Computation, Communication, Rationality and Incentives in Collective and Cooperative Decision Making

Table of contents:

There are no evolutions since the pre-proposal.

1 Summary of the Proposal

Our project aims at studying several classes of collective decision problems under three point of views: (a) the impact of the computational difficulties of the mechanisms involved; (b) the impact of their communication requirements; (c) their vulnerability to strategic behaviour.

The classes of problems we consider concern groups of agents who have to reach a stable state or a common decision: (1) coalition structure formation, where agents have to be partitioned into groups; (2) selection of a common alternative, or a collective set of alternatives, subject to some constraints: voting (single-winner elections, committee elections, multiple referenda), group recommendation, multi-facility location; (3) fair allocation of indivisible resources.

Our project has a two-dimensional structure: classes of problems (1)–(3) on the one hand, classes of tasks or questions (a)–(c) on the other hand. There will be one work package for each of the main questions (a) to (c), and a work package whose role will be to implement some of the methods and solutions from other work packages on a collective decision platform, which will be tested on real users and used for pedagogical purposes.

Personnes impliquées dans le projet / People involved in the project

Organisation	Name	First name	Current position	Field of research	Involvment (PM)	Contribution to the project
LAMSADE	Airiau	Stéphane	MCF[1]	CS[2]	14.4	WP1 WP2 WP4
LAMSADE	Bouyssou	Denis	DR CNRS	E[3]	14.4	WP2 WP3 WP4
LAMSADE	Cazenave	Tristan	Professeur	CS	12	WP1 WP3 WP4
LAMSADE	Cornaz	Denis	MCF	CS	7.2	WP1 WP3
LAMSADE	Galand	Lucie	MCF	CS	14.4	WP1, WP2 WP4
LAMSADE	Gourvès	Laurent	CR CNRS	CS	14.4	WP1 WP3 (co-)coordinator WP3
LAMSADE	Lang	Jérôme	DR CNRS	CS	24	project coordinator coordinator WP2 WP1 WP2 WP3 WP4
LAMSADE	Monnot	Jérôme	DR CNRS	CS	14.4	WP1 WP2 WP3
LAMSADE	Moretti	Stefano	CR CNRS	CS	7.2	WP2 WP3
LAMSADE	Tsoukiàs	Alexis	DR CNRS	CS	4.8	WP2 WP4
LIPN[◇]	Chevaleyre	Yann	Professeur	CS	7.2	WP2 WP4
LIP6	Bampis	Evripidis	Professeur	CS	12	WP1 WP3
LIP6	Escoffier	Bruno	Professeur	CS	7.2	WP1 WP3
LIP6	Maudet	Nicolas	Professeur	CS	9.6	WP2 WP3 WP4
LIP6	Pascual	Fanny	MCF	CS	12	WP3 WP4
LIP6	Perny	Patrice	Professeur	CS	14.4	WP1 WP2 WP4 local resp. LIP6
LIP6	Spanjaard	Olivier	MCF	CS	14.4	coordinator WP1 WP1 WP3
LIP6	Viappiani	Paolo	CR CNRS	CS	14.4	WP2 WP4
LIG[+]	Bouveret	Sylvain	MCF	CS	14.4	WP1 WP2 WP4 (co-)coordinator WP4
ParisTech[+]	Hudry	Olivier	Professeur	CS	4.8	WP1
CREM	Fanelli	Angelo	CR CNRS	CS	12	WP1 WP3 WP4
CREM	Lebon	Isabelle	Professeur	E	7.2	WP3 WP4
CREM	Mbih	Boniface	Professeur	E	7.2	WP3 WP4
CREM	Merlin	Vincent	DR CNRS	E	14.4	(co-)coordinator WP4 WP1 WP3 WP4 local resp. CREM
THEMA[◦]	Courtin	Sébastien	MCF	E	14.4	WP1 WP3 WP4
CES[◦]	Grabisch	Michel	Professeur	CS	7.2	WP1 WP4
PSE[◦]	Laslier	Jean-François	DR CNRS	E	4.8	WP3 WP4
THEMA[◦]	Nuñez	Matias	CR CNRS	E	14.4	WP1 WP3 WP4 (co-)coordinator WP3
CES[◦]	Rusinowska	Agnieszka	DR CNRS	E	12	WP2 WP3 WP4

[1] maître de conférences (assistant professor); [2] computer science; [3] economics;
[◇] attached to LAMSADE for the project; [+] attached to LIP6 for the project; [◦] attached to CREM for the project.

2 Context, Position and Objectives of the Proposal

2.1 Objectives, Originality and Novelty of the Project

Collective decision making ranges from the election of representatives (in political elections and many other contexts) to the choice and location of multiple common facilities, group formation (car sharing, work team formation), the fair allocation of resources, and many other daily-life problems. The wide availability of internet and the increased power of computers makes it possible to implement new decision mechanisms. However, they raise several issues, such as the impact of their computational difficulties, the impact of their informational and communicational requirements, and the vulnerability of the quality of the output to the potentially strategic, selfish behaviour of agents.

These issues have been partly addressed by several communities:

- *classical social choice* (see for instance [11]), focusing mainly on the design and the normative study of collective decision making mechanisms;
- *computational social choice* (see for instance [44, 32, 33]), focusing mainly on the impact of computational complexity on the feasibility of some collective decision mechanism, as well as on their "computational resistance" to various forms of strategic behaviour;
- *algorithmic game theory* (see for instance [125]), focusing on the computation of game-theoretic solution concepts and on the measure of the impact of strategic behaviour on the quality of the solution.

Our project will need the expertise of these three communities, making them interact in a closer way. As a result, it involves participants both from computer science and from economics.

Our project is original in many aspects, which we list now, and discuss in more detail next.

- focus on *constructive approaches* that can lead to a *direct use of algorithms and protocols*;
- evaluate *trade-offs* between computational hardness, communication requirements, impact of strategic behaviour, and other social choice-theoretic properties, requiring expertise from both computer scientists and economists;
- focus on new problems such as *communication requirements* or *verification certificates*;
- bridge *algorithmic game theory* with *(computational) social choice*;
- *implement* most of our mechanisms and *develop a software platform* that will be openly accessible and used for many purposes, including *on-line collective choice experiments* and *pedagogical* purposes.

We now develop these arguments in more detail.

Computational social choice, initiated in the early 90s by the seminal papers of Bartholdi *et al.* [100, 18], has been developing at a rapid pace since the early 2000s. So far the community has focused mainly on theoretical issues, such as assessing the complexity of various collective decision mechanisms or their computational resistance to strategic behaviour. A major novelty of our project is that it focuses on more directly applicable procedures. We are now at a stage where the field is mature enough so as to go beyond impossibility theorems and computation (or communication) complexity barriers. Our project focuses on *constructive* methods. In every subtask of the project we keep in mind the objective to come up with practically implementable mechanisms or protocols, that come with properties or various modes of evaluation.

Rather than simply assessing the computational complexity or inapproximability of mechanisms, we will take a constructive approach and follow two paths: (a) *cope with complexity*: we will design, implement and experiment advanced algorithms, based on combinatorial optimization techniques as well as heuristics, that do their best to cope with computationally hard mechanisms. Remember that hardness is a worst-case concept and that many (if not most) instances do not fall in the worst-case traps; (b) *design approximate mechanisms that are easier to solve, and study their properties*. While polynomial approximation algorithms are new in social choice, viewing these approximations as full-fledged mechanisms and study them from the point of view of normative properties has rarely been dealt with (see the related work section).

While computational social choice has up to now mainly focused on the difficulty of computing mechanisms, we will also focus on other important (perhaps even more important) problems. First, *communication*: not only we will provide algorithms for computing mechanisms but also *communication protocols* that try to minimize interactions between the agents and the central authority (in centralized mechanisms) or between the agents themselves (in distributed mechanisms). We will use tools from communication complexity and protocol design. Again, we take a constructive approach: we will design specific protocols, for existing mechanisms or for new mechanisms. New mechanisms, designed with the purpose of having a low communication complexity, will also be evaluated according to their normative properties. Another novel issue is that of providing, when possible, a *short certificate* of the solution, allowing to verify that the solution computed is the right one, and to some extent to explain it[1], which is important for the social acceptability of collective choice rules.

Another original aspect of the project is its more intricate interleaving of different fields and techniques, such as social choice, algorithmic game theory, preference elicitation and communication complexity, combinatorial optimization, and advanced algorithmics. Especially, we want to bridge more tightly the communities of computational social choice and algorithmic game theory / algorithmic mechanism design, which are among the fastest-growing subfields

[1] We want to stress that 'verification' may be ambiguous: we do not aim at constructing tools for the verification of algorithms or protocols (which also is very relevant in collective decision making, especially in voting, when one wants to allow agents to verify that their vote has been taken into account or that no voter has been allowed to vote twice) but this is a very different facet of voting, studied with very different techniques by a different community.

of both artificial intelligence and theoretical computer science, and which have been evolving rather separately[2], apart from a few recent works that we review in the next subsection. Our goal here is to build constructive approaches: rather than proving impossibility theorems or characterizing fully truthful mechanisms, that usually require very strong assumptions and/or are not very satisfactory (such as, for instance, dictatorships), we will allow and study tradeoffs (again) between the quality of the mechanisms and "how much" they are impacted by strategic behaviours. For that we will use notions and techniques from algorithmic game theory, such as the price of anarchy or the price of stability.

These mechanisms, together with associated algorithms and protocols, will be implemented on a widely accessible platform. Work in computational social choice has been up to now focusing mainly on theoretical issues. We think the field is ready for the next step. Implementation is a major aspect of our project. Since the platform will be openly accessible, it will be possible to test it with real users, in a similar way as Doodle™ proceeds[3]. This point is novel, and the next one is perhaps even more novel: the platform will be used for pedagogical purposes. We feel that social choice is ill-known by lay people and a long-term expected impact of the project is to improve this matter of facts.

In this project we are mainly interested in handling discrete problems and indivisible objects. As a consequence, combinatorial structures and combinatorial optimization will play an important role. The set of possible outcomes (*i.e.*, feasible solutions) will be, for instances, subsets of candidates, collective rankings of candidates, assignments of objects or activities to agents, partitions of agents. Even in single-winner voting, where the output consists of a single candidate, combinatorial structures play often a role in the computation of the solution[4]. Due to its combinatorial structure, the set of feasible solutions grows exponentially with the description of the problem, and makes exhaustive methods prohibitive.

One of the difficulties in applying notions and techniques from algorithmic game theory to collective decision mechanisms is that social choice usually assumes that the input of a mechanism consists of ordinal preferences, while many concepts in game theory (such as power indices in cooperative game theory, and monetary payments in mechanism design) require cardinal, and often transferable, utilities. This also applies to measures of efficiency in algorithmic game theory, such as the price of anarchy or the price of stability. However, we will, as far as possible, try to avoid the use of money, and more generally of cardinal, transferable or even nontransferable utilities, because the applications that we want to address often lack the possibility to propose monetary compensation, and also because many mechanisms in social choice are defined in a purely

[2]We obviously do *not* mean that social choice has not been studied in game-theoretic settings; there are a lot of papers that model group decision using solution concepts from game theory. What we mean is that the specific notions and tools from *algorithmic* game theory, such as the price of anarchy and the price of stability, the computation of various kinds of equilibria, or the communication complexity of mechanisms, has been much less studied there than for auctions and similar problems.

[3]The platform will not start from scratch. The *Whale* platform (`http://whale3.noiraudes.net/`, cf. WP4) is already used by users in real-life situations. It will be developed on a much larger scale.

[4]For instance, some voting rules – such as Kemeny or Slater – are defined via an optimization process over all possible rankings.

ordinal way[5]. There are good reasons for this, widely discussed in the social choice community, such as, in particular, the issue of interpersonal comparison of preferences.

The main objective of the project is to define and study *reasonable* mechanisms for collective and cooperative decision making. By "reasonable" we mean that the mechanisms will be evaluated under the following three points of view:

1. how difficult it is to *compute* the outcome, which trade-offs can we accept between computational tractability and the quality of the solution, and how much space is needed to *verify* the outcome.

2. how much *communication* (and in particular – but not only – how much *elicitation*) is needed for the outcome to be determined, and how we can trade-off communication against the quality of the outcome.

3. how much the mechanism is resistent to *strategic behaviours*: what is the worst-case impact, on the global quality of the outcome, of letting the agents strategize, and how much computation is required for an agent to strategize efficiently.

Thus, we aim at studying collective decision problems from the point of view of the impact, on the quality of the outcome, of the selfish behaviour of agents, of enforcing computational tractability, and of enforcing communication tractability. The classes of collective choice problems that we will consider are:

- *coalition structure formation problems.* Agents have to form groups; each group may be labelled with a task or an activity (for instance: car sharing, common cultural or touristic activity, etc.) or not (team formation, multipartite matching); agents have preferences about the label of the group and/or the identity (or sometimes, only the number) of agents who will be in the same group. The specific framework where every agents' preferences depend only on the identity of the agents in her group is that of *hedonic games* (with bipartite matching as a specific subcase). In *group activity selection*, preferences depend on the activity of the group and the number of agents in the group.

- *selection of a common alternative, a collective set of alternatives, or a collective structured solution.* The agents have to choose alternatives out of a larger set, subject to some constraints. The most natural instances of this problem come from *voting*: single-winner elections (choice of one candidate), committee elections (choice of a subset of candidates, with usually a cardinality constraint on the size of the committee, but also specific constraints such as gender balance), multiple referenda (voters have to agree on yes/no decisions regarding several dependent issues), or computing a collective ranking over alternatives; two other instances we will focus on are group recommendation (the objects can be movies, touristic sites etc.) and multi-facility location (optimal placement of facilities to be

[5] This is especially true in voting, but also in coalition structure formation problems such as matching and hedonic games; this is less true in fair division, where preferences are often cardinal (but yet nontransferable) – the Santa Claus problem [16] is a typical example.

commonly used by the group, in order to minimize transportation costs and maximize the satisfaction of the users).

- *fair resource allocation.* Given a set of indivisible[6] items, and agents who have preferences over the sets of items they can receive, we look for an allocation assigning every item to an agent, that satisfies some fairness or equity criterion as well as some efficiency criteria. The set of allocations may be subject to some feasibility constraints.

There are formal similarities between these classes of problems. In each case, we have to find a solution out of a combinatorial space of solutions that agrees as much as possible with the agents' preferences.

Depending on the kind of problem (voting, selection of a set of items, resource allocation, coalition formation), some of the questions above have been considered to varying extents in the literature, while some are completely new. For instance, computational complexity and approximation have been studied extensively in voting (although less so for multiwinner election contexts), but less so in fair division and coalition formation. The price of anarchy and the price of stability have been studied in resource allocation with monetary transfers (such as auctions) and in facility location problems, but less so in voting, fair division without money, and coalition structure generation. Communication complexity, the design of efficient protocols, and (even more so) "communicational approximation" have been considered only scarcely for the fields mentioned here.

2.2 State of the Art

We give a quick survey of the state of the art related to our project proposal, with a focus on works realized by members of the consortium (marked by 'MC').

2.2.1 Classical social choice

The social choice community has been studying for long formal models for collective decision making (impossibility and characterization theorems, axiomatic study of voting rules, game-theoretical aspects of group decision making) and has more recently started to perform voting experiments, in laboratory or in situ. The social choice theorists participating to the project have a strong expertise in all these topics. France has historically a very strong expertise in social choice, and our senior social choice theorists are among the world leading researchers in the field. Altogether, the economists involved in the project have published more than 50 papers on the evaluation of different voting rules since 2005: axiomatic properties [52, 22, 31], resistance to strategic behaviour [6], analysis of voting equilibria [61], evaluation of the likelihood of paradoxes [59, 60], various forms of experiments [141, 19, 25].

2.2.2 Computational issues in social choice

Computational complexity of voting rules (and to a lesser extent, of fair division mechanisms), and the computational resistance to strategic behaviour

[6] We leave fair division of *divisible* items, such as cake-cutting procedures, out of our project; this field requires very different techniques.

of voting rules, are the core topics of computational social choice. Works on both topics started in the 80's, with seminal works by Bartholdi *et al.* on the one hand [100, 18], and by French researchers (among them Hudry, MC) on the other hand [17]. These topics have generated a lot of publications since the early 2000s. Faliszewski and Procaccia [79] give a survey on the computational barriers to manipulation. Hemaspaandra et al. [78] give a survey of the complexity of winner determination for voting rules. Two more general surveys on computational social choice are [44] (by three MC) and [32]. The sign that computational social choice is now becoming a mature discipline is witnessed by the fact that a handbook is in preparation [33], co-edited by a MC. There is a biannual *International Workshop on Computational Social Choice* (COMSOC) since 2006, and, at the European level, a COST action on Computational Social Choice (the French members of its management committee, and two members of the steering committee, are participants to the project).

The computational aspects of fair division of indivisible resources have been also investigated, although less thoroughly. See [42] for a survey (now a little bit outdated). Some specific fair division problems have received some attention, such as the Santa Claus problem [16]. The computational aspects of coalition structure formation have been studied in a number of recent papers (see Chapter 4 of [40] for a survey), as well as for the specific models of hedonic games [15, 139, 13] and group activity selection [62]. The social-choice theoretic study of voting rules obtained as approximations of other rules has been initiated in [38].

The members of the project have contributed a lot to these research fields; we give only a few recent references: [97, 98, 41, 99, 58] for the computational complexity of voting rules and aggregation functions; [49, 106, 26, 87] for the computation of outcomes of voting rules under incomplete information; [5, 45, 54, 107, 55, 105, 118, 88] for voting and preference aggregation on combinatorial domains; [57, 72, 66] for complexity issues in multi-winner elections and proportional representation; [62] for the complexity of the group activity selection problem; [74, 28, 29, 115, 63, 4, 86, 111, 43] for computational issues in fair division.

2.2.3 Communication complexity and protocols in social choice

Communication complexity has been studied a lot in the context of distributed computation [104], but only little has been done in social choice. Conitzer and Sandholm [56] and Procaccia [133] initiated the communication complexity study of voting rules. Chevaleyre *et al.* [50, 48] (MC) address the compilation complexity of voting and the trade-offs to be made between communication and compilation. Communication protocols for iterative voting have also been considered by MCs [3]. In fair resource allocation, low-communication protocol by picking sequences have been initiated by Brams *et al.* and further studied very recently, in [30] (MC) and [101]. Communication protocols for fully decentralized multi-agent resource allocation has been investigated in a dozen of papers, some of them by some MCs [47, 4, 46]. The communication complexity of coalition structure formation has been initiated in [132].

2.2.4 Algorithmic game theory

Computational complexity was one of the main obstacles in computer science, but the emergence of the Internet, social networks and electronic markets completely changed the situation. Computers are not isolated anymore and one has to cope with the selfish behaviour of their users. Game theory has quickly become a formal basis to understand the problems created by uncoordinated decisions and manipulations made by interacting users. This fast growing field is now known as *algorithmic game theory* (see for instance [125]).

One of the most challenging issues was determining the complexity of computing a Nash equilibrium [84, 76]. Meanwhile, prominent concepts have emerged like the *price of anarchy* [103, 134] and the *price of stability* [9, 134], which are two derivatives of the approximation ratio. Another seminal work is due to Nisan and Ronen who considered *mechanism design* under an algorithmic perspective [126]. Many classical problems have been revisited in a setting where a pool of agents has to disclose a private information.

Since 2005, several MCs have significantly contributed to the domain of algorithmic game theory [7, 51, 24, 90, 73, 89, 81, 8, 91, 92]. They also organized a spring school [1] and two workshops.

The computational aspects of cooperative games has received a lot of attention too. The work was initiated by [131]; the recent book "Computational Aspects of Cooperative Game Theory" [40] gives an extensive survey of the field.

2.2.5 Computational social choice and (strategic or cooperative) games

Voting seen as a strategic game has been analyzed in many papers in classical social choice, such as [120, 119, 122, 138, 124, 82, 68]. and more recently in computational choice [67, 144, 121, 113, 71]. These papers analyze the strategic behaviour of voting rules using *non-cooperative* game theory. These references will be our starting point for WP3 described below. Only a few papers quantify the loss of optimality caused by strategic behaviour in voting [140, 35] and fair division [39], whereas these questions have been addressed in many papers for many classes of strategic games within the algorithmic game theory community.

Analyzing voting manipulation in the setting of *cooperative* games has been considered recently. The coalition formation process in the coalitional manipulation problem has been studied as a cooperative game with transferable utility in[14] and with nontransferable utility in [146].

The members of the project have also contributed to these issues: [108, 69] are two recent articles on the game-theoretic aspects of voting; [145] gives a game-theoretic analysis of sequential voting on multi-issue domains, [4] a game-theoretic analysis of resource allocation with sharable items, [3] study the convergence of iterated majority voting, and [112] gives a mechanism design approach to fair division of indivisible objects. Finally, some members of the project have a strong expertise in cooperative games and coalition formation processes; among many references, see [93, 77, 20, 65, 75].

3 Scientific and Technical Programme, Project Organisation

3.1 Scientific Programme and Project Structure

There will be a work package for the global coordination of the project. There will be four other work packages. Each of these four will have regular meetings, at least once a year. These meetings will sometimes be joined with meetings of other work packages. Because all participants except one are in Paris or in Caen, it will be easy to organize frequent, one-day meetings, which will ensure that the participants will meet often. Beyond these regular meetings, there will be an annual meeting of all participants. At the end of year 3 there will be a summer school, and near the end of the project, an international workshop where we will invite other researchers, mainly from Europe.

3.2 Description by task

3.2.1 WP0: Global coordination of the project

- Coordinator: Jérôme Lang (LAMSADE)
- Participants: Olivier Spanjaard (coordinator, WP1), Laurent Gourvès and Matias Nuñez (coordinators, WP3), Vincent Merlin (local coordinator, CREM; co-coordinator, WP4), Sylvain Bouveret (co-coordinator, WP4), Patrice Perny (local coordinator, LIP6), Jérôme Lang (project coordinator and coordinator, WP2).

This WP will meet every three months to evaluate the current state of the project. Our project has a two-dimensional structure (classes of problems on the one hand, classes of tasks on the other hand); we choose to organize the WPs along the tasks, which seems a natural things to do, because to each task corresponds a specific expertise that can be applied to several classes of problems. There will be regular "transversal" meetings (one every year) dedicated to voting and object selection, fair division, and coalition structure formation. The timing and other organization issues of these transversal meetings will be handled by WP0.

3.2.2 WP1: Computation, approximation, verification

- Coordinator: Olivier Spanjaard (LIP6).
- Main participants: Stéphane Airiau, Evripidis Bampis, Sylvain Bouveret, Tristan Cazenave, Denis Cornaz, Sébastien Courtin, Bruno Escoffier, Angelo Fanelli, Lucie Galand, Laurent Gourvès, Michel Grabisch, Olivier Hudry, Jérôme Lang, Vincent Merlin, Jérôme Monnot, Patrice Perny, Olivier Spanjaard + postdoc(s).

This work package focuses on the computational difficulties of mechanisms for the various tasks considered in this project (voting and object selection, resource allocation, coalition formation). More precisely, given preferences reported by the agents, we consider the problem of computing a solution (for instance, a set of winners, an assignment of resources to agents, or a partition of the agents into coalitions) that satisfies desirable properties.

Now, many collective decision making mechanisms are computationally hard. This is almost always the case in fair division and in coalition formation rules; for voting, although many simple voting rules can be computed in polynomial time, computational hardness often arises when we want to output a winning committee or a collective preference (and even computing a single winner is hard for some rules, such as Kemeny).

The implications of a high computational complexity are manifold. The most obvious one is that when a mechanism has a high complexity, it is not easy to use in practice. In low-stake contexts, especially, we do not want a mechanism that takes hours or days of computation; the solution should be available immediately, and in case a few small modifications have to be done in the input (consider for instance a dynamic environment like Doodle™), we should be able to recompute a new solution from the old solution as quickly as possible, so that the users are able to interact with the mechanism.

Another very important negative impact of high complexity is that it may become practically infeasible to *verify* the solution. Take voting as an example: for some rules, winner determination (more precisely, determining whether a given candidate is a cowinner) is polynomial; for some others, it is NP-complete (for instance, the Banks rule [143]); and for those of a third group (such as Kemeny [94]), winner determination is above NP (under the standard assumptions of complexity theory). A huge difference between the second and the third class of rules is that in class 2 (and 1), a short certificate can be exhibited, which allows voters to verify that the true winner is indeed the one given by the algorithm[7]. However, for class 3, such succinct certificates do not exist, which makes it very difficult to verify the outcome of the election. (Imagine a country using the Kemeny rule; the day after the election, the certificate published in the newspapers would contain thousands of pages.)

We envisage three ways of coping with this intractability of winner determination or verification:

(a) The first way consists in allowing a trade-off between the quality of the output and the amount of computation needed by designing *polynomial approximation algorithms* for rules for which winner determination is NP-hard. Similarly, we will look for polynomial *verification* approximation algorithms, where we do not require polynomial computation time but only polynomial-size certificates (this means that we look for an approximation of a problem above NP by a problem in NP). Note however that, though approximation algorithms seem a good tradeoff in low-stake contexts, it does not hold true in high-stake contexts: it may indeed be unacceptable to allow an approximate mechanism to output something which is not the true winner. Nevertheless, defining an approximation algorithm of (for instance) a voting rule amounts to *defining a new voting rule*, and this voting rule is worth studying on its own. The axiomatic study of such rules has only been addressed very recently [38]. We will do an *deep axiomatic study of these new rules*, so as to assess their social acceptability.

[7]For instance, for the Banks rule, such a certificate is a maximal subset S of candidates such that the restriction of the majority graph to S is transitive.

(b) The second way consists in *designing efficient exact algorithms for instances of reasonable sizes*. This solution calls for heuristics or preprocessing techniques that take advantage of the problem structure in order to speed up the search (as an example, the notion of *similar* candidates for speeding up the computation of a Slater ranking [53]). In order to evaluate the practical performances of heuristics or preprocessing techniques, several real datasets are freely available on the web for carrying out numerical tests, especially the *PrefLib* database [117].

(c) The third way consists in *identifying polynomially solvable classes of instances*. This typically calls for domain restrictions on preference profiles that make winner determination, or verification, easier. Examples of such restrictions are the various approximate notions of single-peakedness, such as [36], or [57, 58] (proposed and studied by members of the project). We will design and experiment algorithms for specific rules under such restrictions, and the plausibility of these restrictions in practical social choice domains will be assessed by theoretical studies (e.g., by determining the probability that a restriction holds for an arbitrary profile) and/or practical tests on real datasets.

Task 1.1 (year 1) Synthetic bibliography on the computational complexity of various tasks associated with the classes of problems considered.

Task 1.2 (years 1-3) Determine the complexity of mechanisms for which it is not known; identify tractable classes of instances. Social choice-theoretic evaluation of the corresponding domain restrictions.

Task 1.3 (years 3-4) Reoptimization algorithms (recomputing a solution close enough to the current solution when the input undergoes a small change).

Task 1.4 (years 2-3) Price of computational tractability: design of polynomial approximation algorithms together with their performance guarantee; possibility and impossibility results about the approximation of some mechanisms.

Task 1.5 (years 3-4) Axiomatic evaluation of rules defined by approximation algorithms.

Task 1.6 (years 2-4) Practical computation of solutions: heuristic search algorithms, implementations, experimentations. study of mechanisms under this point of view; mechanisms with polynomial-size verification certificates.

3.2.3 WP2: Communication

- Coordinator: Jérôme Lang (LAMSADE)
- Main participants: Stéphane Airiau, Sylvain Bouveret, Denis Bouyssou, Yann Chevaleyre, Lucie Galand, Jérôme Lang, Nicolas Maudet, Jérôme Monnot, Stefano Moretti, Patrice Perny, Agnieszka Rusinowska, Alexis Tsoukiàs, Paolo Viappiani + postdoc(s).

This work package is dedicated to the study of the communication cost of the various mechanisms considered in this project (voting, resource allocation, coalition structure formation). Crucially, this cost is irrespective of the computation burden. We are interested in the following questions: what amount of information needs to be exchanged so as to find the winner(s), to determine an optimal allocation of resources, to find a stable coalition structure?

This is a central concern in many settings. Indeed, whereas the cost of computation relies on the computer, and can therefore be bearable when there is no hurry to make a decision, the cost of communication bears on the individuals, and having a mechanism that requires them to spend hours would make it totally infeasible in practice. (Even spending a few *minutes*, in some low-stake settings, is already not acceptable.)

Communication can take place between the agents and a central authority, who is in charge of computing the social outcome and has no prior information on the agents' preferences: in this case we will speak of *elicitation*. But in some other cases it can take place between the agents themselves, who interact in a decentralized way. We will pay a special attention to semi-decentralized approaches (where agents interact between themselves in the presence of a central authority) and fully decentralized protocols (without any central authority).

Interactions are regulated by a *communication protocol*, which prescribes who has to send which information to whom and when (and usually prescribes turn-taking). The *communication cost* of a protocol is usually defined as the (worst-case or average-case) number of bits that have to be sent so that the solution can be determined. In a centralized setting, the *query cost* of a protocol is simply the total number (again, worst-case or average-case) of questions asked by the central authority to the agents. The communication (resp. query) complexity of a mechanism is the communication (resp. query) cost of the best protocol that makes it possible to determine the outcome. Communication complexity, originating in theoretical computer science (see [104]), has already been applied to social choice mechanisms, in particular voting [56]. There are also slightly more specific techniques that can be borrowed from economics [135], provided the problem at hand verifies some properties.

In this project, we will primarily study the trade-off between the communication cost and the quality of the solution obtained. Relaxing the requirement of providing an optimal solution may reduce the communication cost of the mechanism (see [136] for a preliminary study in voting). In a competitive setting, revealing less information also limits the ability of other agents to manipulate. Some of these mechanisms have already been studied (such as picking sequences in resource allocation [30]); we will design *new* mechanisms, whose interest is that they come with a built-in low-communication protocol, possibly highly distributed. these novel mechanisms will be studied from the point of view of normative properties and computation[8]. We will also try to identify protocols reducing manipulation opportunities, whenever possible.

[8] To give a (simplistic) example, consider this protocol : voters communicate (between themselves or with the central authority) in turns; at each step, a designated voter vetoes one alternative; once an alternative has been vetoed by 2 voters, it is eliminated, and the voters are informed of that. As another example, voters may be asked to submit their top alternative, then their second best alternative, etc., until the quality of the outcome is high enough. (Of course, these protocols are simplistic; they are given as examples.)

Finally, we will evaluate the loss of quality induced by the communication gain, *i.e.*, implied by a restriction on the amount of communication or to the commitment to a simple protocol. We will design and study "communication-wise approximation" protocols for collective decision making, that are to communication complexity what polynomial approximation algorithms are to computational complexity.

Task 2.1 (year 1) Provide an *extensive bibliography* on protocols and communication complexity in collective decision making and multi-agent systems.

Task 2.2 (year 2-3) Evaluate the *communication complexity* of various tasks associated with the classes of problems considered in this project (different forms of voting, resource allocation, coalition formation).

Task 2.3 (year 2-3) Design *low-communication mechanisms*. Study them from the point of view of normative properties, computation (joint work with WP1), and strategic behaviour (joint work with WP3).

Task 2.4 (year 2-4) *Price of low communication.* Evaluate the loss of quality induced by the communication gain. Design communication-wise approximation protocols.

3.2.4 WP3: Strategic Models of Collective Behaviour

- Coordinators: Laurent Gourvès (LAMSADE) and Matias Nuñez (THEMA).
- Main participants: Evripidis Bampis, Sylvain Bouveret, Denis Bouyssou, Denis Cornaz, Sébastien Courtin, Bruno Escoffier, Angelo Fanelli, Laurent Gourvès, Jérôme Lang, Jean-François Laslier, Isabelle Lebon, Nicolas Maudet, Boniface Mbih, Vincent Merlin, Jérôme Monnot, Stefano Moretti, Matias Nuñez, Fanny Pascual, Agnieszka Rusinowska, Olivier Spanjaard + postdoc(s).

This work package mainly deals with the impact of strategic behaviour on the feasibility and the quality of the collective decision mechanisms. It first deals with game-theoretic solution concepts in voting contexts. It also focuses on the notion of rationality and its role in voting environments (*alternative notions of rationality*). It finally deals with (in)efficiencies that can be generated by the different notions of rationality.

As far as solutions concepts are concerned, it is clear that some of them are more attractive than others in the games formed by agents reporting their preferences in collective decision mechanisms. For instance, Nash equilibria in standard voting contexts are often too numerous, and almost any outcome is possible in equilibrium, so that this concept has a poor predictive power. Its extensions, such as correlated equilibria (CE) and refinements such as strong equilibria (SE) are appealing concepts [12, 21] that are particularly relevant in collective decision contexts. Likewise, dominance solvability, through the iterated elimination of weakly dominated strategies (IEWDS), has recently been used in strategic voting contexts since it gives a neutral measure of the coordination incentives that each voting rule posits (see for instance [68, 37, 64, 61]). We want to explore in depth these solution concepts. In particular, we would like to address the following questions: in which collective decision games a CE, SE is guaranteed to exist, or when is a game solvable through IEWDS? If the answer is positive, is it computationally hard to compute such an equilibrium? After a sequence of improving deviations made by the players (each one at a time), do the players reach an equilibrium? If an equilibrium is reached, what it the rate of convergence? (See *e.g.* [34] for the standard case of fictitious play.)

Potential applications of this research stream include electoral behaviour. Indeed, the equilibrium dynamics can be useful to describe how the information impacts the (strategic) voters' choices. For instance, in the French presidential election (plurality with runoff), it is far from clear how the voters should vote in the first round. This might trigger coordination failures and poor preference/information aggregation. Can we modify the electoral system so that the voters' opinions are correctly aggregated in equilibrium? The experimental work (to be detailed in WP4) focuses on these "opinion dynamics" and will be useful to complement our theoretical findings.

The second task of this work package is concerned with the idea of rationality. This is a basic axiom which is often interpreted as saying that players are self-interested. However agents (*a fortiori*, human beings) are sometimes not fully rational, at least, not according to the most standard definition of rationality, and it is of paramount importance to understand and formalize this phenomenon. A first interpretation is that the players' actions are affected by fairness, altruism or even spite [110, 114, 83]. A player may optimize a combination of individual cost and social cost. A second interpretation is that a deviation is performed only if it induces a "significant" modification of the player's utility. Finally, the idea of *partial honesty* is also a relevant alternative to the usual postulate of full rationality (see [70], [102] and [116] for applications in Nash implementation and [130] for computational issues). A partially honest player has a strict preference for telling the truth only when truth-telling leads to an outcome which is not worse than the outcome which occurs when he lies. We would like to revisit the collective decision games described in Section 2.1 in which a player's action is guided by one of the above alternative notions of rationality. In particular, the existence, computation and dynamics of the equilibria associated with these notions of rationality should be studied. The recent theory of voting has mostly focused on environments with a large number of voters (as in national elections). In these environments, characterizing several properties of preference/information aggregation becomes more tractable (see [2, 27, 85, 109, 123, 127, 128] among others). In contrast with this strand of the literature, we will mostly focus on elections with few voters. Indeed, one can understand voting with few agents as a specific sort of bargaining that does not fit into the usual Nash approach. Developing concepts of fairness and efficiency in this restricted environment seems to be a fruitful avenue of research. For instance, Approval Bargaining (i.e. using the method of Approval Voting with two agents as discussed by [129]) posits a novel bargaining protocol over a finite set of objects.

The third and final task of the work package deals with the performance degradation of systems due to the selfish behaviour of the users. Central to this area are the concepts of *price of anarchy and the price of stability* (abbreviated PoA/PoS) [103, 9]. They compare the social cost of the worst/best Nash equilibrium with the social optimum, in the worst case. The PoA/PoS were originally defined for Nash equilibria but it is natural to extend them to other solution concepts. The first task deals with refinements/extensions of the Nash equilibrium (CE, SE and IEWDS in particular), and we would like to study the PoA/PoS for these concepts and refinements. In concrete terms, we expect to give tight bounds on the PoA/PoS for the games described in Section 2.1. Besides self-interest, spite and more surprisingly altruism can lead players to socially poor outcomes [114, 96, 10]. Being fair (e.g. envy-free) may also harm the system's performance. The second goal of this task is to bound the PoA/PoS when the players follow the alternative notions of rationality described above. We expect this work to explain to what extent, following an alternative notion of rationality can improve or deteriorate the quality of the system. A similar question is posed in the context of mechanism design. Returning the socially optimal outcome is rarely strategyproof. Therefore one has to sacrifice optimality to ensure strategyproofness. Under an alternate notion of rationality, how bad are strategyproof mechanisms?

Task 3.1. (year 1-2) *Identify* the conditions for existence and for the different game-theoretic solution concepts. *Determine* the hardness of their computation and their dynamics (i.e. convergence to the equilibrium). *Apply* the previous results to understand how the information impacts the voters' choices.

Task 3.2. (year 2-4) *Explore* the implications of replacing the idea of full rationality by its alternatives (altruism, partial honesty, spite). *Develop* game-theoretic models of collective decision based on these weakenings of classical rationality while using the concepts from Task 3.1.

Task 3.3. (year 3-4) *Understand* the PoA/PoS under the solution concepts developed in Task 3.1. *Evaluate* the welfare consequences (quality of the system) in voting scenarios of the alternative notions of rationality.

3.2.5 WP4: Development of a Collective Decision Making Platform, and Promotion of its Uses.

- Coordinators: Sylvain Bouveret (rattached to LIP6) and Vincent Merlin (CREM).
- Main participants: Stéphane Airiau, Denis Bouyssou, Tristan Cazenave, Sébastien Courtin, Angelo Fanelli, Lucie Galand, Jérôme Lang, Jean-François Laslier, Isabelle Lebon, Nicolas Maudet, Boniface Mbih, Vincent Merlin, Matias Nuñez, Fanny Pascual, Patrice Perny, Agnieszka Rusinowska, Alexis Tsoukiàs, Paolo Viappiani + postdoc(s) + master students.

The goal of this work package is to implement and test the algorithms and protocols worked out in WP1, WP2 and WP3 and make them accessible to the users (citizens and decision makers). It aims at promoting new forms of democratic participation in low-stake situations, by enabling citizens to express precisely their preferences, and by giving them tools that should help them make better common decisions. The platform will also be used as a pedagogical tool in events such as “Fête de la science”, and used to perform research experiments (in laboratories or in decentralized contexts).

Task 4.1: Developing and testing the software platform As a result of the advent of computer networks and connected interfaces, people have the opportunity to participate to an increasing number of collective decision making situations. Social choice theory provides a set of principles, rules and methods in this context. However, rather surprisingly, there is an obvious lack of practical tools to help people make collective decisions in everyday low-stake situations.

One of the key outcomes of our project is to bridge this gap by developing a large-scale platform for various tasks of collective decision making. Currently, the most well-known application for everyday life collective decision making is probably the poll application Doodle™, which allows people to express their preferences, and allows the decision maker to have a survey of these preferences. However, these systems soon become unsatisfactory as soon as the individual preferences are more complex and subtle than just approving/disapproving some options (or remaining neutral), and hence fail to give satisfactory answers in most collective decision making situations. To overcome these limitations, Sylvain Bouveret has started to develop an online voting platform, *Whale* (`http://whale3.noiraudes.net/`). With *Whale*, users can currently create open or sealed ballot polls, cast their votes using expressive preference modes (orders, utility functions,...), and use classical voting rules to figure out which issues are the best ones among the candidate issues. Although *Whale* improves upon *Doodle*™, many important decision making mechanisms are not currently implemented.

The idea of this work package is to start from this existing platform and use it as a basis for a more general collective decision making platform. It will be extended in many ways to

- *more general voting problems*, such as multiple referenda, multi-winner elections, voting rules with different kinds of inputs (*e.g.*, Range Voting or *Majority Judgement*), or the selection of most representative voters;
- *other collective decision problems*, such as the fair division of indivisible items and *coalition structure generation* (such as the search for stable or optimal coalition structures in hedonic games or group activity selection);
- new *tasks*, such as multi-stage protocols or automated generation of certificates for verifying the solution (the software should be able to handle interactive decisions, made in several steps).

Moreover, the platform will be open to *new purposes*, such as scientific dissemination purposes, and for voting experiments.

One important aspect of the work on this platform is that it will be developed with several goals in mind.

- Building a *practical tool*, which can be used by inexperienced persons for everyday-life collective decision problems or low-stake voting situations. In practice, this platform could be used as a tool for promoting and implementing participatory democracy.

- Building a *pedagogical tool*, which can be used by teachers to illustrate different concepts from collective decision theory, different voting procedures, fair division methods, and so on. First, a particular attention will be paid to the pedagogical presentation of the different methods that will be proposed to the users. Next *Whale* shall be used as a demonstration platform for popularizing collective decision making during science fairs such as *La Fête de la Science*, for example by getting people to participate to small voting situations and to test several voting procedures.

- Building a *research tool*, which can be used as an experimental platform for testing the algorithms and protocols developed in the other work packages. A crucial functionality of *Whale* will be to store the data collected (for example, the votes) in a database. This database will then be usable by researchers, for instance for mining purposes. This would complement the *PrefLib* effort [117]. The functionalities will be built in order to match as much as possible the usual requirements of experimental economics (anonymity, possibility to build teams and change partners, determination of the gains, etc.)

Whale has already been used for running a few hundreds of polls in real situations (even for recruitment committees in academical contexts). After a first period of development in order to introduce new functionalities we plan to carry out some additional tests on specific real-world applications. The advances on task 4.1 cannot be separated from the dissemination activities of task 4.2, that we describe below.

Task 4.2: Testing and Using the Platform, in Real Life and in Laboratories As said previously, the main objective of this work package is to propose to the citizens tools for decision making in real life context, and create an interaction between users and researchers. The different steps of this task are extremely imbricated with the task 4.1. The collective decision making platform will be tested on real world situations and real users. For example, we plan to test the platform on team formation for experimental works at Ensimag: a coalition formation problem, consisting in constructing groups of students assigned to projects, that currently occurs several times a year and for which no really satisfactory solution is currently implemented. The first tests will be carried out with groups of students, voluntary participants, specific associations, to determine whether the interface is user friendly, the pedagogic aspects easily comprehensible, and its relevance clearly established. After launching the application, the platform will keep track of the comments of the users, and let the possibility for them to let the researcher analyze the (anonymous) data.

Beyond these tests, this platform is also a great opportunity for running voting experiments in laboratory or *in field*. The aim of these experiments is to observe the behaviour of real voters facing a particular voting situation, to elaborate models of human rationality and strategic behaviour for example, and to use observation to validate these models. Here, a decision-making platform like the one that will be developed in this project greatly eases the design and the implementation of such experiments, because (i) researchers can easily define various voting situations, (ii) it will be easier to participate, as the online system does not even require the voters to be physically present, and (iii) the results are directly available numerically.

In particular, we will carry out several voting laboratory experiments which will aim at testing the effect of information on the voters' choices. For several voting procedures, we will run repeated votes in order to figure out to which extent the knowledge about the results in a given round influences and modifies the choices expressed by the voters in subsequent rounds. With this information we will be able to determine how many rounds are needed to end up in a situation where the information about the ballots cast in previous rounds will not modify any voter's preferences in subsequent rounds anymore (if such a stable situation ever exists). This dynamic point of view about the voting procedures tested should allow us to try different models of rationality, and different models of strategic votes as well, which require coordination among the agents on their choices.

The voting platform will not only help us to carry out these experiments in laboratory, but will also allow us to study real decision making situations (*e.g* in real contexts such as associations, sociétés savantes...), and to compare the results observed in these *in field* experiments with the ones obtained in laboratory.

Task planning and human resources

As said previously, the different parts of Task 4.1 and Task 4.2 are imbricated. The current version of *Whale* has been developed with the idea of clearly separating the core application dedicated to the creation of polls and to the expression of preferences by the voters on the one hand, and the voting procedures on the other hand. Keeping this idea in mind for further development, it will thus be easy to work separately on the interface and on the collective decision making algorithms (from WP1, WP2 and WP3). The task lasts the entire project, but the main development will be carried out at the beginning. Sylvain Bouveret will lead the development of the platform, but the platform will receive contributions from various members of the project, especially the ones working on the algorithmic part. The economists involved in the project will interact with the computer scientists on the choice of methods to implement, and help design the pedagogical and experimental aspects. They will use the platform to conduct experiments on the way citizens strategically acts in face of some voting rules or allocation rules. Both communities will participate to pedagogical and dissemination activities.

The development will be organized as follows:

- **Year 1:** State of the art, decisions about the modules to develop first (3 months). Code refactoring and extension of the core of *Whale* to a generic platform capable of embedding various collective decision-making problems (3 months). Extension to more general collective decision problems (6 months).

- **Year 2:** Setting up the platform for users (3 months). Tests on real situations and live demonstrations (6 months). Realization of a voting experiment in laboratory (3 months).

- **Year 3:** Feedback from the first users, and strategic choice about the next modules to implement (3 months) for the remaining of the program. Extension to other collective decision making problems (6 months) for the integration of algorithms from other WPs. Tests on real situations (3 months). During year 3, a summer school on "computational social choice" will be organized in Caen. One of its main objectives will be to present the field to PhD students, to advertise the latest researches on the subject, and to validate the use of the platform as a teaching and experimental tool. Co-funding by the CNRS summer school program will be asked for.

- **Year 4:** Extension to other tasks related to voting (3 months) and modifications due to the previous feedback. Tests on real situations, live demonstration, and second set of laboratory experiments. Comparison between controlled experiments and real life uses (6 months). Final report on the use of the platform and its future.

The development of the core of the platform is guaranteed by the competencies in programming of the computer science team. The master-2 trainees will also play an important role, by developing subparts of the projets. The post doc (ideally an unique candidate who should spend some time in the three leading labs over the 2 years of his/her contract) should play a role in developing and supervising part of the platform. Social networks (Facebook, Twitter, and so on) will also be part of our dissemination strategy.

3.3 Task schedule

We will follow the following schedule. Each work package has a color: purple for WP0, red for WP1, blue for WP2, green for WP3 and brown for WP4. Each block corresponds to a set of tasks. At each end of block we will produce a deliverable summarizing the work done during the block. Two intermediate workshops are planned (year 2 and 3) and a final workshop is scheduled at the end of year 4.

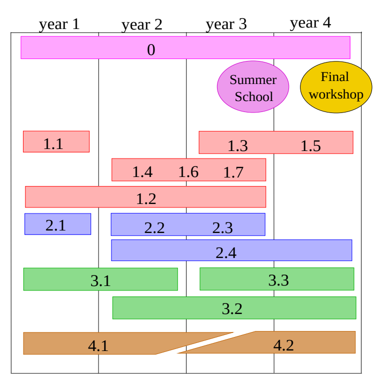

3.4 Consortium Description

One of the most prominent features of our consortium is the distribution of expertise. Our project is ambitious because it bridges (computational) social choice, algorithmic game theory, cooperative game theory, and algorithmic decision theory. For each of these fields we have managed to have top researchers.

social choice theory and game theory Bouyssou, Courtin, Grabisch, Laslier, Lebon, Mbih, Merlin, Nuñez, Rusinowska

computational social choice Airiau, Bouveret, Chevaleyre, Cornaz, Galand, Hudry, Lang, Maudet, Monnot, Spanjaard

algorithmic game theory Bampis, Escoffier, Fanelli, Gourvès, Monnot, Moretti, Pascual, Spanjaard

algorithms and complexity Bouveret, Cazenave, Cornaz, Escoffier, Fanelli, Galand, Gourvès, Monnot, Spanjaard

communication & elicitation protocols Bouveret, Chevaleyre, Lang, Maudet, Monnot, Perny, Tsoukiàs, Viappiani

In addition to the academic researchers, we ask for a two-year postdoc funding (years 2-3), ideally for a single researcher who will stay for two years, or else two different researchers. We think it will be best for the postdoc(s) (and for the project) that (s)he/they should spend some time in all three institutes (LAMSADE, LIP6, CREM), which is why we asked for postdoc funding in all three institutes; but if the ANR thinks it is simpler to have the postdoc in only two places, or even in only one place, we will follow their advice. The postdoc researcher(s) will be (a) computer scientist(s) with a background on computational social choice and/or algorithmic game theory and/or algorithmic decision theory. She/He/They will work on all work packages, and s(he)/they will also spent a part of their time (together with some members of the project) developing some functionalities of the platform. We also ask for master (M2R) trainee grants; these students will take an active part on developing the platform and will be supervised by the members of the project, including the post-doc(s).

The project coordinator, Jérôme Lang, is qualified for coordinating this project for the following two reasons. First, he is one of the founders of the area of computational social choice; he was the program co-chair of the first International Workshop on Computational Social Choice, and is a co-editor of the forthcoming Handbook of Computational Social Choice. He has very strong connections with the rest of the community. Second, he has already been the coordinator of several similar projects. He was the coordinator of the ANR project *Preference Handling and Aggregation over Combinatorial Domains* (PHAC, 2005–2009), the co-leader (with Patrice Perny) of the project "De l'intelligence artificielle aux sciences de la décision : modèles formels et mise en oeuvre de processus de prise de décision individuelle et collective" (2002–2005), within the interdisciplinary program *Société de l'Information* (CNRS) and co-leader (with Andreas Herzig) of the project "La pertinence dans l'interprétation d'actions" within the interdisciplinary program *Cognitique* (2001–2003).

3.5 Scientific Justification of Requested Resources

The requested resources concern the following budget lines: travel expenses for participants; invitation of foreign researchers for joint research and seminars; invitation of colleagues to the final workshop; invitation of colleagues to the summer school; two year post-doc grant; master student grants.

- *travel expenses*: On average, for a participant we count 8,000 € for the the whole duration of the project: one major international conference per year (4*1,500 = 6,000 €), and participation to four smaller events, such as French conferences, regular meetings of the project, short term research visits (4*500 = 2,000 €). Of course, these figures will be higher for the participants who devote more than 25% of their time to the project, and smaller for the other participants. This gives around 88,000 € for LAMSADE, 72,000 € for LIP6 and 72,000 € for CREM. We add a little bit more on top of these amounts (5,000 €) to take into account the travel expenses of the nonpermanent researchers who will be involved in the project and whose activity is not taken into account here: postdoc (funded by the project) and possibly some PhDs (not funded by the project) and/or master students. We end up with the following figures: LAMSADE = 93,000 €, LIP6 = 77,000 €, CREM = 77,000 €.

- *equipment*: The project does not ask for financial help of ANR for equipment. These expenses will be covered by the research centers.

- *postdoc*: We ask for two year post-doc grant (9 months spent at LAMSADE, 9 months spent at LIP6, and 6 months spent at CREM), for a total of 108,000 €.

- *invitation of foreign researchers* (for workshops, regular seminars, or joint research work): 8,000 € for LAMSADE, 4,000 € for LIP6, 9,000 € for CREM. We plan to invite (junior and senior) researchers of the field(s), either in the framework of our workshops, or else to give a seminar and so stay a few days to work with us. Notice that foreign invited researchers can visit both LAMSADE and LIP6 during the same stay, as both are located in Paris.

- *summer school*: One important point of the dissemination strategy is the organization of a summer school on year 3. It will be located in Caen as it is possible to host the participants for a low cost at the student residencies. In parallel, an application to the CNRS summer school program will be made to obtain extra funds.

- *final workshop*. The project will end with a final workshop, where international guests will be invited. It will most likely take place at LAMSADE (10,000 €)

- *grants for master students (M2R)*: we ask for 4 grants at LAMSADE and 6 at LIP6, as both centers have access to master students in computer science. A master student grant amounts to 500 € per month, for six months, that is, 3,000 €.

4 Dissemination, Protection, Exploitation of Results, Global Impact of the Proposal

The project is purely academic, therefore intellectual property is not an issue. The use of the platform developed by the participants will be free and accessible to anyone. The short-term expected scientific impact is of academic nature: the dissemination of results will be made through communications at major conferences and publications in the major journals of the domains involved, both in computer science and in economics: for instance, *Artificial Intelligence Journal*, *Journal of Artificial Intelligence Research*, *Theoretical Computer Science*,

Table 1: Summary of the Budget

Item	LAMSADE	CREM	LIP6	Total
Travels	93,000 €	77,000 €	77,000 €	247,000 €
Invitations	8,000 €	9,000 €	4,000 €	21,000 €
Final Workshop	10,000 €	0	0	10,000 €
Summer School	0	10,000 €	0	10,000 €
Post-doc grant	40,500 €	27,000 €	40,500 €	108,000 €
Master grants	12,000 €	0	18,000€	30,000 €
Experiments	0	9,000 €	0	9,000 €
Total	163,500 €	132,000 €	139,500 €	435,000 €
4%	6,540 €	5,280 €	5,580 €	17,400 €
Grand Total	170,040 €	137,280 €	145,080 €	452,400 €

Games and Economic Behavior, *Mathematical Social Science*, *Social Choice and Welfare*, *Theory and Decision*, *Electoral Studies*, *Internationnal Journal of Game Theory*, *Review of Economic Design*, *ACM Transactions on Computer Systems*, *European Journal of Operations Research*, *International Joint Conference in Artificial Intelligence*, *ACM Conference on Economics and Computation*, *Workshop on Internet Economics*, *International Conference on Autonomous Agents and Multi-Agent Systems*, *International Symposium on Algorithmic Game Theory*, *International Workshop on Computational Social Choice*. The French community has a strong expertise in all the fields concerned, which allows us to expect from the collaborations some research articles of a very high quality with a very high impact.

In the middle and long term, we expect our results to have a significant impact on societal and economic issues, especially through the collective decision platform. Its wide accessibility and its diffusion to a large public (national and international) will first contribute to the visibility of the project, and its use by laypeople will contribute to make the public aware of the principles and methods for collective decision making.

Our project fits completely the axis «Droit, démocratie, gouvernance et nouveaux référentiels» of the challenge «Sociétés innovantes, intégrantes et adaptatives». Its main motivations are to analyze various possible collective decision mechanisms so as to identify those that will be applicable by human groups in practice, and to give citizens access to new methods and paradigms for collective decision making. We therefore expect some of these mechanisms to be used in practice by groups of agents, in various contexts (local communities such as towns, companies, associations, hospitals, universities, schools, co-ownerships, communities formed by social networks etc.) and for various application domains (primarily electronic democracy, but also work team formation, car pooling systems and other facility sharing problems, group recommendation with possible applications to tourism, etc.). A mid-term impact of the project is to *help promoting participatory democracy, through the feasibility study, the implementation and the experimentation of realistic and satisfactory mechanisms and their diffusion to a large public.*

References

[1] http://tja2012.lip6.fr/. Ecole de printemps sur la théorie des jeux algorithmique, Paris, june 2012.

[2] D. Ahn and S. Oliveros. Approval Voting and Scoring Rules with Common Values. mimeo, University of California, Berkeley, 2011.

[3] S. Airiau and U. Endriss. Iterated majority voting. In *ADT*, pages 38–49, 2009.

[4] S. Airiau and U. Endriss. Multiagent resource allocation with sharable items: simple protocols and Nash equilibria. In *AAMAS*, pages 167–174, 2010.

[5] S. Airiau, U. Endriss, U. Grandi, D. Porello, and J. Uckelman. Aggregating dependency graphs into voting agendas in multi-issue elections. In *IJCAI*, pages 18–23, 2011.

[6] N. Andjiga, B. Mbih, and I. Moyouwou. Manipulation of voting schemes with restricted beliefs. *Journal of Mathematical Economics*, 44:1232–1242, 2008.

[7] E. Angel, E. Bampis, and F. Pascual. Truthful algorithms for scheduling selfish tasks on parallel machines. In *WINE*, pages 698–707, 2005.

[8] E. Angel, E. Bampis, and F. Pascual. *Handbook of Approximation Algorithms and Metaheuristics*, chapter Algorithmic Game Theory and Scheduling. Chapman & Hall, 2007.

[9] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *FOCS*, pages 295–304. IEEE Computer Society, 2004.

[10] K. R. Apt and G. Schäfer. Selfishness level of strategic games. In *SAGT*, pages 13–24, 2012.

[11] K. Arrow, A. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*. North-Holland, 2002 (Vol. 1), 2012 (Vol. 2).

[12] R. J. Aumann. Acceptable points in general cooperative n-person games. *Contributions to the theory of games*, IV, 1959.

[13] H. Aziz, F. Brandt, and H. G. Seedig. Computing desirable partitions in additively separable hedonic games. *Artif. Intell.*, 195:316–334, 2013.

[14] Y. Bachrach, E. Elkind, and P. Faliszewski. Coalitional voting manipulation: A game-theoretic perspective. In *IJCAI*, pages 49–54, 2011.

[15] C. Ballester. NP-completeness in hedonic games. *Games and Economic Behaviour*, 49:1–30, 2004.

[16] N. Bansal and M. Sviridenko. The Santa Claus problem. In J. M. Kleinberg, editor, *STOC*, pages 31–40. ACM, 2006.

[17] J.-P. Barthélemy, A. Guénoche, and O. Hudry. Median linear orders: heuristics and a branch and bound algorithm. *EJOR*, 3(42):313–325, 1989.

[18] J. Bartholdi, C. Tovey, and M. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6:227–241, 1989.

[19] A. Baujard, H. Igersheim, I. Lebon, F. Gavrel, and J.-F. Laslier. Who's favored by evaluative voting? an experiment conducted during the 2012 french presidential election. *Electoral Studies*, 34(June):131–145, 2014.

[20] R. Berghammer, A. Rusinowska, and H. de Swart. An interdisciplinary approach to coalition formation. *EJOR*, 195(2):487–496, 2009.

[21] D. Bernheim, B. Peleg, and M. Whinston. Coalition-proof Nash equilibria: I concepts. *Journal of Economic Theory*, 42:1–12, 1987.

[22] S. Bervoets and V. Merlin. Gerrymander-proof representative democracies. *International Journal of Game Theory*, 41(3):473–488, 2012.

[23] N. Betzler, R. Bredereck, J. Chen, and R. Niedermeier. Studies in computational aspects of voting - a parameterized complexity perspective. In *The Multivariate Algorithmic Revolution and Beyond*, pages 318–363, 2012.

[24] V. Bilò, A. Fanelli, M. Flammini, and L. Moscardelli. Graphical congestion games. In *WINE*, pages 70–81, 2008.

[25] A. Blais, K. van der Straeten, and J.-F. Laslier. Vote au pluriel: How people vote when offered to vote under different rules ? *PS: Political Science and Politics*, 46(2):324–328, April 2013.

[26] C. Boutilier, J. Lang, J. Oren, and H. Palacios. Robust winners andwinner determination policies under candidate uncertainty. In *AAAI*, 2014.

[27] L. Bouton and M. Castanheira. One Person, Many Votes: Divided Majority and Information Aggregation. *Econometrica*, 80:43–87, 2012.

[28] S. Bouveret, U. Endriss, and J. Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *ECAI*, pages 387–392, 2010.

[29] S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *J. Artif. Intell. Res. (JAIR)*, 32:525–564, 2008.

[30] S. Bouveret and J. Lang. A general elicitation-free protocol for allocating indivisible goods. In *IJCAI*, pages 73–78, 2011.

[31] D. Bouyssou and M. Pirlot. An axiomatic analysis of concordance-discordance relations. *European Journal of Operational Research*, 199(2):468–477, 2009.

[32] F. Brandt, V. Conitzer, and U. Endriss. *Computational social choice (G. Weiss, ed.)*. MIT Press, to appear. 2nd edition.

[33] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge Press. In preparation.

[34] F. Brandt, F. Fischer, and P. Harrenstein. On the rate of convergence of fictitious play. *Theory Comput. Syst.*, 53(1):41–52, 2013.

[35] S. Brânzei, I. Caragiannis, J. Morgenstern, and A. Procaccia. How bad is selfish voting? In *AAAI*, 2013.

[36] R. Bredereck, J. Chen, and G. Woeginger. Are there any nicely structured preference profiles nearby? In *IJCAI*, 2013.

[37] L. Buenrostro, A. Dhillon, and P. Vida. Scoring Rule Voting Games and Dominance Solvability. *Social Choice and Welfare*, 40:329–352, 2012.

[38] I. Caragiannis, J. Covey, M. Feldman, C. Homan, Ch. Kaklamanis, N. Karanikolas, A. Procaccia, and J. Rosenschein. On the approximability of Dodgson and Young elections. *Artif. Intell.*, 187:31–51, 2012.

[39] I. Caragiannis, Ch. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. The efficiency of fair division. *Theory Comput. Syst.*, 50(4):589–610, 2012.

[40] G. Chalkiadakis, E. Elkind, and M. Wooldridge. *Computational Aspects of Cooperative Game Theory*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2011.

[41] I. Charon and O. Hudry. An updated survey on the linear ordering problem for weighted or unweighted tournaments. *Annals OR*, 175(1):107–158, 2010.

[42] Y. Chevaleyre, P. Dunne, U. Endriss, J. Lang, N. Maudet, and Juan A. Rodriguez-Aguilar. Multiagent resource allocation. *Knowledge Engineering Review. Special Issue on the AgentLink III Technical Forums*, 20(2):143–149, 2005.

[43] Y. Chevaleyre, U. Endriss, S. Estivie, and N. Maudet. Multiagent resource allocation in k -additive domains: preference representation and complexity. *Annals OR*, 163(1):49–62, 2008.

[44] Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. A short introduction to computational social choice. In *Proc. of the 33rd Conference on Current Trends in Theory and Practice of Computer Science*, pages 51–69. Springer-Verlag, 2007.

[45] Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference handling in combinatorial domains: From AI to social choice. *AI Magazine*, 29(4):37–46, 2008.

[46] Y. Chevaleyre, U. Endriss, and N. Maudet. Trajectories of goods in distributed allocation. In *AAMAS (2)*, pages 1111–1118, 2008.

[47] Y. Chevaleyre, U. Endriss, and N. Maudet. Simple negotiation schemes for agents with simple preferences: sufficiency, necessity and maximality. *Autonomous Agents and Multi-Agent Systems*, 20(2):234–259, 2010.

[48] Y. Chevaleyre, J. Lang, N. Maudet, and J. Monnot. Compilation and communication protocols for voting rules with a dynamic set of candidates. In *TARK*, pages 153–160, 2011.

[49] Y. Chevaleyre, J. Lang, N. Maudet, J. Monnot, and L. Xia. New candidates welcome! possible winners with respect to the addition of new candidates. *Mathematical Social Sciences*, 64(1):74–88, 2012.

[50] Y. Chevaleyre, J. Lang, N. Maudet, and Guillaume Ravilly-Abadie. Compiling the votes of a subelectorate. In *IJCAI*, pages 97–102, 2009.

[51] G. Christodoulou, L. Gourvès, and F. Pascual. Scheduling selfish tasks: About the performance of truthful algorithms. In G. Lin, editor, *COCOON*, volume 4598 of *Lecture Notes in Computer Science*, pages 187–197. Springer, 2007.

[52] R. Congar and V. Merlin. A characterization of the maximin rule in the context of voting. *Theory and Decision*, 72(1):131–147, 2012.

[53] V. Conitzer. Computing Slater rankings using similarities among candidates. In *AAAI*, pages 613–619, 2006.

[54] V. Conitzer, J. Lang, and L. Xia. How hard is it to control sequential elections via the agenda? In *IJCAI*, pages 103–108, 2009.

[55] V. Conitzer, J. Lang, and L. Xia. Hypercubewise preference aggregation in multi-issue domains. In *IJCAI*, pages 158–163, 2011.

[56] V. Conitzer and T. Sandholm. Communication complexity of common voting rules. In *ACM Conference on Electronic Commerce*, pages 78–87, 2005.

[57] D. Cornaz, L. Galand, and O. Spanjaard. Bounded single-peaked width and proportional representation. In *ECAI*, pages 270–275, 2012.

[58] D. Cornaz, L. Galand, and O. Spanjaard. Kemeny elections with bounded single-peaked or single-crossing width. In *IJCAI*, 2013.

[59] S. Courtin, B. Mbih, and I. Moyouwou. Susceptibility to coalitional strategic sponsoring: The case of parliamentary agendas. *Public Choice*, 144:133–151, 2010.

[60] S. Courtin, B. Mbih, I. Moyouwou, and T. Senné. The reinforcement axiom under sequential positional rules. *Social Choice and Welfare*, 35:473–500, 2010.

[61] S. Courtin and M. Nuñez. Dominance Solvable Approval Voting Games. mimeo, THEMA- Université de Cergy-Pontoise, 2013.

[62] A. Darmann, E. Elkind, S. Kurz, J. Lang, J. Schauer, and G. J. Woeginger. Group activity selection problem. In *WINE*, pages 156–169, 2012.

[63] B. de Keijzer, S. Bouveret, T. Klos, and Y. Zhang. On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In *ADT*, pages 98–110, 2009.

[64] F. De Sinopoli, G. Iannantuoni, and C. Pimienta. Counterexamples on the superiority of approval vs. plurality. *Journal of Public Economic Theory*, 2013.

[65] H. de Swart, R. Berghammer, and A. Rusinowska. Computational social choice using relation algebra and relview. In *RelMiCS*, pages 13–28, 2009.

[66] C. Delort, O. Spanjaard, and P. Weng. Committee selection with a weight constraint based on a pairwise dominance relation. In *ADT*, pages 28–41, 2011.

[67] Y. Desmedt and E. Elkind. Equilibria of plurality voting with abstentions. In *ACM Conference on Electronic Commerce*, pages 347–356, 2010.

[68] A. Dhillon and B. Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46(1):55–75, 2004.

[69] B. Dutta and J.-F. Laslier. Costless honesty in voting. In *10th International Meeting of the Society for Social Choice and Welfare*, 2010.

[70] B. Dutta and A. Sen. Nash implementation with partially honest individuals. *Games and Economic Behavior*, 74:154–169, 2012.

[71] E. Elkind, U. Grandi, F. Rossi, and A. Slinko. Games Gibbard-Satterthwaite manipulators play. In *COMSOC14*.

[72] E. Elkind, J. Lang, and A. Saffidine. Choosing collectively optimal sets of alternatives based on the condorcet criterion. In *IJCAI*, pages 186–191, 2011.

[73] B. Escoffier, L. Gourvès, and J. Monnot. On the impact of local taxes in a set cover game. In Boaz Patt-Shamir and Tinaz Ekim, editors, *SIROCCO*, volume 6058 of *Lecture Notes in Computer Science*, pages 2–13. Springer, 2010.

[74] B. Escoffier, L. Gourvès, and J. Monnot. Fair solutions for some multiagent optimization problems. *Autonomous Agents and Multi-Agent Systems*, 26(2):184–201, 2013.

[75] B. Escoffier, L. Gourvès, J. Monnot, and S. Moretti. Cost allocation protocols for network formation on connection situations. In *VALUETOOLS*, pages 228–234, 2012.

[76] A. Fabrikant, Ch. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In L. Babai, editor, *STOC*, pages 604–612. ACM, 2004.

[77] U. Faigle and M. Grabisch. Values for Markovian coalition processes. *Economic Theory*, 51:505–538, 2012. DOI: 10.1007/s00199-011-0617-7.

[78] P. Faliszewski, E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. A richer understanding of the complexity of election systems. *CoRR*, abs/cs/0609112, 2006.

[79] P. Faliszewski and A. D. Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.

[80] B. Faltings. Incentive-compatible social choice. In *Proc. IEEE/WIC/ACM International Conference on Web Intelligence*, pages 8–14, 2004.

[81] A. Fanelli, M. Flammini, and L. Moscardelli. The speed of convergence in congestion games under best-response dynamics. *ACM Transactions on Algorithms*, 8(3):25, 2012.

[82] T. Feddersen and W. Pesendorfer. Voting behavior and information aggregation in elections with private information. *Econometrica*, 65(5):1029–1058, 1997.

[83] E. Fehr and K. Schmidt. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114:817–868, 1999.

[84] I. Gilboa and E. Zemel. Nash and correlated equilibria: some complexity considerations. *Games and Economic Behavior*, 1:80–93, 1989.

[85] J. Goertz and F. Maniquet. On the informational efficiency of simple scoring rules. *Journal of Economic Theory*, 146:1464–1480, 2011.

[86] B. Golden and P. Perny. Infinite order Lorenz dominance for fair multiagent optimization. In *AAMAS*, pages 383–390, 2010.

[87] J. Goldsmith, J. Lang, N. Mattei, and P. Perny. Rank-dependent scoring rules. In *AAAI14*.

[88] Ch. Gonzales, P. Perny, and S. Queiroz. Preference aggregation with graphical utility models. In *AAAI*, pages 1037–1042, 2008.

[89] L. Gourvès, J. Monnot, S. Moretti, and N. K. Thang. Congestion games with capacitated resources. In *SAGT*, pages 204–215, 2012.

[90] L. Gourvès, J. Monnot, and O. Telelis. Selfish scheduling with setup times. In *WINE*, pages 292–303, 2009.

[91] L. Gourvès, J. Monnot, and O. Telelis. *Just-in-Time Systems*, chapter Strategic Scheduling Games: Equilibria and Efficiency. Springer, 2012.

[92] L. Gourvès and S. Moretti. *Progress in Combinatorial Optimization*, chapter Combinatorial optimization problems arising from interactive congestion situations. ISTE Ltd and J. Wiley & Sons Inc, 2011.

[93] M. Grabisch and Y. Funaki. A coalition formation value for games in partition function form. *European Journal of Operations Research*, 221:175–185, 2012.

[94] E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of Kemeny elections. *Theor. Comput. Sci.*, 349(3):382–391, 2005.

[95] D. Herreiner and C. Puppe. A simple procedure for finding equitable allocations of indivisible goods. *Social Choice and Welfare*, 19(2):415–430, 2002.

[96] M. Hoefer and A. Skopalik. Stability and convergence in selfish scheduling with altruistic agents. In *WINE*, pages 616–622, 2009.

[97] O. Hudry. A survey on the complexity of tournament solutions. *Mathematical Social Sciences*, 57(3):292–303, 2009.

[98] O. Hudry. On the complexity of Slater's problems. *European Journal of Operational Research*, 203(1):216–221, 2010.

[99] O. Hudry. On the computation of median linear orders, of median complete preorders and of median weak orders. *Mathematical Social Sciences*, 64(1):2–10, 2012.

[100] J. Bartholdi, III, C. Tovey, and M. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6, 1989.

[101] T. Kalinowski, N. Narodytska, T. Walsh, and L. Xia. Strategic behavior in a decentralized protocol for allocating indivisible goods. In *AAAI*, 2013.

[102] N. Kartik and A. Tercieux. Implementation with evidence. *Theoretical Economics*, 7:323–355, 2012.

[103] E. Koutsoupias and Ch. Papadimitriou. Worst-case equilibria. In Ch. Meinel and S. Tison, editors, *STACS*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413. Springer, 1999.

[104] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.

[105] J. Lang, J. Mengin, and L. Xia. Aggregating conditionally lexicographic preferences on multi-issue domains. In *Proceedings of CP-2012*, pages 973–987, 2012.

[106] J. Lang, M. S. Pini, F. Rossi, D. Salvagnin, K. B. Venable, and T. Walsh. Winner determination in voting trees with incomplete preferences and weighted votes. *Autonomous Agents and Multi-Agent Systems*, 25(1):130–157, 2012.

[107] J. Lang and L. Xia. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences*, 57(3):304–324, 2009.

[108] J.-F. Laslier and J. W. Weibull. A strategy-proof Condorcet jury theorem. *Scandinavian Journal of Economics*. Forthcoming.

[109] J.F. Laslier. Strategic approval voting in a large electorate. *Journal of Theoretical Politics*, 21:113–136, 2009.

[110] J. Ledyard. *Handbook of Experimental Economics*, chapter Public goods: A survey of experimental resesarch. Princeton University Press, 1997.

[111] J. Lesca and P. Perny. LP solvable models for multiagent fair allocation problems. In *ECAI*, pages 393–398, 2010.

[112] J. Lesca and P. Perny. Almost-truthful mechanisms for fair social choice functions. In *ECAI*, pages 522–527, 2012.

[113] O. Lev and J. Rosenschein. Convergence of iterative voting. In *AAMAS*, pages 611–618, 2012.

[114] D. Levine. Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1:593–622, 1998.

[115] Ch. Lumet, S. Bouveret, and M. Lemaître. Fair division of indivisible goods under risk. In *ECAI*, pages 564–569, 2012.

[116] H. Matsushima. Role of honesty in full implementation. *Journal of Economic Theory*, 139:353–359, 2008.

[117] N. Mattei and T. Walsh. Preflib: A library for preferences http://www.preflib.org. In *ADT*, pages 259–270, 2013.

[118] N. Maudet, M. S. Pini, K. B. Venable, and F. Rossi. Influence and aggregation of preferences over combinatorial domains. In *AAMAS*, pages 1313–1314, 2012.

[119] R. McKelvey and R. Niemi. A multistage game representation of sophisticated voting for binary procedures. *Journal of Economic Theory*, 18(1):1–22, 1978.

[120] R. McKelvey and R. Wendell. Voting equilibria in multidimensional choice spaces. *Mathematics of Operations Research*, 1(2):144–158, 1976.

[121] R. Meir, M. Polukarov, J. Rosenschein, and N. Jennings. Convergence to equilibria in plurality voting. In *AAAI*, 2010.

[122] H. Moulin. Dominance solvable voting schemes. *Econometrica*, 47(6):1137–51, 1979.

[123] R. Myerson. Comparison of scoring rules in Poisson voting games. *Journal of Economic Theory*, 103:219–251, 2002.

[124] R. Myerson and R. Weber. A theory of voting equilibria. *The American Political Science Review*, 87(1):102–114, 1976.

[125] N. Nisam, T. Roughgarden, É. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[126] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1-2):166–196, 2001.

[127] M. Nuñez. Condorcet consistency of approval voting: A counter example on large Poisson games. *Journal of Theoretical Politics*, 22:64–84, 2010.

[128] M. Nuñez. The Strategic Sincerity of Approval Voting. *Economic Theory*, 56(1):157–189, 2014.

[129] M. Nuñez and J.-F. Laslier. Bargaining through approval. Working paper, 2014.

[130] S. Obraztsova, E. Markakis, and D. Thompson. Plurality voting with truth-biased agents. In *SAGT*, pages 26–37, 2013.

[131] Ch. Papadimitriou and X. Deng. On the complexity of cooperative solution concepts. *Mathematical Operations Research*, 19:257–266, 1994.

[132] A. Procaccia and J. Rosenschein. The communication complexity of coalition formation among autonomous agents. In *AAMAS*, pages 505–512, 2006.

[133] A. D. Procaccia. A note on the query complexity of the Condorcet winner problem. *Inf. Process. Lett.*, 108(6):390–393, 2008.

[134] T. Roughgarden and É. Tardos. *Algorithmic Game Theory*, chapter Introduction to the inefficiency of equilibria. Cambridge University Press, 2007.

[135] I. Segal. The communication requirements of social choice rules and supporting budget sets. *Journal of Economic Theory*, 136:341–378, 2007.

[136] T. Service and J. Adams. Communication complexity of approximating voting rules. In *AAMAS*, pages 593–602, 2012.

[137] P. Skowron, P. Faliszewski, and A. Slinko. Fully proportional representation as resource allocation: Approximability results. In *IJCAI*, 2013.

[138] B. Sloth. The theory of voting and equilibria in noncooperative games. *Games and Economic Behaviour*, 1(5):152–169, 1993.

[139] S. C. Sung and D. Dimitrov. On core membership testing for hedonic coalition formation games. *Operations Research Letters*, 35:155–158, 2007.

[140] D. Thompson, O. Lev, K. Leyton-Brown, and J. Rosenschein. Empirical aspects of plurality election equilibria. In *Proceedings of COMSOC-2012*, 2012.

[141] K. van der Straeten, N. Sauger, J.-F. Laslier, and A. Blais. Sorting out mechanical and psychological effects in candidate elections: An appraisal with experimental data. *British Journal of Political Science*, 44(4):37–944, october 2013.

[142] P. Viappiani and C. Boutilier. Recommendation sets and choice queries: There is no exploration/exploitation tradeoff! In *AAAI*, 2011.

[143] G. J. Woeginger. Banks winners in tournaments are difficult to recognize. *Social Choice and Welfare*, 20(3):523–528, 2003.

[144] L. Xia and V. Conitzer. Stackelberg voting games: Computational aspects and paradoxes. In *AAAI*, 2010.

[145] L. Xia, V. Conitzer, and J. Lang. Strategic sequential voting in multi-issue domains and multiple-election paradoxes. In *ACM Conference on Electronic Commerce*, pages 179–188, 2011.

[146] M. Zuckerman, P. Faliszewski, V. Conitzer, and J. Rosenschein. An NTU cooperative game theoretic view of manipulating elections. In *WINE*, pages 363–374, 2011.