

Collective Decision Making under Incomplete Knowledge: Possible and Necessary Solutions

Jérôme Lang

CNRS, LAMSADE, PSL, Université Paris-Dauphine, France
lang@lamsade.dauphine.fr

Abstract

Most solution concepts in collective decision making are defined assuming complete knowledge of individuals' preferences and of the mechanism used for aggregating them. This is often unpractical or unrealistic. Under incomplete knowledge, a solution advocated by many consists in quantifying over all completions of the incomplete preference profile (or all instantiations of the incompletely specified mechanism). Voting rules can be 'modalized' this way (leading to the notions of possible and necessary winners), and also efficiency and fairness notions in fair division, stability concepts in coalition formation, and more. I give here a survey of works along this line.

1 Introduction

In the last 20 years there has been a significant interest on collective decision making within the AI community, with the aim of designing and supporting new ways for agents to make group decisions, in a more flexible and realistic way than what classical social choice methods would allow.

One reason why collective decision making in modern applications has to be flexible is that preferences can often not be assumed to be elicited entirely and in a single pass. This leads to viewing collective decision making as an interactive process, and ask questions such as: when agents' preferences are partly known at some stage, should we simply wait, or do we know the outcome already, or at least part of it?

Beyond preferences, incomplete knowledge can also bear on the mechanism itself, for various reasons (including reducing the impact of strategic behaviour).

For both sources of incomplete knowledge, a notion that has been used in several subfields of computational social choice is what I will call *possible and necessary solutions*. It has been used in voting under the terminology *possible and necessary winners*, and more recently in fair division and coalition structure formation.

This survey, and the chapter *Incomplete Information and Communication in Voting* [Boutilier and Rosenschein, 2016], in the *Handbook of Computational Social Choice*, complement each other. The handbook chapter has a much wider

scope, as it covers more generally information and communication issues in voting. Its Section 10.3.1 covers possible and necessary winners in voting, with much less details and structure than this survey. Also, the handbook chapter deals only with voting, and not with other social choice domains.

As most of the work I am about to survey bears on incomplete preferences, in each section I start by this setting, and consider incompletely specified mechanisms later.

2 Possible and Necessary Winners in Voting

2.1 Definitions

An *election* consists of a set N of n agents, a set C of m candidates, and a *preference profile* $R = (\succ_1, \dots, \succ_n)$, where each \succ_i is a linear order¹ over C , called a *preference relation* or *vote*. A *property* F (of candidates) maps R to a subset of C . If $F(R)$ is always nonempty then F is a *voting rule*; and if $F(P)$ is always a singleton then it is a *resolute voting rule*.

An example of a property is 'being a Condorcet winner' (CW) : $c \in C$ is a CW is for any $c' \in C \setminus \{c\}$, a majority of voters prefers c to c' , that is, $|\{i, c \succ_i c'\}| > \frac{n}{2}$. An example of a voting rule is *plurality* (F_{plu}): it outputs candidates ranked with largest *plurality score*, defined as the number of votes in which they are ranked on top. A voting rule can be made resolute by composing it with a tie-breaking mechanism, usually a priority relation \triangleright over C . For instance, if $R = (a \succ b \succ c, a \succ b \succ c, b \succ c \succ a, b \succ c \succ a, c \succ a \succ b)$ then the CW property outputs the empty set, the plurality rule outputs $\{a, b\}$, and resolute plurality with tie-breaking priority $a \triangleright b \triangleright c$ (denoted by $F_{plu, a \triangleright b \triangleright c}$) outputs $\{a\}$.

A natural model of incomplete information about votes consists in assuming that only a *partial order* for each voter is known.² A *partial* (or *incomplete*) vote P_i is a strict partial order (an irreflexive and transitive relation) over C . A (complete) vote \succ_i is a *completion* of P_i if it contains P_i . A *partial profile* P is a collection (P_1, \dots, P_n) of incomplete votes. A (complete) profile $R = (\succ_1, \dots, \succ_n)$ is a *completion* of P if for all i , \succ_i is a completion of P_i .

Given a candidate $c \in C$, a voting rule or property F and an incomplete profile P , c is a *possible winner* (for P and F) if there exists a completion R of P such that $c \in F(R)$;

¹We stick here to rules defined from *ordinal preferences*, which is a usual view in voting and more generally in social choice.

²This is nevertheless not the only possibility. See Section 2.4.

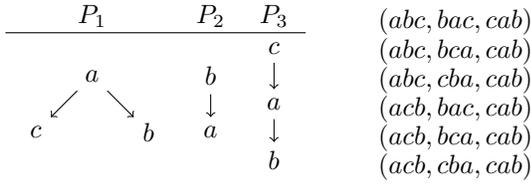


Figure 1: Partial profile P and its six completions

and c is a *necessary winner* if for every completion R of P , $x \in F(R)$ [Konczak and Lang, 2005].

If F is a voting rule then any necessary winner is a possible winner, and there is at least one possible winner. If moreover F is resolute, then there is at most one necessary winner.

Example 1. We take $n = m = 3$, $C = \{a, b, c\}$. Voter 1 only specifies that a is her top candidate. Voter 2 only says that she prefers b to a . Voter 3 reports a complete preference relation. Figure 1 shows the partial profile (left) and its six completions (right). (We note abc instead of $a \succ b \succ c$, etc.)

We show below the winners for all completions of P , for the Condorcet winner property, irresolute plurality, and resolute plurality with tie-breaking priority $b \triangleright a \triangleright c$.

P_1	P_2	P_3	CW	F_{plu}	$F_{plu, b \triangleright a \triangleright c}$
abc	cba	cab	c	c	c
abc	bca	cab	\emptyset	a, b, c	b
abc	bac	cab	a	a, b, c	b
acb	cba	cab	c	c	c
acb	bca	cab	c	a, b, c	b
acb	bac	cab	a	a, b, c	b

- a and c are possible Condorcet winners. There is no necessary Condorcet winner.
- for irresolute plurality: a , b and c are possible winners, and c is a necessary winner.
- for $F_{plu, b \triangleright a \triangleright c}$: b and c are possible winners, and there is no necessary winner.

If we further learn that voter 2 prefers a to c , only completions (abc, bac, cab) and (acb, bac, cab) remain; now a is a necessary CW and b a necessary winner for $F_{plu, b \triangleright a \triangleright c}$.

In the case of a resolute rule, the existence of a necessary winner is a sufficient and necessary condition for the elicitation process to be considered terminated (as there is enough information for the outcome to be known). This question has been considered first in [Conitzer and Sandholm, 2002] and [Walsh, 2008].³ [Ding and Lin, 2013] go further by determining which are the relevant queries voters should be asked.

2.2 Computation

Because the number of completions of a partial profile is exponential, deciding whether a candidate is a possible or a necessary winner is not guaranteed to be easy, and indeed it is often hard. The complexity of these decision problems for a variety of rules is addressed first in [Xia and Conitzer, 2011]

³In the case of irresolute rules, whether the elicitation process is considered terminated once some necessary winner has been found depends on the context.

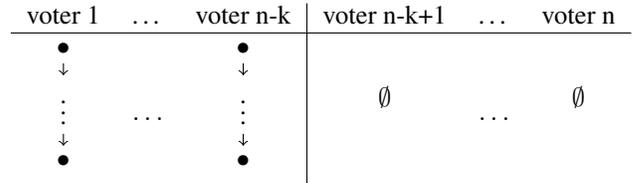
for a variety of rules, and further in [Betzler and Dorn, 2010; Baumeister and Rothe, 2012] for the specific case of positional scoring rules.

The parameterized complexity of the problems is initiated in [Betzler *et al.*, 2009], where the relevant parameters are the number of candidates, number of voters, and degree of incompleteness (number of undetermined pairs in votes). Kernelization complexity results with number of candidates as parameter are given in [Dey *et al.*, 2016]. The impact of the degree of incompleteness is investigated further in [Dey and Misra, 2017], who determine, for several rules, the smallest number of undetermined pairs for which the possible winner problem remains hard.

2.3 Special Cases

Missing Votes

A partial profile $P = (P_1, \dots, P_n)$ is of the *missing voters* (MV) type if for every i , either P_i is either complete or empty. Without loss of generality, let us renumber the voters such that P_i is complete for each $i \leq k$ and empty for $i > k$.



A partial profile of the *missing voters* type

There are at least two interpretations of these MV profiles. Under the *strategic* interpretation, $\{k+1, \dots, n\}$ is a coalition of manipulating voters who, knowing the votes of $\{1, \dots, k\}$, can coordinate and turn the outcome to their advantage. These two problems are defined in [Bartholdi *et al.*, 1989]: for a resolute rule F , given the votes of $\{1, \dots, k\}$, coalition $\{k+1, \dots, n\}$ has a *constructive* (resp. *destructive*) *manipulation* for (resp. against) $c \in C$ if there is a way for them to vote so that c is the winner (resp. so that c is not the winner). Clearly, in a MV profile, c is a possible winner if and only if the missing voters have a coalitional manipulation for c , and a necessary winner if and only if they have no destructive manipulation against c .

These definitions assume that the manipulating voters have full knowledge of the nonmanipulators' votes. This assumption is relaxed in [Conitzer *et al.*, 2011], who assume nonmanipulators' votes to be partial orders and define a dominating manipulation as a collection of votes by the manipulators making the winner preferable (to them) to the winner obtained with truthful votes. In a similar setting, [Dey *et al.*, 2018] study a notion of *weak manipulation* that corresponds to the possible winner problem given k partial votes and $n-k$ empty votes (see Section 2.4 for further discussion).

Under the *temporal* interpretation, votes come in an asynchronous way. At time t , the set of voters who have sent their votes is $V(t)$; we usually assume V to be monotonic⁴ and as soon as we have necessary winner, we can stop waiting for more votes to come. This situation occurs typically

⁴Although one may argue that there may be cases where some voters can cancel their vote after some time.

in date-finding polls. In such a context, during the elicitation process voters can be asked to block the dates mentioned in the poll, and identifying the possible winners before the end of the elicitation process allows them to block less dates.

Example 2. F is the irresolute plurality rule. At time t , let $s_t(x)$ be the plurality score of $x \in C$. The number of voters who have already voted is $|V(t)| = \sum_x s_t(x)$. x is a necessary winner if $s_t(x) \geq \max_{y \neq x} s_t(y) + n - |V(t)|$ and a possible winner if $s_t(x) + n - |V(t)| \geq \max_{y \neq x} s_t(y)$.

This assumes that the set of voters is fixed from the beginning. When it is not, all-or-nothing notions no longer make sense; one can try instead to identify candidates who are winners with a large enough probability [Dey *et al.*, 2017].

Missing Candidates

A partial profile is of the *missing candidates* (MC) type if the partial votes consist of complete rankings over some subset of candidates (the same for all voters), with no information at all on the other candidates. Concrete situations include polls about meeting dates where some dates that were initially considered impossible (for instance because the meeting room was thought to be unavailable) have become possible in the meantime, or applicants to a job that are known only after the first set of applicants have been interviewed and ranked by the members of the committee.

voter 1	voter 2	...	voter n
c	b		b
↓	↓		↓
a	c	...	a
↓	↓		↓
b	a		c
(d?e?)	(d?e?)		(d?e?)

A partial profile of the *missing candidates* type

Clearly, for any reasonable voting rule, these new candidates are possible winners. The question is which of the initial candidates are possible winners, that is, can still win after we learn how the voters compare the new candidates to them.

Example 3. $n = 12$. $\{a, b, c\}$ are the initial candidates and y the only new candidate. The rule is plurality with tie-breaking priority $a \triangleright b \triangleright c \triangleright y$. The initial plurality scores (before preferences about y are known) are $a : 5, b : 4, c : 3$. Clearly, a can still win (for instance if no one ranks y first); b will win provided that at least two (and at most three) voters who used to rank a first now decide to rank y first (third column); c cannot win, because to make the scores of a and b lower than its, we would need too many voters ranking y first, and y would become the winner (rightmost column).

	initial scores	b wins	y wins
a	5	3	2
b	4	4	2
c	3	3	3
y	—	2	5

If there are two new candidates y_1, y_2 instead of one, then c is a possible winner if it has priority over y_1 or over y_2 .

Characterizing possible winners in this specific context has been considered for positional scoring rules [Chevalerey *et al.*, 2012] and a few other rules [Xia *et al.*, 2011].

Truncated Ballots

Asking voters to rank all candidates is often too much of a communication burden on them. To limit the amount of elicitation, we may ask them to rank only their top k candidates, for some k . This leads to a partial profile of the *top- k* type:

voter 1	voter 2	...	voter n
●	●		●
↓	↓		↓
●	●	...	●
↓	↓		↓
?	?		?

A partial profile of the *top- k* ballots type (here $k = 2$)

A possible variation is to allow each voter i to choose their own k_i . Another one consists in *double-truncated ballots*: each voters reports their k^+ top candidates and their bottom k^- . Truncated ballots have been considered in quite many papers; those that focus explicitly on possible and necessary winners are [Baumeister *et al.*, 2012a; Kalech *et al.*, 2011; Ayadi *et al.*, 2019], the latter two in the context of designing low-communication protocols. A protocol studied in [Kalech *et al.*, 2011] starts by eliciting the top-1 ballots, then top-2, and so on, until there is a necessary winner; in practice, this “next-best queries” protocol saves a lot of communication.

Incomplete Lists

A partial profile is of the *incomplete lists* (IL) type if the voters rank only some candidates, namely those they know (the films they have seen, the candidates they have interviewed etc.) The crucial difference with truncated ballots is that here the candidates that do not appear in a ranking can be ranked anywhere and not necessarily below those that appear. The crucial difference with the “missing candidates” setting is that the ranked and the missing candidates are voter-dependent.

voter 1	voter 2	...	voter n
c			c
↓			↓
a	d	...	e
↓	↓		↓
b	a		d
			↓
			f

A partial profile of the *incomplete lists* type

This situation occurs in information retrieval, when one has to merge the ordered lists of documents returned by different search engines, and in systems where individuals rank restaurants, films etc. However, unlike other types of partial profiles, the IL type is less suited to the possible and necessary winner notions: the information is generally too incomplete so that typically, almost all candidates are possible winners. Dealing with IL profiles calls for either the counting version of possible winners (cf. Subsection 2.4), or to specific aggregation methods: for instance, [Caragiannis *et al.*, 2015] aggregate such incomplete lists in the context of peer grading in Massive Online Open Courses, where each student ranks a small number of other students’ exams.

Cloning

Assume we know the voters' preference relations over C , and that now we add p clones a_1, \dots, a_q of candidate $a \in C$. Let $C(a) = \{a_0 = a, a_1, \dots, a_q\}$. The voters' preference relations \succ'_i over the new set of candidates $C \cup \{a_1, \dots, a_q\}$ are such that, for each i , (1) for each $a_j \in C(a)$ and each $x \in C \setminus C(a)$, $a_j \succ'_i x$ if and only if $a \succ_i x$ and $x \succ'_i a_j$ if and only if $x \succ_i a$; (2) for each $x, y \in C \setminus C(a)$, $x \succ'_i y$ if and only if $x \succ_i y$. It is not possible to predict how voters will rank the elements of $C(a)$ with respect to one another. We obtain a partial profile P of the *cloning* type:

voter 1	voter 2	...	voter n
c	$a_0 \ a_1 \ a_2$		c
$\swarrow \downarrow \searrow$ $a_0 \ a_1 \ a_2$	$\swarrow \downarrow \searrow$ b	...	\downarrow b
$\swarrow \downarrow \searrow$ b	\downarrow c		$\swarrow \downarrow \searrow$ $a_0 \ a_1 \ a_2$

A partial profile of the *cloning* type

A resolute rule F is clone-proof if whenever the winner in the initial profile (without clones) is $x \neq a$, the only possible winner in P is x , and whenever the winner in the initial profile is a , the possible winners in P are all in $C(a)$. When F is not clone-proof, determining the possible winners can be challenging; this is studied in [Elkind *et al.*, 2011].

2.4 Variants and Extensions

More quantifier alternations. Possible and necessary winners correspond to simple existential and universal quantification, or, in epistemic logic terms, to possibility and knowledge: x is a necessary winner if x is known to be a winner ($\mathbf{K}win_x$), and x is a possible winner if x is not known not to be a winner ($\neg\mathbf{K}\neg win_x$), where possible worlds corresponds to profile completions. The notion of *strong manipulation* [Dey *et al.*, 2018] uses one more level: given k partial votes, the remaining $n - k$ voters have a strong manipulation for x if there exists a way for them to cast their votes so that for any completion of the first k votes, x is a winner. More formally, given an incomplete profile P and a partition (N_1, N_2) of the voters, x is a $\exists N_1 \forall N_2$ -winner for F if there is a completion $R_1 = (\succ_i, i \in N_1)$ of $(P_i, i \in N_1)$ such that for any completion $R_2 = (\succ_i, i \in N_2)$ of $(P_i, i \in N_2)$, $x \in F(R_1, R_2)$. Other quantifier alternations could be considered.

Probabilistic possible winners. If the incompleteness of the profile is significant, determining possible winners can be disappointing: typically, all candidates, or almost all, are then possible winners. In such a case one may instead want to *count* completions. [Bachrach *et al.*, 2010] define the winning probability of a candidate c as the proportion of completions for which c is a winner; they show that although computing the exact winning probability is hard, giving a good additive approximation of it by sampling is tractable for some rules. [Kenig and Kimelfeld, 2019] obtain results for multiplicative approximation ratios, and beyond the uniform distribution over profile completions, they consider a more realistic family of distributions: the repeated insertion model. Finally, [Hazon *et al.*, 2012] consider a different uncertainty model

about votes: instead of a partial order, each voter comes with a probability distribution over complete votes.

Resoluteness with incomplete preferences. As said above, in some cases there are typically too many possible winners and we may want to define resolute (or at least, more resolute) rules from incomplete profiles. A first way to resoluteness, followed in [Lu and Boutilier, 2011; Benabbou *et al.*, 2016; Lu and Boutilier, 2020] is to select the *minimax regret* winner, that is, the one which is closest to optimal in the worst case. A necessary winner is also a minimax regret winner but, perhaps unexpectedly, a minimax regret winner may not be a possible winner [Lu and Boutilier, 2011]. Another way to resoluteness is to define rules from incomplete profiles as *maximum likelihood estimators* [Xia and Conitzer, 2011].

Other models of uncertainty. [Shiryayev *et al.*, 2013] consider a type of incomplete knowledge about votes which is generally not representable by partial orders: a 'default' (or 'reference') vote of each voter is known, but the voter may report a vote that is within a maximal number of swaps away from her default vote. A *robust winner* is then like a necessary winner, but for this type of incomplete knowledge. Both this model, and possible and necessary winners, are specialisation of *swap bribery* [Elkind *et al.*, 2009].

In the specific case of approval voting, [Barrot *et al.*, 2013] assume that voters' preferences are fully known, and that incomplete knowledge bears on the approval ballots they report. Assuming these ballots are sincere (*i.e.*, a ballot consists of all candidates above some 'threshold' candidate), they characterize and compute the possible (sets of) approval winners. [Benabbou and Perny, 2016] also consider possible approval winners, in a group knapsack context where a common set of objects is to be selected.

Voting and databases. [Kimelfeld *et al.*, 2018] relate possible and necessary winners to possible and necessary answers in incomplete databases, and enrich the possible and necessary winner setting with conjunctive database queries: given an incomplete profile and some attributes describing the candidates, they consider questions such as the existence of a set of possible winners satisfying some conjunctive query.

2.5 Incomplete Knowledge on the Mechanism

Instead of bearing on the votes, incomplete knowledge can bear on some parameters of the voting rule, such as:

1. the voters and candidates that take part to the election [Wojtas and Faliszewski, 2012];
2. the voters' weights [Baumeister *et al.*, 2012b];
3. the order in which voters vote, under a simple strategic behavioural assumption [Gaspers *et al.*, 2014];
4. the initial position (or *seeding*) of candidates in knock-out tournaments (also called voting trees) [Aziz *et al.*, 2018; Mattei *et al.*, 2015];
5. the scoring vector for positional scoring rules [Viappiani, 2018].

3 Beyond Voting

The epistemic principle at work for the possible winner problem (consider all possible completions of partial preferences as possible worlds and quantify over them) naturally extends to more social choice settings, and even beyond social choice.

3.1 Partial Tournaments

In [Aziz *et al.*, 2015a], the input is a *partial tournament* (a directed asymmetric graph) over C , or a *partial weighted tournament* (a weighted directed graph $w(\cdot, \cdot)$, such that for all $a, b \in C$, $w(a, b) + w(b, a) \leq n$, where n is independent from a and b). A partial [weighted] tournament may come from a partial profile⁵, or can be an object on its own, such as a set of partial outcomes in a sports tournament. Given a [weighted] *tournament solution* S , that maps a [weighted] tournament into a subset of C , c is a possible winner for partial tournament T if there is a complete tournament T^* extending T such that $c \in S(T^*)$, and a necessary winner for T if for every complete extension T^* of T , $c \in S(T^*)$. [Aziz *et al.*, 2015a] give computational results for various voting rules. [Brandt *et al.*, 2018] study the social-choice theoretic properties of these extensions of tournament solutions to incomplete profiles. Earlier work [Lang *et al.*, 2012; Pini *et al.*, 2011] studied voting trees and the Schwartz rule defined from incomplete tournaments. [Yang and Guo, 2017] give a parameterized study of the complexity of winner determination for partial tournament solutions.

3.2 Fair Division

A fair division problem (of indivisible items) setting consists of a set of agents $N = \{1, \dots, n\}$, a set of items $O = \{o_1, \dots, o_m\}$, and for each agent i , a preference relation \succeq_i over 2^O . An *allocation* π maps each agent to a set of items, such that $\pi(i) \cap \pi(j) = \emptyset$ for $i \neq j$. It is complete if all items are allocated. The quality of an allocation is usually evaluated through some criteria; here we give only these two:

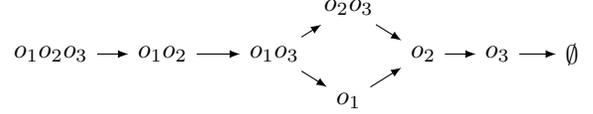
- *envy-freeness*: π is envy-free if for all i, j , $\pi(i) \succeq_i \pi(j)$
- *Pareto-efficiency*: π is Pareto-efficient if there is no π' such that (a) $\pi'(i) \succeq_i \pi(i)$ for all i , and (b) $\pi'(i) \succ_i \pi(i)$ for some i .

Envy-freeness is a fairness criterion while Pareto-efficiency and completeness are efficiency criteria. Efficiency and fairness criteria tend to conflict with each other; for instance, the existence of an allocation that is both envy-free and Pareto-efficient (or even complete) is not guaranteed.

When the number of items grows, it becomes unreasonable to ask agents to report a ranking over all subsets of O . A way to avoid this exponential communication burden consists in letting agents express preferences in a compact way, possibly leaving some incompleteness in the reported preference. Especially, they may rank individual items and then extend the

⁵If a partial tournament T comes from a partial profile P , i.e., if T is the majority graph of P , then the set of possible winners determined from T contains the set of possible winners determined from P (for the same tournament solution). The rule applied to partial tournaments has often a lower complexity than the same rule applied to partial profiles.

preference relation to subsets of items using the *responsive extension*: given an order \triangleright over O , the responsive extension of \triangleright to 2^O is the smallest (complete) preference relation \succeq over 2^O which is monotonic (if $S \subseteq T$ then $T \succeq S$) and weakly separable (if $o, o' \neq S$ then $S \cup \{o\} \succeq S \cup \{o'\}$ if and only if $o \triangleright o'$). The responsive extension of $o_1 \triangleright o_2 \triangleright o_3$ is depicted on the following figure.



Let $P = \langle \triangleright_1, \dots, \triangleright_n \rangle$ where each \triangleright_i is a ranking over O . From P we derive $P^* = \langle \succ_1, \dots, \succ_n \rangle$ where each \succ_i is the responsive extension of \triangleright_i . $T = \langle T_1, \dots, T_n \rangle$ is a completion of P^* if each T_i is a linear order extending \succ_i . Given a fairness or efficiency criterion Γ , an allocation π is

- *possibly* Γ if there exists a completion T of P^* such that π satisfies Γ with respect to T .
- *necessarily* Γ if for every completion T of P^* , π satisfies Γ with respect to T .

For some properties Γ , the characterization of *possibly* Γ and/or *necessarily* Γ is easy. For instance, π is necessarily envy-free if for all agents i, j , and all $k \leq |\pi_i|$, i prefers his k th best item in π_i to the k th best item in π_j [Bouveret *et al.*, 2010]. For some other properties, there is no simple characterization. Such characterizations (when they exist) as well as the computation of allocations that satisfy some properties or combination of properties are given in [Bouveret *et al.*, 2010] and [Aziz *et al.*, 2015b].⁶ Finally, Section 4 of [Aziz *et al.*, 2019] characterizes possibly and necessarily Pareto-optimal allocations under responsive preferences.

Example 4.

$$\begin{aligned} \triangleright_1: & a \triangleright b \triangleright c \triangleright e \triangleright d \triangleright f \\ \triangleright_2: & a \triangleright d \triangleright b \triangleright c \triangleright e \triangleright f \\ \triangleright_3: & b \triangleright a \triangleright c \triangleright f \triangleright d \triangleright e \end{aligned}$$

Because 1 and 2 have the same most preferred item, there is no complete necessarily envy-free allocation, and a fortiori no possibly Pareto-efficient, necessarily envy-free allocation. If 1 leaves, then the allocation π such that $\pi(2) = \{a, d, e\}$ and $\pi(3) = \{b, c, f\}$ is necessarily envy-free. It is also possibly Pareto-efficient, but not necessarily Pareto-efficient, because in a completion where $\{c, d, e, f\} \succ_2 \{a, d, e\}$ and $\{a, b\} \succ_3 \{b, c, f\}$, π is Pareto-dominated by π' defined by $\pi'(2) = \{c, d, e, f\}$ and $\pi'(3) = \{a, b\}$.

Incomplete knowledge can also bear on the mechanism itself. [Aziz *et al.*, 2015c] and [Cechlárová *et al.*, 2017] consider sequential allocation mechanisms where agents pick items in sequence, following a predefined order. In Example 4, if the picking sequence is 123321, 1 picks a , 2 picks d , 3 picks b and c , 2 picks e and 1 picks f . If it is only known that the sequence will be $xyzzyx$ for some permutation xyz of 123, then c is possibly allocated to any agent, a is possibly allocated to 1 or 2, and d is necessarily allocated to 2.

⁶[Bouveret *et al.*, 2010] focus on envy-freeness. [Aziz *et al.*, 2015b] also consider proportional fair share, which needs preferences to be expressed by utility functions.

3.3 Hedonic Games

In a hedonic game, agents form coalitions, and have preferences about who is in their coalition. The classical task consists in finding coalitions structures that satisfy some rationality or stability concepts. However, expressing preferences over coalitions needs exponential space. This lead to studying hedonic games under various compact representations. In [Lang *et al.*, 2015], this is done via a generalization of the responsive extension, where an agent partitions other agents between friends, enemies, and neutral agents, and rank friends and enemies. Since the resulting preferences over sets of agents are incomplete, concepts are modalized: given a concept X , a coalition structure is possibly X if it satisfies X for some completion of the partial profile. The characterizations and complexity results in [Lang *et al.*, 2015] have been complemented by those in [Kerkmann and Rothe, 2019].

Beyond social choice, possible and necessary solution concepts have also been studied in noncooperative games [Brill *et al.*, 2016].

4 Discussion

A quick look at the bibliography shows that while in years 2005-2015 there was a focus on the classical possible winner problem in voting, in the last 5 years the focus moved to variants and extensions of the model and expanded to other social choice settings.

A key question for future research is the interleaving between reasoning about possible outcomes and eliciting further the agents' preferences. To say it differently: how can identifying possible and necessary winners help in making the decision quicker? The existence of a necessary winner is a safe stopping criterion for the elicitation process, but before this point is reached, how can the identification of possible winners help? Removing necessary losers is generally not a good idea, as it could change the result; but sometimes removing some of them is safe (see *e.g.*, [Ayadi *et al.*, 2019] for the special case of the STV rule).

One may also wonder if other meaningful classes of incomplete preferences are yet to be considered. The answer is yes: one should devote some attention to multichotomies (each voter divides candidates in clusters, such as good, acceptable and unacceptable candidates, with all candidates in a cluster being initially incomparable), and multi-attribute preferences with incompletely specified importance over attributes.

Acknowledgements

This work was funded in part by the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program ANR-19-P3IA-0001 (PRAIRIE 3IA Institute).

References

[Ayadi *et al.*, 2019] M. Ayadi, N. Ben Amor, J. Lang, and D. Peters. Single transferable vote: Incomplete knowledge and communication issues. In *AAMAS*, 2019.

- [Aziz *et al.*, 2015a] H. Aziz, M. Brill, F. A. Fischer, P. Harrenstein, J. Lang, and H. G. Seedig. Possible and necessary winners of partial tournaments. *J. Artif. Intell. Res.*, 54:493–534, 2015.
- [Aziz *et al.*, 2015b] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artif. Intell.*, 227:71–92, 2015.
- [Aziz *et al.*, 2015c] H. Aziz, T. Walsh, and L. Xia. Possible and necessary allocations via sequential mechanisms. In *IJCAI*, pages 468–474, 2015.
- [Aziz *et al.*, 2018] H. Aziz, S. Gaspers, S. Mackenzie, N. Mattei, P. Stursberg, and T. Walsh. Fixing balanced knockout and double elimination tournaments. *Artif. Intell.*, 262:1–14, 2018.
- [Aziz *et al.*, 2019] H. Aziz, P. Biró, J. Lang, J. Lesca, and J. Monnot. Efficient reallocation under additive and responsive preferences. *Theor. Comput. Sci.*, 790:1–15, 2019.
- [Bachrach *et al.*, 2010] Y. Bachrach, N. Betzler, and P. Faliszewski. Probabilistic possible winner determination. In *AAAI*, 2010.
- [Barrot *et al.*, 2013] N. Barrot, L. Gourvès, J. Lang, J. Monnot, and B. Ries. Possible winners in approval voting. In *Algorithmic Decision Theory (ADT)*, pages 57–70, 2013.
- [Bartholdi *et al.*, 1989] J. J. Bartholdi, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6:227–241, 1989.
- [Baumeister and Rothe, 2012] D. Baumeister and J. Rothe. Taking the final step to a full dichotomy of the possible winner problem in pure scoring rules. *Inf. Process. Lett.*, 112(5):186–190, 2012.
- [Baumeister *et al.*, 2012a] D. Baumeister, P. Faliszewski, J. Lang, and J. Rothe. Campaigns for lazy voters: truncated ballots. In *AAMAS*, pages 577–584, 2012.
- [Baumeister *et al.*, 2012b] D. Baumeister, M. Roos, J. Rothe, L. Schend, and L. Xia. The possible winner problem with uncertain weights. In *ECAI*, 2012.
- [Benabbou and Perny, 2016] N. Benabbou and P. Perny. Solving multi-agent knapsack problems using incremental approval voting. In *ECAI*, pages 1318–1326, 2016.
- [Benabbou *et al.*, 2016] N. Benabbou, S. Di Sabatino Di Diodoro, P. Perny, and P. Viappiani. Incremental preference elicitation in multi-attribute domains for choice and ranking with the Borda count. In *Scalable Uncertainty Management (SUM)*, pages 81–95, 2016.
- [Betzler and Dorn, 2010] N. Betzler and B. Dorn. Towards a dichotomy for the possible winner problem in elections based on scoring rules. *J. Comput. Syst. Sci.*, 76(8):812–836, 2010.
- [Betzler *et al.*, 2009] N. Betzler, S. Hemmann, and R. Niedermeier. A multivariate complexity analysis of determining possible winners given incomplete votes. In *IJCAI*, pages 53–58, 2009.
- [Boutilier and Rosenschein, 2016] C. Boutilier and J. S. Rosenschein. Incomplete information and communication in voting. In *Handbook of Computational Social Choice*, pages 223–258. Cambridge University Press, 2016.
- [Bouveret *et al.*, 2010] S. Bouveret, U. Endriss, and J. Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *ECAI*, pages 387–392, 2010.
- [Brandt *et al.*, 2018] F. Brandt, M. Brill, and P. Harrenstein. Extending tournament solutions. *Social Choice and Welfare*, 51(2):193–222, 2018.
- [Brill *et al.*, 2016] M. Brill, R. Freeman, and V. Conitzer. Computing possible and necessary equilibrium actions (and bipartisan set winners). In *AAAI*, 2016.

- [Caragiannis *et al.*, 2015] I. Caragiannis, G. A. Krimpas, and A. Voudouris. Aggregating partial rankings with applications to peer grading in massive online open courses. In *AAMAS*, pages 675–683, 2015.
- [Cechlárová *et al.*, 2017] K. Cechlárová, T. Fleiner, and I. Schlotter. Possible and necessary allocations under serial dictatorship with incomplete preference lists. In *Algorithmic Decision Theory (ADT)*, pages 300–314, 2017.
- [Chevaleyre *et al.*, 2012] Y. Chevaleyre, J. Lang, N. Maudet, J. Monnot, and L. Xia. New candidates welcome! possible winners with respect to the addition of new candidates. *Mathematical Social Sciences*, 64(1):74–88, 2012.
- [Conitzer and Sandholm, 2002] V. Conitzer and T. Sandholm. Vote elicitation: Complexity and strategy-proofness. In *AAAI*, pages 392–397, 2002.
- [Conitzer *et al.*, 2011] V. Conitzer, T. Walsh, and L. Xia. Dominating manipulations in voting with partial information. In *AAAI*, 2011.
- [Dey and Misra, 2017] P. Dey and N. Misra. On the exact amount of missing information that makes finding possible winners hard. In *Mathematical Foundations of Computer Science (MFCS)*, pages 57:1–57:14, 2017.
- [Dey *et al.*, 2016] P. Dey, N. Misra, and Y. Narahari. Kernelization complexity of possible winner and coalitional manipulation problems in voting. *Theor. Comput. Sci.*, 616:111–125, 2016.
- [Dey *et al.*, 2017] P. Dey, N. Talmon, and O. van Handel. Proportional representation in vote streams. In *AAMAS*, pages 15–23, 2017.
- [Dey *et al.*, 2018] P. Dey, N. Misra, and Y. Narahari. Complexity of manipulation with partial information in voting. *Theor. Comput. Sci.*, 726:78–99, 2018.
- [Ding and Lin, 2013] N. Ding and F. Lin. Voting with partial information: what questions to ask? In *AAMAS*, pages 1237–1238, 2013.
- [Elkind *et al.*, 2009] E. Elkind, P. Faliszewski, and A. M. Slinko. Swap bribery. In *SAGT*, pages 299–310, 2009.
- [Elkind *et al.*, 2011] E. Elkind, P. Faliszewski, and A. M. Slinko. Cloning in elections: Finding the possible winners. *J. Artif. Intell. Res.*, 42:529–573, 2011.
- [Gaspers *et al.*, 2014] S. Gaspers, V. Naroditskiy, N. Naroditska, and T. Walsh. Possible and necessary winner problem in social polls. In *AAMAS*, pages 613–620, 2014.
- [Hazon *et al.*, 2012] N. Hazon, Y. Aumann, S. Kraus, and M. J. Wooldridge. On the evaluation of election outcomes under uncertainty. *Artif. Intell.*, 189:1–18, 2012.
- [Kalech *et al.*, 2011] M. Kalech, S. Kraus, G. A. Kaminka, and C. V. Goldman. Practical voting rules with partial information. *Autonomous Agents and Multi-Agent Systems*, 22(1):151–182, 2011.
- [Kenig and Kimelfeld, 2019] B. Kenig and B. Kimelfeld. Approximate inference of outcomes in probabilistic elections. In *AAAI*, pages 2061–2068, 2019.
- [Kerkmann and Rothe, 2019] A. M. Kerkmann and J. Rothe. Stability in FEN-hedonic games for single-player deviations. In *AA-MAS*, pages 891–899, 2019.
- [Kimelfeld *et al.*, 2018] B. Kimelfeld, P. G. Kolaitis, and J. Stoyanovich. Computational social choice meets databases. In *IJCAI*, pages 317–323, 2018.
- [Konczak and Lang, 2005] K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *Proceedings of the Multidisciplinary IJCAI-05 Workshop on Advances in Preferences Handling*, pages 124–129, 2005.
- [Lang *et al.*, 2012] J. Lang, M. S. Pini, F. Rossi, D. Salvagnin, K. B. Venable, and T. Walsh. Winner determination in voting trees with incomplete preferences and weighted votes. *Autonomous Agents and Multi-Agent Systems*, 25(1):130–157, 2012.
- [Lang *et al.*, 2015] J. Lang, A. Rey, J. Rothe, H. Schadrack, and L. Schend. Representing and solving hedonic games with ordinal preferences and thresholds. In *AAMAS*, 2015.
- [Lu and Boutilier, 2011] T. Lu and C. Boutilier. Robust approximation and incremental elicitation in voting protocols. In *IJCAI*, pages 287–293, 2011.
- [Lu and Boutilier, 2020] T. Lu and C. Boutilier. Preference elicitation and robust winner determination for single- and multi-winner social choice. *Artif. Intell.*, 279, 2020.
- [Mattei *et al.*, 2015] N. Mattei, J. Goldsmith, A. Klapper, and M. Mundhenk. On the complexity of bribery and manipulation in tournaments with uncertain information. *J. Appl. Log.*, 13(4):557–581, 2015.
- [Pini *et al.*, 2011] M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Possible and necessary winners in voting trees: majority graphs vs. profiles. In *AAMAS*, 2011.
- [Shiryayev *et al.*, 2013] D. Shiryayev, L. Yu, and E. Elkind. On elections with robust winners. In *AAMAS*, pages 415–422, 2013.
- [Viappiani, 2018] P. Viappiani. Positional scoring rules with uncertain weights. In *Scalable Uncertainty Management (SUM)*, pages 306–320, 2018.
- [Walsh, 2008] T. Walsh. Complexity of terminating preference elicitation. In *AAMAS 2008*, pages 967–974, 2008.
- [Wojtas and Faliszewski, 2012] K. Wojtas and P. Faliszewski. Possible winners in noisy elections. In *AAAI*, 2012.
- [Xia and Conitzer, 2011] L. Xia and V. Conitzer. Determining possible and necessary winners given partial orders. *J. Artif. Intell. Res.*, 41:25–67, 2011.
- [Xia *et al.*, 2011] L. Xia, J. Lang, and J. Monnot. Possible winners when new alternatives join: new results coming up! In *AAMAS*, pages 829–836, 2011.
- [Yang and Guo, 2017] Y. Yang and J. Guo. Possible winner problems on partial tournaments: a parameterized study. *J. Comb. Optim.*, 33(3):882–896, 2017.