From AI to Computational Social Choice
New Techniques, new Paradigms, new Applications

Jérôme Lang

27 juillet 2022

Résumé

Computational social choice is an interdisciplinary field at the crossing point of economics, AI, and more generally computer science. It consists in designing, analysing and computing mechanisms for making collective decisions, with various application subfields such as voting, fair allocation of resources, participatory budgeting, selecting representative groups of individuals, or matching with preferences. Since the mid-2000s, The AI community has played a major role in the (computational) social choice research output. Not only it has lead to developing algorithms for collective decision making, but it has also helped reshaping and revitalising the field, by identifying *new paradigms, new problems, new objects of study and applications*. I give below a (non-exhaustive) series of examples showing the role of AI in the study of collective decision making. I also give some thoughts about what should come next.

This paper goes with my IJCAI-22 invited talk. It is written mostly for people outside the computational social choice community.

Collective decision making

(Classical) social choice is the science of *designing and analysing methods for collective decision making.* It is usually seen as a subfield of microeconomics, for reasons that are partly historical. Here are examples of what we mean by collective decision making :

1. electing a president or a parliament.
2. finding a date for a meeting, or a series of meetings.
3. deciding which movie to watch together tonight.
4. choosing a set of public projects to fund.
5. deciding how to divide a public resource such as the budget of a country, or the time of a talk to be divided on a set of topics.
6. in a high school or a university, deciding who gets which class.
7. deciding which Covid patients will get a ventilator.

8. assigning students to university programs.
9. in a company : finding a partition of employees in groups of people who will work together.
10. a jury agreeing on a verdict.
11. in crowdsourcing : aggregating labels given by different individuals.
12. aggregating ranked lists of web pages given by different search engines

In examples 1-9, we aggregate *preferences*; they are completely subjective and there is no ground truth. In examples 10-12 we aggregate *opinions, judgments, beliefs* about a ground truth ; while these opinions are subjective, the ground truth (guilty or not guilty ; location where this picture was taken ; to a lesser extent, relevance of a document to a query) is objective. *As most of social choice is about agregating preferences, in most of these notes, and unless explicitly said otherwise, I will assume this is the case.*

We start with a set of n agents, denoted by N, who have preferences about a set of alternatives. A *profile* is a collection of preferences, one for each agent. A *social choice rule* maps a profile into an alternative, which will thus be the collective decision, ready to be implemented.

input :	agents' preferences over possible alternatives
output :	an alternative

Several keywords have been intentionally left vague : Who are the agents ? What are the alternatives ? How are preferences expressed ?

Agents They are typically individuals involved in the collective decision, and who are generally concerned with the output of the decision. [1] Agents can also be more or less structured groups of individuals (such as companies or countries) ; the early literature on computational social choice often mentioned the possibility autonomous agents (algorithms) acting for the sake of an individual or a group of individuals (this has been more or less abandoned since then ; see the conclusion).

Alternatives They significantly vary with the type of social problem at hand. The alternative space is the set in which a decision has to be selected. Let us come back to our list of typical examples and specify the set of alternatives in each case :

- The president election version of Example 1, and the single-meeting version of Example 2, are instances of *single-winner voting* : alternatives are *candidates*.

1. There may be other individuals, who are concerned as well with the decision, but who do not take part to the collective decision. For instance, laws are collectively decided by members of parliaments or other representatives (who are usually elected by the people concerned with these laws, although there are exceptions to this rule).

- Example 3 is still an instance of voting, but the space of alternatives has a combinatorial structure : alternatives are tuples from a Cartesian product, consisting of a place and a date.
- The parliament version of Example 1, and Example 4, are instances of *multi-winner voting* : alternatives are *sets of candidates* generally satisfying some constraints (the size of the parliament is fixed ; the number of meetings is fixed and should satisfy some temporal constraints like at least one week between two meetings ; the set of selected projects should not go over the budget limit).
- Example 5 is an instance of *portioning* : alternatives are, mathematically speaking, probability distributions that say which proportion of the resource should be spent on each item.
- Example 6 is an instance of *fair division* : alternatives are *assignments from resources to agents*, or equivalently, from agents to sets of resources ; here resources are classes. Each agent typically gets several resources.
- Example 7 is also an instance of fair division, but with the specific contraint that each agent gets at most one resource : it is also an instance of *two-sided matching*.
- Example 8 is also an example of two-sided matching ; unlike in Example 7, agents of one side (students or university programs) have preferences over agents of the other side. An alternative is here a one-to-many matching between students and university programs ; some students may remain unmatched and some universities may not fill their quota.
- Example 9 is an instance of *coalition structure formation* : alternatives are *coalition structures*, that is, *partitions of agents into groups*. Coalition structure formation can be seen as a generalization of matching.

Examples 1-5 are all instances of *public decision making*, or more specifically, in the case of Examples 4 and 5, *decision making with public goods* : all agents are concerned with the global decision. Examples 6-9 are instances of *collective decision making with private goods* : each agent is (generally) concerned with one part of the alternative, that is, her share, her allocation. [2]

Examples 1-9 are all instances of *preference aggregation* : agents have preferences over alternatives, and will be more or less satisfied by the outcome. Examples 10-12 are instances of *epistemic social choice* : what has to be aggregated had nothing to do with preferences but with beliefs, opinions, or judgments about the true state of the world, usually called the *ground truth*.

Preferences Preferences (and also opinions or beliefs, for epistemic problems) can have various forms :

- *dichotomous* : each agent partitions the set of alternatives between good (approved) and bad (disapproved) alternatives (or, in epistemic contexts : alternatives that she finds plausible and those she does not).
- *ordinal* : each agent gives a preference relation over the set of alternatives.

2. There are however exceptions to this ; in some allocation problems, agents may be concerned about how the goods they do not receive are distributed among other agents.

– *cardinal* : each agent evaluates each alternative with a valuation.

Preference relations can be linear orders (strict transitive orders, also called rankings), or more generally weak orders, allowing for indifferences. For cardinal preferences, valuations can be numerical (in which case the function that maps each alternative to a number quantifying the satisfaction of the agent is called a utility function, as in decision theory), or more generally taken from an ordered scale, for instance a qualitative scale such as {very good, good, etc. } Note that dichotomous preferences are a special case of both ordinal and cardinal preferences. [3]

Given a set of agents, each of them with their preferences, the collection of all preferences is called a *preference profile*. *Resolute social choice rules* map each preference profiles to an alternative, which is the collective decision. Sometimes we need to consider *irresolute social choice rules* (also called *social choice correspondences*), mapping each preference profile to a nonempty set of alternatives (typically, to cope with ties) ; and *aggregations functions*, (also called *social welfare functions*), mapping each preference profile to a collective preference relation or a collective valuation function. In voting settings, social choice rules are more commonly called *voting rules*, or simply, *rules*.

How do we study social choice rules ?

1. They must obviously be first *designed*, and sometimes adapted to the problem at hand, in which case I say they will be *engineered*.
2. Next they have to be *studied axiomatically* (which properties do they satisfy ?). Axioms give guarantee on the behaviour of the rule. The classical social choice community has focused on a set of wellk-discussed axioms, that come from economics, such as (Pareto-)efficiency, strategyproofness, and various notions of fairness. They can be seen as a way of ensuring some ethical guarantees.
3. They have to be *implemented* : they should come both with *algorithms* that computes them, but also with *communication protocols* that specify how and when agents report information about their preferences.
4. Finally, they have to be *tested in the real world*, and *analyzed ex post* : how do they perform in practice ? How do people react to them ? How can they be improved ?

I will come back to these four steps but let me take now a small detour and give a very rough and brief history of social choice. A more complete (and more serious) history can be found here.

The early stage of social choice is usually considered to be around the French revolution, with Condorcet and Borda (although, of course, voting had been practised since the Antiquity).

3. However, although a weak order can of course be induced from an evaluation function, ordinal preferences are *not* a special case of ordinal preferences : valuations give information that cannot be ignored !

The second period starts with the birth of *modern*, social choice, around 1950 : mathematicians and economists played the most important role. Social choice was formalized at that time ; the results that came in that period were mainly *axiomatic*. Some of these results — including the most well-known ones — were negative impossibility theorems showing the incompatibility of a small set of seemingly innocuous conditions, such Arrow's theorem. [4]

Some more positive results were characterization theorems, showing that the set of rules satisfying a given set of properties are all rules of a given interesting class. [5] In any case, implementation [6] issues (such as algorithms and communication protocols) were neglected, or considered trivial or uninteresting.

The third stage of social choice came around 1990, with computer scientists taking a deep interest in collective decision making. The idea was at the time to using computational notions and techniques for solving or studying complex collective decision making problems. These computer scientists came from Operations Research, Theoretical Computer Science, and inreasingly from Artificial Intelligence. This was the birth of *computational social choice* (COMSOC), even if the word was probably not used until 2006. [7]

Computational social choice is well defined on this page, which I'm quoting below :

> [It] is a field at the intersection of social choice theory, theoretical computer science, and the analysis of multi-agent systems. It consists of the analysis of problems arising from the aggregation of preferences of a group of agents from a computational perspective.

Its birth and history are described in [17] (from which I'm borrowing a lot here) and, more recently, were perfectly told in Edith Elkind's 2021 IJCAI plenary talk. Interestingly, the first papers published at the turn of the 90's were of two different kinds : a computational study of a voting rule, namely the *Kemeny* rule, which was at the time proven to be NP-hard. [8] ; and a new paradigm : *computational resistance to strategic behaviour*. For the latter, high computational complexity was thought of being positive, since it was thought to be a barrier to strategic behaviour (if a computer has a hard time finding it, humans can only do worse). One of the two papers related to this new paradigm concerned *voting*

4. To give the reader a flavour of what an impossibility theorem looks like, let us state Arrow's theorem — without any explanation :

With at least 3 candidates, an aggregation function satisfies unanimity *and* independence of irrelevant alternatives *if and only if it is a* dictatorship.

5. Of course, Arrow's theorem can also be seen as characterisation theorem since it gives a characterization of dictatorships ; but dictatorships are not good enough so that this result can be considered as a (positive) characterization result.

6. In case an economist reads these notes, a warning : I use the word *implementation* with its computer science meaning, not its game theory meaning.

7. Still, one can argue that the first paper that explicitly described an algorithm for collective decision making was Gale and Shapley's stable matching algorithm, published in 1961. Around the same time, the first formalized interaction protocol for collective decision making was the Dubin-Spanier "algorithm" for cake cutting. See Section 1.2.3 of [17] for a discussion.

8. Its precise complexity was not settled until 2005.

manipulation (strategic voting), and the other one, a notion introduced in the paper, called *voting control* (strategic behaviour of the chair). *This distinction between using computation tools to solve (or assess the computability of) known problems, and the emergence of a new research paradigm, started at that time and never stopped.* Slightly more generally, what I claim here, and want to illustrate on various examples, is that what computer science and especially AI have brought to social choice can be roughly classified in three categories :

- *new techniques* (including algorithms, but not only) :
- *new research paradigms* ;
- *new classes of collective decision problems*, *new application domains*, *new objects ot study.*

In the 1990s and 2000s research in computational social choice was mainly *top-down* : the settings were general and abstract, and the path to real-world implementation was quite long and nontrivial. Now, for the last 10 years we've seen a flurry of *bottom-up* approaches, starting from practical real-world problems and seeing what computational social choice can do for it. Examples among others are participatory budgeting, liquid democracy, bidding-enhancing mechanisms for conferences, construction of citizen assemblies with diversity constraints, crowdsourcing (and more). An exception is matching, where bottom-up approaches came earlier.

In the rest of these notes I will call a *(social choice) context* (or class of problems) simply a class of pairs (input, output) restricting the set of social choice rules of interest. Although new techniques don't necessarily intend to solve problems in new contexts, typically, new paradigms do, and new objects ot study, too.

A selection of topics

The rest of the document considers a selection of topics whose existence is mainly due to the computational social choice community. For each topic, rather than giving formal definitions I prefer to give examples of scenarios.

Caveats

1. The list of topics is non-exhaustive. There are important topics that I omit : some because they are classical social choice topics for which I can hardly claim that they would not have been studied without the COMSOC community ; some because they are too recent and have led to few development until now ; some because I don't know them well ; some simply because I forgot.
2. Obviously, giving exhaustive lists of references would be far too long. These topics, altogether, represent more than 1,000 research articles. I prefer to give *key references* : surveys or books when there are some, paper(s) introducing the topic for the first time, latest paper(s) containing a comprehensive related work section. However, at the risk of not being consistent,

sometimes I will cite a precise reference. Please tell me if some omissions are particularly unfair or inconsistent and I will add these references.

3. There will be a bias towards the research trends I have been active in, and to the papers I was a coauthor of. What you are reading now is not a publication but an informal paper that comes with a talk. To give more convincing arguments I will give prioriy to works that I know well, and often, works I know best coincide with those I took part to (although there are exceptions for both directions !).
4. The order of topics is not random but not clearly explainable either. They come more or less by *chronological* order, adapted according to my argumentation needs.
5. This is a *living paper*. If, in spite of the warnings I expressed above, you think I made a gross omission, a wrong interpretation of a research trend or of a specific work, or more generally if you have suggestions or critiques, please write me and I'll update the text (and will thank you in the end).

Computing voting rules

The context is here the most classical one : the input is a classical profile – collection of rankings – and the output is either a candidate, a nonempty subset of candidates (to handle ties without breaking them), a ranking over candidates, or a nonempty subset of such rankings.

Computing voting rules [9] is the most obvious role that computation can take, and indeed it came first, with [10], the first paper establishing that some voting rules (Kemeny and Dodgson) were NP-hard. Around the same time, and independently, a team of researchers were also exploring algorithms for computing Kemeny's rule and other median orders [7]. This topic remained still until the late 1990s and then considerably took off soon after 2000 and reached its peak during the 2000 decade. Most of it is reviewed in three chapters of the Handbook of Computational Social Choice (referred to "The Handbook" from now on) : [16] for tournament solutions, [38] for weighted tournament solutions, and [21] for the Dodgson and Young rules. The papers on the topic can be roughly classified in the following categories :

1. *Identifying the computational complexity of winner determination.*

   Most of the common rules had a well-identified complexity by the end of the 2000 decade, although the Slater rule in the general case resisted until very recently [48]. *Winner determination* is the decision problem consisting in checking whether a given candidate is among the winners for a given profile (the voting rule being fixed). Roughly speaking, voting rules can be classified in three classes :

9. I will mention here many voting rules but I am not going to define them rules, neither to give references to specific work (except two). Please refer to the Handbook, especially Chapters 2, 3, 4 and 5.

(a) those for which winner determination is polynomial-time computable : this includes all positional scoring rules, many Condorcet rules such as Copeland, minimax, Schulze, as well as plurality with runoff and the "immediate tie-breaking" versions of single transferable vote (STV) and Ranked Pairs.

(b) those for which winner determination is NP-complete : this includes the Banks rule, and the "parallel universe" versions of STV and Ranked Pairs.

(c) those for which winner determination is above NP. For the Slater, Kemeny, Dodgson and Young rules, winner determination is Θ_2^p-complete.

The difference between classes (b) and (c) has a important practical impact : when winner determination is "only" NP-complete, there exists a succinct certificate – given by the computer – that allows the voters to verify the result of the rule. Such a succinct certificate does (likely) not exist for rules of class (c), which means that if the proof of the result was to be published in the newspapers the day following the election, these would possibly be thousands pages long (not a very good idea).

2. *Finding tractable subclasses*

   Such subclasses are typically obtained by assuming standard domain restrictions such as single-peakedness, or generalizations thereof.

3. *Identifying the source of hardness by exploring paratemerized complexity*

   Parameterized complexity allows to identify the source of complexity. Typically, some hard rules (but not all) become easy when the number of candidates is bounded by a constant. The study of voting rules under the lens of parameterized complexity has been surveyed in [32].

4. *Search algorithms*

   An important body of work has been done for the Kemeny rule, less so for other rules. In the 2010 decade, some papers considered using generic solvers for computing voting rules.

5. *Polynomial approximations*

   When defining a polynomial approximation algorithm, the first question to ask is that of the score function to be approximated. For all voting rules defined by the minimisation or maximisation of a score, the choice looks simple enough. This applies in particular to the Kemeny and Dodgson rules.

   A polynomial approximation algorithm for winner determination with respect to a hard rule can also be considered *a genuine rule on its own.* It has been remarked that these approximations are often normatively as desirable as the original rule, *and sometimes even better* : there are for instance approximations of the Dodgson rule that not only are polynomial-time computable but also satisfy an important social choice property, such as monotonicity, that the original rule does not satisfy [22] and Section 6.2 of [21].

The role of computer science for this trend looks obvious. Note however that it goes much beyond simply partitioning rules between easy and hard ones, and providing algorithms for winner determination. The distinction between rules that have succinct certificates and those that are likely not to needed a fine-grained complexity study, referring to complexity classes beyond NP. Algorithm design has benefited from the expertise of several communities such as parameterized complexity, heuristic search, and combinatorial optimization. Finally, computational studies have led to desining new rules, or new variants of rules, such as the parallel universe variants of iteration-based rules, or socially desirable approximations of hard rules.

Computational resistance to strategic behaviour

This is a new paradigm and as such introduces new contexts, that vary slightly with the notion of strategic behaviour at hand. For constructive manipulation by a single voter, for instance, the input is a profile and a distinguished candidate, and the output is the vote of the manipulator (or "failure" if no suitable vote exists).

This trend was also one of the chronologically first. The key idea is that if finding a successful strategic behaviour is computationally hard, then it is is likely that humans won't be better than the computer and will (often) not find it, or not even look for it. The three seminal papers [9, 8, 11] defined two research directions that reached their peak in the 2000s and are still active : a first one on *manipulation by voters*, which is surveyed in [28], and a second one on a new family of strategic behaviours, *control by the chair*, which together with variants and extensions is surveyed in [36].

A notable difference with the previous trend is that it *created a totally new paradigm*, which is often seen as the first new research paradigm launched by the COMSOC community.

Using NP-hardness as a criterion for evaluating the resistance to strategic behaviour has sometimes been criticized as being a worst-case notion : finding a manipulation may be hard in the worst case while easy in most instances, which unfortunately, is typical ; see Section 6.5 of [28]. Still, the computational hardness of finding a manipulation gives an insight on the structure of the problem and seems to be at least partly correlated to the cognitive hardness of the problem : for instance, Single Transferable Vote, which was the first rule to be identified hard to manipulate, is known to be also hard to manipulate by humans, who usually give up trying.

The computational hardness of manipulation should not be confused with the frequency of manipulability, measured as the fraction of profiles for which there exists a manipulation ; the most recent work to date (including an extensive survey) is [33].

Computational resistance to strategic behaviour has also been devoted some attention (though considerably less) for fair division as well as judgment aggregation ; for the latter see the survey [12].

Clearly, notions from computational complexity were crucial for this new research paradigm. It went however much further than identifying the complexity of various forms of strategic behavour : it also contributed to study new forms of strategic behaviour, that had been neglected by economists. The relation between comptational resistance and cognitive resistance to strategic behaviour has not really been studied yet, and it is (I think) a promising topic.

Defining and computing fairness notions for the allocation of indivisible goods

Context : the input contains the preferences of every agent over combinations of goods ; the most common form consists of a valuation matrix where v_i^j is i's valuation for item j. The output is generally an allocation.

Consider a set of indivisible goods $\{g_1, \ldots, g_m\}$ and n agents (generally, $n < m$) who have preferences over bundles of goods they can receive. The two key questions are, how can we define a *fair* (and yet *efficient*) allocation, and how can we compute it ?

The trend took off in the mid-2000s, with the first papers on computing maxmin allocations as well as computing envy-free allocations. Three surveys written ten years later are [14, 49, 53] ; however, the trend considerably expanded since then, with more than 200 papers published since 2016.

There are two families of methods for defining a choice of an allocation : those based on the optimisation of a numerical function, and those on the satisfaction of some qualitative (binary) fairness criteria.

In the first family, we find rules that are defined by maximising social welfare : agents have numerical utility functions that associate a valuation to each possible bundle of goods, and the social welfare of an allocation is the aggregation of the utilities of all agents. In the mid-2000s there was a focus on *egalitarian* social welfare, under the name *Santa Claus problem* : an egalitarian allocation maximizes the satisfaction of the least satisfied agent. The problem is NP-hard ; in the mid-2000s there were a series of papers about approximating it in polynomial time (which turns out not to be possible except if some restrictions on the utility functions are made) ; see Section 12.3.1 of [14]. The *Nash* social welfare, defined as the product of (positively-valued) utilities of the agents, is now recognized as a very good trade-off between fairness and efficiency ; the seminal paper appeared in 2016 [23] and there were lots of follow-up papers in the last few years, including papers addressing its computation and approximation.

In the second family, we start from an efficiency criterion — the most two commons are Pareto-efficiency and its weakening consisting of completeness (all goods are assigned) — and a fairness criterion. In the 2000s the common fairness criterion considered was *envy-freeness* : no agent should prefer another one's share to her own. However, a Pareto-efficient (or even complete) and envy-free allocation is not guaranteed to exist. Early works on this trend addressed the computation of efficient and envy-free allocations, for cardinal or ordinal preferences. It became then soon well-accepted that if one of both criteria had to be weakened, it should be envy-freeness (though there are a few recent papers that explore the other way round). A seminal paper was [52] : considered defining degrees of envy and minimizing them (which was the common way of relaxing envy-freeness at the time) and introduced *envy-freeness up to one good* (EF1), which is weak enough so that it is guaranteed to exist for additive valuations. An intermediate notion was then introduced : *envy-freeness up to any good* (EFX) ; up to now it is an open problem whether its existence is guaranteed for additive valuations, and there is a considerable amount of work around this question. Another relaxation of envy-freeness is *maxmin fair share fairness* (MMS), introduced in 2011 as a discrete counterpart of the the "I cut you choose" principle in the divisible setting. It became soon evident that the maxmin fair share was not guaranteed to exist under arbitrary valuations, guaranteed to exist for two agents or for several agents with identical preferences, but it took some time to find out that it was not guaranteed for additive valuations. Researchers went further and looked for quantative relaxations of MMS and EFX, with the aim of finding maximally strong notions, that is, that are exactly on the edge of being guaranteed to exist.

Example 1

	a	b	c	d	e
Ann	15	3	**2**	**2**	6
Bob	7	**5**	5	5	**7**
Carol	**20**	3	3	3	3

We assume that agents have additive *valuations : the value that an agent gives to a subset of items is the sum of the values that she gives to each of them. For instance,* $v_{Bob}(\{b,e\}) = v_{Bob}(b) + v_{Bob}(e) = 5 + 7 = 12$.

Ann prefers Bob's share $\{b,e\}$ *to her own* $\{c,d\}$ *: the blue allocation is not envy-free. Worse, there is no (complete) envy-free allocation for this instance.*

The blue allocation is however EF1 : Ann no longer envies Bob if we remove one good from Bob's share : $v_{Ann}(\{b,e\} \setminus \{e\}) = 3 \leq v_{Ann}(\{c,d\}) = 4$ *; Ann no longer envies Carol if we remove one good from Carol's share :* $v_{Ann}(\{a\} \setminus \{a\}) = 0 \leq v_{Ann}(\{c,d\}) = 4$ *; and Bob and Carol do not envy anyone.*

This allocation is not EFX because Ann still envies Bob if we remove b from Bob's share : $v_{Ann}(\{b,e\} \setminus \{b\}) = 6 > v_{Ann}(\{c,d\}) = 4$*. Here is one that is EFX :*

	a	b	c	d	e
Ann	15	**3**	2	2	**6**
Bob	7	5	**5**	**5**	7
Carol	**20**	3	3	3	3

On Figure 1 I depict (almost) all considered relaxations of envy-freeness (some of which are described above, some aren't) together with their implication relations. The last survey to date on this precise topic is [2].

It must be stressed that although some work has been done on general valuations, most of the focus has been bearing on additive valuations, perhaps

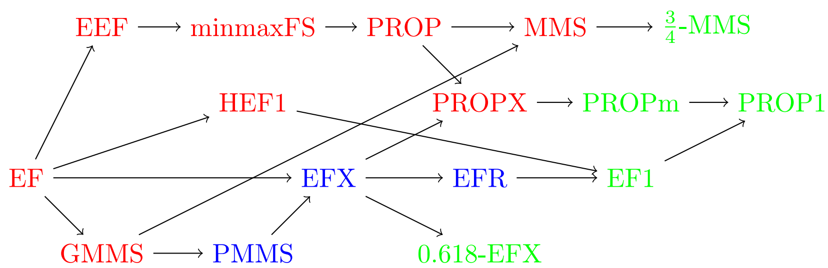

FIGURE 1 – Implications between several relaxations of envy-freeness. In red, relaxations that are not guaranteed to exist for additive valuations ; in green, relaxations that are not guaranteed to exist for additive valuations ; in blue, relaxations for which the question remains open.

because it is so simple and yet raises intriguing and hard mathematical and computational questions. It has to be put in contrast with *combinatorial auctions* [29], where (almost) all the focus bears on nonadditive valuations.

Epistemic voting and crowdsourcing

Epistemic voting is not a novel research paradigm – it started two centuries ago ; and for this reason the context is not new : a classical profile as input, and as output, an alternative, a set of alternatives, or a ranking over alternatives. But it deals with new application domains (such as crowdsourcing) and uses new techniques mostly from machine learning.

The most common view of social choice is *preference aggregation* : there is no ground truth, the state of the world will be decided by the outcome of the votes. The future president is undefined before the election, there is no "true" allocation of teachers to classes, matching between students and programs, or sets of public projects to be built. In contrast, *epistemic social choice* relies on the assumption that there is a ground truth, that votes are noisy reports about it, and the voting rule aims at uncovering it.

Epistemic social choice has a very long history : it started with Condorcet's jury theorem in 1785, which was reinterpreted and formalized by Young in 1988, and came into the computational social community in 2005 [26]. The only survey I know is Section 8.3 of [34]. The work that was done until then was mostly theoretical : some known voting rules were characterized as maximum likelihood estimators for some noise models, and the (sometimes new) rules corresponding to some specific, particularly relevant noise models were identified.

What distinguishes epistemic social choice from preference aggregation is precisely the existence of a ground truth, which allows for using statistical learning methods. In particular, it can be used for *crowdsourcing* : for instance, voters express their beliefs on the true label, under the form of a single label, a set of possible labels, a ranking over labels, or a probability distribution over labels. They may have different reliabilities, which may be known beforehand or inferred from previous data, or sometimes from the profile itself. The ground truth may also have different forms : one true alternative, a true set of alternatives, a true ranking over alternatives. And as a matter of fact, the crowdsourcing literature sometimes uses voting in a non-principled way to aggregate opinions given by different workers. So it did not come as a surprise that epistemic social choice jumped into crowdsourcing and started to confront real data. It started with [35], with linguistic annotations as an application domain. One of the reasons why social choice is useful here is that it comes with a variety of voting formats and voting rules, together with reasons to use them (or not).

As an example, one can think of asking voters not to report a single label but a *set* of labels that they can consider as plausible answers, and then use a voting rule that takes such approval ballots as input. This has been explored in a few recent papers [61, 24, 66, 1].

Example 2 (Approval-based crowdsourcing)

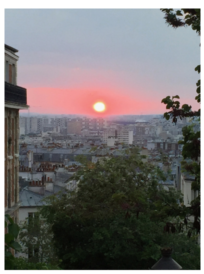

In which of the 20 districts of Paris was this picture taken ? You may give several answers. You will get a reward if your selection contains the true answer, minus a penalty that increases with the size of your selection.

Suppose now the participants answer as follows :

	12	*13*	*14*	*15*	*16*	*17*	*18*	*19*	*20*
Ann							+		
Bob			+		+			+	+
Carol		+		+		+		+	
David							+		+
Eva	+	+	+	+	+	+	+	+	+
Fred	+								
Gloria					+		+	+	+
#	*2*	*2*	*2*	*2*	*3*	*2*	**4**	**4**	**4**

The intuition is that someone who gives few answers (Ann, Fred, to a lesser extent David) is more likely to be knowledgeable (after all, if Ann knows *that the picture was taken in the 18th district, why would she tick other labels ?) than someone who gives many answers. Standard approval voting would declare three winners : 18, 19 and 20 ; but giving more weight to ballots of smaller size leads to output only 18 (which is the true answer !).*

	12	*13*	*14*	*15*	*16*	*17*	*18*	*19*	*20*	*expertise ?*
Ann							+			*high*
Bob			+		+			+	+	*med−*
Carol		+		+		+		+		*med−*
David							+		+	*med+*
Eva			+	+	+	+	+	+	+	*low*
Fred	+									*low !*
Gloria					+		+	+	+	*med−*
#							•			

There is something more we can say from this instance : Fred, who pretends to know the answer (which is likely to be incorrect), is perhaps not very reliable, and his opinion should perhaps be discounted on other instances. Such an interaction between label's plausibilities and voters' reliabilities is studied and experimented on real data in [1]. Other similar recent works : [44, 66, 65].

Epistemic social choice and applications to crowdsourcing use tools from statistical machine learning.

Low-Communication Social Choice and Distortion

New research paradigms, new techniques

In this section I'll focus on voting rules but the questions I discuss are relevant to all of social choice problems, notably fair division.

Communication Complexity and Voting Protocols

A voting rule maps a preference profile to an outcome (a candidate). But it does no say how preferences are elicited from voters so as to arrive to the outcome. Determining the outcome of a voting rule can be seen as a distributed computation problem : voters have private knowledge (their preferences) and have to interact in some way to arrive at the final outcome. Although purely distributed protocols can be envisaged, most studied voting protocols consist in a central authority interacting with voters.

Consider the Borda rule. An obvious communication protocol (Protocol 1) that finds the outcome consists in asking all voters to report their full preference ranking ; this of course works, but required $n \log(m!)$ bits to be transferred from voters to the central authority ; if m is significantly large, this might be time-consuming for voters, and worse than that, some of them may feel at ease or even unable ranking so many candidates, especially when it comes to candidates in the middle of the ranking.

Consider Protocol 2, that asks all voters to name their best candidate, then their second best candidate, and so on until the winner is known. Assume $A = \{a, b, c, d, e\}$, $n = 4$ voters whose best candidates are respectively a, a, b, c. There is not enough information for determining the winner : a may of course win, but b may win too (e.g., if it is ranked second in votes 1,2,4 and a is ranked last in vote 3). Now we ask voters to give their second best candidate. The new partial profile is (ac, ac, bc, ce). c will have a Borda score 13, while a, b, d and e will have a score at most 12, 10, 8, 9, respectively : c is a *necessary winner*, and we can stop the elicitation process at that point. While Protocol 2 has a much better communication complexity than Protocol 1 for many profiles, still there are some profiles for which it does not do better : its worst-case communication complexity is still $n\log(m!)$, but its average communication complexity is lower.

It is known that it is not possible to have a protocol that computes the exact Borda winner with worst-case communication complexity less than $n\log(m!)$ [27]. Let us consider Protocol 3 that asks voters to give only their top 2 candidates and find a winner with best expected Borda score (giving $m-1$ points to a candidate ranked first, $m-2$ to a candidate ranked second, and $\frac{m-1}{2}$ points to any candidate not ranked among the first two). It requires only $O(2n.\log m)$ bits to be sent by the voters ; the winner may not be the correct one, but it can be easily proven that the ratio between the Borda score of the Borda winner and the Borda score of the candidate we compute from such truncated ballots is, in the worst case, is in the worst case in the order of $\frac{m}{2}$.

So there is a trade-off between communication cost and quality : the cheaper the protocol, the lower the optimality guarantee. At one extremity, protocols that possibly need all information, such as Protocols 1 and 2. At the other extremity, the void protocol that elicits nothing and returns a constant outcome, which is extremely bad in the worst case. [10] Inbetween, protocols like Protocol 3, that offer a trade-off between communication cost and the quality of the outcome.

Low-Communication Protocols and Distortion

The protocols we discussed so far were designed so as to determine the outcome, or to approximate it, for a fixed voting rule. One may also want to *design new voting rules that are specifically tailored for coming with low-communication protocols* ; their main interest is that their communication is very low, so that they can be easily implemented in low-stake contexts. Now, these protocols have to be evaluated regarding the quality of their outcome. Unless one aims at approximating a given voting rule, the classical evaluation notion is *distortion*, most usually defined as the loss of utilitarian social welfare (the sum of voters' utilities for the chosen outcome) implied by the use of a low-communication rule. Two key references on voting distortion : the original paper introducing it [60] and the recent survey [3].

A first interpretation of distortion is the *price of ordinality* : if we had access to voters' utilities, one would be able to make an optimal decisions, but eliciting thrse utilities is costly, so if we elicit ordinal information instead, what do we lose in the worst case ? The distortion is infinite if no assumption is made on the voters' utilities (simply consider two candidates a, b and three voters, two

10. Randomized rules allow to do much better : the void protocol that elicits nothing and returns a random candidate with uniform probability has an expected score $O(\frac{nm}{2})$, leading to a ratio in the order of 2.

with utilities ε for a and 0 for b, and one with utilities 1 for b and 0 for a) ; if utilities are normalized (such that their sum over all candidates is constant across all voters) then the optimal worst-case distortion or deterministic voting rules is $O(m^2)$ (and is achieved for instance by plurality). If randomized voting rules are allowed, much better bounds can be obtained.

Things become even better in *metric social choice*, where candidates and voters are located in a common metric space, the disutility of a voter for an alternative being the distance between them. A simple example shows that the worst-case distortion cannot be less then 3 : consider the metric space $[0, 1]$ with distance $d(x, y) = |x - y|$, two candidates a, b positioned respectively at 0 and 1, a fraction $\frac{1}{2} + \varepsilon$ of the voters positioned at $\frac{1}{2} - \varepsilon$ and the others positioned at 1 : the majority winner is a ; the ratio between the social welfares of b and a approaches 3 when ε tends to 0.

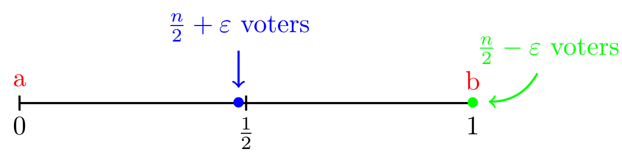

The question whether this upper bound could be reached turned out to be extremely complex ; it has been answered positively in 2021 [42] ; however, the rule used in that paper is quite complex and requires the voters' rankings to be fully elicited in the worst case. A much cheaper rule with the same distortion is given in [45] : *plurality veto*. Each candidate starts with a score equal to his plurality score. These scores are then gradually decreased via an n-round veto process in which a candidate drops out when his score reaches zero. One after the other, voters decrement the score of their bottom choice among the standing candidates, and the last standing candidate wins. This rule makes two queries to each voter, so its communication complexity is only $O(n \log m)$.

Plurality-veto :

– for each alternative x, let $s(x)$ be its plurality score.
– we fix a sequence of $n - 1$ voters
– at each step the designated voter decrements $s(x)$ where x is her worst alternative such that $s(x) > 0$
– the remaining candidate after $n - 1$ steps is the winner

Example 3 *Four candidates, six voters.*

Ann	$a \succ b \succ c \succ d$
Bob	$a \succ c \succ d \succ b$
Carol	$b \succ c \succ d \succ a$
David	$b \succ c \succ a \succ d$
Edith	$c \succ d \succ b \succ a$
Fred	$d \succ c \succ b \succ a$

Suppose the sequence is (Ann, Bob, Carol, David and Edith). Then we have the following

$$\begin{array}{rl} & (a:2, b:2, c:1, d:1) \\ \to_{Ann} & (a:2, b:2, c:1, d:0) \\ \to_{Bob} & (a:2, b:1, c:1, d:0) \\ \to_{Carol} & (a:1, b:1, c:1, d:0) \\ \to_{David} & (a:0, b:1, c:1, d:0) \\ \to_{Edith} & (a:0, b:0, \mathbf{c:1}, d:0) \end{array}$$

but if the sequence is (Fred, Edith, David, Carol, Bob) then

$$\begin{array}{rl} & (a:2, b:2, c:1, d:1) \\ \to_{Fred} & (a:1, b:2, c:1, d:1) \\ \to_{Edith} & (a:0, b:2, c:1, d:1) \\ \to_{David} & (a:0, b:2, c:1, d:0) \\ \to_{Carol} & (a:0, \mathbf{b:2}, c:0, d:0) \end{array}$$

and the winner is b.

This rule belongs to a more general class of rules (deterministic or randomized) whose main interest is that they come with a very cheap communication protocol. I give below a few other examples :

- *Triadic consensus* : at each stage, a triple of voters (i_1, i_2, i_3) is selected randomly ; each of them gives her preferred alternative x_{i_j} , and votes between the other two. If there is a three-way tie, the three voters disappear from the election, otherwise the losing voters are replaced with "copies" of the winning voter. The process is iterated until there remains only one voter.
- *Random pairs* : each voter is presented a small number of randomly chosen pairs of alternatives and chooses between them ; the alternative winning most pairwise contests is declared the winner.
- *Vote until two of you agree* : voters are selected randomly and asked to report their preferred alternative ; as soon as there is an alternative named by two voters, it is declared the winner.
- *Sequential elimination* : let σ be a (fixed or randomized) sequence of voters of length $m-1$; at each step t, the designated voter $\sigma(t)$ eliminates a candidate.

These rules all have reasonable distortion in some setting (not necessarily in the metric setting) and all have a very small communication complexity.

Distortion is heavily based on the notion of competitive ratio, which has been largely studied by the AI and algorithmic game theory communities. Low-communication protocols use notions and techniques from distributed computing and multi-agent systems.

Iterative Voting

Plurality ballots are one of the cheapest ways of expressing preferences ; however, the voting rule that naturally comes with it, plurality, is very poor. Rules based on rankings can be much better, but expressing rankings is expensive in time and cognitive effort, which might deter some voters form participating to the process, or lead other to report partly arbitrary preferences. We have given below several examples of protocols that elicit more than top candidates but less than full rankings. Another middle way is *iterative plurality voting*, whise principle is very simple, intuitive, and close to how people organize informal votes in quite many low-stake contexts : each voter starts by giving their top choice ; the scores of all candidates are made public ; voters can then change their vote ; scores are observed again, and the process (vote, observe) goes on until convergence (or a fixed deadline) is reached. This research trend was started in [55] and there is a survey [56] ; it goes much beyond plurality, but I will focus here on iterative plurality voting. Although the initial motivation for studying iterative voting was to give a game-theoretic study of strategic behaviour under such dynamic voting settings, by studying under which conditions on vote behaviour and the voting rule the convergence to a Nash equilibria is guaranteed, it can also be interpreted as a nice, intutive cheap way of attaining good outcomes (see Section 4.5.4 in [56]), which in the case of plurality voting is furthermore communicationwise cheap, provided convergence is reached quickly.

Here is a simple example ; we have 10 voters : 4 with preferences *abcde*, 3 *edbca*, 2 *cebad* one 2 *bcdae*. Initially, the plurality voting profile is $(a, a, a, a, e, e, e, c, c, b, b)$ and a is the winner.

4 voters	$a \succ b \succ c \succ d \succ e$
3 voters	$e \succ d \succ b \succ c \succ a$
2 voters	$c \succ e \succ b \succ a \succ d$
2 voters	$b \succ c \succ d \succ a \succ e$

Now the two c voters now change their vote to e ; the new profile is $(a, a, a, a, e, e, e, e, e, b, b)$ and the winner is e ; the 4 a voters now change their vote to b : the profile vote is $(b, b, b, b, e, e, e, e, e, b, b)$ and the winner is b (which happens to be the Borda winner) ; under several reasonable behavioural assumptions, convergence has been reached, votes won't change anymore. Of course, unlike low-communication voting rules, iterative plurality voting does not offer guarantees on the communication costs, but will *plausibly* be cheap (the gain is likely to be high if the number of alternatives is large and if there exists a good compromise candidate).

Iterated voting uses notions from algorithmic game theory which have been developed in part by the AI community, as well as simulation tools from multiagent systems.

Complex Alternatives, Complex Preferences

While classical social choice mostly focused on "flat" sets of alternatives, the computational social choice comunity has devoted a lot of energy to collective decision making on *structured* sets of alternatives. This research trend has been studied in several different subdomains, each of them with its specificities : combinatorial domains, multiwinner elections, participatory budgeting, judgment aggregation (and a few additional isolated papers).

Voting over combinatorial domains

The set of alternatives is defined as a Cartesian product. For instance, one can think of multiple referenda (yes/no voting on a number of issues), or a group configuration problem (decide of a common menu consisting of a first course, a main course and a wine). Voting on combinatorial domains is surveyed in [51].

Example 4 *Ann, Bob and Carol have to decide about which topics (among* t_1, t_2, t_3, t_4*) I'm going to speak in the talk. There are two groups of topics :* t_1, t_3 *(odd topics) and* t_2, t_4 *(even topics). I can talk only about two topics. Ann would have a preference for hearing about one odd and one even topic ; in addition, she is especially interested in* t_1*,* t_2 *and* t_3 *; if the selection is* $\{t_1, t_2\}$ *then she will have a satisfaction of 3 (one for* t_1*, one for* t_2*, one because she hears about both topics) ; if the selection is* $\{t_2, t_3\}$ *then she has a satisfaction of 3 as well ; if the selection is* $\{t_1, t_3\}$ *then it is 2, if it is* $\{t_2, t_4\}$*, then it is 1, and so on. Ann's utility function is not additive ; it is said to have* preferential dependencies *: the marginal contribution of a topic depends on the rest of the selection. On the other hand, Bob wants to hear about* t_3 *and that's all – he has no preferential dependencies. Finally, Carol wants to hear about* t_1 *and* t_4*, and in case* t_1 *is not selected then she'd like to hear about* t_2*. If we use the* conditional minisum approval voting *rule [6], the chosen committee will be* $\{t_2, t_3\}$*, with a global satisfaction (*social welfare*) 5. We could also use sequential ballots. The participants first vote on* t_1 *: Ann and Carol vote for, Bob against,* t_1 *is selected ; then they vote on* t_2*, given that they know* t_1 *is selected ; only Ann votes for,* t_2 *is no selected ; then they vote on* t_3*, Ann and Bon vote for,* t_3 *is selected and we return the final selection* $\{t_2, t_3\}$*.*

There are lots of other ways of handling such a collective decision over a combinatorial domain. The three choices to make are (1) the format of ballots (here, conditional approval ballots), (2) the elicitation protocols (e.g., one-shot or sequential), and (3) the rule used for determining the outcome. While everything would be rather easy without preferential dependencies, their presence lead to a lot of complications, and a trade-off must be made between the communication cost of the protocol and the quality of the outcome.

Voting over combinatorial domains makes an important use of preference representation languages developed in the KR community.

Multiwinner elections

The aim of mutiwinner elections is also to make a decision over a combinatorial domain, with one more restriction : the domain of each attribute should be binary. But the main difference is that instead of focusing on preferential dependencies and how to cope wit them, the focus here is on the normative properties and the computation of the rules. Several families of rules have been thoroughly studied ; the choice if rule primarily depends on whether one is interested in *excellence* (such as shortlisting a set of good candidates for an award), *diversity* (such as selecting a set of movies for a group, given that each member of the group will see only one movie), *representation* (such as electing a parliament). For a survey see [37].

Example 5 *We can now select three topics our of five. Here are the votes ;*

	t_1	t_2	t_3	t_4	t_5
4 voters	+	+	+		
3 voters	+	+			
2 voters				+	+
2 voters			+		+
1 voter				+	

If we simply select the topics that gather the highest number of votes (which is referred to as standard multiwinner approval voting, or MAV), then these are t_1*,* t_2 *and* t_3 *; this selection is unfair, as four attendees get their three choices while three get none at all. One may use instead a very egalitarian rule called* Chamberlin-Courant approval voting *which evaluates the score of a committee as the number of attendees who like at least one of the topics in it. Here the only way to satisfy all attendees — and thus to be optimal — is to select* t_4*, one of* t_1 *and* t_2*, and one of* t_3 *and* t_5*. But this is perhaps giving too much weight to the single-minded voter who is only interested in* t_4*...*

Finally, a choice that looks more reasonable is to use proportional approval voting *(PAV), which gives a utility* $1 + \frac{1}{2} + \ldots + \frac{1}{k}$ *to an agent who votes for* k *topics in in the selection, and maximizes the sum of utilities ; then the optimal selection is either* $\{t_1, t_2, t_5\}$*, with global score* $7\frac{3}{2} + 4 = 14.5$*. Proportional approval voting satisfies a strong proportionality property that few rules manage to satisfy, and on which I'll come back later on.*

Participatory budgeting

In participatory budgeting, we also have a set of candidates (topics, projects) and a selection must be made, but the constraint is budget-driven : each project has a cost and there is a budget limit. The greedy method, that ranks projects in decreasing number of votes and selects them one after each other if it is consistent with the budget constraint, is arguably very bad. Here are the votes for the 2018 participatory budgeting in the 4th district of Paris in 2018. [11] We

11. Thanks to Dominik Peters !

see that the three most expensive projects are selected, and needless to say that this is far from being desirable... and still this is the rule they use in Paris. [12]

	cost	votes	select ?
Amélioration de la cour du Collège Charlemagne	350 000	1 066	yes
Rénovation de la fontaine Niki de Saint Phalle	900 000	996	yes
Aménagement du préau de l'école (...)	325 000	825	yes
Aménagement des balcons terrasses de l'école (...)	150 000	807	–
Un arbre de la laïcité dans le 4eme arrondissement	5 000	756	–
Valorisation de la rue de Venise (...)	20 000	740	–
Des vidéoprojecteurs interactifs pour l'école (...)	18 000	591	–
Création de mezzanines de stockage à l'école (...)	150 000	404	–

Maximizing the global score among all budget-feasible selections is usually better— here, it is :

	cost	votes	select ?
Amélioration de la cour du Collège Charlemagne	350 000	1 066	yes
Rénovation de la fontaine Niki de Saint Phalle	900 000	996	–
Aménagement du préau de l'école (...)	325 000	825	yes
Aménagement des balcons terrasses de l'école (...)	150 000	807	yes
Un arbre de la laïcité dans le 4eme arrondissement	5 000	756	yes
Valorisation de la rue de Venise (...)	20 000	740	yes
Des vidéoprojecteurs interactifs pour l'école (...)	18 000	591	yes
Création de mezzanines de stockage à l'école (...)	150 000	404	yes

But it is not perfect either, because it does not take the costs into account when computing scores (only for defining feasible selections). There exist other rules that do better [59], but we will not enter into the details.

Participatory budgeting is a very hot topic at the moment : most papers were published in the last 5 years. Still, there is already a survey [5].

Judgment aggregation

Judgment aggregation is the third main research stream (besides voting over combinatorial domains and multiwinner elections) that is tailored to make decisions over complex, structured domains. The focus here is on the expressivity of the logical relations between the various issues voters have to vote for or against, as *any logical relation* is allowed, and not only cardinality or more generally budget constraints. Participatory budgeting has been generalized this way in [63].

12. They now use a different input with qualitative evaluations instead of yes-no votes, but the problem remains.

Portioning and Apportioning

An old topic in classical social choice is *randomized voting* : the profile is standard, and the output is a *probability distribution over alternatives.* A rule that outputs such a distribution is a *randomized rule*, also sometimes referred to under the strange name *social decision scheme.* The first motivation for using randomized voting is that it gives more possibilities for strategyproofness. [13]

Now, there is another interpretation of randomized rules that have actually nothing to do with randomization ! The output of a randomized rule can be thought of as a division of a common public resource (such as time or money) among the alternatives. We call such a division a *portioning.* So, in summary, a portioning problem consists of

– a total budget to be spent
– a set of budget items
– agents have preferences on how the budget should be divided
– and the goal is to find a budget allocation

This can be seen as (sort of) a continuous version of participatory budgeting. Examples of portioning problems :

dividing time voters are residents of a retirement home ; they have to decide which proportion of the time the common room should be used for various activities

dividing money voters are citizens of a country ; they have to decide which proportion of the budget to spend on education, health, public transport, police etc.

dividing seats voters are citizens ; their votes determine the ideal fraction of the parliament each party should get (portioning) ; these fractions are used to decide the number of seats they each get (apportionment).

Now that we know what the output looks likek we have to talk about the input. What are the possible formats of votes ? The four possibilities that have been explored are :

1. each voter names only one item
2. each voter approves a subset of items
3. each voter provides a ranking of the items
4. each voter gives an ideal division of the budget

And in all cases the output is a portioning p over the set of items X s.t. $\sum_{x \in X} p(x) = 1$.

13. While we know that any deterministic, surjective and strategyproof rule is dictatorial, for randomized rules we have far many pther possibilities : any probabilistic mixture of dictatorships and *duples* (rules whose output is contained in a subset of alternatives of size one or two) is strategyproof. A notable example is when uniformy randomizing over all majority rules between all pairs of candidates (which for simplicity I describe here for an odd number of voters) ; in other words, the rules selects a pair of distinct alternatives $\{x, x'\}$ with probability $m(m-1)/2$ and then output the majority winner.

Portioning : uninominal preferences

If each voter names only one item only, the only reaso,able thing to do is simply to give each alternative x a portion that is the fraction of voters who name it – and this is how scores of political parties are computed in party-list elections. [14]

Portioning : approval preferences

When each voter approves a subset of items, there are various possibilities. Consider the following input :

$$A = (\{p_0\}, \{p_0\}, \{p_0, p_1, p_2\}, \{p_0, p_1, p_3\}, \{p_1\}, \{p_2, p_3\})$$

One can simply assign each item a portion that is proportional to the number of approvals it gets (the *proportional* rule) :

$$p(x) = \frac{|\{i : x \in A_i\}|}{\sum_i |A_i|}$$

This gives .36 to p_0, .27 to p_1, and .18 to p_2 and p_3.

The *split* rule gives weight $\frac{1}{N}$ to each voters and splits this weight equally between all approved alternatives :

$$p(x) = \sum_{i:x \in A_i} \frac{1}{n|A_i|}$$

Here this gives .44 to p_0, .28 to p_2, and .14 to p_3 and p_4.

The *conditional utilitarian rule* [4] decides to give all the weight oa a ballot to the item from this ballot with maximal approval score : if x_i^* be the item in A_i with the largest number of approvals ; then

$$p(x) = \frac{1}{n}|i : x_i^* = x|$$

Here it gives .67 to p_0, .17 to p_1, t to p_2 and $\frac{1}{6} - t$ to p_3, for any $t \in [0, \frac{1}{6}]$. The *majoritarian portioning rule* [30] is similar to CU but computed in a greedy way.

Let $u_i(p) = \sum_{x \in A_i} p(x)$ be the total part of the share that agent i approves. The *Nash* rule selects p maximizing $\prod_{i \in N} u_i(p)$ while the leximin rule [13] selects p with the leximin-optimal vector $(u_i(p))I \in N$. Here the Nash solution gives .53 to p_0, .27 to p_1, .10 to p_2 and p_3, while the leximin solution gives .33 to p_0 and p_1, and .17 to p_2 and p_3.

	P	S	CU	$Nash$	$leximin$
p_0	.36	.44	.67	.53	.33
p_1	.27	.28	.17	.27	.33
p_2	.18	.14	t	.10	.17
p_3	.18	.14	$.17 - t$	.10	.17

14. With the exception, sometimes, that parties that do not reach a bar, typically 5 %, get a portion 0).

As an example, imagine a group of friends who want to share tapas. Ann and Bob like only patatas bravas, Carol likes patatas bravas, tortilla and calamares, David likes patatas bravas, tortilla and chorizo, Edith likes only tortilla, and Fred likes calamares and tortilla. Which share of patatas bravas should there be ? .33 seems to be rather on the low side, and .67 on the high side... The split rule seems here a good compromise.

As another example : no tapas anymore but topics I can speak of during my talk, and the share of a topic is the fraction of the total time of the talk devoted to it. Perhaps here it makes more sense to insist on the "patatas bravas" talk so that the Nash solution makes more sense ?

Portioning : ordinal preferences

Now suppose each voter provide a *ranking* of the items. Suppose there are five topics I can talk about, and five attendees listening (a small IJCAI). Their preferences are as follows :

Frances, George, Helena	$a \succ b \succ c \succ d \succ e$
Ingrid	$e \succ b \succ c \succ d \succ a$
John	$c \succ a \succ e \succ d \succ b$

We could of course pay attention only to the top choices and allocate propotionally to them, and listen to a 60% of the time, and c and e 20% each. But it does not look optimal, as b looks likes a good compromise.

Another possibility would be to define shares as proportional to Borda scores, but this would lead us to listen to d 10% of the time although everyone prefers c to d !

We could also notice that c is a Condorcet winner (it is also a Borda winner, and actually a winner for most commonly used single-winner voting rules). This would lead us to listen to a all the time !

Now we could define the disutility of the agent as the average rank of the topic she's listening to – or equivalently, define her utility as her average Borda score. This has a strong egalitarianistic flavour : no agent will be strongly disavantadged for having nonconventional preferences. Here the Borda-egalitarian solution gives b with proportion 40% and c with proportion 60%. The average Borda score for Frances, George, Helena and Ingrid is $0.4(3) + 0.6\ (2) = 2.4$; and the average Borda score for John is $0.6\ (4) = 2.4$ too. All participants have the same average satisfaction over time.

A last solution, which gives less weight to agents with outlier preferences, is the *Nash* solution, which maximizes the product of average Borda score. Here we get to listen a around 46 % of the time, to b around 29 % and to c around 26 %. The corresponding average scores of the agents are 3.23 for Frances, George and Helena, 1.39 for Ingrid, and 2.42 for John. This is better than the egalitarian solution for 4 four out of the five agents, and much better for three !

Portioning : cardinal preferences

One may think, why couldn't each of the agents simply report their preferred portioning and some aggregation method be used ? Let us suppose the optimal divisions are as follows :

	a	b	c	d	e
F, G, H	40	30	20	10	0
I	5	20	15	10	50
J	30	50	35	15	20

The most obvious idea to aggregate these vectors is the average. This would give

$$(a : 31; b : 22; c : 22; d : 11; e : 14)$$

But now remark that it is very easy for the participants to manipulate the outcome : for instance, Ingrid can report $(a : 0; b : 10; c : 0; d : 5; e : 85)$ leading the the outcome $(a : 30; b : 20; c : 19; d : 10; e : 21)$. Thius is not to say that other methods – using approval or ordinal information – are not manipulable : they are. But averaging is even easier to manipulate, and manipulation may have more effect.

If there were only two items (a and b) and each agent i expresses the ideal portioning $(x_i, 100-x_i)$ then aggregating by the median would be strategyproof. For more than two items this is less easy, but there are solutions that generalize the median [39].

Portioning followed by apportioning

In party list-based elections :

1. candidates are partitioned into *lists*
2. a list is an ordered set of candidates
3. each voter votes for one of the lists ;
4. the score of each list is the fraction of votes it receives ;
5. these fractional scores are then mapped into an integral distribution of seats : *apportionment*
6. each list fills its assigned seats with candidates in the order of the list.

For Step 3 there is a variety of rules ; however Steps 1 and 2 are rarely critically examined A natural idea is to allow inputs consisting of *approvals* or *rankings* over lists.

This has been studied for approval ballots in [20] but one could also do it for ordinal ballots.

Portioning vs. randomized voting

Technically, portioning and randomized look identical : in portioning, the output is a portioning p over the set of items X such that $\sum_{x\in X} = 1$, and in andomized voting it is a probability distribution p over the set of items X. These are the same mathematical objects ! Randomized voting (probabilistic social choice) has been surveyed in [15]. However the motivations in one-shot randomized voting and portioning are completely different : in one-shot voting, the motivation for using randomized rules is that they tend to offer more strategyproofness guarantees than resolute rules while keeping neutrality and anonymity. In portioning, there is no randomness at all, the motivaiton is to find a division of a common resource sich as time or budget. Some properties make sense for one-shot probabilistic voting but not for portioning, and *vice versa*. For instance, Condorcet-consistency, that says that if an alternative is a Condorcet winner then it should get all the resource, makes no sense for portioning !

Proportionality

When you think about proportionality, what you have in mind is probably proportional representation in party-list elections : there are k seats to be filled, there are m parties, each of them with a list of k candidates, each voter casts a vote for a list, each party p_i receives a number of votes s_i and the number of seats for each party is computed by an *apportionment* process, which consists in finding a vector of integer numbers $(t_1, \ldots, t_m)$ summing up to k and being toughly proportional to $(s_1, \ldots, s_m)$; this notion of "being roughly proportional" has several possible interpretations, each of them coming with an apportionment method, but the general principle is independent from the choice of a particular method. In classical social choice and political science, proportionality did not go beyond that.

However, proportionality is a more general principle that applies to many more contexts, and the computational social choice community has done a huge effort in the last 5 years defining proportionality principles that applies to a variety of contexts, defined by their input and their output.

The general principle of proportionality is that *a group of agents that constitutes a fraction α of the electorate should be entitled to decide about how a fraction α of the public resource is used.* In particular, if a solution of the allocation problem is S then there should not be an α-partial solution T and a subelectorate V such that

- V is a fraction at least α of the electorate
- T uses only a part at most α of the resource
- all agents in V (weakly) prefer any extension of T to S, and one of them strictly prefers T to S.

A solution that satisfies this property is *core stable*. However, depending on the setting, there might or might not be a guarantee of existence of a core stable solution.

To start with, proportionality applies to elections where the ballots are not uninominal (that is : one ballot, one party) but consists of *sets of parties* (approval voting — see Section for an example — or *rankings over parties.*

Example 6 *Let the set of candidates be $\{a, b, c, , d, e\}$, from which $k = 3$ candidates have to be selected, and $n = 100$ voters with the following approval preferences*

	a	b	c	d	e
65×	+	+	+		
34×				+	
1×					+

The candidates with highest support are a, b and c. However, the 34 voters who support d can claim that they constitute a fraction $> \frac{1}{3}$ of the electorate and since they all agree on d, they should be entitled $\frac{1}{3}$ of the decision power. Finally, if diversity is the criterion that counts, then probably we should elect a, d and e, that gives each voter one of the candidates they like.

Proportionality naturally applies to *participatory budgeting*, since it is a generalization of multiwinner elections. Here when we say that T uses only a part at most α of the resource, we mean that the cost of T is at most a fraction α of the total budget.

Proportionality also applies to *portioning*. Recall our example

Frances, George, Helena	$a \succ b \succ c \succ d \succ e$
Ingrid	$e \succ b \succ c \succ d \succ a$
John	$c \succ a \succ e \succ d \succ b$

The utilitarian solution a is not core stable because Ingrid, who is entitled to control $\frac{1}{5}$ of the resource, would prefer any extension of $0.2e$ to a. The equal split $0.2(a+b+c+d+e)$ is not core stable either because Frances, George and Helena, who are entitled to control 60% of the resource, prefer eny extension of $0.6a$ to $0.2(a+b+c+d+e)$.

Another setting where proportionality plays an important role is *proportional rankings*. Here the common resource is the set of positions on the list. A fraction α of the electorate is entitled to control $\alpha.t$ seats of of the first t seats, for all t, and prefers a list L to a list L' if the set of positions in L stochastically dominates that in L'.

Porportionality is also a key notion in long-term voting.

Finally, in *fair division of indivisible goods*, the notion of *controlling a fraction α of the resource* is quite complex, since the resource is no longer assigned to all agents (public) but divided among the agents (private). One way of understanding it is *proportional fair share* : if I am one of n agents then my share should be worth at least a fraction $\frac{1}{n}$ of the total utility I would obtained if I was given all goods.

Diversity

While forming committees in the precious section was only guided by the *votes* (expressing the voters' preferences), in some cases they should also be guided by some diversity objectives.

Offline selection

Example 7 *Suppose we have to choose $k = 4$ members for a recruiting committee. We have the following objectives on the composition of the committee :*

- *gender-balanced : ideally 50% male, 50% female*
- *ideally 50%, 25%, 25% of researchers in areas 1, 2, 3*
- *local / external members : ideally (25%, 75 %)*
- *senior / junior members : ideally (75%, 25 %)*

Now we have the following pool of candidates :

Name	*Gender*	*Group*	*Age*	*Affiliation*
Ann	*F*	1	*J*	*L*
Bob	*M*	1	*J*	*E*
Charlie	*M*	1	*S*	*L*
Donna	*F*	2	*S*	*E*
Ernest	*M*	1	*S*	*L*
George	*M*	1	*S*	*E*
Helena	*F*	2	*S*	*E*
John	*M*	2	*J*	*E*
Kevin	*M*	3	*J*	*E*
Laura	*F*	3	*J*	*L*

We can make the following selection

Name	*Gender*	*Area*	*Seniority*	*Affiliation*
Ann	*F*	1	*J*	*L*
Bob	*M*	1	*J*	*E*
Charlie	*M*	1	*S*	*L*
Donna	*F*	2	*S*	*E*
Ernest	*M*	1	*S*	*L*
George	*M*	1	*S*	*E*
Helena	*F*	2	*S*	*E*
John	*M*	2	*J*	*E*
Kevin	*M*	3	*J*	*E*
Laura	*F*	3	*J*	*L*

which reaches the gender-balance objective and the senior/junior balance but only partly the other two objectives. Or this one :

Name	*Gender*	*Area*	*Seniority*	*Affiliation*
Ann	*F*	1	*J*	*L*
Bob	*M*	1	*J*	*E*
Charlie	*M*	1	*S*	*L*
Donna	*F*	2	*S*	*E*
Ernest	*M*	1	*S*	*L*
George	*M*	1	*S*	*E*
Helena	*F*	2	*S*	*E*
John	*M*	2	*J*	*E*
Kevin	*M*	3	*J*	*E*
Laura	*F*	3	*J*	*L*

which is perfect on all objectives except gender parity. We cannot do any better!

Here the diversity on the attributes was considered an *objective* to reach or approach. These target represntation over attribute domains may come from votes (maybe from an portioning problem), or can be exogeneous, that is, fixed by the society, for instance according to some rule or law. In [50] several norms are considered and the committee closest to the target representation according to this norm was computed (which is a hard problem).

In other settings it is rather a *constraint* that must be satisfied, and the objective bears on something else, for instance votes.

Example 8 *Suppose we have to choose $k = 4$ members again. We have the following constraints on the composition of the committee :*

- *gender-balanced : 50% male, 50% female*
- *between 25% and 50 % in area 1, between 40% and 60 % in area 2, and between 10% and 25 % in area 3.*
- *at least 25% of junior and at least 50 % of senior members.*

Now we have the following pool of candidates :

Name	*Gender*	*Area*	*Seniority*
c_1	*F*	1	*J*
c_2	*M*	2	*J*
c_3	*M*	2	*S*
c_4	*F*	3	*S*
c_5	*M*	2	*J*
c_6	*M*	2	*J*
c_7	*M*	2	*J*
c_8	*F*	1	*J*

And now approval votes :

	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8
v_1	+	+	+					
v_2			+		+			
v_3							+	+
v_4			+	+	+		+	
v_5	+							+
v_6						+		
v_7	+	+				+		

The committees maximising the AV-score and the PAV score subject to the constraints are respectively $\{c_1, c_3, c_4, c_7\}$ *and* $\{c_3, c_4, c_6, c_8\}$ *[47].*

Online selection

The problem is similar except that now candidates come online and we have to decide instantaneously whether to select them or not. An example of concrete setting where this is the case is the composition of a diverse citizen's assembly with a representation of gender, age, region of origins, level of education, and occupation that are roughly proportional to their importance in the general population.

If the probability distribution on the arrival of candidates (i.e., vector of attributes) is known, the problem can be expressed as an infinite-horizon *constrained Markov Decision Process*. As the optimal policy is stationary, candidates can be interviewed in parallel.

Example 9 *We have only two attributes : gender and age. The target distribution is* $(\frac{1}{2}, \frac{1}{2})$ *for each of them (half women and half men, half senior and half junior). The distribution on arrivals is* $MS : \frac{1}{2}; SW : \frac{1}{4}; JM : \frac{1}{4})$*, that is, half of the arriving candidates are senior male, one fourth senior female and one fourth junior male. The optimal policy consists in selecting each arriving senior male with probability* $\frac{1}{2}$ *and all other candidates with probability 1.*

Liquid democracy

Quoting from [19] :

> Liquid democracy is a novel paradigm (...) that gives agents the choice between casting a direct vote or delegating their vote to another agent

It can be thought as being a sweet spot between representative democracy and direct democracy. *Direct democracy* has the advantage to give each citizen the freedom to express her opinion on any issue, and the drawback that it may incur a lot of efforts on citizens, or lead citizens to express their views on issues on which they have only little information or opinion. *Representative democracy*, where citizens choose their delegates (members of parliament, steering committee), allows for efficiency gains since citizens have to vote only once ; however, it is generally not possible for them to find someone who agrees with what they think on every single issue, and even if they find one, this person might not be elected — especially so in political elections (for example, in France there is a there is a MP for more than 100,000 persons). Liquid democracy offers the best of both worlds : a citizen who has a clear opinion on an issue will cast a direct vote, and otherwise will delegate the vote to someone they trust *regarding this issue.*

Liquid democracy

Quoting from [19] :

> Liquid democracy is a novel paradigm (...) that gives agents the choice between casting a direct vote or delegating their vote to another agent

It can be thought as being a sweet spot between representative democracy and direct democracy. *Direct democracy* has the advantage to give each citizen the freedom to express her opinion on any issue, and the drawback that it may incur a lot of efforts on citizens, or lead citizens to express their views on issues on which they have only little information or opinion. *Representative democracy*, where citizens choose their delegates (members of parliament, steering committee), allows for efficiency gains since citizens have to vote only once ; however, it is generally not possible for them to find someone who agrees with what they think on every single issue, and even if they find one, this person might not be elected — especially so in political elections (for example, in France there is a there is a MP for more than 100,000 persons). Liquid democracy offers the best of both worlds : a citizen who has a clear opinion on an issue will cast a direct vote, and otherwise will delegate the vote to someone they trust *regarding this issue.*

Liquid democracy can be used for classical preference aggregation. For instance, if I'm asked whether I want to ovte for the steering committee of an organization I'm part of, I may chose between casting a vote myself or delegate to someone I trust and whose interest are close to mine, to avoid spending some time listening to the programs of the candidates and making up my mind.

But liquid democracy can also be used for epistemic social choice, where there is a ground truth and what is to be aggregated is beliefs, opinions about this ground truth. For instance, if I am asked a question about English idioms, I will delegate to Ann ; of this is a prediction about a tennis match, to Bob ; and if this about locating a landmark in a country, I will vote because I feel confident enough. Experiments have shows that delegations improve the ability of a crowd to identify the truth [62].

Liquid democracy is at work in the real world : see `https://en.wikipedia.org/wiki/Liquid_democracy`. It is perfect example of the *reactivity* of the computational social choice community : rather than performing a top-down approach they start with a practical problem and quickly come with solutions. For liquid democracy, the number one problem comes frome the *transitivity of* *delegation* : in he example below, A and B delegate to C, whod elegates to D, who abstains so this delegation path leads to nowhere, meaning that the votes of A, B and C (and D) are lost. On the other hand, theer is a delegation cycle between E, F and H. ? What shall we do with these delegaiton cycles and these paths leading nowhere ?

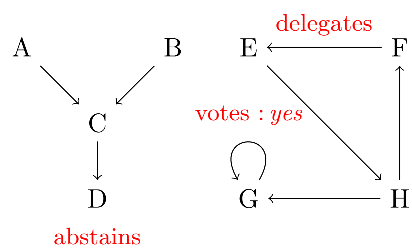

In order to cope with such issues, the COMSOC literature has come up with several solutions, such as *ranked delegations.* In the figure below, dottes edges mean second choice delegations. For instance, B's first choice is to delegate to C, but if this does not lead to anywhere or falls in a cycle then he will delegate to G instead. As for H, her first choice is to delegate to F,whom she trusts more than herself, but she also casts a vote and if there is a problem then her vote will be used.

A delegation function matches every voter to a casting voter. If the voter casts a vote then of course she is matched to herself. A good delegation function should maximize the number of matched voters, not concentrate too much power in the hands of few, and apply a good trade-off between short paths and expressed preferences. A possible delegation function could consists in favouring shortest delegation paths (breadth-first) and another one could be to favor first ranked (depth-first). Such issues are studied especially in [19, 25, 43].

Dynamicity

Long-term voting

The COMSOC research group at the University of Antarctica [15] meets once a week. Everyone enjoys these meetings ; there is an interesting seminar, gossips, coffee and sandwiches, and this is often the possibility in the week to see everybody. Until recently, the group met all Tuesdays at lunchtime ; however, it became apparent that it was heavily unfair to Ann, who has classes precisely on that time and Bob, who for personal reasons works from home on Tuesdays. So the group decided that a different meeting time should be found every single time and use the following procedure :

– every week, the set of possible slots are {M, Tu, W, Th, F}
– in the middle of the week t, members of the group express approval ballots on slots of week $t + 1$

15. `http://antarcticaedu.com/`. Apply for positions, don't wait until it becomes the last liveable place on Earth.

– and then a voting rules is aplied to select the slot of week $t+1$.

Unfortunately, if the voting rules does not take the past into account, something like that may happen : out of 20 members, 18 (all except Ann and Bob) like Tuesdays, and no other day scores better so Tuesdays will be selected each time... and Ann and Bob will never attend. A solution [40, 46] consists in using a *memoryful voting rule*, that weighs voters according to the number of times they have been satisfied in the past. This way, after a few Tuesday, Ann and Bob will have such an important weight that a slot will be selected so that at least one of them is satisfied.

Let us fix a *weighted approval rule* [46] defined by the simple weight function : initially, each voter has weight 1 ; if a voter with weight w_t has been satisfied at round t, then their weight at round $t+1$ is $w_t/2$, otherwise it is $2w_t$. The winning alternative at round t maximizes the sum of weights of voters who approve it. On week 1, the profile is

	M	Tu	W	Th	F
Ann	+	-	+	-	-
Bob	-	-	+	+	+
Carol	-	+	-	-	-
David	+	+	+	+	+
Edith	+	+	-	-	-
Fred	-	+	-	+	-

The winning slot is Tuesday. The updated vector of weights is $(2, 2, 1/2, 1/2, 1/2, 1/2)$. On week 2, the profile is

	M	Tu	W	Th	F
Ann	-	+	+	+	-
Bob	+	-	-	-	-
Carol	-	-	-	+	+
David	+	-	+	+	+
Edith	-	+	-	+	+
Fred	+	+	-	-	+

The winning slot is Thursday (score 5.5). The updated vector of weights is $(1, 4, 1/4, 1/4, 1/4, 1)$. On week 3, the profile is

	M	Tu	W	Th	F
Ann	-	+	+	+	-
Bob	+	-	-	-	+
Carol	+	-	-	+	-
David	-	-	+	+	+
Edith	-	+	-	+	-
Fred	-	+	-	+	-

The winning slot is now Monday. And so on. Perhaps the choice of weights is a bit extreme : Monday would have won even if Bob had been the only one voting for it ! With a smoother weight update function such as $w_{t+1} = {}^{3w_t}/_2$ if the agent was not satisfied at round t and $w_{t+1} = {}^{2w_t}/_3$ otherwise, the weights would have been $(1, {}^9/_4, {}^4/_9, {}^4/_9, {}^4/_9, 1)$, the score of Monday would be ${}^{13}/_4$, that of Thursday, ${}^{10}/_3$, Thursday would have won. With such exponentially increasing/decreasing weights, in the long term all agents are guaranteed to be satisfied roughly the same number of times.

So far we assumed that the preference profile at time t is know only at round t. If it is known in advance then things are very different since we can optimize — for instance maximizing Nash social welfare – off-line. If we had known everything from the beginning we would not have selected Wednesday at round 1, Monday at round 2, and Thursday at round 3 (giving Nash welfare $2.1.2.3.2.2 = 48$).

Finally, an interesting case is when the profile at each round it the same ! In that case, we are back to an apportioning problem from approval ballots – in factn not quite since here not only the *number* of weeks where the meeting will be on Monday counts but also how these weeks are disposed through the year.

Online Selection of Alternatives

A special case of long-term voting is when the decision at time t is binary decision, usually whether or not to include a current candidate into a final selection. Think of a set of human resource experts in company who have to vote for a against a current candidate, given that candidates appear over time, that each time a candidate appears it must be decided immediately whether to select it or not, and that in the end they must hire exactly k candidates. These voters can actually be (proxies for) criteria : voter 1 evaluates the current candidate according to their ability of condcuting research, voter 2 according to their ability of managing a group, voter 3 according to their programming abilities. We must hire two persons ; for each criterion, the reward of the selection is 1 if exactly one of its members satisfies the criterion, $\frac{3}{2}$ if both do, and 0 if none does. (This is the PAV score, see Subsection). This is a voting version of optimal stopping (more precisely of the multiple secretary problem, with nonadditive rewards). Things differ significantly whether we know the distribution on the vector (x_1, x_2, x_3), where $x_i = 1$ of the candidate satisfies voter (criterion) i ; for both contexts see [31].

If we know the distribution : the computation of the optimal policy can be cast as a Markov Decision Process. An example of execution :

– c_1 comes, gets evaluation $(1, 1, 0)$ and is not selected ;
– c_2 comes, gets evaluation $(1, 1, 0)$ and is selected ;
– c_3 comes, gets evaluation $(1, 1, 0)$ and is not selected ;
– c_4 comes, gets evaluation $(1, 0, 1)$ and is selected ;
– stop

Automated Theorem Proving for Social Choice

New techniques

Classical social choice aimed at proving possibility and impossibility theorems, often taking this abstract form, where the context, and thus the definition of a rule, has been defined first.

- Impossibility : If there are at least $m \geq m^*$ alternatives and at least $n \geq n^*$ voters, then no rule satisfies properties $P = \{P_1, \ldots, P_k\}$. For instance, Gibbard and Satterthwaite's theorem has $n^* = 1$, $m^* = 3$, $P =$ {unrestricted domain, surjectivity, strategyproofness, non-dictatorship}.
- Possibility : If there are at most $m \leq m_*$ alternatives and at least $n \geq n_*$ voters, then some rule satisfies properties $P = \{P_1, \ldots, P_k\}$. Gibbard and Satterthwaite's theorem obviously not holds for $m = 2$, because the majority rule is strategyproof.

Sometimes, the smallest known value of n^* and the largest known value of n_* coincide, and similarly for m^* and m_*. But it is not always the case : sometimes there are gaps, as I am explaining on this very classical paradox called the *no-show paradox*, also called the failure of participation. What we knew from [58] :

- for $m \geq 4$ and $n \geq 25$, no voting rule satisfies the Condorcet criterion (of there exists a Condorcet winner then it should be elcted) and the participation property (a voter should never be better off if she abstains than if she votes sincerely). The proof relies on a single counter-example profile with 25 voters.
- if $m = 3$ then there exists a rule that satisfies both properties (for instance, the maximin rule).

So we have a gap : is 25 is the smallest value of m for which the impossibility theorem holds, and if not, what is this minimal value ?

Answering such a question by hand calculations has no guarantee to succeed, and if so, would be extraordinarily tedious. This is where automated theorem proving plays a role. For each value of m, the problem can be formulated as a SAT instance ; if a model is found, then there is a profile that shows impossibility ; if not, then we have a possibility. And indeed it was proved in [18] that the impossibility holds for $m = 4$ and $n \geq 12$, and ceases to hold when $n \leq 11$.

This example is one of many examples showing the usefulness of such a translation to SAT. The reasons why there are so many examples are made clear in [41] :

> [Social choice theory] has three characteristics that suggest applying computer-aided reasoning to it : it uses the axiomatic method, it is concerned with combinatorial structures, and its main concepts can be defined based on rather elementary mathematical notions.

Automated social choice theorem proving is the perfect example of why using AI-based theorem proving techniques such as SAT solving or constraint reasoning can help proving theorems that would otherwise be to hard and/or too tedious to prove or disprove.

Datasets

Perhaps one of the greatest advances made by the COMSOC community is the building and maintenance of collective decision making databases. The largest database for voting data is PrefLib.Org [54].[16] The main two driving ideas that lead to the construction of PrefLib are (quoting from [54]) :

1. How wide is the gap between theoretical intractability results and practical, real-world instances ? If the constructions required to prove theoretical intractability are rare, what does it tell us about the practical applicability of these results ?
2. Models of agent behaviour and rationality seem to be largely driven by intutiive feeling (e.g., a left to right political spectrum) or mathematical expediency (...) How realistic are these assumptions ? Do we ever see them in real-world data ? Can we derive or learn the assumptions we should use from data ?

See `https://www.preflib.org/`. Since the launch of PrefLib some other repositories emerged, especially PabuLib `http://pabulib.org/`, collecting participatory budgeting instances.

A key interest of such libraries is that we can learn interesting facts from these data (see point 2 above).

Online Platforms

Social Choice Engineering

The COMSOC community has not only developed general tools for classes of collective decision making problems, but has also focused on specific application domains that require specific solutions. This is what I call *social choice engineering.*

An important example — especially for IJCAI attendees – is the assignment of papers to reviewers, which can be seen as an allocation problem (more precisely, an allocation of *chores* rather than *goods* : being assigned one more paper usually gives a negative utility to a reviewer). This allocation process needs as input the preferences of reviewers over papers. On common conference management systems, these preferences are indivated by qualitative evaluations such as *eager to review*, *willing to review*, *can review if necessary*, *don't want to review*, *conflict of interest.* Now, everyone who has already chaired the program committee of a conference knows that a difficulty is that (1) a few papers have very few positive bids, sometimes not a single positive bid ; and relatedly, (2) a few reviewers do not place enough positive bids. Incentivizing reviewers to bid for more papers, and especially for papers that don't have many bids, enhances the whole process by avoiding that some papers are assigned arbitrarily to reviewers who did not bid for them [57, 64].

16. There are also specialized databases for matching.

Example 10 *Six papers (P_1, ..., P_6) have been submitted to a conference. Each paper will be reviewed by two reviewers. Since there are six reviewers, each of them should review two papers. Here are the bids :*

	P1	*P2*	*P3*	*P4*	*P5*	*P6*
Ann	+	+				
Bob			+	+		
Carol				+	+	+
David			+	+	+	+
Edith			+	+	+	+
Fred		+		+	+	+
	1	*2*	*3*	*5*	*4*	*4*

P1 is a lonely paper *: it does not have enough bids, and as a consequence will be assigned to at least one reviewer who does not want to review it (and who is likely not to be an expert on this paper) ; this will probably be Bob, who is one the PC members with the lowest number of positive bids the other one is Ann, who placed a bid on P1 and will very likely be assigned to it). See :*

	P1	*P2*	*P3*	*P4*	*P5*	*P6*
Ann	$\boxed{+}$	$\boxed{+}$				
Bob	$\boxed{-}$		$\boxed{+}$	+		
Carol				$\boxed{+}$	$\boxed{+}$	+
David			$\boxed{+}$	+	+	$\boxed{+}$
Edith			+	+	$\boxed{+}$	$\boxed{+}$
Fred		$\boxed{+}$		$\boxed{+}$	+	+

Bob is an underbidder *: he placed only two positive bids ; that's not enough. Ann placed only two positive bids too. Is she an underbidder just like Bob ? Well, yes, to some extent she is an underbidder, but perhaps not as much as Bob, because she is more* cooperative *: she bids for P_1 (the 'most lonely' paper) and P_2 (the second most lonely paper), and we badly need reviewers for these two papers ! On the other hand, Bob's two positive bids are for papers that many people want to review (P_2 and P_4) so he is only moderately helpful. Can we find a way of (a) measuring the cooperation degree of the reviewers, and (b) give an incentive for Bob to bid for more papers and/or papers needing more bids ?*

An answer is given in [57, 64] : each paper P_i is associated with a bonus *b_i that can be interpreted as the perceived probability, by a reviewer who bid for paper i, that it will be assigned to her. If k_i reviewers bid for paper i then paper i has bonus $b_i = \min(1, 2/k_i)$.*

	P1	*P2*	*P3*	*P4*	*P5*	*P6*		
Ann	+	+					2	
Bob			+	+			$16/15$	< 2
Carol				+	+	+	$7/5$	< 2
David			+	+	+	+	$31/15$	
Edith			+	+	+	+	$31/15$	
Fred		+		+	+	+	$12/5$	
bonus	1	1	$2/3$	$2/5$	$1/2$	$1/2$		

The score of a reviewer is the sum of bonus of the papers they bid for. Bob and Carol are underbidding *because their score is lower than the number of papers they should review. Underbidders are at risk of being assigned papers they don't want! And here is what happens indeed :*

	P1	*P2*	*P3*	*P4*	*P5*	*P6*		
Ann	$\boxed{+}$	$\boxed{+}$					2	
Bob	$\boxed{-}$		$\boxed{+}$	+			$16/15$	< 2
Carol				$\boxed{+}$	$\boxed{+}$	+	$7/5$	< 2
David			$\boxed{+}$	+	+	$\boxed{+}$	$31/15$	
Edith			+	+	$\boxed{+}$	$\boxed{+}$	$31/15$	
Fred		$\boxed{+}$		$\boxed{+}$	+	+	$12/5$	
bonus	1	1	$2/3$	$2/5$	$1/2$	$1/2$		

When seeing the danger, Bob and Carol can bid for more papers :

	P1	*P2*	*P3*	*P4*	*P5*	*P6*
Ann	+	+				
Bob		+	+	+	+	
Carol	+			+	+	+
David			+	+	+	+
Edith			+	+	+	+
Fred		+		+	+	+

Now there is a better matching :

	P1	*P2*	*P3*	*P4*	*P5*	*P6*
Ann	$\boxed{+}$	$\boxed{+}$				
Bob		$\boxed{+}$	$\boxed{+}$	+	+	
Carol	$\boxed{+}$			$\boxed{+}$	+	+
David			$\boxed{+}$	$\boxed{+}$	+	+
Edith			+	+	$\boxed{+}$	$\boxed{+}$
Fred		+		+	$\boxed{+}$	$\boxed{+}$

A screenshot from the experiment in [64] :

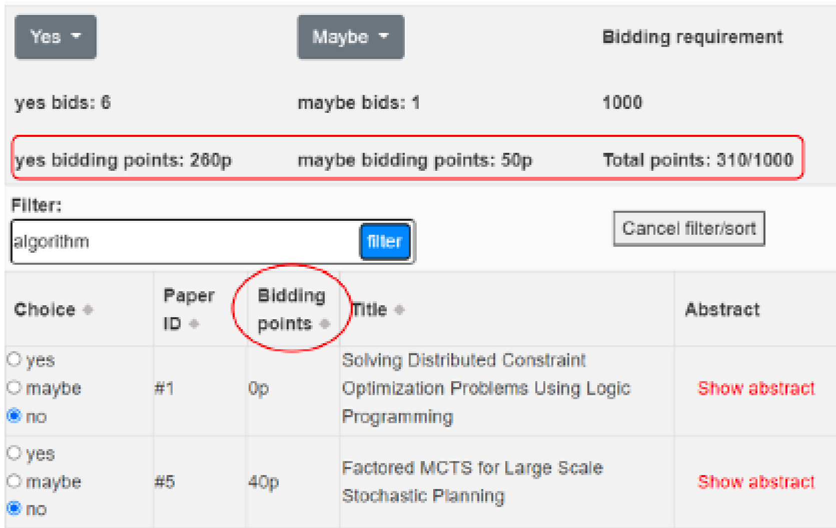

Computational social choice has been applied to various other problems related to conference paper assignment. The notions and techniques involved are from game theory, operations research, fair division, behavioural study.

Acknowledgements

In the end I expect there will be lots of people mentioned here.

Références

[1] Tahar Allouche, Jérôme Lang, and Florian Yger. Truth-tracking via approval voting : Size matters. In *AAAAI-22*, 2022.

[2] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. Fair division of indivisible goods : A survey. *CoRR*, abs/2202.07551, 2022.

[3] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A. Voudouris. Distortion in social choice problems : The first 15 years and beyond. In *IJCAI*, pages 4294–4301. ijcai.org, 2021.

[4] Haris Aziz, Anna Bogomolnaia, and Hervé Moulin. Fair mixing : The case of dichotomous preferences. *ACM Trans. Economics and Comput.*, 8(4) :18 :1–18 :27, 2020.

[5] Haris Aziz and Nisarg Shah. Participatory budgeting : Models and approaches. *CoRR*, abs/2003.00606, 2020.

[6] Nathanaël Barrot and Jérôme Lang. Conditional and sequential approval voting on combinatorial domains. In *IJCAI*, pages 88–94. IJCAI/AAAI Press, 2016.

[7] J.-P. Barthelemy, A. Guenoche, and O. Hudry. Median linear orders : heuristics and a branch and bound algorithm. *Eur. J. Oper. Res.*, 41 :313–332, 1989.

[8] John J. Bartholdi and James Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8 :341–354, 1991.

[9] John J. Bartholdi, Craig Tovey, and Michael Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6 :227–241, 1989.

[10] John. J. Bartholdi, Craig Tovey, and Michael Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6 :157–165, 1989.

[11] John J. Bartholdi, Craig Tovey, and Michael Trick. How hard is it to control an election ? *athematical and Computer Modelling*, 16(8-9) :27–40, 1992.

[12] Dorothea Baumeister, Jörg Rothe, and Ann-Kathrin Selker. Strategic behavior in judgment aggregation. In U. Endriss, editor, *Trends in Computational Social Choice*, pages 145–168. AI Access, 2017.

[13] Anna Bogomolnaia, Hervé Moulin, and Richard Stong. Collective choice under dichotomous preferences. *J. Econ. Theory*, 122(2) :165–184, 2005.

[14] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible goods. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 12. Cambridge University Press, 2015.

[15] Felix Brandt. Rolling the sice : Recent results in probabilistic social choice. In U. Endriss, editor, *Trends in Computational Social Choice*, pages 3–26. AI Access, 2017.

[16] Felix Brandt, Markus Brill, and Paul Harrenstein. Tournament solutions. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 3. Cambridge University Press, 2015.

[17] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. Introduction to computational social choice. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 1. Cambridge University Press, 2015.

[18] Felix Brandt, Christian Geist, and Dominik Peters. Optimal bounds for the no-show paradox via SAT solving. *Math. Soc. Sci.*, 90 :18–27, 2017.

[19] Markus Brill, Théo Delemazure, Anne-Marie George, Martin Lackner, and Ulrike Schmidt-Kraepelin. Liquid democracy with ranked delegations. In *AAAI*, 2022.

[20] Markus Brill, Paul Gölz, Dominik Peters, Ulrike Schmidt-Kraepelin, and Kai Wilker. Approval-based apportionment. In *AAAI*, pages 1854–1861. AAAI Press, 2020.

[21] Ioannis Caragiannis, Edith Hemaspaandra, and Lane A. Hemaspaandra. Dodgson's Rule and Young's Rule. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 5. Cambridge University Press, 2015.

[22] Ioannis Caragiannis, Christos Kaklamanis, Nikos Karanikolas, and Ariel D. Procaccia. Socially desirable approximations for dodgson's voting rule. *ACM Trans. Algorithms*, 10(2) :6 :1–6 :28, 2014.

[23] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. In *EC*, pages 305–322. ACM, 2016.

[24] Ioannis Caragiannis and Evi Micha. Learning a ground truth ranking using noisy approval votes. In Carles Sierra, editor, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 149–155. ijcai.org, 2017.

[25] Rachael Colley, Umberto Grandi, and Arianna Novaro. Unravelling multi-agent ranked delegations. *Auton. Agents Multi Agent Syst.*, 36(1) :9, 2022.

[26] Vincent Conitzer and Tuomas Sandholm. Common voting rules as maximum likelihood estimators. In *UAI*, pages 145–152. AUAI Press, 2005.

[27] Vincent Conitzer and Tuomas Sandholm. Communication complexity of common voting rules. In *EC*, pages 78–87. ACM, 2005.

[28] Vincent Conitzer and Toby Walsh. Barriers to manipulation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 6. Cambridge University Press, 2015.

[29] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.

[30] Pietro Speroni di Fenizio and Daniele A. Gewurz. The space of all proportional voting systems and the most majoritarian among them. *Soc. Choice Welf.*, 52(4) :663–683, 2019.

[31] Virginie Do, Matthieu Hervouin, Jérôme Lang, and Piotr Skowron. Online approval committee elections. In *IJCAI*, 2022.

[32] Britta Dorn and Ildikó Schlotter. Having a hard time ? explore parameterized complexity ! In U. Endriss, editor, *Trends in Computational Social Choice*, pages 209–230. AI Access, 2017.

[33] François Durand. *Vers des modes de scrutin moins manipulables. (Toward Less Manipulable Voting Systems)*. PhD thesis, Paris Diderot University, France, 2015.

[34] Edith Elkind and Arkadii Slinko. Rationalizations of voting rules. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 8. Cambridge University Press, 2015.

[35] Ulle Endriss and Raquel Fernández. Collective annotation of linguistic resources : Basic principles and a formal model. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL-2013)*, pages 539–549, 2013.

[36] Piotr Faliszewski and Jörg Rothe. Control and bribery in voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 7. Cambridge University Press, 2015.

[37] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner voting : a new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, pages 27–48. AI Access, 2017.

[38] Felix Fischer, Olivier Hudry, and Rolf Niedermeier. Weighted tournament solutions. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 4. Cambridge University Press, 2015.

[39] Rupert Freeman, David M. Pennock, Dominik Peters, and Jennifer Wortman Vaughan. Truthful aggregation of budget proposals. *J. Econ. Theory*, 193 :105234, 2021.

[40] Rupert Freeman, Seyed Majid Zahedi, and Vincent Conitzer. Fair and efficient social choice in dynamic settings. In *IJCAI*, pages 4580–4587. ijcai.org, 2017.

[41] Christian Geist and Dominik Peters. Computer-aided methods for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, pages 249–268. AI Access, 2017.

[42] Vasilis Gkatzelis, Daniel Halpern, and Nisarg Shah. Resolving the optimal metric distortion conjecture. In *FOCS*, pages 1427–1438. IEEE, 2020.

[43] Paul Gölz, Anson Kahng, Simon Mackenzie, and Ariel D. Procaccia. The fluid mechanics of liquid democracy. *ACM Trans. Economics and Comput.*, 9(4) :23 :1–23 :39, 2021.

[44] Hadi Hosseini, Debmalya Mandal, Nisarg Shah, and Kevin Shi. Surprisingly popular voting recovers rankings, surprisingly ! In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 245–251. ijcai.org, 2021.

[45] Fatih Erdem Kizilkaya and David Kempe. Plurality veto : A simple voting rule achieving optimal metric distortion. In *IJCAI*, 2022.

[46] Martin Lackner. Perpetual voting : Fairness in long-term decision making. In *AAAI*, pages 2103–2110. AAAI Press, 2020.

[47] Martin Lackner and Piotr Skowron. *Multi-Winner Voting with Approval Preferences*. Springer, 2022.

[48] Michael Lampis. Determining a slater winner is complete for parallel access to NP. In *STACS*, volume 219 of *LIPIcs*, pages 45 :1–45 :14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[49] Jérôme Lang and Jörg Rothe. Fair division of indivisible goods. In *Economics and Computation*, Springer texts in business and economics, pages 493–550. Springer, 2016.

[50] Jérôme Lang and Piotr Skowron. Multi-attribute proportional representation. *Artif. Intell.*, 263 :74–106, 2018.

[51] Jérôme Lang and Lirong Xia. Voting in combinatorial domains. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 9. Cambridge University Press, 2015.

[52] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *EC*, pages 125–131. ACM, 2004.

[53] Evangelos Markakis. Approximation algorithms and hardness results for fair division. In U. Endriss, editor, *Trends in Computational Social Choice*, pages 231–248. AI Access, 2017.

[54] Nicholas Mattei and Toby Walsh. A PREFLIB.ORG retrospective : Lessons learned and new directions. In U. Endriss, editor, *Trends in Computational Social Choice*, pages 289–308. AI Access, 2017.

[55] R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Convergence to equilibria in plurality voting. In *AAAI*, pages 823–828. AAAI Press, 2010.

[56] Reshef Meir. Iterative voting. In U. Endriss, editor, *Trends in Computational Social Choice*, pages 69–86. AI Access, 2017.

[57] Reshef Meir, Jérôme Lang, Julien Lesca, Nicholas Mattei, and Natan Kaminsky. A market-inspired bidding scheme for peer review paper assignment. In *AAAI*, pages 4776–4784. AAAI Press, 2021.

[58] H. Moulin. *Axioms of Cooperative Decision Making.* Cambridge University Press, 1988.

[59] Grzegorz Pierczynski, Piotr Skowron, and Dominik Peters. Proportional participatory budgeting with additive utilities. In *NeurIPS*, pages 12726–12737, 2021.

[60] Ariel D. Procaccia and Jeffrey S. Rosenschein. The distortion of cardinal preferences in voting. In *CIA*, volume 4149 of *Lecture Notes in Computer Science*, pages 317–331. Springer, 2006.

[61] Ariel D. Procaccia and Nisarg Shah. Is approval voting optimal given approval votes ? In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 28 : Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1801–1809, 2015.

[62] Manon Revel. Experiments on liquid democracy.

[63] Simon Rey, Ulle Endriss, and Ronald de Haan. Designing Participatory Budgeting Mechanisms Grounded in Judgment Aggregation. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning*, pages 692–702, 9 2020.

[64] Inbar Rozenzweig, Reshef Meir, and Nicholas Mattei. Mitigating skewed bidding for conference paper matching. In *Match-Up*, 2022.

[65] Tsviel Ben Shabat, Reshef Meir, and David Azriel. Empirical bayes approach to truth discovery problems. *CoRR*, abs/2206.04816, 2022.

[66] Nihar B. Shah and Dengyong Zhou. Approval voting and incentives in crowdsourcing. *ACM Trans. Economics and Comput.*, 8(3) :13 :1–13 :40, 2020.