ^E \{b, c\}$?*
(2) *Can we infer $\{b\} \succ^E \{a, c\}$?*
(3) *Can we infer $\{a, b, d\} \succ^E \{a, c, d\}$?*

S. Barberà, W. Bossert, and P.K. Pattanaik. Ranking sets of objects. In S. Barberà et al. (eds.), *Handbook of Utility Theory*, Vol. 2. Kluwer, 2004.

The Kelly Principle

Let us specify some *axioms* for preference extensions $E : \succcurlyeq \mapsto \succcurlyeq^E$.

The extension axiom:

(EXT) $\{a\} \succ^E \{b\}$ if $a \succ b$

Two further axioms:

(MAX) $\{\max(A)\} \succcurlyeq^E A$ [$\max(A)$ = best element in A w.r.t. $\succcurlyeq$]
(MIN) $A \succcurlyeq^E \{\min(A)\}$ [$\min(A)$ = worst element in A w.r.t. $\succcurlyeq$]

The *Kelly Principle* = (EXT) + (MAX) + (MIN). Thus:

- $A \succ^E B$ if all elements in A are strictly better than all those in B
- $A \succcurlyeq^E B$ if all elements in A are at least as good as all those in B

This seems appropriate for a *cautious* voter considering to manipulate.

J.S. Kelly. Strategy-Proofness and Social Choice Functions without Single-Valuedness. *Econometrica*, 1977.

The Gärdenfors Principle

Two more axioms, encoding some kind of *dominance* principle:

(GF1) $A \cup \{b\} \succ^E A$ if $b \succ a$ for all $a \in A$
(GF2) $A \succ^E A \cup \{b\}$ if $a \succ b$ for all $a \in A$

The *Gärdenfors Principle* = (GF1) + (GF2) thus amounts to this:

If I can get from A to B by means of a (nonempty) sequence of steps, each involving deleting the best element or adding a new worst element, then A is strictly better than B.

Nice: This models the behaviour of a voter who expects ties to be broken by another agent with a fixed but unknown preference order.

The Gärdenfors Principle *entails* the Kelly Principle, but not *vice versa*.

Remark: Sometimes this is called the Fishburn Principle instead.

P. Gärdenfors. Manipulation of Social Choice Functions. *Journal of Economic Theory*, 1976.

An Independence Axiom

We may want to assume that, if you (strictly) prefer A over B, then that preference will not get inverted when we add a new object c:

(IND) $A \cup \{c\} \succcurlyeq^E B \cup \{c\}$ if $A \succ^E B$ and $c \notin A \cup B$

Exercise *What do you think? Makes sense?*

The Kannai-Peleg Theorem

The 1984 paper by Yakar Kannai and Bezalel Peleg is considered *the* seminal contribution to the axiomatic study of *ranking sets of objects*.

Theorem 1 (Kannai and Peleg, 1984) *If $|X| \geqslant 6$, then no weak order $\succcurlyeq^E$ satisfies both the Gärdenfors Principle and independence.*

Probably the first paper treating the problem of preference extension as a *problem in its own right*, from an *axiomatic* perspective.

- For Kelly and Gärdenfors (and others), the problem has been more of a side issue (when studying manipulation in voting).
- Work on the problem of ranking sets of objects itself published before 1984 is descriptive rather than axiomatic.

Y. Kannai and B. Peleg. A Note on the Extension of an Order on a Set to the Power Set. *Journal of Economic Theory*, 1984.

Lemma

Recall the axioms:

$$\left.\begin{array}{ll}\text{(GF1) } A \cup \{b\} \succ^E A & \text{if } b \succ a \text{ for all } a \in A \\ \text{(GF2) } A \succ^E A \cup \{b\} & \text{if } a \succ b \text{ for all } a \in A\end{array}\right\} \text{Gärdenfors Principle}$$

(IND) $A \cup \{c\} \succcurlyeq^E B \cup \{c\}$ if $A \succ^E B$ and $c \notin A \cup B$

Lemma 2 *Gärdenfors + independence* $\Rightarrow A \sim^E \{\max(A), \min(A)\}$

Proof:

- If $|A| \leqslant 2$, then $A = \{\max(A), \min(A)\}$. ✓
- If $|A| > 2$:
  - $A \setminus \{\max(A)\} \succ^E \{\min(A)\}$ by repeated application of (GF1), and thus $A \succcurlyeq^E \{\max(A), \min(A)\}$ by (IND).
  - $\{\max(A)\} \succ^E A \setminus \{\min(A)\}$ by repeated application of (GF2), and thus $\{\max(A), \min(A)\} \succcurlyeq^E A$ by (IND).

  Hence, $A \sim^E \{\max(A), \min(A)\}$. ✓

Proof of the Kannai-Peleg Theorem

Theorem: If $|X| \geqslant 6$, then no weak order $\succcurlyeq^E$ satisfies both the Gärdenfors Principle and independence.

Proof: Suppose $a_6 \succ a_5 \succ a_4 \succ a_3 \succ a_2 \succ a_1$.

Claim: $\{a_2, a_5\} \succcurlyeq^E \{a_4\}$ $(*)$

Proof of claim: if not, then $\{a_4\} \succ^E \{a_2, a_5\}$, as $\succcurlyeq^E$ is complete

$\Rightarrow \{a_1, a_4\} \succcurlyeq^E \{a_1, a_2, a_5\}$ by (IND)

$\Rightarrow \{a_1, a_2, a_3, a_4\} \succcurlyeq^E \{a_1, a_2, a_3, a_4, a_5\}$ by our lemma $\Rightarrow \lightning$ [GP]

Hence: $\{a_2, a_5\} \succ^E \{a_3\}$ from $(*)$ and $\{a_4\} \succ^E \{a_3\}$ [GP]

$\Rightarrow \{a_2, a_5, a_6\} \succcurlyeq^E \{a_3, a_6\}$ by (IND)

$\Rightarrow \{a_2, a_3, a_4, a_5, a_6\} \succcurlyeq^E \{a_3, a_4, a_5, a_6\}$ by our lemma $\Rightarrow \lightning$ [GP]

Done. ✓

Remark: Note that there *are* *preorders* satisfying all axioms, e.g.:

$A \succcurlyeq^E B :\Leftrightarrow \max(A) \succcurlyeq \max(B)$ and $\min(A) \succcurlyeq \min(B)$

Automated Theorem Search

A major challenges in COMSOC is to formally model social choice problems so as to facilitate the *automated verification*—or even the *discovery*—of theorems that help us understand a given domain.

For the relatively simple domain of *ranking sets of objects*, this is indeed possible (MoL thesis by Christian Geist, 2010):

> *Systematic search finds all 84 (minimal) impossibilities in a space of 20 axioms for domains X with $|X| \geqslant k$ (for k up to 8).*

More on this technique in the next lecture.

C. Geist and U. Endriss. Automated Search for Impossibility Theorems in Social Choice Theory: Ranking Sets of Objects. *J. Artif. Intell. Research*, 2011.

The Paradox of Multiple Elections

Suppose 13 voters are asked to each vote *yes* or *no* on three issues:

- 3 voters each vote for YNN, NYN, NNY.
- 1 voter each votes for YYY, YYN, YNY, NYY.
- No voter votes for NNN.

If we use the *simple majority* rule *issue-by-issue*, then NNN wins, because on each issue 7 out of 13 voters go for *no*.

This is an instance of the so-called *paradox of multiple elections:* the winning combination received not a single vote!

Exercise: *Can you spell out when this actually is / is not a problem?*

S.J. Brams, D.M. Kilgour, and W.S. Zwicker. The Paradox of Multiple Elections. *Social Choice and Welfare*, 1998.

Voting in Combinatorial Domains

Elections often have a *combinatorial structure:*

- Electing a committee of k members from amongst m candidates.
- Voting on ℓ propositions (yes/no) in a referendum.

Clearly, the number of alternatives can quickly become *very large*.
So we face both a *choice-theoretic* and a *computational challenge*.

For the remainder of today, let us focus on the case of voting in a *combinatorial domain* defined by ℓ *binary* decisions:

$$D_1 \times \cdots \times D_\ell \text{ with } |D_j| = 2$$

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference Handling in Combinatorial Domains: From AI to Social Choice. *AI Magazine*, 2008.

J. Lang and L. Xia. Voting in Combinatorial Domains. In F. Brandt et al. (eds.), *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

Simplistic Approaches

- *Run ℓ elections in parallel, one for each issue!*

  But such issue-by-issue voting can lead to paradoxical outcomes.

- *Vote on combinations in $D_1 \times \cdots \times D_\ell$ directly!*

  This is like the standard model of voting with 2^ℓ alternatives.
  Hardly feasible for most voting rules, unless ℓ is very small.
  Maybe feasible for the plurality rule, but very indecisive. (*Why?*)

- *Preselect small number of combinations and only vote on those!*

  But *who* selects? Would only work in very special circumstances.

Distance-based Aggregation

Idea: Elicit preferred choices issue-by-issue (as for the paradox), but find a better way to aggregate this information.

One such approach is to use the *minimax rule:* elect outcomes that minimise the *maximal Hamming distance* to the ballots.

Thus, if unhappiness $\sim$ number of issues on which you do not get your way, then this rule maximises the happiness of the unhappiest voter.

Exercise: *What do you think about the "minisum rule", which elects outcomes minimising the sum of Hamming distances to the ballots?*

S.J. Brams, D.M. Kilgour, and M.R. Sanver. A Minimax Procedure for Electing Committees. *Public Choice*, 2007.

Sequential Voting

Idea: Vote separately on each issue, but do so sequentially. So voters can make their vote on one issue dependent on earlier decisions.

Very basic positive results are easy to obtain:

Proposition 3 (Lacy and Niou, 2000) *Sequential majority on binary issues never results in a winning combination that is a Condorcet loser.*

Exercise: *Prove that this is true!*

However, sequential majority cannot guarantee Condorcet consistency. Xia and Lang (2009) have identified conditions under which it can.

D. Lacy and E.M.S. Niou. A Problem with Referendums. *Journal of Theoretical Politics*, 2000.

J. Lang and L. Xia. Sequential Composition of Voting Rules in Multi-issue Domains. *Mathematical Social Sciences*, 2009.

Combinatorial Vote

Idea: Ask voters to report their ballots using a compact preference *representation language* and apply your favourite voting rule to the succinctly encoded ballots received.

Lang (2004) calls this approach *combinatorial vote*.

Discussion: A promising approach, but not too much is known to date about what would be good choices for preference representation languages or voting rules, or what algorithms to use to compute the winners. Also, complexity can be expected to be very high.

Let us look at an example . . .

J. Lang. Logical Preference Representation and Combinatorial Vote. *Annals of Mathematics and Artificial Intelligence*, 2004.

The Language of Prioritised Goals

Consider a binary combinatorial domain $D_1 \times \cdots \times D_\ell$ and associate one propositional variable with each D_j. So the elements of the domain correspond to truth assignments to the variables.

Use propositional formulas to express *goals* and use numbers to indicate their relative *importance*. This induces a weak order:

- Suppose $\varphi : k_1$ has higher *priority* than $\psi : k_2$ if $k_1 > k_2$.
- Under *lexicographic aggregation*, we prefer M to M' if there exists a k such that for all $i > k$ both M and M' satisfy the same number of goals of rank i, and M satisfies more goals of rank k.

Other forms of aggregation are possible.

Combinatorial Vote: Example

Use the language of *prioritised goals* (1 has higher priority than 0) with *lexicographic aggregation* together with the *Borda rule:*

- Voter 1: $\{x : 1,\ y : 0\}$ induces order $xy \succ_1 x\bar{y} \succ_1 \bar{x}y \succ_1 \bar{x}\bar{y}$
- Voter 2: $\{x \vee \neg y : 0\}$ induces order $x\bar{y} \sim_2 xy \sim_2 \bar{x}\bar{y} \succ_2 \bar{x}y$
- Voter 3: $\{\neg x : 0,\ y : 0\}$ induces order $\bar{x}y \succ_3 \bar{x}\bar{y} \sim_3 xy \succ_3 x\bar{y}$

As the induced orders need not be strict linear orders, we use a suitable *generalisation of the Borda rule:* an alternative gets as many points as she dominates other alternatives. So we get these Borda scores:

$$\begin{array}{ll} xy : 3+1+1 = 5 & \bar{x}y : 1+0+3 = 4 \\ x\bar{y} : 2+1+0 = 3 & \bar{x}\bar{y} : 0+1+1 = 2 \end{array}$$

So combinatorial alternative xy wins.

Combinatorial vote *proper* would be to compute the winner *directly* from the goalbases, without the detour via the induced orders.

Computational Complexity

Under the *generalised plurality rule*, a voter gives 1 point to each undominated alternative. Suppose voters specify a *single goal* each:

- The goal $\neg x \wedge y$ induces the order $\bar{x}y \succ xy \sim x\bar{y} \sim \bar{x}\bar{y}$, so only combination $\bar{x}y$ receives 1 point.
- The goal $x \vee y$ induces the order $xy \sim \bar{x}y \sim x\bar{y} \succ \bar{x}\bar{y}$, so combinations xy, $\bar{x}y$, $x\bar{y}$ receive 1 point each.

Things quickly get very complex (proof omitted but easy):

Proposition 4 (Lang, 2004) *Deciding whether a given combination wins for a given profile of preferences is coNP-complete when we use the language of single goals and the generalised plurality rule.*

Recall that coNP is the complement of the complexity class NP (i.e., it is the complexity class of checking validity in propositional logic).

J. Lang. Logical Preference Representation and Combinatorial Vote. *Annals of Mathematics and Artificial Intelligence*, 2004.

Summary

We have seen several approaches for modelling voter preferences in combinatorial domains, e.g., when electing sets of alternatives.

One approach is to elicit preferences on alternatives and consider ways of extending them to preferences on sets of alternatives:

- Kelly, Gärdenfors, optimism, pessimism, ...
- Kannai-Peleg Theorem: dominance + independence impossible

Another family of approaches involves eliciting preferences regarding the combinatorial domain of interest:

- Distance-based voting rules, such as the minimax rule
- Sequential voting, so voters can circumvent paradoxes
- Combinatorial vote with compactly expressed preferences

What next? Automated reasoning for social choice theory.