THÈSE DE DOCTORAT

de l'Université de recherche Paris Sciences et Lettres
PSL Research University

Préparée à l'Université Paris-Dauphine

Interaction among Agents via a Social Network in Computational Social Choice

École doctorale de Dauphine – ED 543

Spécialité INFORMATIQUE

COMPOSITION DU JURY :

Laurent Gourvès
Université Paris-Dauphine
Directeur de thèse

Julien Lesca
Université Paris-Dauphine
Examinateur

Katarína Cechlárová
Pavol Jozef Šafárik University in Košice
Rapporteur

Christoph Dürr
Sorbonne Université
Rapporteur

Ulle Endriss
University of Amsterdam
Examinateur

Jérôme Lang
Université Paris-Dauphine
Président du Jury

Pierre Marquis
Université d'Artois
Examinateur

Agnieszka Rusinowska
Université Paris 1 Panthéon-Sorbonne
Examinatrice

Soutenue par **Anaëlle WILCZYNSKI**
le **13/12/2018**

Dirigée par **Laurent GOURVÈS**

DAUPHINE UNIVERSITÉ PARIS | PSL RESEARCH UNIVERSITY PARIS

Contents

Introduction

Collective decision making appears in many real life situations, ranging from political elections, to the assignment of positions to applicants, including among others the partition of students into working groups. *Social choice theory* consists in the study of collective decision processes. The paternity of this economic field is usually attributed to Kenneth Arrow (1921-2017) with his well-know impossibility theorem [1951], establishing that there is no voting system satisfying together several basic axioms. But actually the origin of social choice can be traced back to the 18th century with the works of Nicolas (marquis) de Condorcet (1743-1794) and Jean-Charles (chevalier) de Borda (1733-1799). A social choice problem aims at aggregating the preferences of different members of a society into a final decision over a set of alternatives. The final decision of a social choice problem can take the form of the selection of one or more alternatives or the (partial) ranking of the alternatives.

The aim of social choice is to design efficient and fair procedures for collective decision making. However, beyond the quality of the procedures, their computation and communication costs must be taken into account for practical reasons. This line of research, born at the end of the 20th century, initially within the computer science community, is known as *computational social choice* [Brandt et al., 2016]. Computational social choice arises as an interdisciplinary field with concrete applications, mixing Economics, Computer science and Mathematics.

Several aspects of computational social choice problems can be investigated. First of all, classically, one can design specific aggregation rules or protocols and establish the properties of the outcome according to some desirable axioms. The computation of a solution is a main issue in order to solve concrete problems. Consequently, an important research area consists in stating the computational complexity of social choice procedures. Moreover, some restrictions, especially about the preferences of the agents, have a strong impact on the outcome of the aggregation rules. It appears relevant to consider realistic preference types for the agents and analyze how the complexity and the outcome of the procedures are affected. In that respect, eliciting the preferences of the agents also appears as a necessary challenge, raising additionally the question of the communication cost of the procedures. Furthermore, in some contexts, agents may lie when they reveal their preferences, in order to strategically orient the outcome of the procedure towards a solution more advantageous for them than the solution they could expect by telling the truth. As the designer of a social choice procedure would generally like to prevent such a strategic behavior, strategy-proofness is an important issue in social choice.

When speaking about social choice, one usually think about voting and political elections. However, the spectrum of social choice problems is much larger than voting. The problems addressed in computational social choice can be classified into several categories, according to the nature of the problems [Chevaleyre et al., 2007b]. It is possible to cite, as main topics, *voting theory*, *resource allocation*, *coalition formation* or *judgment aggregation*.

- **Voting theory**. Agents are voters who have preferences over a set of candidates. The goal is either to select one or several candidates as the winner(s) of the election or to aggregate the preferences of the voters into a final ranking over the candidates.

  As a simple example, take a class of 15 students who want to organize an activity together at the end of the year. The students hesitate among four activities: bowling, climbing, ice skating and hiking. They organize a vote, on an online voting platform, in order to decide which activity they will choose. The students must indicate in the platform the activity that they prefer. The score that each activity obtains after the vote is summarized in the following table.

bowling	climbing	ice skating	hiking
3	5	5	2

  Climbing and ice skating are the two activities which get the most points according to the results of the vote. Between these two winning alternatives, the students finally decide to choose climbing because it is a less expensive activity, compared to ice skating.

  Many different settings fall into voting theory, such as multi-winner elections [Faliszewski et al., 2017] where the goal is to select a subset of candidates as a winning committee, or voting on combinatorial domains [Lang and Xia, 2016] where the preferences of the agents are conditioned by the selection of some other alternatives. Classically, in a social choice perspective, voting procedures can be characterized according to the axioms that they satisfy (see Zwicker [2016] for an overview of the main characterization results). By focusing on the computational aspects of voting procedures, one can investigate the complexity of computing the outcome of a given voting rule, in the line of the seminal paper of Bartholdi et al. [1989a]. Furthermore, designing a procedure that is strategy-proof is a main issue in voting because the nice properties of a given voting rule could fail if the voters adopt a strategic behavior. However, no voting scheme satisfying some natural properties is strategy-proof [Gibbard, 1973, Satterthwaite, 1975].

  Many works focus on the impact of *strategic voting* (see Meir [2018] for a recent overview), where agents can manipulate by casting a ballot that does not correspond to their true preferences. Strategic voting can be investigated through the computation cost of manipulation [Bartholdi et al., 1989b, Bartholdi and Orlin, 1991]. Alternatively, one can model strategic voting as a strategic game [Myerson and Weber, 1993, Sloth, 1993, Desmedt and Elkind, 2010, Meir et al., 2010], by considering that manipulation is unavoidable, even when computing a strategy for manipulation is computationally hard [Faliszewski and Procaccia, 2010, Conitzer and Walsh, 2016]. In a voting game, the outcome of the voting procedure is defined according to the different strategies played by the voters, that are materialized into the ballot that they cast. *Iterative voting* (see Meir [2017] for a recent survey) relies on a dynamic voting game where voters strategically deviate by rounds to a new ballot. This model can have different natural interpretations. In a context of political elections, iterative voting can represent the vote intentions of the voters before the real election. These vote intentions can change progressively during the period preceding the election according to different opinion poll announcements, in a strategic perspective. One could also think about online polls like in the previous example with the choice of activities among the students. If the votes of the students are visible for all the others, a student who has voted first for bowling, can decide to change her vote to ice skating that she prefers to climbing, the current winning alternative. Afterwards, another student who has previously voted for hiking can decide to change her vote to climbing that she prefers to ice skating, the new current winning alternative, and so on. This process of deviations continues until a stable state is reached, if possible, or the horizon for taking a decision is passed.

- **Resource allocation and fair division**. The goal is to fairly allocate a given set of resources to agents [Young, 1995, Brams and Taylor, 1996, Moulin, 2004]. This subject is closely related to real life applications and many challenges can be addressed [Chevaleyre et al., 2006]. Designing procedures that satisfy some guarantee of efficiency or fairness is a main concern in this research area. For instance, *envy-freeness* [Tinbergen, 1946, Foley, 1967, Varian, 1974] is a natural fairness criterion requiring that no agent prefers the share assigned to another agent to her own share. Regarding efficiency, a basic requirement is *Pareto-efficiency*, imposing that it is impossible to improve the satisfaction of some agents without harming another. The procedures should also be easy to implement in terms of computation and communication cost. In this perspective, many works have explored the complexity of finding fair allocations according to different criteria, by analyzing for instance the impact of the preference shape [Bouveret and Lang, 2008, de Keijzer et al., 2009]. Different frameworks exist according to the nature of the resources, divisible or not, shareable or not, goods or chores. When the resources are divisible, the problem is referred to as *cake cutting* (see Procaccia [2016] for a recent overview of the advancements in cake-cutting). The context of fair division with indivisible resources [Brams et al., 2003, Bouveret et al., 2016] is conceptually different from cake-cutting, because it deals with discrete resources instead of working within a continuous space. When each agent must be assigned exactly one indivisible resource, the problem is known as *house allocation* [Hylland and Zeckhauser, 1979, Abdulkadiroğlu and Sönmez, 1998, 1999].

  As an example, take the creation of a working schedule among the employees of a nursery. There are four employees, Adrian, Beatrice, Catherine and Diana, and four time slots to fill in the day, from 7am to 11am (T1), from 10am to 2pm (T2), from 12am to 4pm (T3) and from 3pm to 7pm (T4). The director of the nursery decides how to organize the planning according to the wishes of her employees, represented as follows (the preferred time slot is on top and so on).

Adrian	Beatrice	Catherine	Diana
T2	T3	T2	T1
T3	T4	T1	T3
T1	T2	T4	T2
T4	T1	T3	T4

The director finally assigns to Adrian time slot T2, to Beatrice time slot T1, to Catherine time slot T4 and to Diana time slot T3. Observe that this assignment is not optimal in the sense that Beatrice prefers the time slot of Catherine and vice versa. However, the director is flexible and enables her employees to make some rearrangements. Therefore, Beatrice and Catherine finally decide to exchange their working hours, which is more convenient for them.

A special case in house allocation occurs when, starting from an initial endowment, local improvements can be done by trades among the agents, like in the previous example. This setting is known as a *housing market* [Shapley and Scarf, 1974].

- **Coalition formation**. Coalition formation is a field analyzing how agents can group together, for instance in order to jointly solve some problems. One can investigate which coalitions can form as well as how they will share the profit generated by their cooperation, from the perspective of *cooperative game theory.*

  Hedonic games are especially focused on the coalitions that can form (see Aziz and Savani [2016] for a recent overview on hedonic games). The goal of hedonic games is to partition the agents into disjoint coalitions, in a context where the agents express preferences on which group of agents they would like to form a coalition with. Typical examples are the problem of making working groups among colleagues or students, and the elaboration of a seating plan in a wedding. The main requirement in hedonic games is the *stability* of the solution [Dreze and Greenberg, 1980, Cechlárová and Romero-Medina, 2001, Bogomolnaia and Jackson, 2002]. Different notions of stability can be established. The basic one is *Nash stability* where no agent would prefer to belong to another group.

  When the agents are initially partitioned into two disjoint subgroups according to some characteristics (women/men, applicants/schools or residents/hospitals for instance) and the goal is to pair the agents such that each pair contains an agent coming from each subgroup, the problem is actually a specific matching under preferences [Klaus et al., 2016], known as *two-sided matching.* This problem is famous because it also refers to the seminal work of Gale and Shapley [1962] on the *stable marriage problem.*

- **Judgment aggregation**. The specificity of judgment aggregation is based on the point that the agents do not express preferences but beliefs or judgments about given statements, that they estimate true or false. The goal is to aggregate the different beliefs of the agents in order to take a judgment deciding the truth of the statements, like in a court. The set of statements on which the agents must rule is called the *agenda.* The framework of judgment aggregation has been mainly formalized by List and Pettit [2002] and Dietrich [2007]. Several recent surveys can be found [List and Puppe, 2009, Grossi and Pigozzi, 2014, Endriss, 2016].

The belief of each agent is represented by a formula of propositional logic, that is supposed to be consistent, by individual rationality assumption. The goal is then to aggregate the beliefs of the different agents into a consistent formula.

Another very close framework that has arisen in Artificial Intelligence is *belief merging*, where the goal is to aggregate the beliefs bases of different agents. The belief base of each agent is a finite set of propositional formulas representing her beliefs about the current state of the world. Contrary to judgment aggregation, the belief bases are not limited to a specific agenda. The links between judgment aggregation and belief merging have been for instance examined by Everaere et al. [2015].

Interaction among agents in social choice procedures

The integration of the agents in collective decision making is a main question in computational social choice. In fact, the agents constitute a central and active component of the decision process. Their role differs according to the type of procedure that is chosen for aggregating their preferences.

In a *centralized procedure*, a central authority collects the preferences of the agents, aggregates them following a given rule and outputs a final decision for the society. This is the case in most of the voting procedures, where voters are asked to submit a ballot to a system which then, computes and announces the outcome of the election. In such a procedure, the role of the agents is limited to the submission of their preferences in the format imposed by the system. However, they can still have several options because they may adopt a strategic behavior and misreport their preferences.

Alternatively, in a *decentralized* or *distributed procedure*, the agents iteratively construct the final decision, without the intervention of a central authority, by following a certain protocol. This approach is often adopted in resource allocation. The most classical example is the *cut-and-choose* protocol in *cake-cutting*, the problem of fairly distributing divisible resources among agents, which is traditionally modeled as a problem of cutting a cake. Designed for a context with two agents, the cut-and-choose protocol works as follows: one agent cuts the cake and the other one chooses the allocation.

There also exist procedures mixing the centralized and the distributed approach. Typically, the agents can construct a final decision on the basis of an initial solution coming from a central authority (see the example of the nursery).

For every type of social choice procedure, the agents are active in the process and interact with each other. Indeed, the agents are entirely involved in a distributed procedure where they interact with other agents by definition. In a centralized procedure, the agents are aware of the existence of other agents. Therefore, interaction can for instance occur in the elaboration of a strategy for manipulation, when agents decide to misreport their preferences to the central authority. More precisely, agents can collaborate with other agents or get some information from them in order to design efficient strategies. Moreover, the evaluation and validation by the agents of the centralized procedure can appear as necessary in order to ensure the sustainability of the process, and avoid that the agents dispute the procedure. In this perspective, the efficiency and fairness of the procedure, as perceived by the agents, are essential but imply communication and comparison among the agents and thus, more generally, some interaction among them.

The interaction among the agents in social choice problems can take different forms: communication, collaboration, cooperation, influence, information gathering, information diffusion, and others. Take the previous example of students wanting to organize an activity at the end of the year. Now suppose that the online voting platform on which the students report their preferences is closed, in the sense that nobody can observe the votes of the others. Then, no student has incentive to cast false preferences unless she knows that her preferred activity, for example hiking, has no chance to win. If such a student collects information about the preferences of some other students, for example by communicating with them or by listening their claims about what they want, she can realize that hiking is not enough supported among the students and thus she could vote for ice skating, her second choice, instead of hiking. Alternatively, in the example of the nursery, Beatrice and Catherine need to collaborate in order to exchange their working hours.

Traditionally, in social choice, it is assumed that an agent is able to have social interactions with any other agent. However, in real life, due to communication or distance burdens, the interaction with some agents may not be possible. An *accessibility relation* can define the ability of interaction among some agents. In large scale instances, like for example in political elections, the accessibility relation among the agents is quite poor. On the contrary, in small scale instances, like for example in a problem of task assignment in a department of a company, the accessibility relation is very dense but can be limited by questions of affinities among the employees. For instance, in the example of the nursery, the collaboration between Beatrice and Catherine, leading to the exchange of their assigned working hours, can occur because Beatrice and Catherine are friends and thus they can easily exchange some services. Suppose that Beatrice and Adrian are also friends but that Diana is a new employee unknown to Beatrice. Then, even if Beatrice and Diana would also profitably exchange their working hours regarding their preferences, they do not do so. Following the same idea, in the example of the choice of activities, a student can discover the preferences of some other students only if she has some links with them, making the communication possible.

In general, the accessibility relation may be very different regarding the context, avoiding the possibility of assuming a unique accessibility relation for every instance of a social choice problem. This observation leads to consider the possibility of social interaction, given by the accessibility relation, as a part of the input of the problem. There exist many ways to represent the accessibility relation. Due to the important role of the social networks nowadays in the social relations, it appears meaningful to consider that the accessibility relation is modeled by a social network, where interaction is possible between two agents if they are connected in the social network. In general, a social network can represent affinities among people, geographical closeness, peer / colleague relations, or even an online social network. A social network is classically represented by a graph over the agents, more precisely a set of nodes and links which model the connection between two agents in the social network. For instance, in the example of the nursery, the relations between Beatrice and the other employees, making possible some interaction, can be represented by the following graph structure.

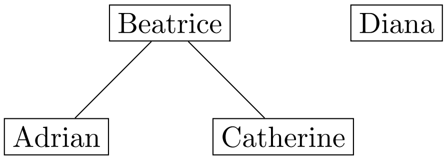

The friendship relation between Beatrice and Adrian, and between Beatrice and Catherine are represented by links between the nodes associated with these individuals.

Social networks and social choice

Nowadays, online social networks such as Facebook, Twitter or Instagram, are everywhere and pervasive. They have an ever-growing impact on the behaviors in the society, making relevant the study of the social mechanisms on which they lie. More generally, the study of the structure of social relations, that can be represented by a network, is referred to as *social network analysis* [Wasserman and Faust, 1994, Scott, 2017]. Such social relations can refer to the family circle, to the friends, to the colleagues, or even to individuals from the same geographical area, in addition to the recent online social networks. Main questions in social network analysis are how the networks form, what are the characteristics and properties of social networks, and how the agents are tied by social relations.

Social network analysis is actually an active research area, and a hot topic for many researchers [Freeman, 2004]. A typical example of the attraction of the researchers for this line of research is the creation of the journal *Social Networks* [Freeman et al., 1978]. Many aspects of the social networks can be investigated, leading to a huge literature dealing for instance with models of networks [Erdös and Rényi, 1959, Barabási and Albert, 1999], structural properties [Newman et al., 2011], experimental and behavioral studies [Travers and Milgram, 1967, 1977], or the measurement of the influence of some nodes in the network [Rusinowska et al., 2011]. Moreover, the influence of social networks on the behavior of individuals deserves attention, especially in the study of social decisions and interactions. In this context, their impact on economic questions like strategic interactions, markets and game theory issues is central and has been notably addressed by Jackson [2008] and Easley and Kleinberg [2010].

Assuming a social network, or more generally a graph structure, over the agents has already been investigated in cooperative game theory by assuming a *cooperation graph* over the agents [Aumann and Dreze, 1974, Myerson, 1977]. In computational social choice, the social networks have been quite recently introduced in order to study some social interactions among the agents. A large part of these works deals with coalition formation problems. For instance, in hedonic games, some stability questions where the possible coalitions of agents arise from a social network have been examined [Igarashi and Elkind, 2016]. Following the same principle, coalitional games and problems of coalition formation with assignment of tasks to groups, have been investigated in a context where the coalitions are determined by a social network [Elkind, 2014, Igarashi et al., 2017]. Moreover, many works aim at capturing, in two-sided matchings, several types of social behavior thanks to a social network, such as altruism [Anshelevich et al., 2013a], information gathering via social contacts [Arcaute and Vassilvitskii, 2009], peer effects [Bodine-Baron et al., 2011], or collaboration [Hoefer, 2013].

A few works also used social networks in judgment aggregation. For instance, Colombo Tosatto and van Zee [2014] have translated the judgment aggregation framework into a social network analysis setting, by modeling the social relations between the agents, such as agreement, with a graph.

A recent survey [Grandi, 2017] shows many relevant lines of research integrating a social network over the agents in social choice issues like voting or opinion diffusion. First of all, one could ask whether the social network in which the voters are embedded has an impact on their vote. Some works consider voting as a way to aggregate opinions in order to recover a ground truth. Conitzer [2012], by assuming that the probability that a voter estimates the right alternative is independent from the probability to be influenced by her neighbors in the social network, answers that the network does not matter in the votes of the agents. Based on a similar model, the *independent conversations model* considers voters that aim at recovering a ground truth thanks to the aggregation of the opinions of their connected agents on the social network [Conitzer, 2013, Procaccia et al., 2015, Tsang et al., 2015]. This model is very close to the models of *opinion diffusion* in social networks [DeGroot, 1974] (see Jackson [2008] for a recent survey), which is a very wide topic in social network analysis, related to the propagation of some ideas or opinions through the network. Opinion diffusion in networks has been recently examined within the social choice community, in a strategic perspective [Grandi et al., 2015, Brill et al., 2016, Grabisch et al., 2017, Grandi et al., 2017]. Concerning strategic voters, some recent works use a social network, or a *knowledge graph*, for modeling the information that the agents have when they act strategically in voting [Chopra et al., 2004, Sina et al., 2015, Tsang and Larson, 2016, Tsang et al., 2018].

Another topic for which the social networks offer a new vision of the problems is resource allocation. In this respect, some fairness criteria can be adapted in order to take into account the social constraints induced by the social network. For instance, one can define a *local* notion of envy-freeness according to the links of the social network, that can be used either in cake-cutting [Abebe et al., 2017, Bei et al., 2017], or in the case of indivisible resources [Aziz et al., 2018, Bredereck et al., 2018, Flammini et al., 2018]. Beyond the evaluation of the allocation, trades of resources that are defined according to a social network can be examined [Chevaleyre et al., 2007c, Kleinberg and Tardos, 2008].

Contributions and organization

The goal of the thesis is to relax the classical assumption that any agent is accessible from anyone else. More precisely, in the continuity of some recent works in computational social choice, we model the possibility of interaction among agents by a social network, modeled as a graph structure over the agents. We aim at understanding the impact of this generalization on the social interactions that occur in social choice problems. We study two types of social interactions that are particularly important in collective decision making: *collaboration* among agents and *information gathering* about the situation of other agents.

We focus on two classical problems of computational social choice: *strategic voting* and *resource allocation*. In particular, strategic voting is examined through iterative voting, whereas in resource allocation we investigate the specific house allocation problem, which is related to the problem of matching a set of indivisible items with a set of agents. In iterative voting and house allocation, how does the social network influence the interaction among the agents? How does the limitation on accessible information or possible collaborating agents impact the execution and the outcome of the procedures for these settings?

We examine these questions according to the two types of social interactions under study: collaboration and information gathering. The results concerning the first aspect are presented in Part I, through coalitional manipulation in iterative voting and swaps in house allocation, where the only possible collaborators are given by the social network. The second aspect is explored in Part II, through local envy-freeness in house allocation and uncertainty in iterative voting, where the most part of the information available to the agents is given by their connections in the social network. A preliminary part in Chapter 1 presents the two settings under study, that are strategic voting and house allocation. Our contributions are grouped in parts I and II, whereas the related work is concentrated in the preliminary chapter (Chapter 1), dedicated to the context.

Chapter 1 introduces the general notations that are used all along the thesis, as well as some key notions about the context. In fact, we present the main components of a social choice problem where the agents are embedded in a social network that is a graph over the agents. Moreover, we introduce the specific frameworks that we study, strategic voting and house allocation, by providing some basic notations and state-of-the-art main concepts. We also recall some useful concepts of computational complexity theory.

Let us describe more in details the contributions of the thesis.

We first investigate the collaboration among the agents, conditioned by a social network, in social choice problems. In Chapter 2, we consider an iterative voting framework where groups of agents can manipulate. We assume that the only groups of agents that can form, in order to establish common strategies of manipulation, are the subsets of agents that are fully connected in the social network. Moreover, an altruistic condition, imposing that the deviating coalitions of voters must not harm other connected agents in the network, limits the possible deviations. The existence of equilibria in the voting game that are immune to such coalitional manipulation is investigated, as well as the convergence of the deviation dynamics to such equilibria. We provide guarantees of existence for equilibria according to several voting rules and essentially negative results regarding the convergence of the dynamics.

Chapter 3 is dedicated to a house allocation problem where, starting from an initial endowment (each agent has one indivisible item), two connected agents can swap their object if the trade is mutually profitable. Consequently, a dynamics of local improvements defined according to the social network is examined. We study in particular the complexity of several decision problems related to the reachability of some objects for the agents through sequences of trades. We provide hardness results and tractable cases for these natural questions, according to the structure of the graph.

Secondly, we study how the social network can model the information that the agents can acquire from the other agents. In this context, Chapter 4 focuses on a house allocation setting where the evaluation of an allocation of indivisible goods to the agents is conditioned by the visibility of the agents. In particular, we examine a local envy-freeness notion, where the agents can only envy agents who are connected to them in the social network. We investigate the complexity of some decision problems based on the existence of locally envy-free allocations and optimization problems where we aim at finding allocations that minimize local envy. We provide hardness results, approximation algorithms and tractable cases according to the topology of the network.

In Chapter 5, we analyze an iterative voting setting under uncertainty, where the voters only have two types of information: a global information given by a public opinion poll and a local one thanks to their connections in the social network. These two sources of information enable the voters to develop a certain belief about the current vote intentions of the other agents. We define a dynamics of manipulation where the voters deviate to a best possible ballot according to their belief about the current state. The convergence properties of this dynamics as well as the quality of its outcome are explored. Moreover, we investigate the question of manipulation by the polling institute.

Finally, we conclude in the last chapter by summarizing the key points of our contributions and by giving some avenues for future works.

The contributions of the thesis are published in the proceedings of international conferences of Artifical Intelligence, namely ECAI-16 [Gourvès et al., 2016], IJCAI-17 [Gourvès et al., 2017], AAMAS-18 [Beynier et al., 2018], SAGT-18 [Saffidine and Wilczynski, 2018] and AAAI-19 [Wilczynski, 2019].

Chapter 1

Preliminaries and notation: Presentation of the problems

1.1 Introduction

This preliminary chapter introduces the notions that will be used all along the document. The main components of a social choice problem, where the agents are embedded in a social network, are presented in Section 1.2. We will also give the basic notations and some key concepts from the state-of-the-art for the two specific settings that are investigated, namely voting and more precisely strategic voting in Section 1.3, and resource allocation with a focus on house allocation in Section 1.4. Finally, since computational aspects of social choice problems are studied, some notions on computational complexity are provided in Section 1.5.

1.2 General framework

The focus of our analysis is on *social choice* problems. Social choice issues can gather very different problems, from voting to judgment aggregation through resource allocation or coalition formation, but a common denominator is *collective decision making*. This section is devoted to the presentation of the common and basic notations of a social choice problem.

Generally, we are given a finite set of n agents $N = \{1, \dots, n\}$, and a finite set of $m \geq 2$ alternatives $M = \{a, b, \dots\}$. Notation $[p]$ stands for the set $\{1, \dots, p\}$ for any integer p, and $[p..q]$ for the set $\{p, \dots, q\}$. Each agent has preferences over the set of alternatives[1], and the goal is to reach a collective decision defined on the set of alternatives. In addition to these general components, we assume that the agents are embedded in a *social network*, represented by a graph $G = (N, E)$ over the agents.

[1]We do not consider the special case of judgment aggregation problems where the words belief and judgment are more appropriate.

1.2.1 Preferences

There are several competing models for expressing preferences of agents. In this subsection, the most common models are presented.

- The preferences can be *cardinal* and expressed via *utility functions*. Each agent $i \in N$ has her own utility function $u_i : M \to \mathbb{R}$ where $u_i(x)$ corresponds to the value that agent i attributes to alternative $x \in M$. This is an absolute evaluation of the alternatives from the agents.

- Another very common model is to consider *ordinal* preferences, where agents relatively evaluate the alternatives according to an order.

- A natural restriction of both models relies on *dichotomous* preferences, where agents can only approve a chosen subset of alternatives.

Eliciting cardinal preferences may require a significant cognitive effort for the agents but, most of all, this raises the problem of normalization for interpersonal comparisons: the agents do not necessarily have the same scale in mind for evaluating the alternatives. Consequently, *we decide to focus on ordinal complete preferences*, by not handling the case of cardinal preferences, nor the case of partial preferences.

We assume that the preferences of agent $i \in N$ are expressed via a weak order $\succsim_i$ over M. The preference profile of the whole society is given by $\succsim = (\succsim_1, \ldots, \succsim_n)$. We mostly restrict ourselves to the case of strict ordinal preferences where the agents are not allowed to express indifference between alternatives. In this restriction, we consider that the preferences of agent i are represented via linear order $\succ_i$ and the preference profile is denoted by $\succ$.

The set of all possible weak orders over M is denoted by $\mathcal{W}(M)$ and the set of all possible linear orders by $\mathcal{L}(M)$. For a weak order $\succsim_i \in \mathcal{W}(M)$, associated with the preferences of agent i, $\sim_i$ denotes the symmetric part of relation $\succsim_i$, and $\succ_i$ its asymmetric part. In other words, $a \succ_i b$ if and only if agent i strictly prefers alternative a to alternative b, i.e., $a \succsim_i b$ holds but not $b \succsim_i a$, and $a \sim_i b$ if and only if agent i is indifferent between alternative a and alternative b, i.e., both $a \succsim_i b$ and $b \succsim_i a$ hold. Consequently, $a \succsim_i b$ means that either $a \succ_i b$ or $a \sim_i b$. Subrelation $\sim_i$ induces equivalence classes of indifference. For a linear order $\succ_i$, let us denote by $\mathfrak{r}_i : M \to [m]$ the mapping induced by $\succ_i$ that indicates the position of the alternatives in the preference ranking, i.e., alternative x is ranked at the $\mathfrak{r}_i(x)^{\text{th}}$ position in $\succ_i$.

Observe that a unique weak order $\succsim_i$ can be derived from cardinal utilities u_i with respect to the natural order over $\mathbb{R}$, that is $a \succ_i b$ if and only if $u_i(a) > u_i(b)$ and $a \sim_i b$ if and only if $u_i(a) = u_i(b)$. Moreover, dichotomous preferences are ordinal preferences where there are at most two indifference classes. Usually, they are represented cardinally by assuming that utility function u_i of agent i is such that $u_i : M \to \{0, 1\}$.

When the entire preference profile is explicitly given, we adopt the general formalism that is presented in the following example.

Example 1.1 *Let us consider an instance with three agents and three alternatives, where* $N = \{1, 2, 3\}$ *and* $M = \{a, b, c\}$*. Take a preference profile* $\succ$ *such that:* $a \succ_1 b \succ_1 c$, $a \succ_2 b \succ_2 c$ *and* $c \succ_3 b \succ_3 a$*. The preference profile is represented as follows.*

$$\begin{array}{ll} 1: & a \succ b \succ c \\ 2: & a \succ b \succ c \\ 3: & c \succ b \succ a \end{array}$$

The name of each agent is followed by the preferences of the agent where, for the sake of simplicity, the subscripted index referring to the agent is omitted. When different agents have the same preferences, they are sometimes grouped, i.e., in this example we could have written $1,\ 2: \quad a \succ b \succ c$*, for the preferences of the first two agents.*

1.2.1.a Restricted preferences

There can exist some *structure* on the global preferences $\succsim \in \mathcal{W}(M)^n$ for the whole society. This structure restricts the preference domain to $\mathcal{D} \subset \mathcal{W}(M)$. Many preference restrictions can be considered but we focus on two types: one related to an order over the alternatives, namely *single-peakedness* [Black, 1948], and one related to an order over the agents, namely the *single-crossing condition* [Karlin, 1968, Mirrlees, 1971, Roberts, 1977].

Definition 1.1 (Single-peakedness) *A preference profile* $\succsim$ *is single-peaked if and only if there exists an order* $>^M$ *over the alternatives such that for every agent* i*, there exists a unique peak alternative* x^* *such that for every pair of alternatives* a *and* b*,* $x^* >^M a \geq^M b$ *implies that* $x^* \succ_i a \succsim_i b$ *and* $a \geq^M b >^M x^*$ *implies that* $x^* \succ_i b \succsim_i a$*.*

The idea of single-peakedness is that every agent shares a common axis over the alternatives, implying that all the agents prefer alternatives that are closer to their peak alternative, i.e., their best alternative, in the axis. Classical examples are axis left-right in politics or temperature scales. As its denomination induces, a single-peaked preference relation with respect to a given axis produces only one peak in its graphical representation. This notion is then graphically well-understandable, as we can observe in Example 1.2.

Example 1.2 *Consider an instance with three agents and four alternatives, where* $N = \{1, 2, 3\}$ *and* $M = \{a, b, c, d\}$*. The preference profile given by the following graphical representation is single-peaked. The common axis over the alternatives, given by the abscissa axis, is* $a >^M b >^M c >^M d$*. Each curve represents the preference order of a given agent. In the ordinate axis, the higher is a point of the curve associated with agent* i*, the more preferred the associated alternative in abscissa according to the preference order of* i*. Observe that, in this representation, the curve associated with the preference order of each agent has only one peak.*

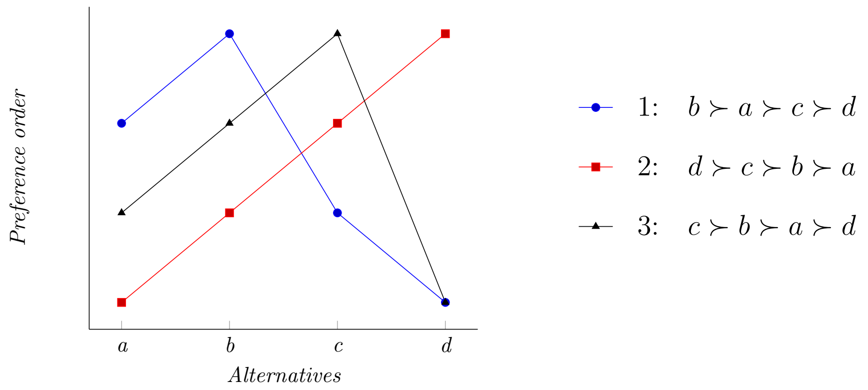

Definition 1.2 (Single-crossing condition) *A preference profile $\succsim$ is single-crossing if and only if there exists an order $>^N$ over the agents such that for every pair of alternatives a and b, there exists a unique agent i^* such that all the agents j with $j <^N i^*$ (respectively, $j >^N i^*$) all share the same preferences $a \succsim_j b$ or $b \succsim_j a$.*

The single-crossing condition has similarities with the idea of single-peakedness but the axis is over the agents. In a single-crossing preference profile, there exists an order $>^N$ over the agents such that for each ordered pair of alternatives a and b, the set of agents preferring a to b is a unique interval of $>^N$. As an example, one may think about agents ranked along a left-right axis according to their political orientations. Generally, in such a context, the preferences over two alternatives that concerns societal or economical questions, such as different types of support for welcoming migrants or different levels of State intervention, will be different for two consecutive agents in the axis at one point of the axis, at most. This type of preference restriction is illustrated in Example 1.3.

Example 1.3 *Consider an instance with five agents and three alternatives, where $N = \{1, 2, 3, 4, 5\}$ and $M = \{a, b, c\}$. The following preferences are single-crossing with respect to order $>^N$, given by the abscissa axis, that we assume to be the natural order over the numbers $\{1, 2, \dots, 5\}$. Each curve corresponds to the positions of a given alternative in the preference orders of the agents. Observe that the curves associated with each pair of alternatives cross at most once.*

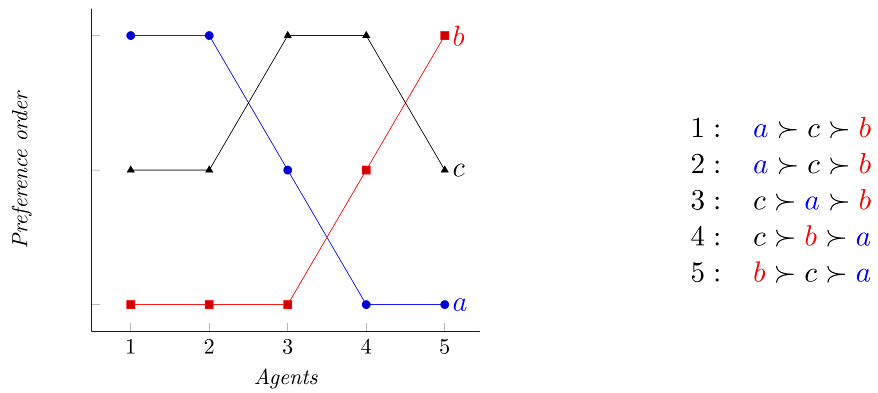

Note that there exist preference profiles that are both single-peaked and single-crossing.

1.2.1.b Generating preference profiles

How to generate preference profiles in experiments is a typical concern in simulation of social choice problems. First of all, note that, by simplicity, it is common to focus on generating strict preferences. A widespread and simple procedure for generating strict preferences for a fictitious population is to uniformly pick linear orders among all the possible rankings. The culture associated with this protocol is called *impartial culture (IC)*.

However, in practice, the preferences of agents are not uniformly distributed [Regenwetter et al., 2006]. Therefore, it can be meaningful to introduce some correlation within the preference profiles. A typical way to incorporate correlation is to consider an urn with all possible rankings and every time an order is picked to represent the preferences of an agent, then q rankings identical to this order are added to the urn, with a predefined value q inducing a specific correlation ratio. For instance, $q = m!/9$ induces a 10%-correlation ratio, culture that we denote by CR-10%. This model is inspired from *Pólya-Eggenberger urn models* [1923], and has for instance been used for generating preference profiles in voting [Berg, 1985, Lepelley and Valognes, 2003, Grandi et al., 2013].

Moreover, it is also possible to consider single-peaked preferences profiles that we can generate according two different distributions:

- either following a uniform distribution over all possible single-peaked profiles [Walsh, 2015], culture that we call *single-peaked uniform (SP-U)*,

- or by choosing, following a uniform distribution, a peak alternative for the preference order, and then choosing the next preferred alternatives with equal probability either on the left of the peak in the axis, or on the right [Conitzer, 2009]. We call this culture *single-peaked uniform peak (SP-UP)*.

Furthermore, it may be sometimes more appropriate to use real data sets. This is now possible in social choice via the *Preflib* platform [Mattei and Walsh, 2013], which currently contains 3,000 data sets of preferences coming from different experimental studies. This online library continues to be populated thanks to the contribution of many social choice researchers who accept to share their data with the community.

Globally, the preferences of the agents mainly impact the collective decision of a social choice problem. Other key points are the relations among the agents and how they can interact. These social connections can be modeled via a social network over the agents.

1.2.2 Social network

The social network in which the agents are embedded is a graph over the agents, denoted by $G = (V, E)$, where $V = N$ is a finite set of *vertices* (or *nodes*) representing the agents, and $E \subseteq V \times V$ translates an irreflexive binary relation R^E on V. Graph G captures the possibility of interaction between the agents. For an overview of basic notions about graphs, see for instance Golumbic [1980].

1.2.2.a Basic notions of graph theory

In general, E is a set of ordered pairs of distinct vertices such that $(u, v) \in E$ if and only if $uR^E v$. An element (u, v) in E, whose endpoints are $u \in V$ and $v \in V$, is called an *arc* and is graphically represented by a directed arrow from u to v. A graph $G = (V, E)$ is also called a *directed graph* (or *digraph*).

When binary relation R^E which characterizes E is *symmetric*, i.e., for every pair of vertices u and v, $uR^E v$ implies that $vR^E u$, the graph is called an *undirected graph.* In an undirected graph, the elements of E are called *edges* and are denoted by $\{u, v\}$ for $u, v \in V$. More precisely, in case of an undirected graph, for every pair of distinct vertices u and v, $\{u, v\} \in E$ if and only if $uR^E v$ and $vR^E u$. An edge $\{u, v\} \in E$ is said to be *incident* to vertex $u \in V$ and to vertex $v \in V$. An edge $\{u, v\}$ is graphically represented by a single line without arrow joining u and v.

Graph G is said to be *complete* if for every ordered pair of distinct vertices u and v, it holds that $uR^E v$. Alternatively, a graph is *empty* if $E = \emptyset$. When the binary relation represented by E is transitive, the graph is said to be *transitive* (see for instance the third graph in Figure 1.1).

Definition 1.3 (Transitive graph) *A graph $G = (V, E)$ is transitive if and only if for every ordered triple of distinct vertices (u, v, w), it holds that if $uR^E v$ and $vR^E w$ then $uR^E w$.*

The *complement* of graph G, denoted by $\overline{G} = (V, \overline{E})$, is such that, for every ordered pair of vertices u and v, $uR^{\overline{E}} v$ if and only if $u \neq v$ and $\neg(uR^E v)$. A graph $G' = (V', E')$ is a *subgraph* of $G = (V, E)$ if $V' \subseteq V$ and $E' \subseteq E \cap (V' \times V')$. A *spanning subgraph* of graph $G = (V, E)$ is a subgraph $G' = (V', E')$ of G such that $V' = V$. An *induced subgraph* of graph $G = (V, E)$, denoted by $G[V']$ for a subset of vertices $V' \subseteq V$, is the restriction of G to the members of V', that is the subgraph $G' = (V', E')$ of G such that $E' = E \cap (V' \times V')$. A set of vertices $V' \subseteq V$ forms a *clique* in G if and only if $G[V']$ is a complete graph.

Definition 1.4 (Clique) *A subset of vertices $V' \subseteq V$ forms a clique of graph $G = (V, E)$ if and only if, for every ordered pair of vertices u and v in V' such that $u \neq v$, it holds that $uR^E v$.*

A graph that is composed of a set of disjoint cliques is called a *cluster graph.* When the graph is undirected, the notion that is complementary to the clique is the *independent set.*

Definition 1.5 (Independent set) *When $G = (V, E)$ is an undirected graph, a set of vertices $V' \subseteq V$ forms an independent set in G if and only if for all vertices u and v in V', $\{u, v\} \notin E$.*

The independent set is generally defined in the context of undirected graphs. However, to extend it to directed graphs, an independent set is simply a subset of vertices $V' \subseteq V$ such that, for every pair of vertices u and v in V', neither (u, v) nor (v, u) belong to E. A graph $G = (V_1 \cup V_2, E)$ is said to be *bipartite* if and only if V_1 and V_2 are independent sets. In a bipartite graph, which can be either directed or undirected, if uR^Ev then either $u \in V_1$ and $v \in V_2$, or $u \in V_2$ and $v \in V_1$.

Another interesting subset of vertices in an undirected graph is the *vertex cover*, defined as a subset of vertices covering all the edges of the graph.

Definition 1.6 (Vertex cover) *A set of vertices $V' \subseteq V$ forms a vertex cover in an undirected graph $G = (V, E)$ if and only if for every edge $\{u, v\} \in E$, $\{u, v\} \cap V' \neq \emptyset$.*

Observe that in undirected graph G, a subset $V' \subseteq V$ is an independent set of G if and only if $V \setminus V'$ is a vertex cover of G.

A *path* in an undirected graph G is a sequence of edges of the form $(\{u_1, u_2\}, \{u_2, u_3\}, \ldots, \{u_{p-1}, u_p\})$, where the vertices $u_1, \ldots, u_p$ are all different. The *distance*, in an undirected graph G, between two vertices u and v, denoted by $dist_G(u, v)$, refers to the number of edges in a shortest path between u and v. For the sake of simplicity, we may denote a path as a sequence of distinct vertices $[u_1, u_2, \ldots, u_p]$ such that $\{u_i, u_{i+1}\} \in E$ for every $i \in [p-1]$. More generally, a path in a directed graph refers to a sequence of vertices $[u_1, u_2, \ldots, u_p]$ such that $u_iR^Eu_{i+1}$ or $u_{i+1}R^Eu_i$ for every $i \in [p-1]$. A graph is *connected* if it contains a path between any pair of vertices u and v. A *cycle* in an undirected graph G is a closed path, that is a sequence of edges of the form $(\{u_1, u_2\}, \{u_2, u_3\}, \ldots, \{u_{p-1}, u_p\}, \{u_p, u_1\})$, where the vertices $u_1, \ldots, u_p$ are all different. Classes of undirected graphs can be defined according to an acyclicity property. This is the case of the *trees*.

Definition 1.7 (Tree) *Undirected graph G is a tree if and only if G is connected and does not contain cycles.*

Similar notions exist when only one direction matters, by using arcs in directed graphs. A *directed path* (or *dipath*) is a sequence of arcs of the form $((u_1, u_2), (u_2, u_3), \ldots, (u_{p-1}, u_p))$, where the vertices $u_1, \ldots, u_p$ are all different. Note that, in an undirected graph G, given a path $[u_1, u_2, \ldots, u_p]$ of G, one can consider a directed path from u_1 to u_p as a subgraph $G' = (V', E')$ of G that is directed such that $V' = \{u_1, \ldots, u_p\}$ and $E' = \{(u_i, u_{i+1}) : i \in [p-1]\}$. A directed graph is *strongly connected* if it contains a dipath between any ordered pair of vertices u and v. A *directed cycle* (or *dicycle*) is a sequence of arcs of the form $((u_1, u_2), (u_2, u_3), \ldots, (u_{p-1}, u_p), (u_p, u_1))$, where the vertices $u_1, \ldots, u_p$ are all different. An acyclicity property based on directed cycles can be derived in order to define specific classes of graphs, such as the directed acyclic graphs.

Definition 1.8 (Directed acyclic graph (DAG)) *A directed acyclic graph (DAG) is a directed graph with no directed cycle.*

For examples of a DAG, see the two first graphs in Figure 1.1.

In directed graphs, a *successor* of a vertex $v \in V$ is a vertex $u \in V$ such that $(v, u) \in E$. Alternatively, a *predecessor* of a vertex $v \in V$ is a vertex $u \in V$ such that $(u, v) \in E$. The *out-degree* $\delta^+_G(v)$ of a vertex v is the number of successors of v, whereas the *in-degree* $\delta^-_G(v)$ of v is the number of vertices for which v is a successor (or equivalently the number of predecessors of v). A vertex with null in-degree is called a

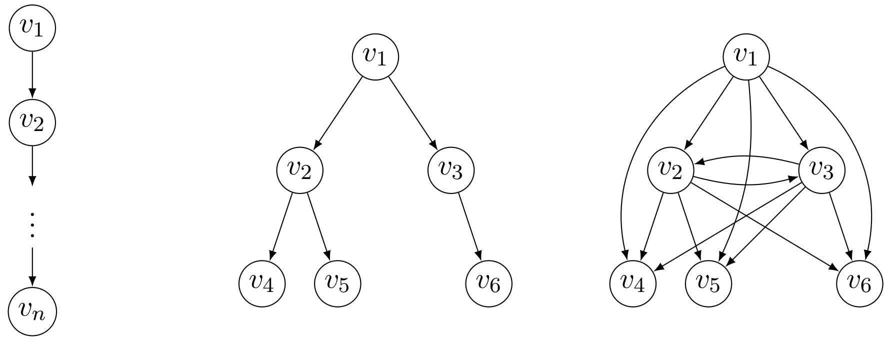

Figure 1.1: Examples of a dipath (on the left), a DAG (in the middle) and a transitive graph (on the right)

source and a vertex with null out-degree is called a *sink*. Observe that a dipath is a specific DAG containing exactly one source and one sink (see Figure 1.1).

For the case of undirected graphs, vertex u is called a *neighbor* of v if $\{u, v\} \in E$. The *neighborhood* of a vertex $v \in V$, that we denote by $\mathcal{N}_G(v)$, is the set of all the neighbors of v, i.e., $\mathcal{N}_G(v) := \{u \in V : \{u, v\} \in E\}$. The degree $\delta_G(v)$ of a vertex $v \in V$ is the number of neighbors of v, i.e., $\delta_G(v) := |\mathcal{N}_G(v)|$. A vertex with degree equal to one is called a *leaf*.

The out-degree (respectively, degree) of a graph G refers to the maximum out-degree (respectively, degree) of a vertex in G and is denoted by Δ_G^+ (respectively, Δ_G). The diameter d_G of an undirected graph is the maximum value of a shortest path between any pair of vertices, i.e., $d_G = \max_{u,v \in V} dist_G(u, v)$. Similarly, the diameter d_G of a directed graph is the maximum value of a shortest directed path between any ordered pair of vertices. The degree and the diameter of a graph are commonly used to characterize some classes of graphs, especially when these parameters are bounded.

For instance, the *paths* and *stars* (see Figure 1.2) are particular trees that can be defined according to the degree of some of their vertices, as well as the *spider graph* [Bahls et al., 2010], which is a generalization of both a path and a star.

Definition 1.9 (Path) *A path is a tree with exactly two leaves.*

Definition 1.10 (Star) *A star is a tree with exactly $n-1$ leaves and one center vertex of degree $n-1$.*

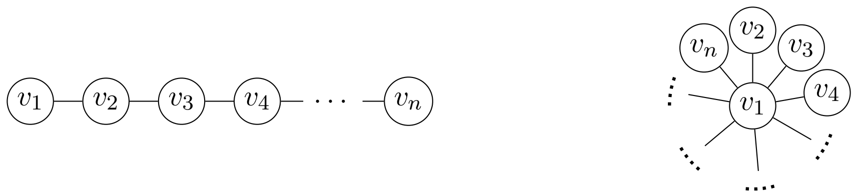

Figure 1.2: Examples of a path graph (on the left) and a star graph (on the right)

Definition 1.11 (Spider) *A spider is a tree with at most one vertex, called the center, of degree greater than 2.*

Other typical graphs defined according to the degree of the vertices are the *regular graphs*, both in undirected and directed version.

Definition 1.12 (Regular graph) *An undirected graph $G = (V, E)$ is said to be regular of degree q if and only if $\delta_G(v) = q$ for every vertex $v \in V$.*

Definition 1.13 (Regular directed graph) *A directed graph $G = (V, E)$ is said to be regular of out-degree q if and only if $\delta_G^+(i) = q$ for every vertex $v \in V$.*

In an undirected graph, a *matching* is a subset E' of edges such that for any edges e and e' in E, $e \cap e' = \emptyset$. Undirected graph $G = (V, E)$ is said to be a matching if E is a matching. In other words, a matching graph is a regular undirected graph of degree equal to one, where the vertices are connected in pairs.

Beyond the degree and the diameter, it is also possible to measure the *density* of a graph, denoted by D_G, which is related to the proportion of edges/arcs that are present in the graph G. The density of graph G is defined as follows.

$$D_G = \begin{cases} \frac{2.|E|}{|V|\cdot|V-1|} & \text{if } G \text{ is undirected} \\ \frac{|E|}{|V|\cdot|V-1|} & \text{otherwise} \end{cases}$$

This parameter is important for instance in the generation of random graphs.

1.2.2.b Random graphs and real networks

An interesting question for experimental studies is how to generate random graphs, and especially graphs that are supposed to represent social networks. Various models for generating social networks have been proposed in the literature, based on the observation of some characteristics of real large social networks, such as the *small-world* phenomenon or the *scale-free* property.

Small-world networks and Erdös-Rényi random graphs. The experimental studies of Travers and Milgram [1967, 1977] have pointed out that individuals seem to be connected within a distance of around 6 intermediates to any other person, leading to the concept of *small-world* phenomenon. The idea is that the distance between any two nodes tends to be small, compared to the size of the network. In a small-world network, any agent is connected with a high probability to any other agent via a short path. A small-world network contains many short paths, in other words the diameter of the graph tends to be small, typically in the order of $\mathcal{O}(\log(n))$.

This is the case in expectation for *Erdös-Rényi random graphs* [Erdös and Rényi, 1959] for which the expected diameter is $\mathcal{O}(\log(n))$. To generate an Erdös-Renyi graph, given a parameter $p \in [0, 1]$, each edge is added with independent probability p. This produces a random graph of density tending to p.

Scale-free networks and Barabási-Albert random graphs. Another important characteristic observed in real social networks is the *scale-free* property, formulated by Barabási and Albert [1999] as follows. A network is scale-free if the degree of its vertices follows a power-law distribution, that is the fraction of vertices of degree k is proportional to $k^{-\gamma}$, for some constant γ. The main observable feature on a scale-free network is that it contains many *hubs*, that are nodes with high degree. Consequently, a scale-free network looks like an aggregation of stars. The idea is that there exist some points of attraction in the network, which are highly connected nodes to which other nodes are more likely to connect.

The *Barabási-Albert random graphs* are scale-free. In particular, the degree of the nodes of a Barabási-Albert graph follows a power-law distribution with degree exponent γ equal to three. In such random graphs, the network is iteratively constructed by adding to a subgraph a new node which is connected with higher probability to high degree nodes, following a preferential attachment mechanism. More precisely, given a subgraph G' defined on a subset of vertices $V' \subseteq V$, we construct a new graph G'' by adding a new node $v \in V \setminus V'$ that is connected to any node $u \in V'$ with probability $p_u = \frac{\delta_{G'}(u)}{\sum_{v' \in V'} \delta_{G'}(v')}$. Note that, in our experiments, we will not impose any average degree in the graph, but we will generate Barabási-Albert graphs that are undirected.

Homophily in networks. A network respects *homophily* if two "similar" nodes tend to be connected in the graph. In the context where agents are embedded in a social network, two agents can be considered similar if they have close preferences over the set of alternatives. How to measure this closeness in preferences is nevertheless subjective and several models can be imagined to take into account this similarity. For instance, Tsang and Larson [2016] integrate homophily in their model for generating random graphs by considering Erdös-Rényi graphs and Barabási-Albert graphs where they take into account the cardinal preferences of the agents. They multiply the probability of connection between two agents i and j by the proximity of the personal preferred values of the agents, which are specific points in a common scale.

In an ordinal perspective, we choose, in the experiments of the thesis, another way to generate a network that respects homophily. In our protocol, based on strict ordinal preferences, the more the agents agree on pairwise comparisons of the alternatives, the more likely they are connected. More precisely, two agents i and j are connected in G with probability equal to $q_{ij} = |\{(a,b) \in M^2 : a \succ_i b \text{ and } a \succ_j b\}|/(m \cdot (m-1)/2)$. It is noteworthy that the probability of connection between two agents is inversely proportional to the Kendall-Tau distance between their respective preference rankings. In particular, in this model, two agents with exactly the same preferences are necessarily connected. Note that for a directed graph, we decide independently whether (i,j) and (j,i) exist but with the same probability.

Real networks. In order to run experiments, it is always attractive to work with real data sets. Interestingly, it is possible for networks thanks to online libraries such as *Stanford Large Network Dataset Collection (SNAP)* [Leskovec and Krevl, 2014] or *Social Computing Data Repository* [Zafarani and Liu, 2009] that collect anonymized data from real networks, and notably social networks like Facebook or Twitter. However, these data are more adapted to graph mining studies because they contain very large graphs (which are computationally difficult to handle for the type of experiments we run in the thesis) and they do not provide any linked preference data. Therefore, for our purpose, since there is no data set (yet) with real preferences of agents and real connections among them, we need to randomly generate the preferences or the network or both.

1.2.3 Solution of a social choice problem

The goal of a social choice procedure is defined according to the set of alternatives M: choosing an alternative or a subset of alternatives, ranking the alternatives or a subset of alternatives. Note that the set of alternatives may have a combinatorial structure, such as the space of possible allocations in resource allocation problems or the space of all partitions of agents in hedonic games. In all the social choice problems that we will consider, the goal will always be the selection of one alternative within the set of alternatives M, implying that M is the set of possible solutions. The objective of a social choice procedure is then to select an alternative according to the preferences of the agents over M.

This objective can be reached through a *distributed* (also called *decentralized*) procedure or a *centralized* procedure, or procedures mixing both characteristics. In a centralized procedure, the agents report their preferences, according to a given protocol for elicitation, to a central authority which collects and aggregates them into a final decision that is finally communicated to the agents. On the contrary, in a distributed process, the agents iteratively construct a final decision by interacting among them without the intervention of a central authority [Sandholm, 1999].

The social choice procedures are usually evaluated according to their properties, i.e., the axioms that they satisfy [Plott, 1976], and according to their computational cost in terms of communication and elicitation [Conitzer and Sandholm, 2005], and in terms of the computation of the solution [Hemaspaandra, 2018].

A general property that is commonly desirable for a social choice procedure is *strategy-proofness*. A procedure is strategy-proof if the agents never have incentive to misreport their preferences. While no voting system, reasonable according to some basic axioms, is strategy-proof [Gibbard, 1973, Satterthwaite, 1975], some procedures based on the idea of *serial dictators* are immune to manipulation. These procedures may appear more attractive in the context of allocating goods [Svensson, 1999] since this setting relies on private goods.

A solution is rarely unanimously recognized as the best alternative by every agent, due to preferences that may be conflicting. Therefore, it appears necessary to design measures for evaluating the quality of solutions, such as for instance the *efficiency* of a solution. In the context of ordinal preferences, the most natural criterion is *Pareto-efficiency*.

Definition 1.14 (Pareto-efficiency) *A solution $s \in M$ is said to be Pareto-efficient if there is no other solution $s' \in M$ which Pareto-dominates it, that is if there is no solution $s' \in M$ such that for every agent $i \in N$, $s' \succsim_i s$ and there exists at least one agent $j \in N$ such that $s' \succ_j s$.*

Other well-known quality measures are based on the cardinal efficiency of the solutions, like the evaluation of the social welfare, in case the preferences are expressed via utility functions. One can notably cite the *utilitarian*, the *egalitarian* or the *Nash social welfare.*

Definition 1.15 (Utilitarian social welfare) *A solution $s \in M$ maximizes the utilitarian social welfare if $s \in \arg\max_{s' \in M} \sum_{i \in N} u_i(s')$.*

Definition 1.16 (Egalitarian social welfare) *A solution $s \in M$ maximizes the egalitarian social welfare if $s \in \arg\max_{s' \in M} \min_{i \in N} u_i(s')$.*

Definition 1.17 (Nash social welfare) *A solution $s \in M$ maximizes the Nash social welfare if $s \in \arg\max_{s' \in M} \prod_{i \in N} u_i(s')$.*

Other quality measures can be derived for more specific settings. In particular, we focus our study on two social choice problems, voting and resource allocation, for which specific requirements can be imposed on the solution and the procedure.

1.3 Voting theory

Voting is present in everyday life. Besides political elections, one could cite voting for agreeing on a date or a place for a meeting, for instance in online voting via Doodle (`www.doodle.com`) or Whale platform (`https://whale.imag.fr/`), elections of the members of a laboratory council, the choice of the menu in a restaurant, and others.

In an election model, agents are also called voters and the alternatives in M are candidates. The voters have preferences over the candidates in M and the goal is to aggregate their preferences in order to determine the *winner* (or the *winning set* in case of multi-winner elections) of the election. Whereas *social welfare functions* aim at aggregating the preferences of the voters into a collective ranking over the alternatives, we focus on mechanisms that output the winning alternative(s) of the election through a *social choice function*, also called a *voting rule*.

Let us denote a voting rule by $\mathcal{F}$. A certain type of admissible ballot, denoted by $\mathcal{B}_\mathcal{F}$, is required for a given voting rule $\mathcal{F}$. Each voter i is asked to submit a ballot $\sigma_i \in \mathcal{B}_\mathcal{F}$ to a central system which computes the outcome of the voting rule on this set of ballots. The voting rule takes as input a set of ballots, one for each voter (under the assumption that the model does not take into account abstention), which is named a voting profile, denoted by $\sigma \in \mathcal{B}_\mathcal{F}{}^n$. Therefore, a voting rule is a mapping associating with each voting profile a set of winners, i.e., $\mathcal{F} : \mathcal{B}_\mathcal{F}{}^n \to 2^M \setminus \{\emptyset\}$. When the ballot σ_i of voter i coincides with her preferences $\succsim_i$, this voter is said to be *sincere*. If all the voters are sincere in the ballot σ_i they submit to the system, then the associated voting profile $\sigma \in \mathcal{B}_\mathcal{F}{}^n$ is said to be *truthful*.

We are concerned with *single-winner elections*, where the winner must be a unique candidate. The rule is said to be *resolute* if it returns a unique winner, that is if the voting rule $\mathcal{F}$ is a mapping $\mathcal{F} : \mathcal{B}_\mathcal{F}{}^n \to M$. A *tie-breaking rule* is usually used in order to obtain a single winner, in case the initial rule is not by essence resolute. We consider in this work a deterministic tie-breaking that relies on a linear order over the candidates, denoted by $\rhd$, and we assume that this is by default the alphabetical order over the alternatives, i.e., $a \rhd b \rhd c \rhd \ldots$. The resolute voting rule, resulting from the composition of a voting rule $\mathcal{F}$ and a tie-breaking rule based on $\rhd$, is denoted by $\mathcal{F}_\rhd$. The rule $\mathcal{F}_\rhd$ outputs, for a given voting profile σ, the candidate in $\mathcal{F}(\sigma)$ that comes first in order $\rhd$.

1.3.1 Voting rules

Several voting rules have been designed in voting theory. Most of the voting rules that we consider in this document assign a score to the candidates and output as winner(s) the candidate(s) maximizing this score. More formally, for such rules, given a profile of ballots $\sigma \in \mathcal{B_F}^n$, the score of candidate $x \in M$ under rule $\mathcal{F}$ is given by $\mathrm{Sc}^{\sigma}_{\mathcal{F}}(x)$ (we omit the reference to σ and $\mathcal{F}$ when the context is clear), and $\mathcal{F}(\sigma) \subseteq \arg\max_{x \in M} \mathrm{Sc}^{\sigma}_{\mathcal{F}}(x)$. We always consider the resolute version $\mathcal{F}_{\triangleright}$ of the voting rule (even if the mention to $\triangleright$ is often omitted for the sake of simplicity), where $\triangleright$ selects the winner in case of ties in $\arg\max_{x \in M} \mathrm{Sc}^{\sigma}_{\mathcal{F}}(x)$.

In general, we assume that the voters are asked to submit as a ballot the whole ranking of their preferences with no indifference, that is $\mathcal{B_F} \in \mathcal{L}(M)$. The order over the candidates associated with ballot $\sigma_i \in \mathcal{L}(M)$ is denoted by $\succ^{\sigma_i}$, and r^{σ}_i is the function giving the rank of the candidates in σ_i, i.e., candidate x is ranked at the $r^{\sigma}_i(x)^{\text{th}}$ position in ballot $\sigma_i \in \mathcal{L}(M)$. In the case where there is indifference in the preferences of a voter i, a sincere ballot σ_i refers to an order $\succ^{\sigma_i} \in \mathcal{L}(M)$ where $x \succ^{\sigma_i} y$ implies that $x \succsim_i y$. Consequently, in presence of indifference in the preferences, several ballots can be sincere.

We focus our study on some particular voting rules, which belong to three families of rules: *Positional Scoring Rules (PSRs)*, *run-off voting rules* and *pairwise comparison voting rules*.

The PSRs compute the score of the candidates according to their position in the ranking ballot of the voters. Indeed, the score $\mathrm{Sc}^{\sigma}_{\mathcal{F}}(x)$ of alternative x under voting profile σ depends on the absolute position of x in each ballot. Concretely, we are given a vector $\alpha = (\alpha_1, \ldots, \alpha_m)$ such that $\alpha_1 \geq \ldots \geq \alpha_m$ and $\alpha_1 > \alpha_m$. If x is placed at position k in a ballot, then x receives α_k points. The score of x is defined as the sum of these points over all ballots. Thus, for all $x \in M$, $Sc^{\sigma}_{\mathcal{F}}(x) = \sum_{i \in N} \alpha_{r^{\sigma}_i(x)}$. Each PSR is characterized by its vector α. We can mention in particular:

- *Borda* : $\alpha = (m-1, m-2, \ldots, 0)$. More generally, any vector which is an affine transformation of $(m-1, m-2, \ldots, 0)$ is an acceptable α-vector for Borda. In other words, Borda gives points to the candidates in each ballot inversely to their rank in the ballot, with more points to the candidate ranked first and exactly the same difference of points between two consecutive positions in the ranking.

- *k-approval* ($k \in [m-1]$): $\alpha = (1, \ldots, 1, 0, \ldots, 0)$ with k consecutive ones. This rule assigns one point to all the candidates ranked within the first k candidates of the ballot.

- *Plurality*: $\alpha = (1, 0, \ldots, 0)$. This rule gives one point only to the candidate at the first position of each ballot. It corresponds to k-approval when k is equal to 1.

- *Antiplurality* (also known as *Veto*): $\alpha = (0, \ldots, 0, -1)$. The principle of this rule is that the voters do not approve candidates in their ballot but disapprove a candidate, by expressing a *veto* against this candidate. Up to an affine transformation of the scores, it corresponds to the k-approval rule when k is equal to $m-1$.

Note that for Plurality and Veto, it is sufficient to submit only one candidate as a ballot: the only approved candidate for Plurality and the only disapproved candidate for Veto. Therefore, for Plurality and Veto, we can assume that $\mathcal{B_F} = M$. In such cases, for the sake of simplicity, we may sometimes write a voting profile σ as an n-tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ where $\sigma_i \in M$ is the ballot of voter i.

Besides the absolute evaluation of the candidates like in PSRs, it is possible to evaluate a candidate relatively to the other candidates, like in *pairwise comparison voting rules*. In such rules, it is useful to compute the number of voters who support some candidate against another one. Given a voting profile σ with the associated vector of linear orders $\succ^\sigma \in \mathcal{L}(M)^n$ and two alternatives $a, b \in M$, let $W^\sigma(a,b)$ and $\omega^\sigma(a,b)$ be the set of voters who prefer a to b and the number of voters who prefer a to b, respectively. That is, $W^\sigma(a,b) = \{i \in N : a \succ_i^\sigma b\}$ and $\omega^\sigma(a,b) = |W^\sigma(a,b)|$. The information contained in $\omega^\sigma : M^2 \to [n]$ enables to compute the *weighted majority graph* of the voting profile σ. A weighted majority graph is a complete directed graph over M where every arc (a,b) for candidates a and b is assigned a weight equal to the number of voters who prefer a to b, that is $\omega^\sigma(a,b)$. For the sake of simplicity, we directly refer to ω^σ to designate the weighted majority graph of profile σ. Observe that $\omega^\sigma(a,b) + \omega^\sigma(b,a) = n$ for any candidates a and b such that $a \neq b$. We focus on two pairwise comparison rules where the score $Sc^\sigma_{\mathcal{F}}(x)$ of a candidate x in a voting profile σ is defined according to the number of voters who prefer x to the other candidates:

- *Maximin*: $Sc^\sigma_{\mathcal{F}}(x) = \min_{y \in M \setminus \{x\}} \omega^\sigma(x,y)$. The Maximin score of each candidate is the minimum number of voters supporting it, considering any pairwise comparison in the voting profile with another candidate.

- *Copeland*: $Sc^\sigma_{\mathcal{F}}(x) = |\{y \in M \setminus \{x\} : \omega^\sigma(x,y) > \frac{n}{2}\}|$. The Copeland score of each candidate is the number of alternatives that it beats with absolute majority in a pairwise comparison in the voting profile[1].

Finally, the run-off voting rules proceed by rounds. Among them, we consider *Single Transferable Vote (STV)* and *Plurality with run-off* which are elimination-based rules. These voting rules are well-known and are actually used for political elections in several countries: STV is notably used in Australia and Ireland, and Plurality with run-off is used in France.

- *STV*[2] is an iterated process where at each round, the loser of Plurality, i.e., the candidate with the smallest Plurality score, gets eliminated (use $\rhd$ to break ties). The input voting profile is then updated by removing every occurrence of this candidate from the ballots of the voters. The process continues until an alternative obtains an absolute majority of votes under Plurality and thus gets elected.

- *Plurality with run-off* proceeds in two rounds. Only the first two candidates according to Plurality (use $\rhd$ to break ties) remain for the second round, and the voting profile is updated by removing all the occurrences of the other candidates. The winner of Plurality in the remaining profile is chosen as the winner of Plurality with run-off. Note that only the first round is necessary if an alternative gets a strict majority of the votes.

[1] Other versions of the Copeland score also subtract the number of defeats by pairwise comparisons.

[2] The name used in political elections for single-winner elections is *Instant Runoff Voting*, whereas the term STV refers to multi-winner elections. However, according to the social choice literature, we use the common terminology of STV even for single-winner elections.

1.3.2 Evaluation of a voting rule

The traditional evaluation of a voting rule is axiomatic: one tries to identify some desirable properties that are satisfied by the voting rule. Since there is no perfect voting rule satisfying natural requirements [Arrow, 1951], the choice of a voting rule is determined by the properties that one would like to satisfy in a given context, according to the fact that some desired axioms are incompatible.

For instance, a condition that appears reasonable to satisfy is the fact that, a priori, no alternative is favored compared to another, property called *neutrality*. The neutrality axiom implies symmetric treatment of the alternatives, or in other words, impartiality among the candidates. Suppose $\mu : M \to M$ is a permutation over the candidates. Then, for a given voting profile σ, let us denote by $\sigma.\mu$ the voting profile σ where any occurrence of alternative x has been replaced by $\mu(x)$.

Definition 1.18 (Neutrality) *A voting rule $\mathcal{F}$ is neutral if and only if for every voting profile $\sigma \in \mathcal{B_F}^n$ and every permutation $\mu : M \to M$, $\mathcal{F}(\sigma) = x$ implies that $\mathcal{F}(\sigma.\mu) = \mu(x)$.*

All the voting rules cited in the previous subsection are basically neutral but we lose this property by combining them with a deterministic tie-breaking rule based on a linear order over the candidates.

Another natural axiom is the *unanimity* requirement, imposing that the voting rule must elect an alternative that is unanimously supported by all the voters.

Definition 1.19 (Unanimity) *A voting rule $\mathcal{F}$ is unanimous if and only if for every voting profile $\sigma \in \mathcal{B_F}^n$ such that there exists a candidate x ranked first by every voter, $\mathcal{F}$ elects x, i.e., if $|\{i \in N : x \succ_i^\sigma y,\ \forall y \in M \setminus \{x\}\}| = n$, then $\mathcal{F}(\sigma) = x$.*

Despite the fact that the unanimity condition appears simple to satisfy, some reasonable voting rules do not satisfy this basic axiom. Among the voting rules evoked in the previous subsection, k-approval (for $k > 1$) and Veto are not "strictly" unanimous, in the sense that they do not necessarily elect, as a single-winner, an alternative ranked first by every voter. However, such an alternative always belongs to the set of winners ex aequo under these rules, before a tie-breaking rule applies.

The unanimity condition can be generalized by requiring that the outcome of the voting rule must be a candidate ranked first by at least a given portion of the voters. By considering more than half of the voters as a specific quota, this boils down to the definition of the *majority rule*. The associated axiom, that also implies unanimity, is *majority consistency*.

Definition 1.20 (Majority consistency) *A voting rule $\mathcal{F}$ is majority consistent if and only if for every voting profile $\sigma \in \mathcal{B_F}^n$ such that there exists a candidate x ranked first by a strict majority of voters, $\mathcal{F}$ elects x, i.e., if $|\{i \in N : x \succ_i^\sigma y,\ \forall y \in M \setminus \{x\}\}| > n/2$, then $\mathcal{F}(\sigma) = x$.*

Among the voting rules cited in the previous subsection, k-approval (for $k > 1$), Veto and Borda are not majority consistent.

Note that unanimity and majority consistency can be interpreted in two different ways, according to either a relative evaluation of the alternatives, a point of view supported by Nicolas (marquis) de Condorcet, or to an absolute evaluation of the alternatives, a vision defended by Jean-Charles (chevalier) de Borda. For instance, according to the former point of view, majority consistency forces to elect a candidate that a majority prefers to any other candidate, whereas according to the latter conception, majority consistency must elect a candidate that is ranked first by a majority of voters. In fact, in the conception of Borda, this is the rank of the candidates in the preference profile which matters, whereas for Condorcet this is the quality of the candidates in comparison to the other candidates which must be taken into account. Following the idea of Condorcet, the *Condorcet winner* is a candidate which beats any other candidate in a pairwise comparison. By definition, there is at most one Condorcet winner.

Definition 1.21 (Condorcet winner) *A candidate x is the Condorcet winner in a voting profile σ, denoted by $CW(\sigma)$, if and only if $|\{y \in M \setminus \{x\} \mid \omega^{\sigma}(x,y) > \frac{n}{2}\}| = m-1$.*

However, a Condorcet winner is not guaranteed to exist for any instance of an election [de Caritat marquis de Condorcet, 1785]. This implies that a voting rule defined only by the election of the Condorcet winner cannot output a winner for any preference profile. Nevertheless, under some restrictions such as single-peaked preferences, there always exists a Condorcet winner ("Median voter theorem" [Hotelling, 1929, Black, 1948, Downs, 1957]). The restriction of possible preference profiles to those admitting a Condorcet winner is referred to as the *Condorcet domain*. A voting rule which always elects the Condorcet winner alone, when restricted to the Condorcet domain, is said to be *Condorcet consistent*.

Definition 1.22 (Condorcet consistency) *A voting rule $\mathcal{F}$ is Condorcet consistent if and only if for every voting profile σ, $CW(\sigma) \neq \emptyset$ implies that $\mathcal{F}(\sigma) = CW(\sigma)$.*

Among the rules we have previously presented, only Copeland and Maximin are Condorcet consistent. Note that a Condorcet consistent rule is necessarily majority consistent and thus unanimous.

When a voting rule is not Condorcet consistent, it is interesting to empirically observe how often it elects the Condorcet winner, when it exists. The *Condorcet efficiency* of a voting rule is the probability that the rule elects the Condorcet winner when restricted to the Condorcet domain [Gehrlein, 1997]. Some theoretical and, especially, experimental studies can be conducted in order to evaluate the Condorcet efficiency of a given voting rule.

In the same idea but regarding the absolute evaluation of the candidates, one could refer to the frequency of electing a candidate which corresponds to the winner of the Borda rule as the *Borda efficiency*. However, evaluating a voting rule following the Borda efficiency is less meaningful than Condorcet efficiency. Indeed, since the Borda scores are an absolute evaluation of the quality of a candidate, it is more relevant to observe how good is the Borda score of the winner of a given voting rule [Grandi et al., 2013]. In order to normalize the Borda scores and more easily compare them, we use the *Borda closeness* measure, denoted by $BC_{\mathcal{F}}(\sigma)$ for a given voting rule $\mathcal{F}$ and a voting profile σ, where Sc_B^σ is the vector of Borda scores on voting profile σ and $\mathcal{F}_B$ denotes the Borda rule.

$$BC_{\mathcal{F}}(\sigma) = \begin{cases} \frac{\mathrm{Sc}_B^\sigma(\mathcal{F}(\sigma)) - \min_{x \in M} \mathrm{Sc}_B^\sigma(x)}{\mathrm{Sc}_B^\sigma(\mathcal{F}_B(\sigma)) - \min_{x \in M} \mathrm{Sc}_B^\sigma(x)} & \text{if } \min_{x \in M} \mathrm{Sc}_B^\sigma(x) \neq \max_{x \in M} \mathrm{Sc}_B^\sigma(x) \\ 1 & \text{otherwise} \end{cases}$$

More precisely, the Borda closeness measures how much the Borda score of the winner of profile σ under $\mathcal{F}$ is close to the Borda score of the Borda winner in σ compared to the Borda score of the Borda loser, i.e., the minimum possible Borda score in profile σ.

Finally, another important point to take into account in the evaluation of a voting rule is computational. The thesis only focuses on voting rules for which the computation of the output is polynomial in the number of voters and the number of candidates. All the voting rules that are cited in the previous subsection satisfy this property.

1.3.3 Strategic voting: voting game and iterative voting

Tactical voting appears very frequently in real life elections, and in political elections in particular. The voters adopt a strategic behavior in voting when they submit a ballot which is not in accordance with their real preferences. By this way, they hope to avoid some unfortunate outcomes and try to maximize their satisfaction. The strategic behavior in voting is usually called *manipulation* [Taylor, 2005].

Despite the fact that manipulation can be undesirable, the Gibbard-Satterthwaite theorem [1973, 1975] establishes that no voting rule that satisfies basic axioms is immune to manipulation. An approach to circumvent the problem of manipulation in voting is to design voting rules that are computationally hard to manipulate, in order to prevent voters to efficiently design strategies for manipulation. A rich literature has developed on this point, analyzing the computation cost of a manipulation for different voting rules [Bartholdi et al., 1989b, Bartholdi and Orlin, 1991]. However, this approach is not sufficient to avoid manipulation, as pointed out by more recent works [Conitzer and Sandholm, 2006, Faliszewski and Procaccia, 2010]. The main drawback of this approach relies on the fact that the complexity of manipulation is analyzed in the worst case. Therefore, even if there exists some cases for which a rule is computationally hard to manipulate, this does not prevent the manipulation of the voters in practice, because these hard cases may never occur.

Another perspective consists in exploring the natural game-theoretical properties of strategic voting. In this idea, manipulation is not necessarily considered as a bad behavior, but as a behavior to be taken into account in the voting process.

1.3.3.a Voting game

Strategic voting can be seen as a strategic game. Let us recall some basic notions about game theory (see for instance Osborne and Rubinstein [1994]).

A strategic game aims at modeling a situation where several individuals must simultaneously choose an action, according to the payoff (or the cost) that this action, combined with the actions chosen by the other individuals, will generate. Strategic games are *non-cooperative games* [Nash, 1951], in the sense that the individuals are self-interested and have their own goals and beliefs, even if they may associate with others.

Alternatively, in a *cooperative game*, individuals can form coalitions, in which the share, among the members of the coalition, of the payoff (or the cost) generated by their joint action is also an issue of the game.

In a strategic game, the individuals are referred to as *players*, whose set is denoted by $\mathscr{N}$. The non-empty set of possible actions, or *strategies*, for any player $i \in \mathscr{N}$ is denoted by $\mathscr{S}_i$. The game is said to be *finite* if the set $\mathscr{S}_i$ of actions of every player $i \in \mathscr{N}$ is finite, and the game is *symmetric* if the set of actions is the same for every player, i.e., $\mathscr{S}_i = \mathscr{S}_j$ for all players $i \neq j$. A state of the game is a strategy profile $s \in \mathscr{S} := \mathscr{S}_1 \times \ldots \times \mathscr{S}_{|\mathscr{N}|}$ where all the players $i \in \mathscr{N}$ simultaneously choose an action or a strategy. Such a state is a profile of *pure strategies*. The satisfaction of a player $i \in \mathscr{N}$ is defined according to a *payoff* function, $v_i : \mathscr{S} \to \mathbb{R}$ (or a *cost* function $c_i : \mathscr{S} \to \mathbb{R}$), or ordinal preferences $\succcurlyeq_i$ over $\mathscr{S}$. By considering that the choice of the players can be non-deterministic, a *mixed strategy* for player $i \in \mathscr{N}$ is a probability distribution over $\mathscr{S}_i$, defining in consequence a profile of mixed strategies. In such a case, the players have preferences over the set of lotteries on $\mathscr{S}$.

A *solution concept* of the game defines specific formal rules in order to predict a possible outcome of the game. An *equilibrium* is a stable state with respect to a set of possible moves, also called *deviations*, predefined by the solution concept. A solution concept defines which types of deviations are valid, according to the type of move and the deviating players taken into account. For instance, a *Nash equilibrium* [Nash, 1951] is a state that is immune to unilateral deviations with strict improvement from the players.

A voting game is given by an instance $\langle N, M, \succsim, \mathcal{F} \rangle$, where the players correspond to the voters, i.e., $\mathscr{N} = N$. The strategies of the voters are the ballots they can submit to the voting system. A voting game is symmetric in the sense that the set of strategies $\mathscr{S}_i$ for each voter $i \in N$ is equal to the set $\mathcal{B}_\mathcal{F}$ of all ballots acceptable by voting rule $\mathcal{F}$. A state of a voting game is a voting (or strategy) profile $\sigma \in {\mathcal{B}_\mathcal{F}}^n$. The states in ${\mathcal{B}_\mathcal{F}}^n$ are not directly evaluated by the voters, contrary to a general strategic game, because the voters $i \in N$ have preferences $\succsim_i$ over the set of candidates M. Therefore, instead of directly evaluating $\sigma \in {\mathcal{B}_\mathcal{F}}^n$, the voters evaluate the winner(s) $\mathcal{F}(\sigma)$ of the state according to voting rule $\mathcal{F}$. We focus on voting games that are non-cooperative games with pure strategies.

Let us formulate some solution concepts of strategic games in the context of voting games. A strategy profile $\sigma \in {\mathcal{B}_\mathcal{F}}^n$ that is restricted to the agents of $N \setminus \{i\}$ for a given player $i \in N$ is denoted by σ_{-i}. We can then write strategy profile σ as $\sigma = (\sigma_i, \sigma_{-i})$. The Nash equilibrium is defined as follows in a voting game.

Definition 1.23 (Nash equilibrium) *A strategy profile $\sigma \in {\mathcal{B}_\mathcal{F}}^n$ is a Nash equilibrium if and only if for every voter i and every ballot $\sigma'_i \in \mathcal{B}_\mathcal{F}$, it holds that $\mathcal{F}(\sigma) \succsim_i \mathcal{F}(\sigma_{-i}, \sigma'_i)$.*

In a Nash equilibrium, no single voter has incentive to change her ballot. The existence and identification of Nash equilibria in voting games is a main concern (see for instance Pattanaik and Sengupta [1982], Feddersen et al. [1990], Desmedt and Elkind [2010]). However, there are usually many Nash equilibria in voting games, and they are not all relevant for describing plausible outcomes of the game [Farquharson, 1969].

Nevertheless, in voting, it can be unrealistic and optimistic to consider that only single agents can deviate, drawing a weakness of the Nash equilibrium. This leads to consider other solution concepts, which are based on coalitional deviations. A strategy profile σ that is restricted to a coalition $C \subseteq N$ is denoted by σ_C, and the strategy profile σ that is restricted to the agents out of C is denoted by σ_{-C}. Thus, (σ'_C, σ_{-C}) denotes σ in which σ_i is replaced by σ'_i if and only if $i \in C$.

A stable state with respect to deviations, for any possible coalition of voters, with strict improvement for each member of the coalition, refers to the *strong (Nash) equilibrium* [Aumann, 1959]. A strong equilibrium is defined as follows in a voting game.

Definition 1.24 (Strong (Nash) equilibrium) *A strategy profile $\sigma \in \mathcal{B_F}^n$ is a strong equilibrium if and only if there is no coalition $C \subseteq N$ of voters and no strategy profile $\sigma'_C \in \mathcal{B_F}^{|C|}$ such that $\mathcal{F}(\sigma_{-C}, \sigma'_C) \succ_i \mathcal{F}(\sigma)$ for every voter $i \in C$.*

In a strong equilibrium, no coalition of voters has incentive, for all the members of the coalition, to change her ballot. Some conditions of existence of a strong equilibrium have been established for voting games, by determining certain types of games that always admit a strong equilibrium [Peleg, 2002, Peleg and Peters, 2010]. Strong equilibria have also been characterized according to the notion of Condorcet winner [Sertel and Sanver, 2004] and for specific voting rules [Messner and Polborn, 2007].

Besides classical solution concepts from game theory, some solution concepts have been designed especially for the voting game framework. For instance, the so-called *voting equilibrium* [Myerson and Weber, 1993] is a state consistent with a preliminary poll.

1.3.3.b Iterative voting

It is possible to consider a dynamic version of a voting game, called *iterative voting* (see Meir [2017] for a recent survey). Initially, all the voters simultaneously cast a ballot, which leads to the election of a given candidate according to voting rule $\mathcal{F}$. From this initial state, voters (or coalitions of voters) can manipulate by successively changing their strategy while the rest of the voters keeps her current ballot. Actually, the voters deviate by rounds according to the type of moves that is allowed by a given solution concept. These deviations as well as the initial state define a *dynamics* of the game. Iterative voting can be seen as the responses to a succession of polls where the voters can observe the previous votes and strategize consequently [Reijngoud and Endriss, 2012]. Alternatively, it can simply describe the changes in the vote intentions of the voters during a pre-election period, where the vote intentions can potentially evolve regarding the others' opinions.

More formally, let us denote by $\sigma^0 \in \mathcal{B_F}^n$ the initial state of the game, where by assumption of no abstention, every voter gives a ballot to the central system. This initial state is a voting profile which is usually supposed to be truthful in the literature[1]. In general σ^t denotes the current voting profile at step t, and the voters deviate according to the outcome of the state that is given by $\mathcal{F}(\sigma^t)$. For the sake of simplicity, when the context is clear, we denote by Sc^t, instead of Sc^{σ^t}, the score of voting profile σ^t under $\mathcal{F}$. A dynamics is a sequence of states $(\sigma^0, \sigma^1, \ldots, \sigma^T)$ such that each pair of consecutive

[1] This restriction is rather natural: the voters start by giving their true opinion, especially if they do not have any previous information about the game, and if they are not satisfied with the outcome then they reconsider their vote.

steps $t, t+1$ is associated with the deviation of some voter (or coalition of voters) that turns σ^t into σ^{t+1}. A deviation at step t is supposed to occur between the steps t and $t+1$. The dynamics is said to *converge* when the associated sequence of deviations is finite. The dynamics fails to converge when a state appears more than once in the sequence of deviations.

It is commonly assumed that only one voter (or coalition of voters) is deviating at each step[1]. Different sequences of deviations can occur, depending on who deviates at a step if different agents can do so. It is possible to define rules to choose the voter (or coalition of voters) who deviates at a given step, among the voters (or coalitions of voters) having incentive to manipulate. These rules are formalized through a *turn function* (or scheduler) [Apt and Simon, 2012], that we denote by τ. In general, we do not assume a specific fixed turn function for the iterative voting game.

In the initial iterative voting formulation of Meir et al. [2010], a single voter is deviating at each step of the game in order to strictly improve the outcome of the election according to her preferences. Consequently, if the dynamics converges, then it reaches a stable state corresponding to a Nash equilibrium. The convergence properties of the iterative voting game have been widely studied [Meir et al., 2010, Lev and Rosenschein, 2012, Reyhani and Wilson, 2012, Grandi et al., 2013, Obraztsova et al., 2015]. Up to our knowledge, only Nash equilibria have been studied in the iterative voting setting, i.e., only unilateral deviations with strict improvement are made at each step of the game. In this context, it is possible to define a *better response* as a strict improvement for the voter who deviates.

Definition 1.25 (Better response) *A ballot $\sigma'_i \in \mathcal{B}_\mathcal{F}$ is a better response for agent i from strategy profile $\sigma \in \mathcal{B}_\mathcal{F}{}^n$ if $\mathcal{F}(\sigma_{-i}, \sigma'_i) \succ_i \mathcal{F}(\sigma)$.*

Players are said to be *rational* if they aim to maximize their payoff according to the information they get about the strategies of the other players. Under such a classical assumption in game theory, it seems natural to restrict to best response deviations.

Definition 1.26 (Best response) *A ballot $\sigma'_i \in \mathcal{B}_\mathcal{F}$ is a best response for agent i from strategy profile $\sigma \in \mathcal{B}_\mathcal{F}{}^n$ if σ_i is a better response for agent i and there is no other strategy ballot σ''_i such that $\mathcal{F}(\sigma_{-i}, \sigma''_i) \succ_i \mathcal{F}(\sigma_{-i}, \sigma'_i)$.*

Concerning convergence of the voting game to a Nash equilibrium, one of the most important results is the guarantee of convergence of the dynamics defined by *direct best responses* under Plurality and Veto [Meir et al., 2010, Lev and Rosenschein, 2012, Reyhani and Wilson, 2012][2].

Definition 1.27 (Direct best response (Plurality)) *From voting profile $\sigma \in M^n$, a deviation to a ballot $\sigma'_i \in M$ is a direct best response for agent i under Plurality, if σ'_i is a best response for i from σ and $\sigma'_i = \mathcal{F}(\sigma_{-i}, \sigma'_i)$.*

[1] Trivial cases of non convergence can occur otherwise (see for instance Proposition 2 of Meir et al. [2010]).

[2] These references do not explicitly use the term of *direct best responses* which will be introduced later, more precisely under the term *direct best replies*, by Meir et al. [2017].

Definition 1.28 (Direct best response (Veto)) *From voting profile $\sigma \in M^n$, a deviation to a ballot $\sigma'_i \in M$ is a direct best response for agent i under Veto, if σ'_i is a best response for i from σ and $\sigma'_i = \mathcal{F}(\sigma)$.*

A direct best response in Plurality consists in approving the candidate that will become the new winner. Alternatively, a direct best response in Veto consists in disapproving the candidate that is the current winner. Clearly, if there exists a better response for agent i from strategy profile σ, then there exists a best response for i from σ, and then there also exists a direct best response for i from σ. Therefore, the restriction to direct best responses does not weaken the Nash equilibrium that is reached by the dynamics, but only chooses a certain type of best response.

However, the generalization of this convergence result to other voting rules does not hold. Nevertheless, by further restricting the deviations in such a way that the associated solution concept is not a Nash equilibrium but a weaker notion, it is possible to reach convergence [Reijngoud and Endriss, 2012, Grandi et al., 2013, Obraztsova et al., 2015]. These restricted manipulation moves aim at modeling simple heuristic deviations that can be cognitively manageable for the agents, especially for some rules like Borda where the ballot is a linear order over the candidates. See Meir [2017] for an overview on the restricted manipulation moves.

An interesting question in this framework is the quality of the equilibria that are reached by the dynamics in case of convergence. This question has also been investigated, for the case of unilateral deviations, in a theoretical point of view [Reijngoud and Endriss, 2012, Grandi et al., 2013], via for instance Condorcet consistency, and through an empirical analysis [Reijngoud and Endriss, 2012, Grandi et al., 2013, Koolyk et al., 2017]. These empirical studies highlight the fact that, in practice, the dynamics cycles very rarely. Moreover, the iterative voting process can be viewed as a new voting rule whose outcome is the equilibrium reached by the dynamics. Therefore, the quality of the equilibria in terms of Condorcet efficiency and Borda scores can be investigated [Reijngoud and Endriss, 2012, Grandi et al., 2013, Koolyk et al., 2017]. It turns out that iterative voting produces rather good outcomes according to these criteria, that are often better than the quality of the outcome without iteration.

Note that some parameters are important for choosing a best response or a heuristic for best response [Reijngoud and Endriss, 2012, Grandi et al., 2013]. For instance, the voters are almost always assumed to be *rational* and *myopic*. Indeed, the lookahead horizon of the voters should also be taken into account. The usual myopia of the voters means that their strategic horizon is only the next step and they do not establish long-term strategies. Some works relax this assumption. Obraztsova et al. [2016] follow an idea close to local-dominance [Meir et al., 2014] (further explained in a subsequent paragraph), by considering that some thresholds determine the optimism horizon of the voters and they focus their study on Plurality and Veto. They define as an *NM-Plurality response* (referring to *non-myopic* voters) a deviation to a ballot approving the preferred candidate among the possible winners at a particular optimism horizon, i.e., among the candidates whose current score is not less than the score of the winner minus the optimism horizon value of the voter.

Another important point to take into account within the deviations of the voters is the information that the voters have for computing a manipulation at a given step. This raises the question of uncertainty about the current voting profile.

1.3.3.c Voting under uncertainty

In iterative voting, an important parameter to take into account in the dynamics of deviations is the *information available to the voters*. In the classical framework that we have depicted, the voters are assumed to know the current votes of all the other voters, allowing them to design best response strategies. However, this assumption appears highly unrealistic, especially in large scale instances or political elections.

Starting from this observation, an important literature begins to develop in order to face with uncertainty in voting. Let us briefly present some of them[1].

- The Bayesian approach assumes that the voters think in terms of probabilities over the possible voting profiles [Myerson and Weber, 1993, Messner and Polborn, 2005, Hazon et al., 2008]. Then, they try to maximize their expected utility function.

- The knowledge-based approach uses modal logic frameworks in order to formalize manipulation in a context of incomplete information [Chopra et al., 2004, Van Ditmarsch et al., 2012]. Different types of manipulation defined according to the knowledge of the agents are enumerated.

- In the *local-dominance* perspective, some uncertainty thresholds are added to the classical framework of iterative voting [Meir et al., 2014, Meir, 2015]. The focus is given on Plurality. A local-dominant strategy is a ballot that dominates all the possible ones within the set of possible profiles considered by the voter. The possible current profiles are all the profiles that are at most at a certain distance from the real current profile. The distance is given by the uncertainty thresholds and a certain metric.

- Another approach analyzes the possibility of manipulation according to an *information function*, giving a specific type of information, such as the winner, the scores or the weighted majority graph of the current voting profile, which is communicated to the voters [Reijngoud and Endriss, 2012, Endriss et al., 2016]. According to the type of information that is given to the voters, the susceptibility to manipulation of different voting rules is investigated. The voters adopt a risk-averse behavior in the sense that the played strategy must be undominated over all the possible voting profiles that coincide with the information they have.

- Another way to deal with uncertainty is to consider a set of possible profiles according to partial available votes [Conitzer et al., 2011, Dey et al., 2016]. The voters infer as possible voting profiles all those which coincide with the partial votes. While Conitzer et al. [2011] choose to study the dominant strategy within this set as a best response, Dey et al. [2016] focus on different types of manipulation based on possible or necessary winners within the set of possible profiles.

Recently, some works assume a *social network structure* over the agents in order to understand certain behaviors in strategic voting [Grandi, 2017]. In this line of research, uncertainty in voting has been studied by assuming that the agents directly infer from their links in the social network the current votes of each voter [Chopra et al., 2004, Sina

[1] For a more complete survey, see Meir [2017].

et al., 2015, Tsang and Larson, 2016, Tsang et al., 2018]. This naturally leads to a bias in the deviations and the convergence of an iterative voting process can be analyzed in consequence.

The collective decision that is made via voting systems sets a configuration which will be the same for every agent. Other interesting questions arise in social choice problems where an alternative defines different situations according to the identity of the agents. This is the case in resource allocation problems.

1.4 Resource allocation of indivisible goods

Resource allocation of *indivisible* goods [Chevaleyre et al., 2006, Bouveret et al., 2016] is a main issue in Artificial Intelligence, which is at the intersection of Economics and Computer Science, and has many real life applications. One can cite for instance divorce settlements, allocation of articles to reviewers, assignment of courses to students, and many others.

In a resource allocation problem, given a set O of r resources $\{o_1, \ldots, o_r\}$ (also called items or objects), and a set of agents who have preferences over the subsets of resources, the goal is to allocate the items among the agents in an efficient and fair manner [Young, 1995, Brams and Taylor, 1996, Moulin, 2004]. We assume that the set of resources is finite and that the resources are neither shareable nor divisible among the agents, that is each resource must be assigned to at most one agent. The divisibility of the resources is assumed in a particular resource allocation setting called *cake cutting* [Steinhaus, 1948, Robertson and Webb, 1998, Procaccia, 2016]. In the context of indivisible goods, the set of alternatives M is the set of all possible assignments of objects to agents. An allocation $\pi \in M$ is a mapping $\pi : N \to 2^O$ such that for any pair of agents i and j, $\pi(i) \cap \pi(j) = \emptyset$, where $\pi(i)$ denotes the bundle assigned to agent i [1]. If $\bigcup_{i \in N} \pi(i) = O$, all the objects have been assigned, and the allocation is said to be *complete*. All the allocations that we consider are assumed to be complete, therefore the set of possible assignments M is of size $m = n^r$.

Remark that, contrary to the general framework of social choice presented in Section 1.2.1, we have stated that the agents have preferences over the resources, instead of preferences over the set of alternatives that are the possible allocations. Actually, this simplification makes sense by making the classical assumption that the agents only care about their own share. Under such an assumption, we can simplify the problem by only considering preferences over all subsets of items, instead of preferences over all the possible assignments. Moreover, in case of cardinal preferences, if one assumes that the preferences are additive, then the value that an agent attributes to a subset of objects is equal to the sum of the values that she assigns to the isolated objects. Therefore, in such a case, it suffices to know the valuations of the agents over O to deduce the preferences over the shares and then the preferences over the assignments.

In the thesis, we assume that no monetary compensation is allowed.

[1] In order to simplify the notations and not confuse objects with allocations, we use π to denote an allocation instead of the general notations for alternatives in M like a, b or x.

1.4.1 Evaluation of an allocation of goods

Two dimensions are investigated to evaluate the quality of an assignment in resource allocation: *efficiency* and *fairness.*

Efficiency usually refers to *Pareto-efficiency* or to the maximization of the *social welfare.* An allocation is Pareto-efficient if no other allocation is at least as good for every agent and strictly better for at least one agent. The maximization of the social welfare, which applies when the preferences are cardinal, can involve different notions of welfare, such as the utilitarian, the egalitarian or the Nash social welfare.

Various notions of fairness can also be used, depending on the interpretation of fairness, like for instance *proportionality* or *envy-freeness.* Used in a context where the agents have utilities over the subsets of resources, $u_i : 2^O \to \mathbb{R}$, and initially proposed for divisible resources [Steinhaus, 1948], the proportionality requirement prescribes that every agent should value her share at least her value for the entire set of resources divided by the number of agents. From the definition of proportionality, it appears that this notion is of interest especially when the utilities of the agents are *additive*, that is each bundle of objects is evaluated by an agent as the sum of the utilities of its different items.

Definition 1.29 (Proportionality) *An allocation $\pi \in M$ is proportional if and only if $u_i(\pi(i)) \geq u_i(O)/n$ for every agent i.*

Envy-freeness [Tinbergen, 1946, Foley, 1967, Varian, 1974] requires that no agent finds her share less valuable than the share of another agent. This notion can be defined either in a context of cardinal preferences with utilities $u_i : 2^O \to \mathbb{R}$, or ordinal preferences over all bundles of resources.

Definition 1.30 (Envy-freeness) *An allocation $\pi \in M$ is envy-free if and only if for every agent i, $u_i(\pi(i)) \geq u_i(\pi(j))$, or more generally, $\pi(i) \succsim_i \pi(j)$, for every agent j.*

Note that, in order to make the notion of envy-freeness meaningful, the allocation needs to be complete. Otherwise, the trivial allocation allocating no object is always envy-free.

When the preferences are cardinal, several measures for evaluating the degree of envy of an agent, and more generally, of the society, can be designed. A rich literature has developed on this point [Lipton et al., 2004, Chevaleyre et al., 2007a, Nguyen and Rothe, 2014]. A general classification for measuring the degree of envy of an allocation has been done by Chevaleyre et al. [2017], in a model where monetary compensations are allowed. Let us present their classification by restricting to the case without payment.

Given an allocation π, the positive envy that an agent i feels towards agent j is given by $e_{i,j}$, defined as follows.

$$e_{i,j} = \max\{0, u_i(\pi(j)) - u_i(\pi(i))\}$$

One can either define the degree that agent i envies j by considering how much i prefers the bundle of j compared to hers, denoted by e^{raw}, or only focus on the fact that she is envious of j, i.e., a boolean degree of envy denoted by e^{bool}.

$$e^{raw}(i,j) = e_{i,j} \qquad e^{bool}(i,j) = \begin{cases} 1 & \text{if } e_{i,j} > 0 \\ 0 & \text{otherwise} \end{cases}$$

Now, let us state the global degree of envy of a single agent i, by considering all the possible agents j that i can envy. The sum and the maximum are two natural operators for aggregating the individual degrees of envy, defined by e^{op} where $op \in \{raw, bool\}$, into a collective one.

$$e^{max,op}(i) = \max_{j \in N} e^{op}(i,j) \qquad e^{sum,op}(i) = \sum_{j \in N} e^{op}(i,j)$$

Similarly, these operators enable to define a degree of envy for the whole society, where the operator for the degree of envy of a single agent is $op_1 \in \{sum, max\}$ and the envy between two agents is defined according to $op_2 \in \{raw, bool\}$.

$$e^{max,op_1,op_2} = \max_{i \in N} e^{op_1,op_2}(i) \qquad e^{sum,op_1,op_2} = \sum_{i \in N} e^{op_1,op_2}(i)$$

One can then investigate the optimization of the global degree of envy of the society, by searching for an allocation minimizing the envy of the society, such as the minimization of the maximum pairwise envy [Lipton et al., 2004], for which the objective is to minimize $e^{max,max,raw}$.

Note that, instead of defining envy by the difference of the utilities, one can define an envy-ratio $e'_{i,j}$ and then derive the previous degrees of envy for a single agent and the whole society based on this new definition of envy [Lipton et al., 2004, Nguyen and Rothe, 2014].

$$e'_{i,j} = \begin{cases} \max\left\{1, \frac{u_i(\pi(j))}{u_i(\pi(i))}\right\} & \text{if } u_i(\pi(i)) \neq 0 \\ +\infty & \text{otherwise} \end{cases}$$

In general, fairness and efficiency are difficult to combine. For instance, deciding whether an allocation that is both envy-free and Pareto-efficient exists is computationally hard [Bouveret and Lang, 2008, de Keijzer et al., 2009]. Even with a weaker notion of efficiency, by only imposing that the allocation is complete, the probability of existence of an envy-free allocation can be very low [Dickerson et al., 2014].

1.4.2 Specific problems: house allocation and housing market

Some restricted versions of resource allocation with indivisible goods are interesting to investigate, for their applications or their specific properties. This is the case of the setting where exactly one resource must be assigned to each agent.

1.4.2.a House allocation

When exactly one object has to be assigned to each agent, the problem of resource allocation of indivisible goods refers to an *assignment problem* [Gardenfors, 1973], called *house allocation* [Hylland and Zeckhauser, 1979, Abdulkadiroğlu and Sönmez, 1998, 1999]. In the economic literature, this problem is also known as a *one-sided matching* [Zhou, 1990], a specific type of matching under preferences [Manlove, 2013, Klaus et al., 2016], where the agents have preferences over the resources but the resources do not have any preference over the agents. Note that another problem exists, called *two-sided matching*, where the resources also have preferences over the agents. This is the classical setting of the *stable marriage* problem [Gale and Shapley, 1962] which can have some applications in the assignment of students to universities for instance. Real life applications of house allocation can be, for example, the allocation of tasks to employees, or the assignment of time slots in a time schedule.

In this context, the set O of resources is of size $r = n$, and the set M of alternatives, i.e., all possible assignments, is of size $m = n!$. A possible allocation $\pi \in M$ is therefore a mapping $\pi : N \to O$ where $\pi(i) \neq \pi(j)$ for every agents i and j. The allocations are sometimes written as n-tuples where i^{th} coordinate corresponds to $\pi(i)$, the object assigned to agent i in π. Moreover, we denote by $\pi_{|N'}$, for $N' \subseteq N$, the partial allocation from π restricted to the agents in N'. Note that there is no need of defining preferences over bundles of objects in this framework. So, we simply assume that the agents have ordinal preferences over the set O of resources.

The questions of Pareto-efficiency and fairness have also been investigated for this particular problem [Bogomolnaia and Moulin, 2001, Abraham et al., 2005]. The definition of envy is very simple in house allocation: an agent envies another one if she prefers the object assigned to the other agent to her own item. However, it is a very strong requirement because an envy-free allocation exists if and only if the agents' top objects (the objects ranked first in their preferences) are all different.

A classical and efficient mechanism in house allocation that is widely used for real applications is the *Serial dictatorship* protocol [Satterthwaite and Sonnenschein, 1981]. Serial dictatorship specifies an order over the agents, and each agent is asked to choose an object among the available resources at her turn. The order is either fixed or randomly generated (*random serial dictatorship* [Abdulkadiroğlu and Sönmez, 1998]). This mechanism is proved to be *strategy-proof* [Svensson, 1999], that is no agent has incentive to choose an object that is not her preferred one among the remaining items.

1.4.2.b Reallocation and housing market

In house allocation, a subtlety is added when the agents are initially endowed with a resource. Concretely, there exists an initial allocation $\pi^0 : N \to O$ such that $\pi^0(i)$ denotes the initial object of agent i. The goal is then to *reallocate* the resources among agents. This variant of the problem is called *housing market* [Shapley and Scarf, 1974]. In such a context, the agents typically trade their resources.

A trade is said to be *rational* if every agent involved in the exchange obtains an object that she prefers to her previous resource. The housing market setting as well as the notion of rationality is presented in the following example.

Example 1.4 *Let us consider an instance of a housing market problem with four agents, where $N = \{1, 2, 3, 4\}$ and $O = \{o_1, o_2, o_3, o_4\}$. The preferences of the agents are the following where the initial endowment of each agent is framed.*

$$\begin{array}{rl} 1: & o_1 \succ o_2 \succ o_4 \succ \boxed{o_3} \\ 2: & o_2 \succ o_3 \succ o_1 \succ \boxed{o_4} \\ 3: & o_3 \succ o_1 \succ \boxed{o_2} \succ o_4 \\ 4: & o_2 \succ o_4 \succ o_3 \succ \boxed{o_1} \end{array}$$

An exchange between agent 1 and agent 3 would be rational: they both prefer the object assigned to the other agent.

Two approaches can be adopted in housing market: a *centralized approach* where a central authority decides how to rearrange the objects among the agents by guiding their exchanges, or a *distributed approach* where the agents directly exchange their objects, without any external intervention, in order to obtain a better assignment.

Efficient centralized algorithms have been designed for housing market, namely the well-known *Top-Trading Cycle (TTC)* algorithm [Shapley and Scarf, 1974] and some variants [Abdulkadiroğlu and Sönmez, 1999, Aziz and De Keijzer, 2012]. Here is the global idea of the TTC algorithm. Consider a bipartite directed graph over two sets of nodes that respectively represent the agents and the resources. Draw an arc from each object to its initial owner agent, and an arc from each agent to her most preferred object. Then, by performing the exchanges corresponding to the directed cycles in this directed graph and by repeating this process, we obtain an allocation that is, among other properties, Pareto-efficient. This simple and efficient algorithm is presented in the following example considering the same instance as in Example 1.4.

Example 1.5 *Let us consider the same instance as in Example 1.4 and construct the bipartite directed graph of the TTC algorithm (the dashed arcs stand for the current allocation of object, whereas the plain arcs connect the agents to their most preferred object).*

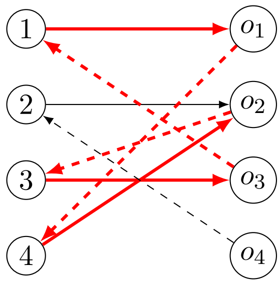

Let us consider the directed cycle in red. By performing the exchanges along this dicycle, we reach an allocation $\pi' = (o_1, o_4, o_3, o_2)$ *which is Pareto-efficient.*

The housing market framework as well as TTC-like algorithms are actually of wide interest for a very important application in the real world, namely the *kidney exchange problem* [Roth et al., 2004, Cechlárová and Lacko, 2012]. In such a context, additional constraints like the necessity of considering small cycles of exchanges are needed.

From the TTC seminal algorithm, centralized algorithms have been largely investigated, notably for the case of multiple resources per agent [Konishi et al., 2001, Todo et al., 2014, Sonoda et al., 2014, Fujita et al., 2015, Aziz et al., 2016a]. The main drawback of the centralized approach is the requirement for the agents to communicate a part of their preferences to an external authority they need to trust.

Alternatively, the agents may perform the reallocation by trading and negotiating among them, following a distributed process of exchanges. Less demanding in terms of communication, this approach has the benefit of the independence of the agents from an external authority but may lead to less satisfactory final allocations because of the agents' myopia. Indeed, the agents only perform local improvements without global vision of the whole allocation, contrary to an external central authority. Some works examine this distributed process of exchanges, in order to find conditions for realistic trades and analyze the quality of the possible outcomes. Widely studied in a general resource reallocation problem [Sandholm, 1998, Chevaleyre et al., 2017], this approach has been recently investigated in housing market [Damamme et al., 2015]. This latter work analyzes dynamics of exchanges where the agents, starting from the initial allocation, perform rational trades among limited set of agents until a stable allocation is reached. In such a stable allocation, no rational exchange is possible within a limited set of agents. The rationality assumption in the trades is essential in order to ensure the convergence of the dynamics. Damamme et al. [2015] notably shows that the dynamics associated with *swap deals*, i.e., trades between only two agents, eventually converges when the preferences are single-peaked.

An important dimension to study in resource allocation, as well as in other social choice problems, is the computational cost of the procedures. Indeed, real world applications require algorithms which solve every instance in a reasonable amount of time.

1.5 Computational complexity background

Since our goal is to analyze the problems within a computational social choice perspective, we investigate the computational cost of the procedures and protocols that we study. This section is devoted to the presentation of some basics about classical and parameterized complexity.

A problem $\mathcal{P}$ is defined according to a general description of its parameters, and a certain set of constraints characterizing the feasible solutions of the problem. An instance $\mathcal{I}$ of $\mathcal{P}$ is then a specification of particular values for the parameters of $\mathcal{P}$. An algorithm is a step by step method for solving a problem. An algorithm is said to be *deterministic* if its output for a given instance $\mathcal{I}$ is fixed, and *non-deterministic* if its output may be different for several runs of the algorithm on the same instance $\mathcal{I}$.

Problems can be classified into two categories: the *decision problems* and the *search problems*. They differ with respect to the type of answer that must be returned by an algorithm solving the problem. A search problem asks for a particular solution that satisfies the conditions of the problem, whereas a decision problem formulates a question, requiring a binary answer, *yes* or *no*. Among the search problems, one can distinguish *optimization problems* that aim to find a particular solution optimizing a given objective function f, defined over the feasible solutions of the problem. In order to analyze the complexity of search problems, it is more convenient to consider their decision version. The decision version of a search problem asks for the existence of a particular solution that satisfies the conditions stated in the search problem with, in addition, for the case of an optimization problem, a condition on the value that must take a feasible solution according to f. For instance, in the Minimum Vertex Cover optimization problem, the goal is to find a vertex cover $V' \subseteq \mathcal{V}$ in an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with minimum size. The decision version of this problem is Vertex Cover, asking whether there exists a vertex cover $V' \subseteq \mathcal{V}$ in an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ such that $|V'| \leq k$ for a given integer k given as input. In general, we will formalize the decision problems as follows.

Decision Problem:	
Instance:	Parameters that define an instance of the problem
Question:	Problem to solve that must be answered by *yes* or *no*

The analysis of the complexity of an algorithm is commonly asymptotic, according to the input size $|\mathcal{I}|$ (in our case the number of agents n and the number of alternatives m), for a given instance $\mathcal{I}$ of a problem $\mathcal{P}$. We follow a classical approach of *worst-case complexity.*

1.5.1 Classical complexity

For a more complete background about computational complexity theory, see for instance Arora and Barak [2009] and Garey and Johnson [1979].

In general, a *complexity class* defines a set of problems that can be solved by an algorithm using a given resource in terms of time or space for its execution. We will restrict ourselves to complexity classes focusing on time resources. Complexity classes are usually linked by inclusion relations, inducing a hierarchical structure over them. A complexity class C_2 is said to be higher than class C_1 in the hierarchy if it is known that $\mathsf{C}_1 \subseteq \mathsf{C}_2$. *Completeness* of a problem for a given complexity class means that the problem belongs to the hardest problems of this class. For two classes C_1 and C_2 such that $\mathsf{C}_1 \subseteq \mathsf{C}_2$, the problems that are C_1-complete are supposed to be easier to solve that problems C_2-complete, unless $\mathsf{C}_1 = \mathsf{C}_2$.

Completeness of a problem for a given complexity class is proved by showing both *hardness* and *membership* of the problem according to this complexity class. Proving that a problem $\mathcal{P}$ is C-hard, for a given complexity class C, can be done by showing that any instance $\mathcal{I}'$ of a problem $\mathcal{P}'$, that is complete for C, can be transformed into an instance $\mathcal{I}$ of $\mathcal{P}$, in a way that does not consume more resources than those required in the definition of C. This informally means that $\mathcal{P}$ is at least as hard as a problem that is complete for class C, proving a *lower bound* for the complexity of $\mathcal{P}$, which cannot be complete for a lower class in the hierarchy, unless the classes are equal. On the contrary, proving membership of a problem $\mathcal{P}$ in a given complexity class C can be done by exhibiting an algorithm solving $\mathcal{P}$ with the resources defined by C. This proves an *upper bound* for the complexity of $\mathcal{P}$, which cannot be complete for a higher class in the hierarchy, unless the classes are equal.

Traditionally, an algorithm is said to be computationally efficient if its running time is polynomial in the size of the instance, even in the worst case scenario. In other words, such an algorithm terminates within $\mathcal{O}(|\mathcal{I}|^{\mathcal{O}(1)})$ steps, for any instance $\mathcal{I}$ of the problem, where the number of steps refers to the number of basic operations. All the problems that are solvable in polynomial time belong to the complexity class P. A problem $\mathcal{P}$ is in P if there exists a deterministic algorithm solving any instance $\mathcal{I}$ of $\mathcal{P}$ in polynomial time.

1.5.1.a NP

For decision problems, there exists a hierarchy of complexity classes beyond P. We restrict our attention to the first level of this hierarchy, namely the complexity class NP. Class NP gathers all the problems that can be solved by a non-deterministic algorithm in polynomial time. Clearly, $\mathsf{P} \subseteq \mathsf{NP}$. Class NP is the best known and used complexity class because it suffices to prove a lower bound, i.e., hardness, to this class to prove that there is no efficient algorithm for the problem, unless $\mathsf{P} = \mathsf{NP}$. Knowing whether the inclusion of P in NP is strict is a very famous and challenging open problem, even if there is a consensus within the computer science community for conjecturing that $\mathsf{P} \neq \mathsf{NP}$.

Actually, the difference between P and NP relies on the difference between the construction and the verification of a solution. Indeed, P gathers all the problems for which the answer of the problem can be computed efficiently, whereas NP gathers all the problems for which a solution can be verified efficiently. A decision problem $\mathcal{P}$ is in NP if for every instance $\mathcal{I}$ of $\mathcal{P}$, this is possible to verify in polynomial time whether a given solution x of $\mathcal{I}$ is a *certificate* for asserting that $\mathcal{I}$ is a *yes*-instance of $\mathcal{P}$.

The first problem that has been proved to be NP-complete is the SATISFIABILITY problem, which became the reference problem for this class. SATISFIABILITY is defined as follows for formulas in conjunctive normal form (CNF-formulas) (recall that a CNF-formula is a propositional formula that is a conjunction of clauses, and a clause is a disjunction of literals).

SATISFIABILITY (SAT):	
Instance:	A set X of variables, a CNF-propositional formula φ over X
Question:	Does there exist a truth assignment of the variables in X such that φ is satisfied?

Theorem 1.1 (Cook-Levin (1971)) *SAT is* NP*-complete.*

After the establishment of the NP-completeness of SATISFIABILITY, Karp [1972] provided a list of 21 NP-complete problems with proofs based on *polynomial reductions*. Some years later, Garey and Johnson [1979] presented an extensive list of NP-complete problems, and now the size of this class is still growing.

The proof of NP-completeness relies on two facts: membership to NP and NP-hardness. A problem $\mathcal{P}$ is said to be NP-hard if there exists a polynomial reduction from a known NP-complete problem to $\mathcal{P}$.

Definition 1.31 (Polynomial reduction) *A polynomial reduction from problem $\mathcal{P}$ to problem $\mathcal{P}'$ is a polynomial algorithm that transforms any instance $\mathcal{I}$ of $\mathcal{P}$ to an instance $\mathcal{I}'$ of $\mathcal{P}'$ such that $\mathcal{I}$ is a* yes*-instance of $\mathcal{P}$ if and only if $\mathcal{I}'$ is a* yes*-instance of $\mathcal{P}'$, and $|\mathcal{I}'| = \mathcal{O}(|\mathcal{I}|^{\mathcal{O}(1)})$.*

Informally, a problem is NP-hard if it is at least as hard as an NP-complete problem. This proves that the exact complexity class of the problem is either NP or a higher class in the hierarchy, and thus shows a lower bound of hardness. Conversely, the membership to NP shows an upper bound of hardness: that means that the exact complexity class of the problem is either NP or a lower class in the hierarchy. Consequently, a problem is NP-complete, that is the exact complexity class of the problem is NP, if it is in NP and NP-hard.

In this document, we will use polynomial reductions from some variants of SATISFIABILITY, that are known to be NP-complete. These variants are based on particular CNF-formulas. The best known is 3-SAT, that belongs to the original list of the 21 NP-complete problems presented by Karp [1972].

3-SAT:	
Instance:	A set X of variables, a CNF-formula φ over X where each clause contains exactly three literals
Question:	Does there exist a truth assignment of the variables in X such that φ is satisfied?

Theorem 1.2 (Karp [1972]) 3-SAT *is* NP-*complete.*

Interestingly, the 3-SAT problem delimits the tractability frontier within the satisfiability problems conditioned by the size of the clauses. Indeed, whereas 2-SAT, the restriction to clauses of size two, is solvable in polynomial-time [Cook, 1971b,a, Even et al., 1976], and even in linear time [Aspvall et al., 1979], considering only three literals per clause is sufficient to get intractability.

We will also consider more restrictive CNF-formulas where the occurrences of the variables are limited as well.

2P1N-SAT:	
Instance:	A set X of variables, a CNF-formula φ over X where each variable in X occurs exactly three times: twice as a positive literal and once as a negative literal
Question:	Does there exist a truth assignment of the variables in X such that φ is satisfied?

Theorem 1.3 (Yoshinaka [2005]) 2P1N-SAT *is* NP-*complete.*

Note that in the reduction proving the NP-completeness of 2P1N-SAT [Yoshinaka, 2005], the clauses of the constructed formula can be reorganized such that the negative occurrence of each variable always occurs at first, or at second or at third position among the occurrences of the variable in the formula. Therefore, considering instances of 2P1N-SAT with such properties also leads to an NP-complete problem.

We also consider another NP-complete variant of the Satisfiability problem, where the occurrences of the variables are limited as well as the size of the clauses.

(3, B2)-SAT:	
Instance:	A set X of variables, a CNF-formula φ over X where each clause contains exactly three literals and each variable in X occurs exactly four times: twice as a positive literal and twice as a negative literal
Question:	Does there exist a truth assignment of the variables in X such that φ is satisfied?

Theorem 1.4 (Berman et al. [2004]) (3, B2)-SAT *is* NP-*complete.*

Besides the satisfiability problems, NP-complete problems based on graphs are also widely used in polynomial reductions. Among them, one can cite for instance the Independent Set problem.

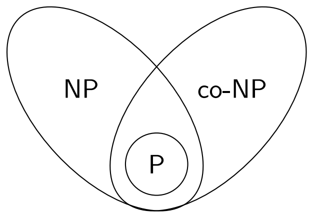

Figure 1.3: Relation between complexity classes P, NP and co-NP

INDEPENDENT SET:	
Instance:	Undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, integer k
Question:	Is there an independent set $I \subseteq \mathcal{V}$ in $\mathcal{G}$ such that $\|I\|= k$?

Theorem 1.5 (Garey and Johnson [1979]) INDEPENDENT SET *is* NP-*complete.*

1.5.1.b co-NP

Let us define the problem $\overline{\mathcal{P}}$ as the complement of decision problem $\mathcal{P}$, i.e., for any instance $\mathcal{I}$ of $\mathcal{P}$, $\mathcal{I}$ is a *yes*-instance of $\mathcal{P}$ if and only if $\mathcal{I}$ is a *no*-instance of $\overline{\mathcal{P}}$. If we define as co-C the set of decision problems whose complement is in C, then co-P = P, because P is closed under complementation. Indeed, if it is possible to answer in polynomial time to any instance of a problem $\mathcal{P}$, then it is also possible to answer in polynomial time to the dual question.

Let us introduce the class co-NP containing the decision problems that are complementary to those belonging to NP. Observe that the classes NP and co-NP are not complementary because they have P as a non-empty intersection (see Figure 1.3). A problem $\mathcal{P}$ is in co-NP if, for any instance $\mathcal{I}$ of $\mathcal{P}$, it is possible to verify in polynomial time whether a given solution x of $\mathcal{I}$ is a certificate for asserting that $\mathcal{I}$ is a *no*-instance of $\mathcal{P}$.

Following the same principle as for NP, a problem is co-NP-complete if it is in co-NP and co-NP-hard. The co-NP-hardness of a problem is proved via a polynomial reduction from a co-NP-complete problem. The reference problem for co-NP-completeness is UNSATISFIABILITY, the complementary problem of SATISFIABILITY.

UNSATISFIABILITY:	
Instance:	A set X of variables, a propositional formula φ over X
Question:	Is φ unsatisfiable for any truth assignment of the variables in X?

Intuitively, problems in NP refer to a unique existential quantifier whereas problems in co-NP refer to a unique universal quantifier. For example, in a *yes*-instance of SATISFIABILITY, there *exists* a truth assignment of the variables satisfying the formula, whereas in a *yes*-instance of UNSATISFIABILITY, *for all* truth assignments of the variables, the formula is not satisfied. Note that the question whether NP = co-NP is still open.

1.5.1.c Approximation

Let us briefly talk about *approximation*, which is related to optimization problems. When the decision version of an optimization problem is proved to be computationally hard, an alternative option is to design an efficient algorithm that finds a feasible solution that is close to the optimum with respect to objective function f.

A ρ-approximation algorithm, for a given optimization problem $\mathcal{P}$ with objective function f, outputs in polynomial time, in any instance $\mathcal{I}$ of $\mathcal{P}$, a feasible solution whose value with respect to f is guaranteed to be within a factor of ρ of the optimum of f in $\mathcal{I}$. Note that the approximation factor ρ, i.e., the performance guarantee, may depend on the input size. If $\mathcal{P}$ is a maximization problem, then $\rho < 1$, otherwise $\rho > 1$.

1.5.2 Parameterized complexity

The aim of parameterized complexity is to refine the classical complexity analysis of NP-hard problems by determining the parameters that make the problems computationally hard (for some basic notions about parameterized complexity, see for instance Flum and Grohe [2006] and Downey and Fellows [2013]). In fact, the goal is to find a parameter k of the problem, for instance the size of the solution, which is preferably small in practice, such that all the computational hardness of the problem is contained in k. This idea is formalized via the complexity class FPT, analogous to P in the parameterized world, i.e., the problems in FPT are solvable efficiently by assuming that k remains small.

Definition 1.32 (FPT) *A problem $\mathcal{P}$ is fixed-parameter tractable (*FPT*) with respect to parameter k if every instance $\mathcal{I}$ of $\mathcal{P}$ is solvable in time $\mathcal{O}(g(k).|\mathcal{I}|^{\mathcal{O}(1)})$ for a computable function g.*

Note that g must be computable, in the sense that there exists an algorithm guaranteed to terminate that can compute the output of g. Assuming the problem we are trying to solve is NP-hard (if the problem is solvable in polynomial time, then there is no need to use a parameterized complexity approach), the algorithm computing g cannot be executed in polynomial time, unless P = NP. However, if parameter k is small compared to the size of the instance $|\mathcal{I}|$, we achieve that the blow-up is limited to the small value k.

Another important parameterized complexity class for delimiting the tractability frontier of the parameterized problems is XP. Indeed, if k is a constant, then a problem in XP is solvable in polynomial time.

Definition 1.33 (XP) *A problem $\mathcal{P}$ is in* XP *with respect to parameter k if every instance $\mathcal{I}$ of $\mathcal{P}$ is solvable in time $\mathcal{O}(|\mathcal{I}|^{g(k)})$ for a computable function g.*

Clearly, FPT $\subseteq$ XP. Another noticeable class that contains FPT is para-NP. Intuitively, an NP-complete problem that remains hard whatever the value of the parameter k is in para-NP.

Similarly as for classical complexity, there exist hierarchies of parameterized complexity classes for decisions problems beyond FPT, which are linked by inclusion relations. Actually, there are two different hierarchies of parameterized complexity classes which include FPT and are included in XP: the W-hierarchy and the A-hierarchy. The two hierarchies are linked via inclusion relations, more precisely at each level $t \geq 1$, $W[t] \subseteq A[t]$, and $W[1] = A[1]$, as we can observe in Figure 1.4.

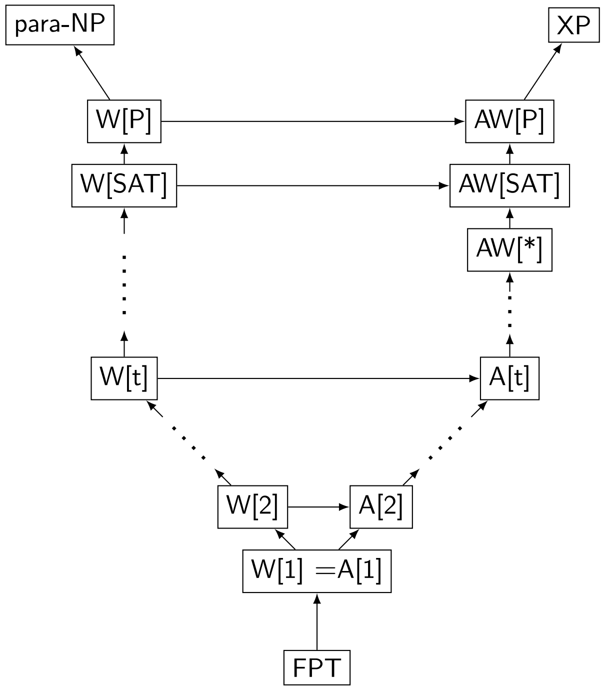

Figure 1.4: Parameterized complexity classes beyond FPT: A-hierarchy and W-hierarchy. The arrows refer to the inclusion relation (adaptation of a Flum and Grohe [2006]'s figure)

Like in classical complexity, hardness for a given parameterized complexity class is proved via reductions from problems known to be complete for the given class, but this time they are FPT-reductions. An FPT-reduction from a problem $\mathcal{P}$ of parameter k to a problem $\mathcal{P}'$ of parameter k' is an FPT algorithm according to parameter k that transforms any instance $\mathcal{I}$ of $\mathcal{P}$ to an instance $\mathcal{I}'$ of $\mathcal{P}'$ such that $\mathcal{I}$ is a *yes*-instance of $\mathcal{P}$ with parameter k if and only if $\mathcal{I}'$ is a *yes*-instance of $\mathcal{P}'$ with parameter $k' = g(k)$ for a computable function g, and $|\mathcal{I}'| = \mathcal{O}(|I|^{\mathcal{O}(1)})$. Note that membership of problem $\mathcal{P}$ to a given parameterized complexity class can also be proved thanks to FPT-reductions, by using the reduction in the reverse direction, that is there exists an FPT-reduction from $\mathcal{P}$ to problem $\mathcal{P}'$ known to be complete for this class.

1.5.2.a The W-hierarchy

The reference problem of the W-hierarchy is about weighted satisfiability in boolean circuits where the level in the hierarchy refers, by simplifying, to the complexity of the gates and paths in the circuit. Parameterized classes FPT and W[1] can be thought of as corresponding to P and NP in the parameterized world. Therefore, proving that a problem is W[1]-hard excludes, under classical complexity assumptions, that there exists an FPT algorithm solving this problem.

In this document, we will use some parameterized problems, known to be complete within the W-hierarchy. We notably use parameterized problems defined on graphs, such as the following CLIQUE problem, known to be W[1]-complete, as well as a colored parameterized version of the independent set problem, namely MULICOLORED INDEPENDENT SET.

CLIQUE:	
Instance:	Undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, parameter k
Question:	Is there a clique $K \subseteq \mathcal{V}$ in $\mathcal{G}$ such that $\|K\|= k$?

Theorem 1.6 (Downey and Fellows [1995a]) CLIQUE *is* W[1]*-complete.*

MULTICOLORED INDEPENDENT SET:	
Instance:	Undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, parameter k, partition $\mathcal{V} = \mathcal{V}_1 \cup \ldots \cup \mathcal{V}_k$
Question:	Is there an independent set $I \subseteq \mathcal{V}$ in graph $\mathcal{G}$ of size $\|I\|= k$ such that $\|I \cap \mathcal{V}_i\|= 1$ for all $i \in [k]$?

Theorem 1.7 (Fellows et al. [2009]) MULTICOLORED INDEPENDENT SET *is* W[1]*-complete.*

We will also use a typical problem of the W-hierarchy, which is the parameterized version of a variant of the SATISFIABILITY problem, known to be complete for a higher class than W[1] in the W-hierarchy, unless the whole hierarchy collapses (see Figure 1.4). Recall that a propositional formula is monotone if it only contains Boolean operators $\vee$ and $\wedge$, but no negation ($\neg$).

MONOTONE WEIGHTED SATISFIABILITY:	
Instance:	Set X of variables, monotone propositional formula φ over X, parameter k
Question:	Does there exist a truth assignment of the variables in X of weight k, i.e., exactly k variables are assigned to true, such that φ is satisfied?

Theorem 1.8 (Abrahamson et al. [1995]) MONOTONE WEIGHTED SATISFIABILITY *is* W[SAT]*-complete.*

1.5.2.b The A-hierarchy and model checking

The A-hierarchy is a hierarchy of parameterized complexity classes that include FPT and are included in XP, which is also linked to the W-hierarchy by inclusion relations (see Figure 1.4). The parameterized complexity classes of the A-hierarchy are closely related to first-order logic since they are defined according to model checking for first-order formulas. We will use in this document some reductions based on model checking.

Before formalizing the model checking problem, let us briefly state some first-order logic notions. We define a vocabulary τ as a finite set of relation symbols, each relation $R \in \tau$ having an associated arity $ar(R)$. A finite structure $\mathcal{A}$ of vocabulary τ (or τ-structure) consists of a finite set U of elements, called the universe, and an interpretation $R^{\mathcal{A}} \subseteq U^{ar(R)}$ of each relation symbol R in τ over U, with corresponding arity.

We assume a countably infinite set of variables. Atomic formulas over vocabulary τ are of the form $x_1 = x_2$ or $R(x_1, \dots, x_p)$ where $R \in \tau$, $ar(R) = p$, and $x_1, \dots, x_p$ are variables. The class FO of all first-order formulas over τ consists of formulas that are constructed from atomic formulas over τ using standard Boolean connectives $\neg, \wedge, \vee$, as well as quantifiers $\exists, \forall$ followed by a variable. Let φ be an FO formula. The variables of φ that are not in the scope of a quantifier are called *free variables*. Let $\varphi(\mathcal{A})$ denote the set of all assignments of elements of U to the free variables of φ such that φ is satisfied. We say that $\mathcal{A}$ is a *model* of φ, or $\mathcal{A}$ satisfies φ, if $\varphi(\mathcal{A})$ is not empty.

The class Σ_ℓ contains all first-order formulas of the form $\exists x_{11}, \dots, \exists x_{1k_1} \forall x_{21}, \dots, \forall x_{2k_2} \dots \mathcal{Q} x_{\ell 1}, \dots, \mathcal{Q} x_{\ell k_\ell} \varphi$, where φ is a quantifier free FO-formula, and $\mathcal{Q} = \exists$ if ℓ is odd and $\mathcal{Q} = \forall$ otherwise. In particular, the class Σ_1 contains all FO formulas of the form $\exists x_1, \dots, \exists x_k \varphi$, and Σ_2 contains all FO formulas of the form $\exists x_1, \dots, \exists x_k \forall y_1, \dots, \forall y_k \varphi$, where φ is a quantifier free FO formula.

Let Φ be a class of formulas. The model checking problem takes as input a finite structure $\mathcal{A}$ and a formula $\varphi \in \Phi$, and asks whether $\mathcal{A}$ is a model of φ, i.e., $\varphi(\mathcal{A}) \neq \emptyset$. A natural parameterized version of this problem considers as a parameter the size of (a reasonable encoding of) φ, denoted by $|\varphi|$.

Model Checking (MC(Φ)):	
Instance:	Finite structure $\mathcal{A}$, formula $\varphi \in \Phi$, parameter $k = |\varphi|$
Question:	Does $\mathcal{A}$ satisfy φ, i.e., is $\varphi(\mathcal{A}) \neq \emptyset$?

The A-hierarchy relies on parameterized model checking problems where the level in the hierarchy refers to the number of alternations between existential and universal quantifiers in the formula. Indeed, the reference problem which is complete for class A[t] is MC(Σ_t). The next two theorems establish the completeness of the model checking problem restricted to formulas in Σ_1 and in Σ_2, for classes A[1] and A[2], respectively, the two lower parameterized complexity classes in the A-hierarchy (see Figure 1.4).

Theorem 1.9 (Flum and Grohe [2006]) *The* Model Checking *problem for the existential fragment of first-order logic,* MC(Σ_1)*, is* A[1]*-complete (=* W[1]*-complete).*

Theorem 1.10 (Flum and Grohe [2006]) *The* Model Checking *problem for the second-level alternation of first-order logic,* MC(Σ_2)*, is* A[2]*-complete.*

Note that the parameterized complexity classes W[1] and A[1] are equivalent. This is explained by the fact that Σ_1 is built on simple existential first-order formulas.

1.6 Conclusion

In this chapter, we have presented the notation that will be useful for the rest of the document, as well as key concepts on the two problems under study, strategic voting and house allocation. Moreover, we have provided some notions on complexity theory, helpful for understanding the proofs of the thesis. The reader is invited to refer to this chapter all along the document for more details about some definitions or state-of-the-art concepts.

The contributions of the thesis are presented in the next two parts. The first one is dedicated to the study of social choice problems where a social network models the possibility of collaboration among the agents, whereas the second one focuses on problems where the information possessed by the agents is conditioned by the social network.

Part I

The Social Network as a Collaboration Tool

Introductory comments

Let us first investigate how the social network can model the possibility of *collaboration* among the agents. Collaboration within a group of agents is an interaction where all the members of the group, treated on an equal footing, work together by combining their capabilities, within a free association, in order to achieve a common objective, or compatible goals. By definition, this interaction requires a reciprocity relation among the agents. Note that collaboration is slightly different from *cooperation*, because cooperation can induce participation to a common project within a hierarchical structure where the agents can perform complementary tasks.

Collaboration must be mutually profitable for all the involved agents, who must actively participate to the common action. This necessarily implies a symmetry in the accessibility relation which represents the possibility of collaboration among the agents. Consequently, as we model the possibility of collaboration by a social network, we assume for this part that the social network is represented by an undirected graph over the agents.

We focus on two problems of social choice, namely strategic voting and house allocation. We study the former, in Chapter 2, under the prism of coalitional manipulation, where the possible coalitions are fully connected components of the social network. The members of a coalition collaborate by commonly elaborating a strategy in order to achieve a common goal: the election of a candidate that they prefer to the current winner. We examine house allocation, in Chapter 3, via the specific framework of housing market, where exchanges of items occur between connecting agents in the network. In such a setting, two agents collaborate by participating to a trade that is profitable for both agents.

Chapter 2

Coalitional Manipulation in Iterative Voting

Abstract

Strategic voting is investigated via a voting game where the strategic deviations are performed by coalitions of voters. We examine coalitional equilibria that are immune to manipulation performed by realistic coalitions, based on the social network. Concretely, the coalitions of voters are given by the cliques of the network, which form groups of agents that can collaborate and agree on common strategies. We assume that the agents are not fully selfish as they have consideration for their relatives. The corresponding solution concept, introduced by Hoefer et al. [2011], is called a *considerate equilibrium*. We study its existence in strategic voting games, and the ability of the agents to converge to such an equilibrium using well-known voting rules: Plurality, Veto, Plurality with run-off, Borda, k-approval, STV, Maximin and Copeland.

Résumé

On s'intéresse dans ce chapitre au vote stratégique, que l'on étudie sous l'aspect d'un jeu de vote dynamique, où des coalitions d'agents peuvent s'entendre pour dévier vers de nouveaux bulletins. Notre travail repose sur l'étude d'états stables par rapport à des déviations faites par des coalitions réalistes, basées sur le réseau social. Plus précisément, les coalitions possibles de votants sont données par les cliques du réseau social, représentant des groupes d'agents pouvant communiquer et établir des stratégies communes. De plus, les agents sont supposés être altruistes et avoir de la considération pour leurs proches dans le réseau. La notion d'équilibre correspondante est nommée *équilibre de considération* [Hoefer et al., 2011]. On se propose d'étudier son existence dans des jeux de vote, ainsi que la capacité du jeu à atteindre un tel équilibre pour différentes règles de vote classiques : Pluralité, Véto, Pluralité à deux tours, Borda, k-approbation, STV, Maximin et Copeland.

This chapter is an extension of [Gourvès, Lesca, and Wilczynski, 2016].

Contents

2.1 Introduction

Manipulation occurs in voting, especially in political elections, when some electors may behave strategically by misreporting their preferences [Blais, 2004, Alvarez et al., 2006, Daoust, 2015]. In this perspective, considering an election as a strategic game where voters need to identify the best ballot that they should submit to the electoral system is meaningful.

2.1.1 Deviations, coalitions and social network

Manipulation from single agents, where stable outcomes of the voting game are Nash equilibria, has been widely studied [Feddersen et al., 1990, Desmedt and Elkind, 2010, Meir et al., 2010]. The main benefit of focusing on Nash equilibria is the guarantee of its existence for numerous voting systems. Typically, voting profiles where all the agents play the same ballot are immune to unilateral deviations for many reasonable voting rules. As a consequence, the number of Nash equilibria in a given instance of a voting game can be very large [Dhillon and Lockwood, 2004], raising the question of the relevance of Nash equilibrium, as a solution concept for capturing a plausible outcome of the game. Moreover, especially in the context of voting, the significance of only considering unilateral manipulation can be questionable. Indeed, in real elections, coalitions can form to agree on some voting strategies. For instance, it is very frequent that some political parties give voting instructions to their members. A drawback of the Nash equilibrium is its weakness against coalitional deviations, as we can observe in the following example.

Example 2.1 *Take as an example a group of nine friends who plan to go to a restaurant. They have four different options: an Italian restaurant (I), a Japanese restaurant (J), a Lebanese restaurant (L) or a Portuguese restaurant (P). They have the following preferences over the different restaurants:*

$$\begin{aligned} \textit{Allan:} &\quad L \succ I \succ J \succ P \\ \textit{Bob, Claire:} &\quad J \succ L \succ P \succ I \\ \textit{Damian, Elise:} &\quad P \succ L \succ J \succ I \\ \textit{Flora, Gamal, Henri, Imen:} &\quad I \succ J \succ L \succ P \end{aligned}$$

They organize a vote via the Plurality rule where they all give their preferred alternative. Observe that the winning option is the Italian restaurant with four points while the other alternatives do not have more than two points. This state, where everybody tells her real preferences is a Nash equilibrium because no single individual can change the outcome or has incentive to do it. However, if the group formed by Bob, Claire, Damian and Elise agrees on a deviation to the vote approving the Lebanese restaurant, which is the second best alternative of every member of the group, then this option will win. Therefore, the state where everybody says the truth is not immune to coalitional manipulation. Since these four individuals are friends and have close preferences, they may be able to make contact and establish a common strategy, but the Nash equilibrium concept cannot prevent such a deviation.

As a consequence, it appears interesting to consider coalitional manipulation in voting games. In the context of voting, the characterization of strong equilibria – a state immune to strict improvement deviations performed by any possible coalition – has for instance been investigated [Sertel and Sanver, 2004, Messner and Polborn, 2007, Peleg and Peters, 2010]. Despite its conceptual attractiveness, the main disadvantage of the strong equilibrium is that it rarely exists. But actually, considering that any coalition of voters can form may be highly pessimistic, because it supposes that all the agents in any subgroup are able to coordinate their moves in order to manipulate. Such a coordination requires an important level of communication and trust for the manipulators. In practice, one can reasonably exclude some coalitions from the definition of an equilibrium that is supposed to be stable against group deviations. For example, this is the case of the *partition equilibrium* introduced by Feldman and Tennenholtz [2009] where only the coalitions belonging to a prescribed partition of the voter set have the ability to deviate.

More generally, one can determine which coalitions can form by using social networks, graph structures representing social relations among agents, as depicted in the following example.

Example 2.2 *Take the example of a computer science laboratory that aims to recruit a new professor. Five candidates apply for the position: Amira, Boris, Clara, Dimitri and Elias. Three teams exist in the lab, with different interests, namely the logic team, the complexity team, and the machine learning team. Let us focus on a particular researcher, Robert, belonging to the recruitment committee and who works both in the logic and the complexity teams. Note that the three teams of the lab can be represented as three cliques in a social network where Robert both belongs to the cliques representing the logic team and the complexity team.*

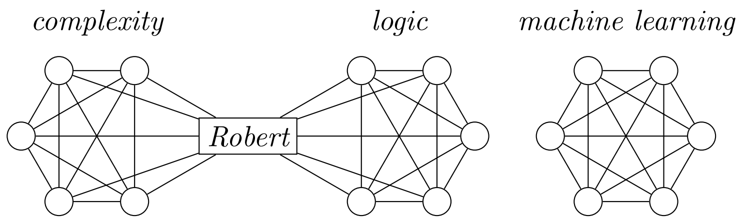

The preferences of Robert are the following: Boris $\succ$ Amira $\succ$ Dimitri $\succ$ Clara $\succ$ Elias. Robert learns that his first choice, Boris, has no chance to win because nobody supports him. But he learns that Elias has a large support within the machine learning group. This candidate is not appreciated by Robert as well as the whole logic team. Consequently, the members of the logic team, which is a cohesive group, agree to vote for Dimitri because he works on related topics. The complexity team does not succeed to find a consensus and thus Robert decides to participate to the coalition formed by the logic team, which decides to vote in favor of Dimitri that the whole team prefers to Elias.

Since elaboration of a common strategy within a group of agents necessitates communication, agreement and coordination among all its members, cliques of the social network, which are fully connected components of the graph, are good candidates for representing plausible coalitions for manipulation, as underlined in Example 2.2. By assuming that the agents are embedded in a social network and that their relations are characterized by their links in the network, one can further suppose that a voter is tied by social relations which force her to have consideration for other agents. Exploiting the social relations which bind the members of a group has received much attention [Jackson, 2008]. In a more altruistic perspective, one can assume that a voter is not only guided by her own preferences over the candidates but she pursues the goal of optimizing the welfare of the communities where she belongs. A natural solution concept which combines notions of altruism and cohesive groups based on a social network is the *considerate equilibrium*. Introduced by Hoefer et al. [2011], a considerate equilibrium is a state robust to deviations by coalitions of agents who form a clique in the social network, but a coalition only deviates when it is not harmful for her relatives (neighbors in the graph). As for the partition equilibrium that it extends, the considerate equilibrium has only been studied, as far as we know, for a special case of congestion game.

2.1.2 Related work

Coalitional manipulation in voting has been widely investigated via the complexity of computing such a manipulation [Conitzer et al., 2007, Xia et al., 2009, Zuckerman et al., 2009]. Nevertheless, some works has examined the game-theoretical properties of coalitional manipulation in strategic voting, especially through a cooperative game [Bachrach et al., 2011]. The cooperative game model notably enables to discuss the conditions of coalition formation for manipulating in voting. These questions have also been addressed via another model by Slinko and White [2008]. However, to the best of our knowledge, the social context in which the voters are embedded has not been explored so far for coalitional manipulation in voting.

Nevertheless, the integration of a social network in voting games has been recently proposed in iterative voting in a context of uncertainty where the links of the graph provide some information to the agents about the current vote of their relatives [Sina et al., 2015, Tsang and Larson, 2016, Tsang et al., 2018]. However, there is no notion of possible collaboration given by the graph.

Moreover, so far, iterative voting has been especially analyzed via unilateral deviations. Consequently, we aim to investigate coalitional manipulation in iterative voting where the coalitions are not arbitrary but given, more realistically, by the social network.

2.1.3 Contributions and organization

The global idea is to use the solution concept of *considerate equilibrium* in order to refine the game-theoretical analysis of the voting game. Concretely, we explore the existence of a considerate equilibrium in strategic voting games and the ability of the dynamics of the game to converge to such an equilibrium for different voting rules. Our main contributions are existence proofs of a considerate equilibrium in a voting game under well-established voting rules, namely Plurality, Veto, Plurality with run-off, STV and Maximin. We also investigate the possibility for the voters to reach a considerate equilibrium in a natural iterative process. In this respect, our results are rather negative because convergence to a partition equilibrium, or convergence to a Nash equilibrium, which are less demanding goals than convergence to a considerate equilibrium, fail.

We first introduce in Section 2.2 the notion of coalitional deviation and considerate equilibrium. Then, we study the existence of a considerate equilibrium for different voting rules in Section 2.3. Section 2.4.1 is devoted to the special case of the Veto rule, which always admits, in addition to a considerate equilibrium, a strong equilibrium. The existence of a strong equilibrium under Veto enables to design a new quality measure for evaluating voting rules. Then, we investigate the ability of the dynamics of the voting game to converge to a considerate equilibrium in Section 2.5. We present some experiments in Section 2.6 on the quality of the considerate equilibria that are reached by the dynamics. Finally, we conclude in Section 2.7 with a discussion on the global results and the impact of consideration within deviating moves.

2.2 Coalitional deviations in strategic voting games

Let us consider a voting game $\langle N, M, \succsim, \mathcal{F}_\rhd \rangle$ with pure strategies, as described in Section 1.3.3, where voters can deviate following an iterative voting process. The voters in $N = \{1, \ldots, n\}$ have preferences, which are not necessarily strict, over the m candidates in $M = \{a, b, c, \ldots\}$. The winner of the election is determined by voting rule $\mathcal{F}$, where ties are broken thanks to linear order $\rhd$ over the candidates, assumed to be the alphabetical order, i.e., $a \rhd b \rhd \ldots$.

We focus on coalitional manipulation: a coalition of voters can agree to simultaneously deviate together to a new strategy, i.e., to a new ballot in the context of voting. Note that the new ballot is not necessarily the same for all the members of the deviating coalition.

2.2.1 Classical solution concepts

First of all, let us define how individual preferences can be extended to collective preferences for each group $C \subseteq N$ of voters. For any x and $y \in M$, we use the following notations:

- $x >_C y \Leftrightarrow [\forall i \in C,\ x \succ_i y]$,
- $x \geq_C y \Leftrightarrow [\forall i \in C,\ x \succsim_i y]$ and $[\exists j \in C,\ x \succ_j y]$,
- $x \geqq_C y \Leftrightarrow [\forall i \in C,\ x \succsim_i y]$,
- $x \sim_C y \Leftrightarrow [\forall i \in C,\ x \sim_i y]$.

Let us provide a "local" Pareto-efficiency definition, which corresponds to the restriction of Pareto-efficiency (see Definition 1.14) to a coalition of voters.

Definition 2.1 *An alternative $x \in M$ is said to be* undominated *(respectively,* strictly undominated*) over coalition $C \subseteq N$ if there is no alternative $y \neq x$ such that $y \geq_C x$ (respectively, $y \geqq_C x$).*

Clearly, a Pareto-efficient alternative is undominated over N. An undominated alternative always exists for any non-empty coalition $C \subseteq N$ but strictly undominated alternatives may be absent. The two notions are illustrated in Example 2.3.

Example 2.3 *Let us consider an instance with four voters and four candidates, where $N = \{1, 2, 3, 4\}$ and $M = \{a, b, c, d\}$. The preference profile is the following.*

$$\begin{array}{ll} 1: & (a \sim b) \succ c \succ d \\ 2: & (a \sim b) \succ d \succ c \\ 3: & d \succ (a \sim b) \succ c \\ 4: & c \succ b \succ (a \sim d) \end{array}$$

Let us focus on the dominance properties of the candidates with respect to coalition $C = \{1, 2, 3\}$. Candidate d is strictly undominated over C because it is the best alternative of agent 3. Alternatives a and b are undominated over C because there is no candidate that is at least as preferred as them for all the members of C and strictly preferred to them for at least one member of C. However, a and b are not strictly undominated because all the members of C are indifferent between them. Finally, candidate c is dominated over C by alternatives a and b.

However, candidate c is strictly undominated over N, and thus Pareto-efficient, because c is the best candidate of agent 4. Candidate d is trivially Pareto-efficient since it is strictly undominated over a coalition of agents. Finally, the situation of a and b changes according to the whole set of voters. Indeed, b is strictly undominated over N, and thus Pareto-efficient, because agent 4 prefers b to a and b is undominated over C with $a \sim_C b$. For the same reasons a is dominated by b over N and thus a is not Pareto-efficient.

Let $\mathcal{C}$ be a non-empty family of coalitions of N, i.e., $\mathcal{C} \subseteq 2^N \setminus \{\emptyset\}$. This set represents the collection of coalitions that can possibly deviate. Different solution concepts can be designed, based on $\mathcal{C}$, but also on the type of *deviation* that is performed by the coalition.

Definition 2.2 (Improving move (IM)) *For coalition $C \in \mathcal{C}$, an improving move from state σ is a joint strategy $\sigma'_C \in \mathcal{B_F}^{|C|}$ such that $\mathcal{F}(\sigma'_C, \sigma_{-C}) >_C \mathcal{F}(\sigma)$.*

In an improving move, each member of the coalition must be better off by the deviation, that is each member must strictly prefer the outcome given by voting rule $\mathcal{F}$ of the new state. This type of deviation can be relaxed by considering *weak improving moves.*

Definition 2.3 (Weak improving move (WIM)) *For coalition $C \in \mathcal{C}$, a weak improving move from state σ is a joint strategy $\sigma'_C \in \mathcal{B_F}^{|C|}$ such that $\mathcal{F}(\sigma'_C, \sigma_{-C}) \geq_C \mathcal{F}(\sigma)$.*

Weak improving moves are appealing because they allow the participation of some voters who do not benefit, but who are not harmed by the deviation.

Note that an ordered pair $(\mathcal{C}, \mu)$ for $\mu \in \{\text{IM,WIM}\}$ can be used to define some notions of equilibrium.

Definition 2.4 *State σ is a $(\mathcal{C}, \mu)$-equilibrium if and only if no coalition $C \in \mathcal{C}$ can deviate from σ with a move of type μ.*

With this general definition, one can redefine classical well-known solution concepts. In this idea, a *Nash equilibrium* (NE) (Definition 1.23) is a ($\mathcal{C}$,IM)-equilibrium where $\mathcal{C}$ is the set of all singletons of agents, i.e., $\mathcal{C} = \{\{i\} : i \in N\}$. Similarly a *strong equilibrium* (SE) (Definition 1.24) is a ($\mathcal{C}$,IM)-equilibrium where $\mathcal{C} = 2^N \setminus \{\emptyset\}$. In this context, one can also evoke a refinement of the strong equilibrium that is called a *super strong equilibrium.* A super strong equilibrium (SSE) [Voorneveld, 1999, Feldman and

Tennenholtz, 2009] is a ($\mathcal{C}$,WIM)-equilibrium where $\mathcal{C} = 2^N \setminus \{\emptyset\}$. These three solutions concepts are linked as follows for a given state $\sigma \in \mathcal{B_F}^n$:

$$[\sigma \text{ is a } SSE] \Rightarrow [\sigma \text{ is a } SE] \Rightarrow [\sigma \text{ is a } NE] \tag{2.1}$$

2.2.2 Considerate equilibrium

Let us consider a special collection of coalitions $\mathcal{C}$ that takes into account the social context in which the voters are embedded. Given a social network of voters, represented by an undirected graph $G = (N, E)$, an edge $\{u, v\} \in E$ indicates that voters u and v are related, e.g., they have the possibility to communicate, allowing them to collaborate together within a deviating coalition. It makes sense to define the cliques of graph G as the possible coalitions. In that case $\mathcal{C} = \{C \in 2^N \setminus \{\emptyset\} : \forall i, j \in C,\ \{i, j\} \in E\}$ and each coalition $C \in \mathcal{C}$ has a set of neighbors $\mathcal{N}(C) = \{i \in N \setminus C : \exists j \in C \text{ such that } \{i, j\} \in E\}$.

Therefore, we consider from now on an instance of a *linked voting game* as a tuple $\langle N, M, \succsim, \mathcal{F}_\triangleright, G\rangle$, which is an instance of the classical voting game where agents are tied by their social relations in social network G.

Interestingly, the fact that the voters are related implies that each individual is not only guided by her own preferences, but she can also care about how deviating can negatively impact her relatives. In order to take into account the social context and the fact that a player can have consideration for other voters, more precisely her neighbors in the graph, an appropriate notion of deviation, called *considerate improving move*, can be introduced [Hoefer et al., 2011].

Definition 2.5 (Considerate improving move (CIM)) *For coalition $C \in \mathcal{C}$, a considerate improving move σ'_C from state σ is a weak improving move from σ for C where, in addition, $\mathcal{F}(\sigma'_C, \sigma_{-C}) \geqq_{\mathcal{N}(C)} \mathcal{F}(\sigma)$.*

In a considerate improving move performed by coalition C, at least one player in C is better off and no player of $C \cup \mathcal{N}(C)$ can be worse off.

Consequently, the pair ($\mathcal{C}$,CIM), where $\mathcal{C}$ contains all the cliques of the social network $G = (N, E)$, leads to a new type of equilibrium called *considerate equilibrium* [Hoefer et al., 2011].

Definition 2.6 (Considerate equilibrium) *A state σ is a considerate equilibrium if and only if there is no coalition in $\mathcal{C} := \{K \subseteq N : \forall i, j \in K \text{ such that } i \neq j,\ \{i, j\} \in E\}$ that can perform a considerate improving move from σ.*

In what follows, we will say that a linked voting game always admits a considerate equilibrium if for any instance of the voting game, and any social network $G = (N, E)$, there exists at least one state σ which is a considerate equilibrium. The concept of considerate equilibrium is illustrated in the following example.

Example 2.4 *Let us consider an instance with three voters and two candidates, where $N = \{1, 2, 3\}$ and $M = \{a, b\}$. The social network and the profile of preferences are:*

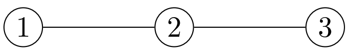

$$\begin{array}{ll} 1: & b \sim a \\ 2: & b \succ a \\ 3: & a \succ b \end{array}$$

The social network is a path, therefore the possible coalitions are $\mathcal{C} = \{\{1,2\},\{2,3\},\{1\},\{2\},\{3\}\}$. *Consider a state* σ *where* a *is elected. Coalitions* $\{1,2\}$ *and* $\{2\}$ *are the only coalitions having incentive to move, they want to make* b *win since* $b \geq_{\{1,2\}} a$ *and* $b >_{\{2\}} a$. *Because* $\{3\} \in \mathcal{N}(\{1,2\}) \cap \mathcal{N}(\{2\})$ *and* $a \succ_3 b$, *coalitions* $\{1,2\}$ *and* $\{2\}$ *cannot perform a CIM (without even taking into account the ability of the coalitions to change the outcome), otherwise they would harm their neighbor* 3. *Thus,* σ *is a considerate equilibrium.*

Following the same idea, if the social network $G = (N, E)$ is composed of a set of disjoint cliques and only maximal cliques of G are considered in $\mathcal{C}$, then an equilibrium associated with ($\mathcal{C}$,WIM) is called a *partition equilibrium* [Feldman and Tennenholtz, 2009]. Observe that in such a case, a CIM corresponds to a WIM.

Definition 2.7 (Partition equilibrium) *Given a partition* $\mathcal{P}$ *over the agents, a state* σ *is a partition equilibrium if and only if there is no coalition in* $\mathcal{C} := \mathcal{P}$ *that can perform a weak improving move from* σ.

Therefore, if a considerate equilibrium is guaranteed to exist for any social network G, then a partition equilibrium exists. Furthermore, if $E = \emptyset$, then a considerate equilibrium corresponds to a Nash equilibrium. Thus, a Nash equilibrium is a special case of a considerate equilibrium where the social network is empty. Hence, the existence of a considerate equilibrium for any social network implies the existence of a Nash equilibrium.

$$\exists \text{ considerate equilibrium}^1 \Rightarrow \exists \text{ partition equilibrium}^2 \Rightarrow \exists \text{ Nash equilibrium} \qquad (2.2)$$

2.3 Existence of considerate equilibria

We investigate in this section the existence of a considerate equilibrium in the linked voting game under different voting rules. As we will show, many rules guarantee the existence of a considerate equilibrium. However, a considerate equilibrium fails to exist for simple and reasonable rules such as Borda or Copeland.

2.3.1 Strict majority susceptible rules

Consider the set of coalitions that are able to change the outcome, according to $\mathcal{F}$, of any voting profile via a manipulation move. Such coalitions are called *powerful coalitions.* Let us denote by $\mathcal{C}_p^{\mathcal{F}}$ the set of powerful coalitions according to voting rule $\mathcal{F}$.

Definition 2.8 (Powerful coalition) *Coalition* $C \subseteq N$ *is powerful if and only if for any voting profile* $\sigma \in \mathcal{B}_{\mathcal{F}}{}^n$ *there exists a joint strategy* σ'_C *such that* $\mathcal{F}(\sigma_{-C}, \sigma'_C) \neq \mathcal{F}(\sigma)$.

The notion of powerful coalitions is similar, by forgetting the dynamic aspect, to the idea of *winning coalitions* in the context of *simple games* [Shapley, 1962, Taylor and Zwicker, 1999], a specific cooperative game in voting, and to the notion of *decisive*

[1] For every social network $G = (N, E)$
[2] For every partition $\mathcal{P}$ over N

coalitions [Ferejohn, 1977], that are coalitions imposing their preferences to the outcome of the voting rule.

A voting profile σ is said to be *unanimous*[1] if σ_i is the same ballot for every voter i. For instance, under Plurality, since $\mathcal{B}_{Plurality} = M$, every alternative $x \in M$ induces a unanimous profile $(x, \ldots, x)$. Clearly, for a unanimous voting rule, the winner of a unanimous profile $\sigma = (\sigma_1, \ldots, \sigma_1)$, where $\sigma_1 \in \mathcal{L}(M)$, is the alternative ranked on top of ballot σ_1. Let $s_x \in \mathcal{B_F}^n$ be a unanimous voting profile favoring candidate x. More precisely, if $\mathcal{B_F} = \mathcal{L}(M)$ or $\mathcal{B_F} = \mathcal{W}(M)$, then $(s_x)_i = [x \succ \text{rest}]$, for every voter i, where "rest" denotes an arbitrary ranking over $M \setminus \{x\}$ that is the same for every voter i. If $\mathcal{B_F} = M$, like in Plurality or Veto, then $s_x = (x, \ldots, x)$ for Plurality and $s_x = (y, \ldots, y)$ where $y \neq x$ for Veto. A particular unanimous voting profile that deserves attention is s_a, where a is the top candidate in tie-breaking $\rhd$.

Let us denote by $\mathcal{C}^{\mathcal{F}}_{wp}$ the set of coalitions that are able to change the outcome, according to $\mathcal{F}$, of a unanimous voting profile, that is different from s_a, with a manipulation move. Such coalitions are called *weakly powerful.*

Definition 2.9 (Weakly powerful coalition) *A coalition $C \subseteq N$ is weakly powerful if there exists a unanimous voting profile $\sigma \in \mathcal{B_F}^n$ for which there exists a joint strategy σ'_C such that $\mathcal{F}(\sigma_{-C}, \sigma'_C) \neq \mathcal{F}(\sigma)$, and there is no joint strategy σ''_C such that $\mathcal{F}((s_a)_{-C}, \sigma''_C) \neq \mathcal{F}(s_a)$.*

Let us define *strict majority susceptible rules* as voting rules for which the powerful coalitions are exactly all the coalitions with a strict majority of voters, and the weakly powerful coalitions are exactly all the coalitions with half of the voters.

Definition 2.10 (Strict majority susceptibility) *A voting rule $\mathcal{F}$ is strict majority susceptible if and only if $\mathcal{C}^{\mathcal{F}}_p = \{C \subseteq N : |C| > n/2\}$ and $\mathcal{C}^{\mathcal{F}}_{wp} = \{C \subseteq N : |C| = n/2\}$.*

By definition, the voting game where $\mathcal{F}$ is a strict majority susceptible rule always admits a Nash equilibrium, e.g., the unanimous profile s_a. However, it is known that a strong equilibrium is not guaranteed to exist (see e.g., Sertel and Sanver [2004]), as we can see in the following example showing the well-known Condorcet paradox.

Example 2.5 *Let us consider an instance of a voting game with three voters and three candidates, where $N = \{1, 2, 3\}$, $M = \{a, b, c\}$ and $\mathcal{F}$ is a strict majority susceptible rule. The profile of preferences is:*

$$\begin{array}{ll} 1: & a \succ b \succ c \\ 2: & c \succ a \succ b \\ 3: & b \succ c \succ a \end{array}$$

If a is elected then coalition $\{2, 3\}$ has incentive to deviate in order to elect c. If c is elected then coalition $\{1, 3\}$ has incentive to deviate in order to make b elected. Finally, if b is elected then coalition $\{1, 2\}$ has incentive to deviate in order to make a elected.

[1] This notion must not be confused with the unanimity axiom (see Definition 1.19) that applies to voting rules.

Note that when a Condorcet winner exists (see Definition 1.21), such as under the single-peakedness restriction, there always exists a strong equilibrium when $\mathcal{F}$ is a strict majority susceptible rule, e.g. the unanimous profile s_x where x is the Condorcet winner. Actually, when there exists a Condorcet winner x, a strong equilibrium under a strict majority susceptible rule necessarily elects x, otherwise a strict majority of voters has incentive to make x the winner and by definition such a coalition can enforce the election of x.

Note that Example 2.5 also rules out the existence of a super strong equilibrium, observation trivially implied by (2.1). However, we shall see that a considerate equilibrium must exist.

Theorem 2.1 *A considerate equilibrium is guaranteed to exist for any instance of the linked voting game where $\mathcal{F}$ is a strict majority susceptible rule.*

Proof: If there is no clique $K \subseteq N$ in graph G such that $|K| > n/2$, then $\mathcal{C}_p^{\mathcal{F}} = \emptyset$ and unanimous profile s_a is a considerate equilibrium. Suppose from now on that $\mathcal{C}_p^{\mathcal{F}} \neq \emptyset$.

Let us denote by $\mathcal{Q}$ the set of all agents belonging to a coalition which can change the outcome of a unanimous state, that is $\mathcal{Q} = \bigcup_{C \in \mathcal{C}_{wp}^{\mathcal{F}} \cup \mathcal{C}_p^{\mathcal{F}}} C$. For any coalitions $C \in \mathcal{C}_p^{\mathcal{F}}$ and $C' \in \mathcal{C}_{wp}^{\mathcal{F}} \cup \mathcal{C}_p^{\mathcal{F}}$, C and C' have at least one member in common, therefore $C \subseteq (\mathcal{N}(C') \cup C')$ and $\mathcal{Q} \subseteq (C \cup \mathcal{N}(C))$, for all $C \in \mathcal{C}_p^{\mathcal{F}}$. Hence, C' cannot deviate so that a worse candidate, from the viewpoint of at least one member of C, is elected. Consequently, if an alternative x strictly undominated over $C \in \mathcal{C}_p^{\mathcal{F}}$ exists, then unanimous profile s_x is a considerate equilibrium.

Suppose from now on that for every coalition $C \in \mathcal{C}_p^{\mathcal{F}}$, there is no strictly undominated alternative over C. However, an undominated alternative over C must exist. For a coalition C and an alternative x, let I_x^C be the *indifference set* of x within C, i.e. $I_x^C := \{y \in M \setminus \{x\} : y \sim_C x\}$. Let ND be the subset of alternatives which are both undominated over at least one coalition of $\mathcal{C}_p^{\mathcal{F}}$ and undominated over $\mathcal{Q}$. The set ND is never empty because $\mathcal{C}_p^{\mathcal{F}} \neq \emptyset$. We denote by x the best alternative of ND w.r.t $\rhd$ and by $C \in \mathcal{C}_p^{\mathcal{F}}$ a coalition for which x is undominated. We will analyze deviations from the unanimous profile s_x and every time a deviation is performed, then we will consider for the next step the unanimous profile of the corresponding winner.

Since we start from unanimous profiles, the deviations are only performed by coalitions in $(\mathcal{C}_p^{\mathcal{F}} \cup \mathcal{C}_{wp}^{\mathcal{F}})$. The winner of every deviation belongs to $I_x^C \cup \{x\}$ because the deviating coalitions have consideration for C. A coalition $C' \in \mathcal{C}_p^{\mathcal{F}}$ cannot deviate from s_x, because $\mathcal{Q} \subseteq (C' \cup \mathcal{N}(C'))$, x is undominated over $\mathcal{Q}$ by definition of ND and $C' \subseteq \mathcal{Q}$. However, a coalition $C' \in \mathcal{C}_{wp}^{\mathcal{F}}$, i.e., of size $|C'| = n/2$, can deviate if there exists $y \in I_x^C$ such that $y \geq_{C'} x$ and $y \rhd x$. If $y \in$ ND then $y \rhd x$ contradicts the fact that x is the best alternative of ND w.r.t. $\rhd$, so $y \notin$ ND. Since $y \geq_{C'} x$ and x is undominated over $\mathcal{Q}$, there exists $j \in \mathcal{Q} \setminus (C \cup C')$ such that $x \succ_j y$. If there exists $C'' \in (\mathcal{C}_p^{\mathcal{F}} \cup \mathcal{C}_{wp}^{\mathcal{F}})$ for which $j \in C''$ with $C'' \cap C' \neq \emptyset$, then C' cannot deviate to y by consideration for j because $j \in \mathcal{N}(C')$, therefore $\overline{C'} = N \setminus C'$ is a coalition of size $n/2$, i.e., belongs to $\mathcal{C}_{wp}^{\mathcal{F}}$, and $j \in \overline{C'}$. It follows that $\mathcal{Q} = N$. Obviously, $C' \cap C \neq \emptyset$ and $\overline{C'} \cap C \neq \emptyset$. These voters cannot be the beneficiaries of a deviation since they are indifferent among the alternatives of I_x^C. Let $I = C' \setminus C$ and $J = \overline{C'} \setminus C$ denote respectively the voters of C' and $\overline{C'}$ who do not belong to C. Then, C' deviates to y and we consider for the next step the unanimous profile s_y. Observe that for a given succession of such deviations performed by coalitions in $\mathcal{C}_{wp}^{\mathcal{F}}$, with size $n/2$, the winners are not in ND and the rank of the winner in $\rhd$ strictly improves. Thus, there is only a finite number of such deviations.

Now, every time a coalition $C'' \in \mathcal{C}_p^{\mathcal{F}}$ can deviate to $z \in I_x^C$ from a unanimous profile s_{y^t} where y^t is the winner at step t (one can verify that C'' cannot deviate to x), it follows that $z \geq_N y^t$ because $\mathcal{Q} \subseteq (C'' \cup \mathcal{N}(C''))$. There exists a voter i such that $z \succ_i y^t$, thus $i \in I \cup J$, say $i \in J$, implying that $j \in \mathcal{N}(C')$. So, no coalition can make j less satisfied. Hence, the improvements within the sequence of the defined deviations follow the preferences of every such j and, during a succession of deviations performed by coalitions of size $n/2$ (i.e., in $\mathcal{C}_{wp}^{\mathcal{F}}$), the rank of the winner in $\rhd$ is improved. Since we have a finite number of alternatives, we finally reach an alternative for which the unanimous associated profile is a considerate equilibrium. □

Let us remark that several well-known voting rules satisfy the condition mentioned in the previous theorem, among them Plurality, Plurality with run-off, STV and Maximin (see Section 1.3.1 for the definitions of these rules).

Lemma 2.2 *Plurality, Plurality with run-off, STV and Maximin are strict majority susceptible rules.*

Proof: Let $\mathcal{F}$ belong to {Plurality, Plurality with run-off, STV, Maximin}. Any strict majority of voters can enforce the election of a candidate $y \neq \mathcal{F}(\sigma)$ from a voting profile σ, by unanimously deviating to the same ballot with y on top. This is due to the majority consistency of these voting rules (see Definition 1.20). Therefore, $\{C \subseteq N : |C| > n/2\} \subseteq \mathcal{C}_p^{\mathcal{F}}$.

We will prove that the strict majority coalitions actually characterize the set of powerful coalitions. Let us consider the unanimous voting profile s_a, with a first ranked candidate in tie-breaking $\rhd$, and a coalition $C \subseteq N$ such that $|C| \leq n/2$. Whatever the deviation of C from s_a leading to a new profile σ', a still obtains a Plurality score at least equal to $n/2$ in σ'. No other candidate can have a better Plurality score, so candidate a is still winning in σ' under Plurality, thanks to the tie-breaking rule. For Plurality with run-off and STV, candidate a remains until the last round of candidate eliminations. At this point, a still has a score at least equal to $n/2$, so a remains the winner under Plurality with run-off and STV thanks to $\rhd$. Concerning Maximin, the minimum number of voters preferring a to another alternative in σ' is at least $n/2$ since a is still ranked first by at least $n/2$ voters. Because for any candidate x, a is ranked before x in σ' by at least $n/2$ voters, no candidate different from a has a Maximin score larger than $n/2$. It follows that a remains the winner under Maximin thanks to the tie-breaking rule. In conclusion, a coalition C of size $|C| \leq n/2$ cannot change the outcome of the election from any profile, since this is not the case from s_a. Thus, $\mathcal{C}_p^{\mathcal{F}} = \{C \subseteq N : |C| > n/2\}$.

Let us now consider the weakly powerful coalitions. Take a unanimous voting profile s_x such that $x \neq a$, and a coalition $C \subseteq N$ of size $|C| = n/2$. The four rules that we consider are unanimous so x is the winner of s_x. Coalition C is able to change the outcome by deviating to a joint strategy where every agent in C deviates to a ballot ranking a on top. Under Plurality, Plurality with run-off, STV and Maximin, a and x both get a score of $n/2$ in the new profile (for the run-off rules, this is the score at the last elimination round). Since a is better than x in $\rhd$, a becomes the winner of the new profile for all these rules. However, as already mentioned, no coalition of size $n/2$ is able to change the outcome of unanimous profile s_a. Thus, $\{C \subseteq N : |C| = n/2\} \subseteq \mathcal{C}^{\mathcal{F}}_{wp}$.

Observe that no other coalition is weakly powerful. Indeed, for any unanimous voting profile s_x, the deviation of a coalition of size strictly less than $n/2$ does not change the fact that x is ranked first by a strict majority of voters in the new profile. Therefore, by majority consistency of Plurality, Plurality with run-off, STV and Maximin, x remains the winner in the new profile. Hence, $\mathcal{C}^{\mathcal{F}}_{wp} = \{C \subseteq N : |C| = n/2\}$ and finally these four rules are strict majority susceptible rules. □

Combining Theorem 2.1 and Lemma 2.2 leads to the following corollary.

Corollary 2.3 *Every instance of the linked voting game where $\mathcal{F} \in \{$Plurality, Plurality with run-off, STV, Maximin$\}$ admits a considerate equilibrium.*

Note that when $n > 2$, the considerate equilibria that we have constructed in the proof of Theorem 2.1 are also Nash equilibria because they are unanimous profiles. However, these two requirements may be conflicting. Indeed, there exist instances for $n = 2$ where the set of considerate equilibria and the set of Nash equilibria do not intersect, as we can see in the next example.

Example 2.6 *Let us consider an instance with two voters and three candidates, where $N = \{1, 2\}$, $M = \{a, b, c\}$, $a \rhd b \rhd c$, and $\mathcal{F}$=Plurality. The social network is a complete graph, that is agents 1 and 2 are connected. The profile of preferences is*

$$\begin{array}{ll} 1: & (a \sim c) \succ b \\ 2: & b \succ c \succ a \end{array}$$

The cliques of the network are $\{\{1, 2\}, \{1\}, \{2\}\}$. The next table gives the winner for all possible strategies/ballots of the voters. The considerate equilibria are marked with pink (light) cells and the Nash equilibria with blue (dark) cells.

		Voter 2		
		a	*b*	*c*
	a	*a*	*a*	*a*
Voter 1	*b*	*a*	*b*	*b*
	c	*a*	*b*	*c*

Nevertheless, with n strictly larger than 2, unanimous considerate equilibria are also Nash equilibria.

2.3.2 The Copeland rule

The Copeland rule (see Section 1.3.1) is a pairwise comparison voting rule that is Condorcet consistent, like Maximin. However, contrary to Maximin, Copeland is not a strict majority susceptible rule.

Lemma 2.4 *Copeland is not strict majority susceptible.*

Proof: Let us consider an instance with an even number of voters. Take a coalition $C \subseteq N$ of size $|C|= n/2$ and any voting profile σ. Suppose that x is the winner of σ under Copeland. Let q_z be the best rank that a candidate z has within a ballot σ_j of a voter j in $N \setminus C$. Choose as a candidate y a candidate different from x that has the best q_y (there can be several such candidates). Observe that q_y is either the second rank if every agent in $N \setminus C$ ranks x first in σ, or the first rank otherwise. Let us consider a ballot $\sigma'_i \in \mathcal{L}(M)$ where y is ranked at the first position of the ballot and x at the last, and denote by σ'_C the joint strategy where every agent i in C plays ballot σ'_i. Observe that the winner of strategy profile $\sigma'' = (\sigma_{-C}, \sigma'_C)$, where every agent not in C plays the same ballot as in σ and coalition C plays the joint strategy σ'_C, is y. Indeed, since y is ranked second or first in σ'' for some voter not in C and all the voters in C rank y first, candidate y beats at least $n-2$ candidates via pairwise comparisons in σ'', leading to a Copeland score of at least $n-2$ for y. If q_y is a first rank, then the Copeland score of y is $n-1$ and y is the Condorcet winner in σ'', leading to the election of y in σ'' under Copeland. Otherwise, if q_y is a second rank, then the Copeland score of y in σ'' is $n-2$, where the only candidate not beaten by y is x, because this implies that x is ranked first by all the voters in $N \setminus C$. Therefore, no candidate z can beat neither x nor y because they are both ranked first by $n/2$ voters, leading to a Copeland score for z smaller than the score of y. Since x is ranked last by every voter in C, x cannot beat any candidate at the absolute majority and its Copeland score is zero. Thus, y is winning in σ''.

To summarize, any coalition of size $n/2$ can change the outcome of the election from any voting profile. Hence, by definition, any such a coalition is a powerful coalition, and $\{C \subseteq N : |C|= n/2\} \subseteq \mathcal{C}_p^{\mathcal{F}}$ when $\mathcal{F}$=Copeland. This implies that Copeland is not a strict majority susceptible rule. □

It follows that the proof of Theorem 2.1, for the existence of a considerate equilibrium, does not hold when $\mathcal{F}$=Copeland.

Indeed, it turns out that a considerate equilibrium is not guaranteed to exist in such a voting game. We show by a simple counterexample that even a Nash equilibrium may not exist in the voting game when using the Copeland rule, ruling out the existence of a considerate equilibrium for any instance of the linked voting game by (2.2) (recall that when $E = \emptyset$ the two solution concepts coincide).

Proposition 2.5 *A Nash equilibrium is not guaranteed to exist in the voting game where $\mathcal{F}$=Copeland, even if the preferences are strict and single-peaked.*

Proof: Let us consider an instance with two voters and three candidates, where $N = \{1, 2\}$, $M = \{a, b, c\}$, $a \rhd b \rhd c$, and $\mathcal{F}$=Copeland. The profile of preferences is:

$$\begin{array}{ll} 1: & a \succ b \succ c \\ 2: & c \succ b \succ a \end{array}$$

The preferences are single-peaked with respect to the axis $a >^M b >^M c$. According to the Copeland rule, an alternative x beats y in a pairwise election if x is ranked before y in the two ballots. Let us consider any voting profile σ.

If alternative a wins in σ, then agent 2 deviates by placing a at the last position of her new ballot and placing on top of her new ballot the first ranked candidate in the ballot of agent 1 between b and c. By this way, either b or c, that agent 2 prefers to a, wins.

If alternative c wins in σ, then it suffices for agent 1 to place c at the last position of her new ballot. Indeed, by this way, the Copeland score of c is zero so c can never win, since it is ranked last in the tie-breaking. Therefore, candidate a or b necessarily wins and agent 1 prefers both a and b to c.

Finally, consider the case where b wins in σ. If c is not ranked first in agent 2's ballot, then agent 1 plays $a \succ c \succ b$. By this way, the Copeland score of b is always zero and the Copeland score of c cannot be more than one. Actually, the only possibility for c to have one point is that the ballot of agent 2 is the same as agent 1's ballot. But in this case, a gets a Copeland score equal to two and thus wins. In all the other cases where b and c have zero point, a necessarily wins thanks to the tie-breaking. Otherwise, i.e., if c is ranked first in the ballot of agent 2, then agent 1 plays $a \succ b \succ c$. Therefore, the Copeland score of both b and c must be zero, implying that candidate a, that agent 1 prefers to b, necessarily wins, thanks to the tie-breaking.

In conclusion, there is no Nash equilibrium in this instance. □

Let us remark that the previous counterexample is quite specific since one voter represents half the electorate. Indeed, for more than two voters, a Nash equilibrium is guaranteed to exist.

Proposition 2.6 *A Nash equilibrium exists for any instance of the voting game where $\mathcal{F}$=Copeland and $n > 2$.*

Proof: Let us consider a unanimous voting profile s_x. Clearly, candidate x wins the election since Copeland respects the unanimity axiom (see Definition 1.19). Then, even if one voter deviates to a new ballot, x is the Condorcet winner since more than half of the voters rank it first. Because Copeland is Condorcet consistent, x remains the winner and no single voter has the power to change the outcome of the election. □

One could think that this positive result when $n > 2$ could also lead to an existence result for a considerate equilibrium. However, the counterexample provided in the proof of Proposition 2.5 can be extended to more than two voters, in order to show that even a partition equilibrium may not exist.

Proposition 2.7 *A partition equilibrium may not exist where $\mathcal{F}$=Copeland even when $n > 2$, and the preferences are strict and single-peaked.*

Proof: Let us consider an instance where n is even and $M = \{a, b, c\}$ (recall that $a \rhd b \rhd c$). The preferences of the voters are the following: $a \succ_i b \succ_i c$ if $1 \leq i \leq n/2$, and $c \succ_i b \succ_i a$ otherwise. They are single-peaked with respect to the axis $a >^M b >^M c$. The partition over the voters is defined as $\mathcal{P} = \{\mathcal{C}_1, C_2\}$ with $C_1 := \{i : 1 \leq i \leq n/2\}$ and $C_2 := \{i : n/2 < i \leq n\}$, where the voters in each coalition have the same preferences.

Let us consider any voting profile σ. If candidate a wins in σ, then all the members of coalition C_2 deviate with the same new ballot where a is ranked at the last position and the first ranked candidate is a candidate between b and c, which is ranked at least once among the two first candidates of a ballot within C_1. By this way, candidate a has a Copeland score equal to zero while the chosen candidate between b and c has at least a Copeland score equal to one. Therefore, either b or c wins and all the members of C_2 prefer them to a.

If alternative c wins in σ, then all the members of coalition C_1 deviate with the same new ballot where c is ranked at the last position. It is sufficient because candidate c is ranked last in the tie-breaking.

Finally, consider the case where candidate b is winning in σ. It suffices for coalition C_1 to deviate to a new ballot where alternative a is ranked at the first position and the relative position of b and c is reversed compared to the current ballots in C_2. That is, if $b \succ_{n/2+i} c$, for $1 \leq i \leq n/2$, then voter i submits the new ballot $a \succ c \succ b$, otherwise i submits $a \succ b \succ c$. Globally, in this new voting profile, there is no absolute majority of voters that prefers b over c or vice versa, and no absolute majority of voters that prefers b or c over a. Therefore, candidate a has at least the same Copeland score as candidates b and c, but in any case a wins thanks to the tie-breaking. □

Nevertheless, restricting to instances where n is odd, which is a classical assumption in social choice, enables to ensure the existence of a considerate equilibrium.

Proposition 2.8 *A considerate equilibrium is guaranteed to exist in any instance of the linked voting game where n is odd and $\mathcal{F}$=Copeland.*

Proof: It suffices to observe that Copeland is a strict majority susceptible rule when n is odd. When n is odd, one can forget the second part of the definition of a strict majority susceptible rule (Definition 2.10), involving coalitions of size exactly $n/2$. Under Copeland, a coalition of size strictly larger than $n/2$ can enforce the election of any candidate, from any voting profile σ, by majority consistency of Copeland (see Definition 1.20). By the same property, these coalitions are the only coalitions that are able to change the outcome of a unanimous profile. Hence, by using Theorem 2.1, we can conclude to the existence of a considerate equilibrium. □

In a nutshell, the existence of a considerate equilibrium is ensured in a linked voting game under Copeland when the number of voters is odd. However, even a Nash equilibrium may not exist, as soon as the number of voters is even.

2.3.3 The Borda rule and other Positional Scoring Rules (PSRs)

Each PSR is characterized by its score vector α (see Section 1.3.1). After normalization, the score vector of every PSR can be written as $\alpha = (1, \alpha_2, \ldots, \alpha_{m-1}, 0)$ where $\alpha_i \in [0, 1]$ and the α_is remain non increasing, for $i \in \{2, \ldots, m-1\}$. The PSR is neither Plurality nor Veto if and only if $\alpha_2 > 0$ and $\alpha_{m-1} < 1$. If $\alpha_i \in \{0, 1\}$ for all $i \in \{2, \ldots, m-1\}$, then the associated PSR is k-approval for a given k.

We know that for Borda and k-approval (k is a constant), a dynamics of the voting game may not converge to a Nash equilibrium [Lev and Rosenschein, 2012], even if the initial state is truthful and only best responses are used (see Definition 1.26). Therefore, this result holds too for a considerate equilibrium because the two solution concepts coincide when graph G is empty. Furthermore, we can prove that a Nash equilibrium (and thus, a considerate equilibrium) is not guaranteed to exist for Borda and k-approval where k is fixed according to either the number of consecutive ones, or to the number of consecutive zeros in score vector α.

Proposition 2.9 *Let ℓ be an integer such that $\ell > 1$ and consider the k-approval rule where $k = \ell$ or $k = m - \ell$. A Nash equilibrium is not guaranteed to exist in the voting game, even if the preferences of the voters are strict and single-peaked.*

Proof: Let us consider an instance with two voters and 2ℓ candidates, where $N = \{1, 2\}$ and $M = \{x_1, x_2, \ldots, x_{2\ell}\}$, and the tie-breaking rule is such that $x_1 \rhd x_2 \rhd \cdots \rhd x_{2\ell}$. Thus, $m = 2\ell$. Since $k = \ell$ or $k = m - \ell$, the voters need to approve, in their k-approval ballot, exactly half of the candidates. The profile of preferences is:

$$\begin{array}{rl} 1: & x_1 \succ x_2 \succ \ldots \succ x_m \\ 2: & x_m \succ x_{m-1} \succ \ldots \succ x_1 \end{array}$$

The preferences are single-peaked with respect to the axis $x_1 >^M x_2 >^M \cdots >^M x_m$. Observe that, with such a k and m, there are in total m approvals that are distributed over the m candidates, by considering the ballots of the two voters. Trivially, a candidate cannot have more than two approvals since there are only two voters and a voter must approve exactly k different candidates. Therefore, in case the winner is not approved by both voters, the winner must be x_1, since x_1 is ranked first in the tie-breaking.

By construction of the preferences, there is always one agent who prefers at least k alternatives to the current winner. Let us consider any voting profile σ. Suppose alternative x_i is elected in σ.

If $i \leq k$, then a better response for agent 2 is to disapprove candidates $x_1, \ldots, x_i$ in her new ballot and approve her preferred candidate within the candidates that agent 1 approves in her current ballot (the rest of the approved candidates can be arbitrary). Observe that this candidate cannot be a candidate of $\{x_1, \ldots, x_i\}$ by construction of the preferences and the fact that $k = m/2$. By this way, no candidate within $\{x_2, \ldots, x_i\}$ can win because it is only approved by one voter. Moreover, there is at least one candidate with two approvals since agent 2 approves a candidate also approved by 1. As already mentioned, this candidate does not belong to $\{x_1, \ldots, x_i\}$, thus a candidate x_j such that $j > i$, that is preferred by agent 2 to x_i, wins.

If $i > k$, a better response for agent 1 is to disapprove in her new ballot candidates $x_i, \ldots, x_m$. By this way, no candidate within $\{x_i, \ldots, x_m\}$ is approved by both voters, thus a candidate x_j such that $j < i$, that agent 2 prefers to x_i, must win. □

A similar negative result regarding the existence of a Nash equilibrium can be stated for specific PSRs.

Proposition 2.10 *A Nash equilibrium is not guaranteed to exist in the voting game where $\mathcal{F}$ is a PSR with $0 < \alpha_2 \leq \frac{1}{2}$, even if the preferences of the voters are strict and single-peaked.*

Proof: Let us consider an instance with two voters and three candidates where $N = \{1, 2\}$, $M = \{a, b, c\}$, $a \rhd b \rhd c$, and $\mathcal{F}$ is a PSR whose score vector is $\alpha = (1, \alpha_2, 0)$ for $0 < \alpha_2 \leq \frac{1}{2}$. The profile of preferences is:

$$\begin{array}{ll} 1: & a \succ b \succ c \\ 2: & c \succ b \succ a \end{array}$$

The preferences are single-peaked with respect to the axis $a >^M b >^M c$.

Let us consider any voting profile σ. If alternative a is elected in σ, then agent 2 puts a last in her new ballot and puts first the best ranked alternative within agent 1's ballot between alternatives b and c. This implies that a has at most one point, and the chosen candidate between b and c has a score strictly better than one, making it the new winner.

If alternative c is elected in σ, then agent 1 puts c last in her new ballot and puts first the best ranked alternative within agent 2's ballot between alternatives a and b. Therefore, candidate c has at most one point and the score of the chosen candidate between a and b is strictly better than one, making it the new winner.

Finally, if b is elected in σ, then agent 1 can make a the winner. If c is ranked first by agent 2 then agent 1 plays $a \succ b \succ c$, implying that c obtains a score equal to one, b has at most one point and a at least one point. So, candidate a is winning in any case thanks to the tie-breaking. Otherwise, agent 1 plays $a \succ c \succ b$, inducing that candidates b and c obtain at most one point while a gets at least one point. Thus, candidate a is winning in any case thanks to the tie-breaking. □

Note that Borda is covered by Proposition 2.10 when there are three candidates (in this case, the normalized vector score of Borda is $\alpha = (1, 1/2, 0)$). Therefore, a Nash equilibrium (and thus, a considerate equilibrium) is not guaranteed to exist with the Borda rule.

Corollary 2.11 *A Nash equilibrium may not exist in the voting game where $\mathcal{F}$=Borda.*

Observe that the counterexample provided in the proof of Proposition 2.10 is an instance where there are only two voters. It turns out that this type of instance is very specific because for more voters, a Nash equilibrium is guaranteed to exist.

Proposition 2.12 *A Nash equilibrium exists for any instance of the voting game where $\mathcal{F}$=Borda and $n > 2$.*

Proof: In this proof, we assume that the score vector α of Borda is not normalized, i.e., $\alpha := (m-1, \ldots, 1, 0)$. Let us construct a voting profile σ where $\sigma_1 :=\succ_1$ (if there was initially indifferences in the preferences of agent 1, then we use an arbitrary tie-breaking in order to obtain a linear order). Say, without loss of generality, that $\sigma_1 = x \succ c_1 \succ c_2 \succ \ldots \succ c_{m-1}$. Now let us set $\sigma_2 := x \succ c_{m-1} \succ \ldots \succ c_2 \succ c_1$. For the rest of the voting profile, we construct σ_i as follows for every agent $i > 2$.

$$\sigma_i := \begin{cases} \sigma_1 & \text{if } i \text{ is odd and } i \neq n \\ \sigma_2 & \text{if } i \text{ is even, or if } i \text{ is odd and } i = n \end{cases}$$

Clearly, since Borda is a unanimous voting rule, candidate x is the winner. Agent 1 has no incentive to deviate, since her preferred candidate is elected. Observe that the sum of the scores of any candidate c_i in the ballots σ_1 and σ_2 is $m - 2$.

Assume that there is an agent i (for $i > 1$) who deviates in order to make another candidate, say c_i, the new winner. The best strategy for agent i is to rank c_i first and x last in the new submitted ballot. Let us denote by σ' the deviating voting profile. The new score of candidate x is $\text{Sc}^{\sigma'}(x) = (n-1).(m-1)$. Since c_i is the new winner, it follows that $\text{Sc}^{\sigma'}(c_i) \geq \text{Sc}^{\sigma'}(x)$.

If n is even, then the best possibility for c_i is to obtain a score of $\text{Sc}^{\sigma'}(c_i) = (m-1)+\frac{n}{2}.(m-2)$. Therefore, by simplifying the inequality $\text{Sc}^{\sigma'}(c_i) \geq \text{Sc}^{\sigma'}(x)$, we get $4-\frac{4}{m} \geq n$. Since $n > 2$ and n is even, $n \geq 4$ and we have $1 - \frac{1}{m} \geq 1$. We reach a contradiction because the number of candidates m is finite.

Now consider the case where n is odd. If $\sigma_i = \sigma_2$, that is if i is even, or i is odd and $i = n$, then the best possibility for c_i is to obtain a score of $\text{Sc}^{\sigma'}(c_i) = (m-1)+\lfloor\frac{n}{2}\rfloor.(m-2)$. Therefore, by simplifying the inequality $\text{Sc}^{\sigma'}(c_i) \geq \text{Sc}^{\sigma'}(x)$, we get $2-\frac{1}{m} \geq \lceil\frac{n}{2}\rceil$. Since $n \geq 3$, we obtain $1 - \frac{1}{2m} \geq 1$, leading to a contradiction.

Otherwise, if $\sigma_i = \sigma_1$, that is if i is odd and $i \neq n$, then the best possibility for c_i is to obtain a score of $\text{Sc}^{\sigma'}(c_i) = (m-1) + \lceil\frac{n}{2}\rceil.(m-2)$. Therefore, by simplifying the inequality $\text{Sc}^{\sigma'}(c_i) \geq \text{Sc}^{\sigma'}(x)$, we get $2 - \frac{3}{m} \geq \lfloor\frac{n}{2}\rfloor$. Since n is odd and there is an agent i such that $i > 1$ is odd but $i \neq n$, then we have $n \geq 5$ and thus we obtain $1 - \frac{3}{2m} \geq 1$, leading to a contradiction. □

Now, the natural question is whether the case $n > 2$ can also ensure the existence of a considerate equilibrium. However, this is not the case because we can show that even a partition equilibrium may not exist when $n > 2$ and n is odd. Note that the counterexample for a partition equilibrium under Copeland when $n > 2$ and n is even, provided in the proof of Proposition 2.7, also works for Borda. But we can also exhibit a particular counterexample that covers the case where n is odd, ruling out the possibility for the game to admit a considerate equilibrium when n is odd, contrary to Copeland.

Proposition 2.13 *A partition equilibrium may not exist when $\mathcal{F}$=Borda, even when n is odd, and the preferences are strict and single-peaked.*

Proof: Let us consider an instance with five voters and three candidates where $N = \{1,\ldots,5\}$ and $M = \{a,b,c\}$ (recall that $a \rhd b \rhd c$). The partition over the voters is defined as $\mathcal{P} = \{C_1, C_2\}$, where $C_1 := \{1,2\}$ and $C_2 := \{3,4,5\}$. The preferences, single-peaked with respect to the axis $a >^M b >^M c$, are given as follows.

$$\begin{aligned} 1,\ 2:&\quad a \succ b \succ c \\ 3,\ 4,\ 5:&\quad c \succ b \succ a \end{aligned}$$

Let us consider any voting profile σ. If candidate a is elected in σ, then all the members of coalition C_2 deviate to ballot $c \succ b \succ a$, implying that the score of a is at most 4 in the new profile while the score of c is at least equal to 6. Therefore, either b or c wins in the new profile and these two options are preferable for C_2 than the election of a.

If candidate c is the winner of σ, then all the members of coalition C_1 deviate to a ballot where c is ranked last and the best candidate between a and b within C_2 in profile σ is ranked first. In this new profile, candidate c obtains at most 6 points while the best candidate between a and b in σ_{C_2} obtains at least 2 points in σ_{C_2}, and thus at least 6 points in total in the new profile. Thanks to the tie-breaking rule, either a or b wins in the new profile, situation that C_1 prefers to the election of c.

Finally, let us consider the case where b wins in σ. In such a case, the members of coalition C_2 deviate in order to make c the new winner, and thus they all give a new ballot where c is ranked first, leading to a new profile where c gets at least 6 points. If, (i), c is ranked first in σ by at least one agent between 1 and 2, then there is no need to elaborate a more sophisticated strategy because c obtains at least 8 points in the new profile, while another candidate can obtain at most 6 points. So, candidate c becomes the new winner. If, (ii), both voters 1 and 2 rank first the same alternative, say a, in σ, then all the members of C_2 must rank a last in their new ballot. This leads to a new profile where c gets at least 6 points, a gets 4 points and b at most 5 points, making c elected. Otherwise, (iii), each of the candidates a and b is ranked first in σ by an agent in C_1. In this case, two agents in C_2 submit the new ballot $c \succ a \succ b$ while the remaining agent submit ballot $c \succ b \succ a$. Therefore, candidate c obtains at least 6 points in the new profile while both a and b cannot obtain more than 5 points, leading to the election of c. □

It follows that no natural restriction on the number of voters leads to the existence of a considerate equilibrium for the Borda rule.

2.4 Special case of the Veto rule

Even though Veto is a PSR like Plurality or Borda, this voting rule has a particular behavior regarding coalitional manipulation, since we prove stronger results than for the other classical rules. Indeed, we show that a strong equilibrium and a considerate equilibrium are guaranteed to exist in any instance of the linked voting game using this rule. We deduce from the strong equilibria under Veto a quality measure for voting rules.

2.4.1 Existence of equilibria

In order to prove the existence of equilibria in the voting game, we use *feasible elimination procedures (f.e.p.)*, a mechanism introduced by Peleg [1978]. The initial definition of an f.e.p. takes into account strict preferences, however it can be directly adapted for non-strict preferences.

Definition 2.11 (Feasible elimination procedure (f.e.p.)) *For a mapping* $\beta : M \to \mathbb{N}$ *such that* $\sum_{x \in M} \beta(x) = n + 1$*, an f.e.p. is a sequence* $(x_1, C_1; x_2, C_2; \ldots; x_{m-1}, C_{m-1}; x_m)$*, where* $\forall i \in [m-1], C_i \subseteq N$*, that satisfies the following conditions:*

(i) $|C_i| = \beta(x_i)$ *and* $\forall j \in [m-1] \setminus \{i\}, C_i \cap C_j = \emptyset$

(ii) $M = \bigcup_{k \in [m]} x_k$

(iii) $\forall \ell \in C_i$ *and* $\forall k \in \{i+1, \ldots, m\}, x_k \succsim_\ell x_i$

An f.e.p. is guaranteed to exist for any preference profile in $\mathcal{L}(M)^n$ and for any mapping β such that $\sum_{x \in M} \beta(x) = n+1$ (Peleg [1984]'s Lemma 5.3.9)[1]. The induction proof does not rely on the assumption of strict preferences, hence this existence result holds for any preference profile in $\mathcal{W}(M)^n$.

Let us define more precisely mapping β for the case of the Veto rule. For a candidate x, let the value $\beta(x)$ correspond to the minimum amount of vetos required to ensure that x cannot be chosen by Veto (whatever the other ballots). Let q and r be respectively the quotient and the rest of the euclidean division of n by m, i.e., $n = m \cdot q + r$. It is easy to check that β must take the following values:

$$\beta(x) = \begin{cases} q+1 & \text{if } x \text{ is ranked among the } r+1 \text{ first alternatives in tie-breaking } \rhd \\ q & \text{otherwise} \end{cases} \tag{2.3}$$

By the definition of q and r, $\sum_{x \in M} \beta(x) = n+1$ holds.

The following theorem is an adaptation of Peleg and Peters [2010]'s Theorem 9.2.6 to the case of the Veto rule.

Theorem 2.14 *A strong equilibrium exists for any instance of the voting game where $\mathcal{F}$=Veto.*

Proof: We show that it is possible to construct, from a given f.e.p., a state σ which is a strong equilibrium for Veto. Let $(x_1, C_1, x_2, C_2, \ldots, C_{m-1}, x_m)$ be an f.e.p. for the mapping β described in (2.3), and let $\sigma \in M^n$ be a state such that for every index $i \in [m-1]$ and every voter $j \in C_i$, $\sigma_j = x_i$ (recall that $\mathcal{B}_{Veto} = M$). Since we have chosen a mapping β corresponding to the number of vetos for each candidate that are necessary to avoid its election, $\mathcal{F}(\sigma) = x_m$.

Let us prove that σ is a strong equilibrium. By contradiction, if σ is not a strong equilibrium, then there exists a coalition $C \subseteq N$ and a joint strategy σ'_C such that $\mathcal{F}(\sigma'_C, \sigma_{-C}) >_C x_m$. Let y denote the candidate $\mathcal{F}(\sigma'_C, \sigma_{-C})$. By the definition of β, there must be an index $i \in [m-1]$ and a voter $\ell \in C_i$ such that $y = x_i$ and $\ell \in C$, because otherwise y is vetoed by at least $\beta(y)$ voters and cannot be chosen by $\mathcal{F}$. But by condition (*iii*) of Definition 2.11, this implies that $x_m \succsim_\ell x_i = y$, a contradiction with $y >_C x_m$. □

As stated in the previous proof, an f.e.p. associated with mapping β induces a strong equilibrium under Veto. However, the reverse is not true: any strong equilibrium does not induce an f.e.p., as we can observe in the following example.

Example 2.7 *Let us consider an instance with three voters and three candidates, where $N = \{1,2,3\}$, $M = \{a,b,c\}$ and $\mathcal{F}$=Veto. The preference profile is described as follows.*

$$\begin{aligned} 1: &\quad b \succ \boxed{a} \succ \underline{c} \\ 2: &\quad b \succ \boxed{\underline{a}} \succ c \\ 3: &\quad b \succ \boxed{c} \succ \underline{a} \end{aligned}$$

[1]Note that the condition $n+1 \geq m$ appears in the existence result of Peleg [1984] but it is easy to fulfill this condition by only keeping the $n+1$ first candidates in the tie-breaking rule. Indeed, only these alternatives can be elected.

The voting profile σ*, that is framed in the previous preference profile, where* $\sigma_1 = \sigma_2 = a$ *and* $\sigma_3 = c$*, is a strong equilibrium that elects candidate* b*. However, there is no associated f.e.p.* $(x_1, C_1; x_2, C_2; x_3)$ *where the ballot of the agents in* C_i *is* x_i *and* $x_m = b$*, because condition* (iii) *of Definition 2.11 cannot hold. Nevertheless, there exists a strong equilibrium electing* b *with an associated f.e.p., i.e., the voting profile* σ'*, underlined in the preference profile, where* $\sigma'_1 = c$ *and* $\sigma'_2 = \sigma'_3 = a$ *and the associated f.e.p. is* $(c, \{1\}; a, \{2, 3\}; b)$*.*

We will show that identifying the candidates that can be elected in a strong equilibrium is polynomial, as stated in the following proposition.

Proposition 2.15 *Verifying whether a given candidate* x *can be elected under* $\mathcal{F}$=*Veto in a strong equilibrium can be done in polynomial time.*

Proof: We will reduce our verification problem to a network flow problem [Ahuja et al., 1993]. Let us consider the directed network $H = (\{s\} \cup N \cup (M \setminus \{x\}) \cup \{t\}, F)$, where s denotes a source node and t a sink node. We describe as follows the arcs (i, j) in F with their associated lower bounds l_{ij} and capacities c_{ij}:

- arcs (s, i) for every $i \in N$ with $l_{si} = 0$ and $c_{si} = 1$
- arcs (i, y) for $i \in N$ and $y \in M \setminus \{x\}$ if $x \succsim_i y$, with $l_{iy} = 0$ and $c_{iy} = 1$
- arcs (y, t) for every $y \in M \setminus \{x\}$ with $l_{yt} = \beta(y)$ (as defined in (2.3)) and $c_{yt} = n$

We claim that there exists a strong equilibrium under Veto electing x if and only if there exists a feasible flow in network H.

Assume that there exists a feasible flow in H and let us consider a maximum flow solution $f : F \to \mathbb{N}$. Then, consider the strategy profile σ where $\sigma_i = y$ for $y \in M \setminus \{x\}$ if $f((i, y)) > 0$, i.e., if flow f circulates in arc (i, y). Let us denote by L the voters that were not assigned a strategy thanks to the network flow solution. The nodes $i \in L$ have necessarily an outdegree equal to zero in H because f is a maximum flow, therefore the voters in L rank x last in their preferences. We arbitrarily assign them strategies within $M \setminus \{x\}$. Clearly, the winner of voting profile σ under Veto is candidate x because it is the only alternative that is not vetoed. We will prove that no improving move can be performed from σ, and thus that σ is a strong equilibrium.

Since there is a feasible flow in network H, every candidate $y \in M \setminus \{x\}$ obtains at least $\beta(y)$ vetos in σ within the voters in $N \setminus L$. Observe that $|L| < \beta(x)$ because $\sum_{y \in M} \beta(y) = n + 1$ and there is a feasible flow. It follows that there cannot be any improving move from a subset of voters in L: they are not sufficiently numerous to veto x with $\beta(x)$ votes whereas all the other candidates $y \in M \setminus \{x\}$ are vetoed at least $\beta(y)$ times by the voters in $N \setminus L$. Let us denote by V_y the voters in $N \setminus L$ that are vetoing candidate y in σ. Every voter $i \in V_y$ prefers x at least as much as y, so there is no improvement move for such a voter if candidate y becomes the new winner. Assume that there exists a set of candidates $S \subseteq M \setminus \{x\}$ such that a subset of voters $A_S \subseteq \bigcup_{y \in S} V_y$ agree to deviate (possibly with some voters in L) from σ. Then, $\beta(z)$ vetos are still active for each candidate $z \notin (S \cup \{x\})$ within the voters $N \setminus (A_S \cup L)$, while the voters in A_S remove some vetos against candidates in S. Therefore, this deviation will necessarily result in the election of a candidate in S, and this is not an improving move for A_S.

Now let us suppose that there is no feasible flow in network H. It follows that there exists a subset $S \subseteq M \setminus \{x\}$ of alternatives such that $|\bigcup_{y \in S}\{i : x \succsim_i y\}| < \sum_{y \in S} \beta(y)$, and thus $|\bigcap_{y \in S}\{i : y \succ_i x\}| > n - \sum_{y \in S} \beta(y)$. Since $\sum_{z \in M} \beta(z) = n + 1$, then $|\bigcap_{y \in S}\{i : y \succ_i x\}| \geq \sum_{z \in M \setminus S} \beta(z)$. Therefore, the voters in $\bigcap_{y \in S}\{i : y \succ_i x\}$ are sufficiently numerous to veto any alternative z in $M \setminus S$ (x included), with the minimum number of vetos $\beta(z)$ that is needed to avoid the election of each $z \in M \setminus S$, whatever the rest of the strategy profile. By performing such a deviation move from any possible strategy profile electing x, the voters in $\bigcap_{y \in S}\{i : y \succ_i x\}$ obtain the election of a candidate within S, that they all prefer to x, and thus x can never be elected in a strong equilibrium under Veto. □

By the previous proposition, one can deduce that even the construction of a strong equilibrium under Veto can be done in polynomial time.

Although Theorem 2.14 holds for non-strict preferences, it does not enable to ensure the existence of a super strong equilibrium when the preferences are not strict. Indeed, in the solution concept defined by the super strong equilibrium, a voter participates to a deviating coalition if she is not worse off by this deviation so, she may participate even if the deviation is not an improvement move for her. This behavior is illustrated in the following counterexample.

Example 2.8 *Let us consider an instance with three voters and two candidates, where* $N = \{1, 2, 3\}$*,* $M = \{a, b\}$ *and* $\mathcal{F}$*=Veto. The profile of preferences is:*

$$\begin{aligned} 1 :&\quad a \succ b \\ 2 :&\quad b \succ a \\ 3 :&\quad a \sim b \end{aligned}$$

Let us consider any voting profile σ*. If* a *is elected in* σ*, then coalition* $\{2, 3\}$ *can deviate from* σ *by vetoing* a*, in order to make* b *elected. If* b *is elected in* σ*, then coalition* $\{1, 3\}$ *can deviate from* σ *by vetoing* b*. Both cases correspond to weak improving moves (WIM) because the situation of agent* 1 *or* 2 *is strictly improved by the deviation but agent* 3 *is indifferent and agree to participate in any deviation. Therefore, there is no super strong equilibrium in this instance.*

Nevertheless, the following theorem shows that a considerate equilibrium is guaranteed to exist even with non-strict preferences.

Theorem 2.16 *A considerate equilibrium exists in any instance of the linked voting game where* $\mathcal{F}$*=Veto.*

Proof: The proof relies on a refinement of the f.e.p. to the concept of considerate equilibrium. The *considerate f.e.p.* (c.f.e.p.) is defined in a similar fashion as f.e.p., except that condition (*iii*) is replaced by (*iii'*) $\forall \ell \in C_i$ and $\forall k \in \{i+1, \ldots, m\}, x_k \succ_\ell x_i$, or $[x_k \sim_\ell x_i$ and $x_i \not\geq_{\mathcal{N}_G(\ell)} x_k]$, where $\mathcal{N}_G(\ell)$ is the set of neighbors of ℓ in G. First of all, let us show that a c.f.e.p. exists for any preference profile $\succsim \in \mathcal{W}(M)^n$. To this end, for any voter $i \in N$, we construct a strict preference $\succ'_i$ which is consistent with the strict part of $\succsim_i$, and where ties are broken by $\geq_{\mathcal{N}_G(i)}$ (or arbitrarily in case of incomparability for $\geq_{\mathcal{N}_G(i)}$). This construction leads to a profile of strict preferences $\succ'$. By Remark 9.2.1

of Peleg and Peters [2010], we know that an f.e.p. exists for $\succ'$. We state that such an f.e.p. is also a c.f.e.p. for $\succsim$. Indeed, conditions (i) and (ii) trivially hold because they are similar for both f.e.p. and c.f.e.p., and they do not depend on the considered profile of preferences. Furthermore, the construction of $\succ'$ ensures that for any $\ell \in N$, $x \succ'_\ell y$ implies $x \succ_\ell y$, or $[x \sim_\ell y$ and $y \not\geqq_{\mathcal{N}_G(\ell)} x]$. Therefore, (iii') holds and the f.e.p. for $\succ'$ is also a c.f.e.p. for $\succsim$.

The remainder of this proof follows the same line as the proof of Theorem 2.14. Let $(x_1, C_1; x_2, C_2; \ldots, C_{m-1}; x_m)$ be a c.f.e.p. for the mapping β described in (2.3), and let σ be a state such that $\sigma_j = x_i$, for every index $i \in [m-1]$ and every voter $j \in C_i$. Let us show that σ is a considerate equilibrium. By contradiction, if σ is not a considerate equilibrium then there exists a coalition $C \in \mathcal{C}$ (where $\mathcal{C}$ refers to the set of cliques of G) and a joint strategy σ'_C such that $\mathcal{F}(\sigma'_C, \sigma_{-C}) \geq_C x_m$ and $\mathcal{F}(\sigma'_C, \sigma_{-C}) \geqq_{\mathcal{N}(C)} x_m$. Let y denote the candidate $\mathcal{F}(\sigma'_C, \sigma_{-C})$. There must be an index $i \in [m-1]$ and a voter $\ell \in C_i$ such that $y = x_i$ and $\ell \in C$. Furthermore, any neighbor of ℓ in G belongs to C or $\mathcal{N}(C)$, and no other voter belongs to C. Therefore, $C \cup \mathcal{N}(C) = \mathcal{N}_G(\ell) \cup \{\ell\}$. But, by condition (iii') of the definition of a c.f.e.p., this implies that $x_m \succ_\ell x_i = y$ or, $x_m \sim_\ell x_i$ and $y = x_i \not\geqq_{C \cup \mathcal{N}(C)} x_m$, a contradiction with $y \geq_C x_m$ and $y \geqq_{\mathcal{N}(C)} x_m$. □

In consequence, a strong equilibrium as well as a considerate equilibrium are guaranteed to exist for any instance of the linked voting game where $\mathcal{F}$=Veto.

2.4.2 Design of a quality measure

The main results of the previous subsection are the existence of a strong equilibrium under the Veto rule for any instance, as well as the possibility to check in polynomial time whether a given candidate can be elected in a strong equilibrium. Strong equilibria are robust and stable voting profiles since no coalition of voters can agree and be able to deviate to another ballot. Since Veto is a "reasonable" voting rule (basically Veto is not dictatorial and every candidate can be elected a priori in the non-resolute version of Veto[1]), candidates elected in a strong equilibrium under Veto can appear as consensual alternatives, in the sense that there is no group of voters preferring another alternative that is sufficiently large to make this alternative the new winner.

Consequently, it would be natural to consider a social choice function based on the selection of the candidates that can be elected in a strong equilibrium under Veto. This perspective has notably been studied for general feasible elimination procedures [Peleg, 1978, 1984, Peleg and Peters, 2010]. However, this approach is limited by the fact that the number of such alternatives can be large. Indeed, we can observe in Figure 2.1, where 10,000 preferences profiles are generated under different cultures of preferences (impartial culture (IC), uniform single-peaked (SP-U) or single-peaked uniform peak (SP-UP), see Section 1.2.1.b for more details) that, in average, half of the alternatives are candidates elected in a strong equilibrium under Veto. The number is even much larger when the preferences are single-peaked via uniform-peak generation.

This implies that the role of the tie-breaking rule may be too important in the procedure. One could think about combinations of voting rules in order to face this

[1]See for instance Baharad and Nitzan [2005] and Kurihara [2018] for a characterization of the Veto rule.

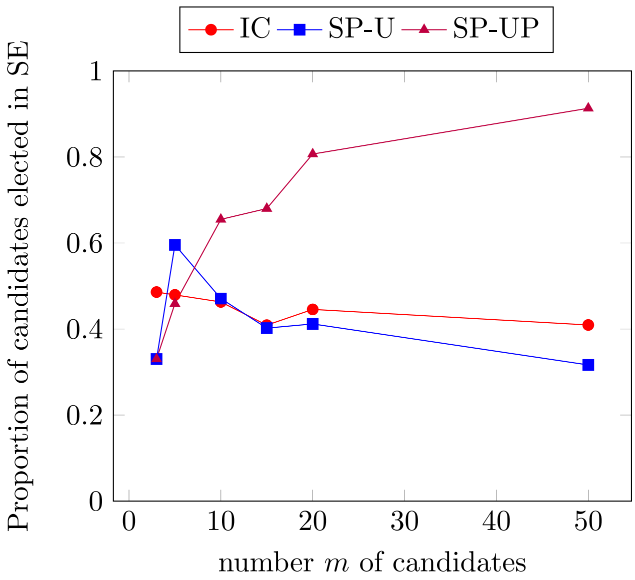

Figure 2.1: Average proportion of the candidates that can be elected in a strong equilibrium under Veto for $n = 100$ and 10,000 runs

problem, like for instance choosing the Borda winner among the candidates that can be elected in a strong equilibrium under Veto.

We choose to use the possibility of being elected in a strong equilibrium under Veto as a characteristic for a "good" candidate. Contrary to the well-known Condorcet winner, such a candidate is guaranteed to exist (Theorem 2.14) for any preference profile. This property is called *Veto-SE*.

Definition 2.12 (Veto-SE) *An alternative x is a Veto-SE if and only if there exists a voting profile $\sigma \in M^n$ such that $\mathcal{F}(\sigma) = x$ where $\mathcal{F}$=Veto, and σ is a strong equilibrium in the voting game under Veto.*

Observe that the Veto-SE property is conceptually different than the concept of Condorcet winner. Whereas in the latter the majority has an important impact, the Veto-SE notion takes into consideration the whole electorate, in the sense that a minority of voters can have a significant power. More precisely, if there is a sufficient number of voters that do not like an alternative at all, then this alternative would never be elected in a strong equilibrium under Veto. On this point, this concept may even be opposed to the Condorcet winner idea. Actually, although the concept of Condorcet winner seems to be a strong requirement, a Condorcet winner is not necessarily a Veto-SE alternative, as we can notice in the following example.

Example 2.9 *Let us consider an instance with three voters and four candidates, where $N = \{1, 2, 3\}$, $M = \{a, b, c, d\}$ and $\mathcal{F}$=Veto. The profile of preferences is:*

$$\begin{array}{ll} 1: & b \succ a \succ c \succ d \\ 2: & c \succ d \succ a \succ b \\ 3: & b \succ d \succ c \succ a \end{array}$$

Alternative b is the Condorcet winner. However, from a state where b is elected, voter 2 has always incentive to deviate by vetoing b, and this veto is sufficient to avoid the election of b. Hence, there is no strong equilibrium electing b.

In order to get a clearer picture on the Veto-SE property, we run some experiments. The output of different classical voting rules is analyzed according to the Condorcet efficiency and the Veto-SE property. The results are presented in Figure 2.2. In the experiments, we generate instances with 10 voters and 5 candidates, where the preferences are drawn from impartial culture (IC), 10%-correlated culture (CR-10%) or uniform single-peaked culture (SP-U) (see Section 1.2.1.b). Seven voting rules are evaluated: Plurality, 2-approval, Borda, Maximin, Copeland, STV and Veto. Under the three different cultures of preferences and for each of the seven chosen voting rules, we display the Condorcet efficiency (blue/dark bars) and the Veto-SE efficiency (red/light bars) of the voting rules. For each voting rule, the Condorcet efficiency refers to the proportion of instances for which the voting rule elects the Condorcet winner, when it exists. If the current instance is not a Condorcet domain, that is there is no Condorcet winner, then we pass to another instance until we have seen 10,000 instances where there exists a Condorcet winner. Concerning the *Veto-SE efficiency*, it refers to the proportion of instances for which the given voting rule elects a Veto-SE candidate, that is a candidate for which there exists a strong equilibrium profile under Veto electing it. Since a Veto-SE is guaranteed to exist in any instance (Theorem 2.14), we do not need to restrict to specific profiles nor to Condorcet domains. Therefore, we stop to look at the veto-SE efficiency as soon as we have seen 10,000 profiles.

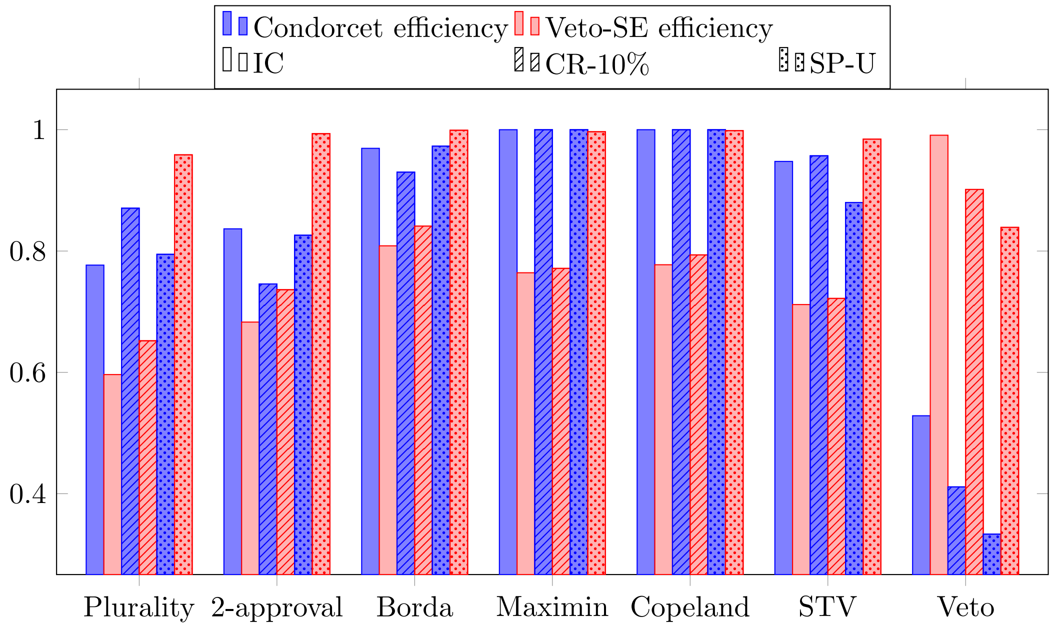

Figure 2.2: Quality of different voting rules for $n = 10$ and $m = 5$

Observe that, globally, Veto-SE efficiency is always above 60%, which is not so weak. However, Veto-SE efficiency is almost always lower than Condorcet efficiency, except for

Veto and uniform single-peaked profiles. Therefore, despite the fact that around half of the candidates are Veto-SE candidates (see Figure 2.1), it seems that electing a Veto-SE candidate is a requirement rather difficult to satisfy for voting rules other than Veto.

In general the two measures may be conflicting: there are only a few configurations where Veto-SE efficiency and Condorcet efficiency are close. The voting rule that appears rather balanced on the two criteria is Borda.

Note that the Veto-SE efficiency is very high when the profiles are single-peaked. This may be explained by the fact that a single-peaked profile, uniformly generated among all single-peaked profiles (SP-U), is strongly correlated, even more than a profile generated with 10% of correlation ratio (CR-10%). Therefore, the voting rules are more likely to select a candidate that is sufficiently consensual to be elected in a strong equilibrium under Veto.

2.5 Convergence of the dynamics

Let us now study, in an iterative voting process, the convergence properties of the dynamics of the linked voting game. Recall that, in iterative voting, a dynamics refers to a sequence of deviations $(\sigma^0, \sigma^1, \ldots, \sigma^T)$ where a deviation is performed by a voter (or a group of voters) between each two consecutive states σ^t and σ^{t+1} for $t \in [T-1]$. Whereas the classical iterative voting setting presented in Section 1.3.3.b only deals with unilateral deviations, we consider that coalitions of voters in $\mathcal{C}$ can perform deviations in this dynamics.

Some deviations are more relevant than others for responses in iterative voting. Indeed, by assuming that the voters are rational, it makes sense to focus on the best possible deviations for coalitions.

A *μ-response* for a coalition $C \in \mathcal{C}$, according to a type of move $\mu \in \{\text{IM,WIM,CIM}\}$ (see Section 2.2) from a state σ, is a deviation of type μ performed by coalition C from state σ. This corresponds to the coalitional version of a *better response* (see Definition 1.25) in the classical iterative voting setting where only unilateral deviations are taken into account. Let us define in consequence a *best μ-response* for coalition C from state σ as one of the best possible deviations of type $\mu \in \{\text{IM,WIM,CIM}\}$ that C can perform from σ.

Definition 2.13 (Best μ-response) *A deviation ballot $\sigma'_C \in \mathcal{B_F}^{|C|}$ is a best μ-response from state σ for coalition $C \subseteq N$ if σ'_C is a μ-response for C from σ and there is no other μ-response σ''_C such that $\mathcal{F}(\sigma_{-C}, \sigma''_C) \geq_C \mathcal{F}(\sigma_{-C}, \sigma'_C)$.*

This definition extends the notion of *best response* (see Definition 1.26), which is used for unilateral improving moves in the classical iterative setting, to coalitional deviations. These restricted deviations are useful in the study of the convergence of the dynamics of the voting game.

Let us study the convergence of the dynamics of the game to a considerate equilibrium when a considerate equilibrium is guaranteed to exist. More precisely, according to the results of Section 2.3, we investigate the convergence of the dynamics under different voting rules such as Plurality, Veto, Maximin, STV, Maximin and Plurality with run-off, which ensure the existence of a considerate equilibrium.

2.5.1 Plurality and Veto

Theorem 2.1 and 2.16 give the existence of a considerate equilibrium and therefore, existence of a partition equilibrium, in any instance of the linked voting game under Plurality and Veto. However, if we let the voters deviate, in an iterative voting perspective, do we reach this equilibrium? By considering unilateral deviations, we know that under the restrictions to direct best responses (see Definitions 1.27 and 1.28), the dynamics of the voting game is guaranteed to converge to a Nash equilibrium [Meir et al., 2010, Lev and Rosenschein, 2012, Reyhani and Wilson, 2012] for both voting rules, whereas it does not hold for general manipulation moves. As we will show in this subsection, the same positive result for convergence is not possible for a considerate equilibrium since it is not the case even for the partition equilibrium and additional natural restrictions about the deviations.

2.5.1.a Plurality

In the deviations under Plurality, only the support for a candidate matters. Indeed, there is no possibility to block or to limit the scores of other candidates via a deviation under Plurality and the only strategy for making elected a given candidate is to support it as much as possible. Therefore, if there exists a joint strategy σ'_C for coalition $C \subseteq N$ from state σ such that y is elected, then unanimous joint strategy $(s_y)_C$, where every voter $i \in C$ gives ballot $y \in M$, also make y elected from σ. This idea applies to any type of deviation $\mu \in \{\text{IM,WIM,CIM}\}$ that has been defined in Section 2.2.

We will extend the definition of *direct best response* under Plurality (see Definition 1.27) to a unanimous direct μ-response, for a type of deviation $\mu \in \{\text{IM,WIM,CIM}\}$, by considering coalitional deviations instead of simply unilateral deviations.

Definition 2.14 (Unanimous direct best μ-response) *A deviation ballot σ'_C is a unanimous direct best μ-response under Plurality from state σ for coalition $C \subseteq N$, if σ'_C is a best μ-response for C from σ and $\sigma'_i = \mathcal{F}(\sigma_{-C}, \sigma'_C)$ for every $i \in C$.*

Remark that when the coalitions are singletons, a unanimous direct best {IM,WIM,CIM}-response corresponds to a direct best response.

Let us now provide a counterexample for the convergence of the dynamics of the game defined according to a partition equilibrium. Although originally, the partition equilibrium is defined for weak improving moves, we consider an instance with strict preferences therefore in this case WIM=IM.

Proposition 2.17 *The dynamics associated with the partition equilibrium may not converge in the voting game with $\mathcal{F}$=Plurality, even if the initial voting profile is truthful, each deviation is a unanimous direct best IM-response, and the preferences are strict, single-peaked and single-crossing.*

Proof: Let us consider an instance with twelve voters and four candidates, where $N = \{1, 2, \ldots, 12\}$, $M = \{a, b, c, d\}$, $a \rhd b \rhd c \rhd d$, and $\mathcal{F}$=Plurality. The profile of preferences is:

$$\begin{aligned} 1,\ 2,\ 3: &\quad d \succ c \succ b \succ a \\ 4,\ 5,\ 6: &\quad c \succ d \succ b \succ a \\ 7,\ 8,\ 9: &\quad a \succ b \succ c \succ d \\ 10,\ 11,\ 12: &\quad b \succ c \succ d \succ a \end{aligned}$$

The preferences are single-peaked with respect to the order over the candidates $a >^M b >^M c >^M d$ and single-crossing with respect to the order over the voters $1 >^N 2 >^N 3 >^N 4 >^N 5 >^N 6 >^N 10 >^N 11 >^N 12 >^N 7 >^N 8 >^N 9$. The partition over the voters is $\{\{1\},\{2\},\{3\},\{4\},\{5\},\{6\},\{7,8,9\},\{10,11\},\{12\}\}$. The next table gives a sequence of states where the first and last ones coincide. The deviations are indicated with arrows and bold letters, and step 0 is the truthful profile.

Steps	σ^0		σ^1		σ^2		σ^3		σ^4		σ^5		σ^6
1, 2, 3:	d		d		d		d		d		d		d
4:	c	$\rightarrow$	**d**		d		d	$\rightarrow$	**c**		c		c
5, 6:	c		c		c		c		c		c		c
7, 8, 9:	a		a		a	$\rightarrow$	**b**		b		b	$\rightarrow$	**a**
10, 11:	b		b	$\rightarrow$	**c**		c		c	$\rightarrow$	**b**		b
12:	b		b		b		b		b		b		b
$\mathcal{F}(\sigma^t)$	a		d		c		b		c		b		a

One can easily verify that each deviation is a unanimous direct best IM-response. □

The previous proposition rules out the possibility of convergence of dynamics defined by a considerate equilibrium since the coalitions given by the members of a partition can be transposed into the maximal cliques of a cluster graph.

Interestingly, Proposition 2.17 can be mitigated if we consider a special partition of N where all coalitions have the same size. The result follows from a simple extension of a proof given by Meir et al. [2010].

Proposition 2.18 *If $\mathcal{P}$ is a partition of N such that all coalitions of $\mathcal{P}$ have the same size, then the dynamics associated with a $(\mathcal{P}, WIM)$-equilibrium, i.e., a partition equilibrium, converges for any instance of the voting game and any initial profile σ^0, if the deviations are unanimous direct best WIM-responses.*

Proof: We consider each group of the partition as a meta-agent. The preferences of such a meta-agent i representing group $C \in \mathcal{P}$ are such that $a \succ_i b$ only if $a \geq_C b$ (we do not express the rest of the preferences which could be partial because we only care about the best response deviations). Then, it suffices to follow the proof of Meir et al. [2010]'s Theorem 3. □

Plurality is not Condorcet consistent, however a unanimous voting profile s_x, where x is the Condorcet winner, is always a considerate equilibrium (as stated by Sertel and Sanver [2004] for strong equilibria) since no absolute majority of voters prefers another alternative and Plurality is a strict majority susceptible rule. Nevertheless, when a

Condorcet winner exists, the game may not converge to a state electing it, because even dynamics associated with a Nash equilibrium (corresponding to a considerate equilibrium where the graph is empty) may not converge to a state electing it, even from a truthful initial profile and for direct best responses.

Example 2.10 *Let us consider an instance with three voters and four candidates, where* $N = \{1, 2, 3\}$*,* $M = \{a, b, c, d\}$*, and* $\mathcal{F}$*=Plurality. The profile of preferences is:*

$$\begin{aligned} 1: &\quad a \succ b \succ c \succ d \\ 2: &\quad c \succ b \succ d \succ a \\ 3: &\quad d \succ b \succ a \succ c \end{aligned}$$

By considering unilateral deviations conditioned by direct best responses, the only possible deviation from truthful profile $\sigma^0 = (a, c, d)$ *electing* a*, is the deviation of agent 2 to ballot* d*. We then reach state* (a, d, d)*, electing* d*, which is a Nash equilibrium whereas* b *is the Condorcet winner.*

In a nutshell, while the dynamics associated with unilateral deviations that are direct best responses always converges to a Nash equilibrium, the straightforward generalization to coalitional deviations that are direct and unanimous for coalitions of voters defined by a partition over the voters make the convergence of the dynamics fail.

2.5.1.b Veto

Contrary to Plurality, a ballot under Veto does not describe a support for a candidate but a veto against a candidate. Consequently, the fact that all the members of a same coalition vetoes the same alternative may not be the best strategy for the coalition. When the coalition is a single agent, this is clear that a best response can be to directly veto the current winner (see Definition 1.28 for the direct best response under Veto). However, for a larger coalition, it could be more advantageous, for instance, to veto the current winner but also other candidates, as we can see in the following example.

Example 2.11 *Let us consider an instance with seven voters and four candidates, where* $N = \{1, \ldots, 7\}$*,* $M = \{a, b, c, d\}$*,* $a \rhd b \rhd c \rhd d$*, and* $\mathcal{F}$*=Veto. The preferences of the agents are described as follows.*

$$\begin{aligned} 1: &\quad b \succ d \succ c \succ a \\ 2: &\quad d \succ a \succ c \succ b \\ 3: &\quad a \succ b \succ d \succ c \\ 4,\ 5: &\quad c \succ b \succ a \succ d \\ 6,\ 7: &\quad a \succ b \succ c \succ d \end{aligned}$$

Let initial state σ^0 *be the truthful profile. In* σ^0*,* a *is elected thanks to the tie-breaking rule. Let us focus on coalition of voters* $C := \{4, 5\}$*. If coalition* C *deviates within a unanimous strategy, then the members of* C *must both veto candidate* a*, in order to make the outcome of the election change. This deviation leads to the election of* b *(thanks to* *the tie-breaking rule). Agents* 4 *and* 5 *both prefer* b *to* a*, therefore this deviation is an* IM*-response for coalition* C *from* σ^0*.*

However, by removing the constraint of a unanimous strategy for coalition C*, this coalition can agree on a joint strategy where agent* 4 *vetoes candidate* a *while agent* 5 *vetoes candidate* b*. Such a deviation leads to the election of* c*, which is the best candidate of agents* 4 *and* 5*. Therefore, this is a best* IM*-response for coalition* C *from* σ^0*.*

Hence, restricting to unanimous strategies under Veto may not allow the voters to design best deviations.

Nevertheless, as for Plurality, the dynamics associated with the considerate equilibrium is not guaranteed to converge, even if we restrict ourselves to the dynamics associated with the partition equilibrium and we consider best IM-responses.

Proposition 2.19 *The dynamics associated with the partition equilibrium may not converge in the voting game with* $\mathcal{F}$=*Veto, even if the initial profile is truthful, each deviation is a best IM-response, and the preferences are strict, single-peaked and single-crossing.*

Proof: Let us consider an instance with six voters and four candidates, where $N = \{1, 2, \ldots, 6\}$, $M = \{a, b, c, d\}$, $a \rhd b \rhd c \rhd d$, and $\mathcal{F}$=Veto. The preferences are single-peaked with respect to the order over the candidates $c >^M d >^M a >^M b$ and single-crossing with respect to the order over the voters $2 >^N 1 >^N 6 >^N 3 >^N 4 >^N 5$. The profile of preferences is:

$$\begin{aligned} 1:&\quad d \succ c \succ a \succ b\\ 2:&\quad c \succ d \succ a \succ b\\ 3:&\quad a \succ d \succ b \succ c\\ 4, 5:&\quad b \succ a \succ d \succ c\\ 6:&\quad a \succ d \succ c \succ b \end{aligned}$$

The partition over the voters is $\{\{1\}, \{2\}, \{3\}, \{4, 5\}, \{6\}\}$. The next table gives a sequence of states where the steps 1 to 7 form a cycle. The ballots are written in the form $-x$, for a candidate x, in order to represent the fact that this ballot is a veto against x. The deviations are marked with arrows and bold letters. Step 0 corresponds to the truthful profile.

Steps	σ^0		σ^1		σ^2		σ^3		σ^4		σ^5		σ^6		σ^7
1:	$-b$		$-b$		$-b$		$-b$		$-b$		$-b$		$-b$		$-b$
2:	$-b$	$\rightarrow$	$-\mathbf{a}$		$-a$		$-a$		$-a$	$\rightarrow$	$-\mathbf{d}$		$-d$	$\rightarrow$	$-\mathbf{a}$
3:	$-c$		$-c$		$-c$		$-c$		$-c$		$-c$		$-c$		$-c$
4, 5:	$-c$		$-c$		$-c$	$\rightarrow$	$-\mathbf{a}$		$-a$		$-a$	$\rightarrow$	$-\mathbf{c}$		$-c$
6:	$-b$		$-b$	$\rightarrow$	$-\mathbf{d}$		$-d$	$\rightarrow$	$-\mathbf{b}$		$-b$		$-b$		$-b$
$\mathcal{F}(\sigma^t)$	a		d		a		b		d		c		a		d

One can easily verify that each deviation is a best IM-response. In particular, these deviations are also unanimous and direct, in the sense that all the members of the coalitions deviate to the same ballot vetoing the current winner. □

While general unanimous μ-responses, where every member of a coalition deviates to the same ballot, are not necessarily best μ-responses (Example 2.11), the counterexample provided in the previous proof exhibits deviations that are unanimous and direct (by vetoing the current winner). Therefore, the dynamics associated with a partition equilibrium may fail to converge even for a straightforward generalization of the direct best response under Veto (Definition 1.28), which guarantees the convergence of the dynamics for unilateral deviations.

2.5.2 Non-convergent rules for single-agent deviations

Besides Plurality and Veto, we have the existence of a considerate equilibrium for other voting rules, such as Maximin, STV and Plurality with run-off. However, the dynamics of the game may not converge to a considerate equilibrium for these voting rules, since it may not even converge to a Nash equilibrium, as we will show in this subsection. Moreover, this is still the case for restricted best responses.

Focusing, in this subsection, on unilateral deviations, we consider dynamics based on best responses as defined in Section 1.3.3.b (Definition 1.26). Nevertheless, we strengthen our negative results by using specific best responses minimizing a certain distance to the sincere ballot.

Proposition 2.20 *The dynamics associated with the Nash equilibrium is not guaranteed to converge in the voting game with $\mathcal{F} \in \{STV$, Plurality with run-off$\}$, even if the initial profile is truthful, the voters' preferences are strict, single-peaked and single-crossing, and each move is a best response minimizing the distance to the sincere ballot in terms of number of differences in pairwise comparisons.*

Proof: Let us consider an instance with four voters and four candidates, where $N = \{1, 2, 3, 4\}$, $M = \{a, b, c, d\}$, $a \rhd b \rhd c \rhd d$, and $\mathcal{F} \in \{$STV, Plurality with run-off$\}$. The preferences are single-peaked with respect to the order over the candidates $a >^M c >^M d >^M b$ and single-crossing with respect to the order over the voters $2 >^N 3 >^N 4 >^N 1$. The next table gives a sequence of states where the first and last ones coincide. Deviations are marked with bold letters. Step 0 represents the truthful profile. Each linear order is the ballot of a voter and the last line of the table specifies the winner at each step.

	σ^0	σ^1	σ^2	σ^3	σ^4
1:	$c \succ a \succ d \succ b$	$\mathbf{a} \succ \mathbf{c} \succ \mathbf{d} \succ \mathbf{b}$	$a \succ c \succ d \succ b$	$\mathbf{c} \succ \mathbf{a} \succ \mathbf{d} \succ \mathbf{b}$	$c \succ a \succ d \succ b$
2:	$b \succ d \succ c \succ a$	$b \succ d \succ c \succ a$	$\mathbf{d} \succ \mathbf{b} \succ \mathbf{c} \succ \mathbf{a}$	$d \succ b \succ c \succ a$	$\mathbf{b} \succ \mathbf{d} \succ \mathbf{c} \succ \mathbf{a}$
3:	$d \succ b \succ c \succ a$	$d \succ b \succ c \succ a$	$d \succ b \succ c \succ a$	$d \succ b \succ c \succ a$	$d \succ b \succ c \succ a$
4:	$c \succ d \succ a \succ b$	$c \succ d \succ a \succ b$	$c \succ d \succ a \succ b$	$c \succ d \succ a \succ b$	$c \succ d \succ a \succ b$
$\mathcal{F}(\sigma^t)$	b	a	d	c	b

This counterexample works for STV and Plurality with run-off. □

It is known that the strategic voting game under the Maximin rule is not guaranteed to converge to a Nash equilibrium with an arbitrary deterministic tie-breaking [Lev and

Rosenschein, 2012]. However, even with a deterministic tie-breaking that is a linear order, the game under Maximin is not guaranteed to converge, as we can see in the next proposition. Moreover, we do not use a general best response but a restricted one.

Proposition 2.21 *The dynamics associated with the Nash equilibrium is not guaranteed to converge in the voting game with $\mathcal{F}$=Maximin, even if the initial profile is truthful, the voters' preferences are strict, single-peaked and single-crossing, and each move is a best response minimizing the distance to the sincere ballot in terms of number of differences in pairwise comparisons.*

Proof: Let us consider an instance with five voters and four candidates, where $N = \{1, 2, 3, 4, 5\}$, $M = \{a, b, c, d\}$, $a \rhd b \rhd c \rhd d$ and $\mathcal{F}$=Maximin. The preferences are single-peaked with respect to the order over the candidates $d >^M b >^M c >^M a$ and single-crossing with respect to the order over the voters $1 >^N 4 >^N 3 >^N 2 >^N 5$. The next table gives a sequence of states where the steps 1 to 5 form a cycle. Deviations are marked with bold letters. Step 0 corresponds to the truthful profile. Each linear order is the ballot of a voter and the last line of the table specifies the winner at each step.

	σ^0	σ^1	σ^2	σ^3	σ^4	σ^5
1:	$d \succ b \succ c \succ a$	$d \succ b \succ c \succ a$	$\mathbf{c} \succ \mathbf{d} \succ \mathbf{b} \succ \mathbf{a}$	$c \succ d \succ b \succ a$	$\mathbf{d} \succ \mathbf{b} \succ \mathbf{c} \succ \mathbf{a}$	$d \succ b \succ c \succ a$
2:	$a \succ c \succ b \succ d$	$\mathbf{a} \succ \mathbf{d} \succ \mathbf{c} \succ \mathbf{b}$	$a \succ d \succ c \succ b$	$a \succ d \succ c \succ b$	$a \succ d \succ c \succ b$	$a \succ d \succ c \succ b$
3:	$c \succ b \succ a \succ d$	$c \succ b \succ a \succ d$	$c \succ b \succ a \succ d$	$c \succ b \succ a \succ d$	$c \succ b \succ a \succ d$	$c \succ b \succ a \succ d$
4:	$b \succ d \succ c \succ a$	$b \succ d \succ c \succ a$	$b \succ d \succ c \succ a$	$b \succ d \succ c \succ a$	$b \succ d \succ c \succ a$	$b \succ d \succ c \succ a$
5:	$a \succ c \succ b \succ d$	$a \succ c \succ b \succ d$	$a \succ c \succ b \succ d$	$\mathbf{a} \succ \mathbf{b} \succ \mathbf{d} \succ \mathbf{c}$	$a \succ b \succ d \succ c$	$\mathbf{a} \succ \mathbf{c} \succ \mathbf{b} \succ \mathbf{d}$
$\mathcal{F}(\sigma^t)$	c	a	c	a	b	a

□

Considering the very restrictive conditions of the previous propositions, it seems that no convergence guarantee could be established based on best response dynamics for these voting rules. Consequently, an option would be to restrict to specific manipulation moves that are not necessarily best responses [Reijngoud and Endriss, 2012, Grandi et al., 2013, Obraztsova et al., 2015].

2.6 Experiments

We run different experiments in order to evaluate the quality of considerate equilibria. The quality of the equilibria that are reached via a dynamics in iterative voting is analyzed as well as the number of considerate equilibria for a given instance.

Let us clarify the deviations of the voters that we have considered. Best μ-responses (Definition 2.13) are used for both Plurality and Veto, for $\mu \in$ {IM,WIM,CIM} defined according to the type of solution concept that is considered. More precisely, we have used unanimous direct best μ-responses for Plurality in order to choose a natural move among the possible best μ-responses. However, we do not make this restriction for Veto since the best unanimous direct responses are not necessarily best μ-responses (Example 2.11), and we consider any best μ-response. However, for the other voting rules, we restrict ourselves to unanimous μ-responses that are the best, for the sake of computational simplicity, leading to take into account equilibria that are not necessarily immune to any form of deviation.

2.6.1 Number of equilibria

First of all, we focus on the number of considerate equilibria according to the density of the social network. We run a first set of experiments where, given 7 voters and 3 candidates, we generate all possible voting profiles (in total $(3!)^7$ voting profiles). Then, by randomly generating 100 instances with strict preference profiles drawn from impartial culture and Erdös-Rényi graphs of density 0 (=empty), 0.25, 0.5, 0.75 and 1 (=complete), we count how many voting profiles are considerate equilibria under different voting rules. We choose seven voting rules that are theoretically analyzed in the previous sections, namely Plurality, Borda, Maximin, Copeland, STV, Veto and 2-approval. The proportion of voting profiles that are in fact considerate equilibria are presented in Figure 2.3.

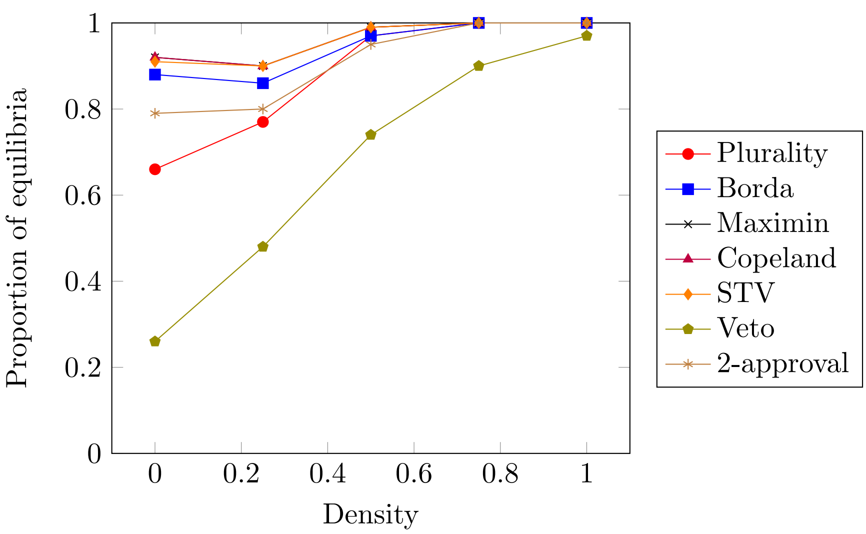

Figure 2.3: Proportion of considerate equilibria when $n = 7$ and $m = 3$ under impartial culture

The main observation is that considerate equilibria have the same drawback as Nash equilibria, that is the number of equilibria is terribly huge. Indeed, for all the voting rules except Veto, more than half of the possible profiles is a considerate equilibrium. The proportion is even higher than 0.8 as soon as the density of the graph is more than 0.25. Globally, the number of considerate equilibria increases with the density of the graph, which makes sense since the coalitions of manipulators have more neighbors to take care. Indeed, when the graph is complete, almost all the voting profiles are considerate equilibria. Regarding the voting rules, it is interesting to remark that Maximin, Copeland and STV seem to behave the same way: they have almost the same proportion of equilibria for each density. Together with Borda, they have a large number of considerate equilibria. Nevertheless, given the theoretical results regarding existence of equilibria, especially for Borda and Copeland, it seems that these results are strongly biased by the fact that we only consider unanimous deviations for these voting rules. However, note that the Nash equilibria, when the density is equal to zero, are not concerned by the restriction to unanimous responses, but that the number of Nash equilibria for STV, Copeland, Maximin and Borda are nevertheless important. Observe that there are significantly less considerate equilibria under Veto than for the other rules whereas, contrary to Borda, Veto always guarantees the existence of a considerate equilibrium. By taking into account any possible deviation, we might expect a behavior for Borda closer to the behavior of the Veto rule. In the same idea, the lower number of equilibria for Veto can be explained by the fact that the voters can elaborate more sophisticated strategies than in Plurality where the only option is to give points to a target candidate. Observe also that the number of considerate equilibria in general is even higher than the number of Nash equilibria (case where the density is equal to zero in the results), although known to be problematically high.

These results highlight a main drawback of the considerate equilibrium, that is due to the consideration assumption. Indeed, this assumption forces the possible deviating coalitions to take into consideration the preferences of their neighbors by avoiding any harmful deviation. This blocks the deviation of many coalitions and thus creates many equilibria.

In order to understand the importance of the consideration assumption, we run the same experiments but this time we count the proportion of "considerate equilibria" without consideration assumption, that are coalitional equilibria where the coalitions are given by the cliques of the network but with no other restriction. The results of this new set of experiments are given in Figure 2.4.

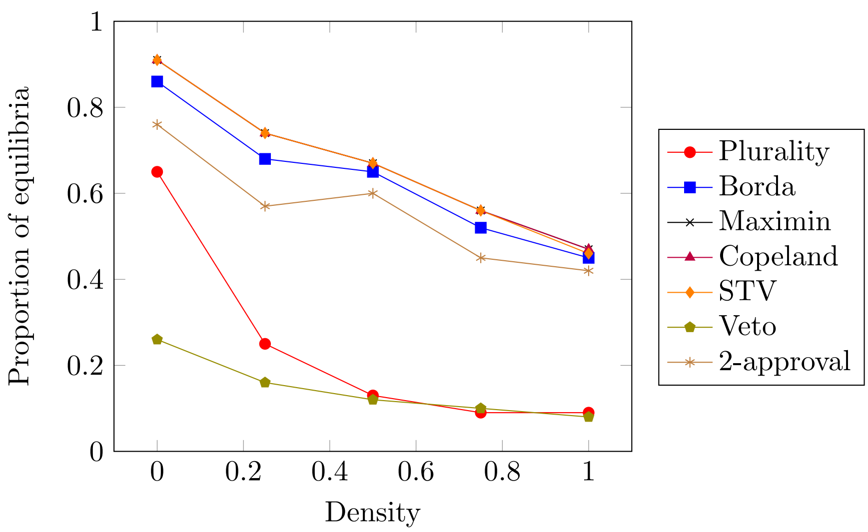

Figure 2.4: Proportion of coalitional equilibria (determined by cliques of the network) without consideration when $n = 7$ and $m = 3$ under impartial culture

Starting from the same point, i.e., the number of Nash equilibria, when the density is null, the curves are totally inversed. Whereas the number of considerate equilibria increases with the density of the graph, the number of coalitional equilibria without consideration decreases when the density increases. This also seems natural since cliques, and consequently coalitions, are bigger when the graph is dense. For Plurality and Veto, when the density is higher than 0.5, the number of equilibria is extremely low. However, for the other voting rules, the number of equilibria is around the half of the states even when the graph is complete (note that the curves associated with Maximin, Copeland and STV are confused). Once again, this is clearly due to the restriction to unanimous deviations, which surely leads us to conclude that some states are equilibria whereas they are not.

Globally, the number of equilibria seems more interesting for filtering the good states of the games when the consideration assumption is removed. However, by removing the consideration assumption, we lose the existence property, see for instance Example 2.8, which is also a main disadvantage.

A good compromise could be the partition equilibrium. Indeed, every time a considerate equilibrium is guaranteed to exist for any social network, a partition equilibrium is also guaranteed to exist. Moreover, there is no overlap between the coalitions and no consideration assumption, so one could expect a more reasonable number of equilibria. Therefore, we also count the number of partition equilibria for 7 voters and 3 candidates where the partitions are generated randomly. We compare these results with the number of considerate equilibria and coalitional equilibria without consideration in graphs with homophily (see Section 1.2.2.b), which is a graph constructed according to the preferences of the agents. The results are presented in Figure 2.5.

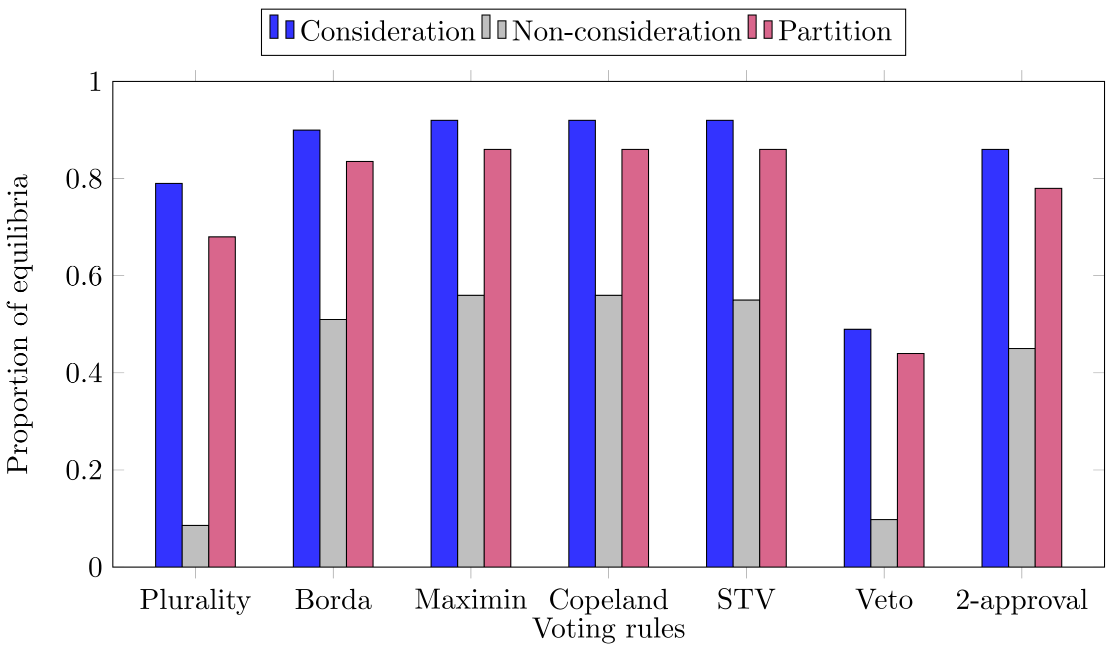

Figure 2.5: Proportion of equilibria when $n = 7$, $m = 3$ and coalitions are given by the cliques of a graph respecting homophily or by a partition

As expected, the number of partition equilibria is in between the number of considerate equilibria and the number of coalitional equilibria without consideration. However, although smaller than the number of considerate equilibria, this number remains too large to discriminate the states.

Nevertheless, investigating the quality of the states that can be actually reached by an iterative voting process enables us to focus on more plausible outcomes, despite the huge number of equilibria.

2.6.2 Convergence to equilibria

Let us consider now best response dynamics of iterative voting and analyze their convergence rate to equilibria.

We assume that the initial state is truthful and we use a random turn function (see Section 1.3.3.b) for choosing the coalition which deviates if several coalitions have incentive to manipulate at a given step of the game.

First of all, in order to have a better idea of the actual possibility of considerate deviation, we count the instances for which a considerate deviation can occur from the initial truthful state ("effective" considerate manipulation), i.e., the proportion of instances for which the truthful state is not a considerate equilibrium. Concretely, for 5 candidates and different number of voters we look at the effective considerate manipulation in an Erdös-Rényi graph of density 0.25 (where there are less considerate equilibria than for higher densities) for 10,000 generated preference profiles under impartial culture. We focus on the seven voting rules used in the previous subsection: Plurality, Borda, Maximin, Copeland, STV, Veto and 2-approval. The results are presented in Figure 2.6.

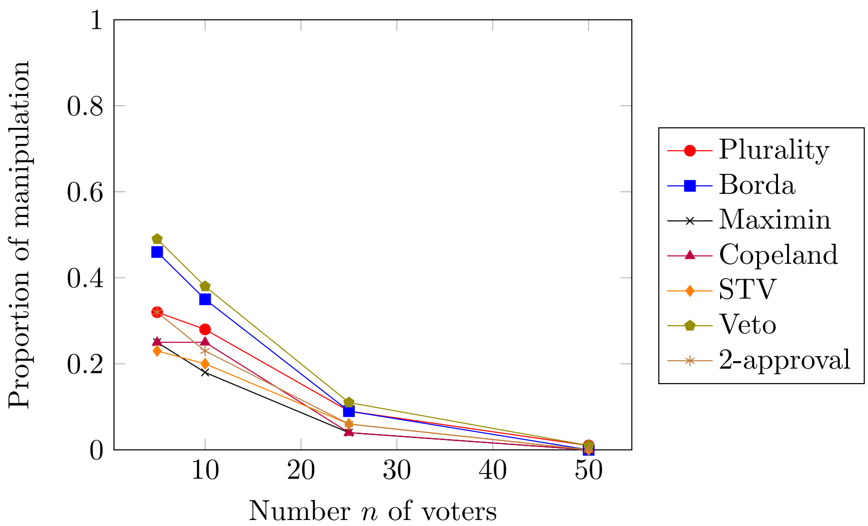

Figure 2.6: Effective considerate manipulation when $m = 5$ under impartial culture and Erdös-Rényi graphs of density 0.25

We observe that the effective manipulation, i.e., the proportion of instances for which there is actually a considerate deviation from the initial truthful state, is extremely low and sharply decreases when the number of voters increases. Indeed, for a number of voters larger than 25, effective considerate manipulation is already almost null and is inexistent for 50 voters. This brings us to consider a small number of voters for our experiments, because if there is no manipulation then the reachable considerate equilibrium under evaluation would be in fact the truthful state, and thus the study of the convergence of the dynamics would be biased by this fact.

Consequently, from now on, we focus on a setting with 10 voters. According to the large number of equilibria when the graph is dense, the question of effective manipulation according to the density of the graph is also meaningful. Therefore, for 10 voters and 5 candidates, and Erdös-Rényi graphs of different densities, we explore the effective considerate manipulation of 10,000 generated instances of a linked voting game for preferences drawn from impartial culture. The results are presented in Figure 2.7.

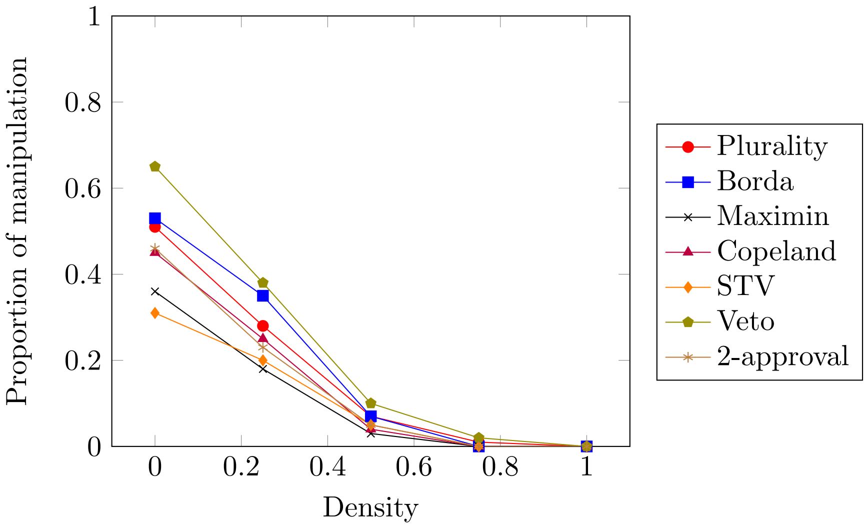

Figure 2.7: Effective considerate manipulation when $n = 10$ and $m = 5$ under impartial culture

In accordance with the results on the number of considerate equilibria, effective manipulation drastically decreases with the density of the graph since it is already almost null for a graph of density 0.5. Because of that, from now on, we will focus on graphs of density 0.25 or graphs that respect homophily. Recall that the empty graph corresponds to unilateral deviations defining a Nash equilibrium.

Since the low level of effective manipulation seems to be due to the consideration assumption, we also explore, like in the previous subsection, dynamics associated with coalitional deviations without consideration assumption, where the coalitions are the cliques of a social network, or where the coalitions are given by a partition, defining a partition equilibrium. We thus investigate the convergence rate of the dynamics for a Nash equilibrium ("NE"), a partition equilibrium ("PE"), and a coalitional equilibrium with and without consideration for graphs of density 0.25 ("CE (ER 0.25)" and "$\overline{\text{C}}$E (ER 0.25)", respectively) or graphs that respect homophily ("CE (homophily)" and "$\overline{\text{C}}$E (homophily)", respectively). We present the results in Figure 2.8 for the rules Plurality, Borda, Maximin, STV and Veto, where we run 10,000 instances with 10 voters and 5 candidates, in which the preferences of the agents are drawn from impartial culture. The partitions over the agents for the study of the partition equilibrium are generated randomly, and the graphs that respect homophily are generated according to the preferences of the agents, following the protocol described in Section 1.2.2.b.

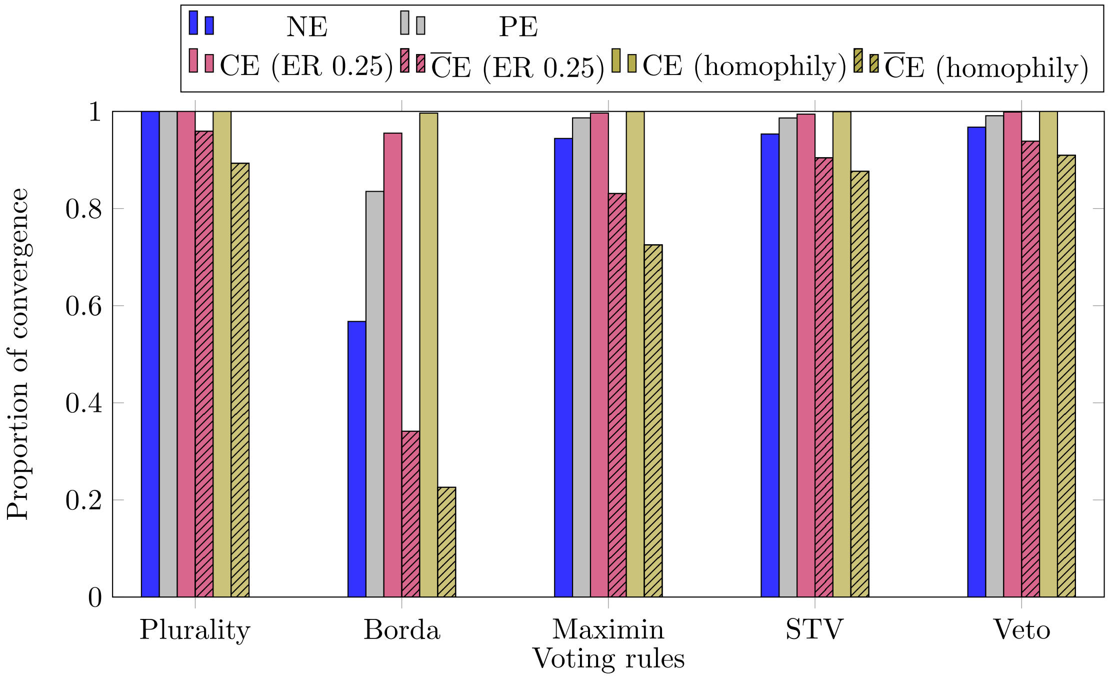

Figure 2.8: Proportion of convergence to equilibria when $n = 10$ and $m = 5$ under impartial culture

Globally, the convergence rate is very good. Indeed, almost all the dynamics converges for more than 80% of the cases. Only the dynamics under Borda has a low convergence rate. This may be due to the high susceptibility to manipulation of the Borda rule [Durand, 2015], for which even a partition equilibrium is not guaranteed to exist (Proposition 2.13). A good point for the consideration assumption is that the convergence rate is always clearly better than without consideration. Similarly, convergence occurs more frequently for considerate dynamics than for partition dynamics which also occurs more frequently than for Nash dynamics.

A question that arises from the analysis of the convergence ratio is the number of steps that is needed before reaching convergence. For the same set of experiments we provide the average number of steps before convergence when convergence occurs. The results are presented in Figure 2.9.

One can observe that the number of steps before reaching convergence is very low. In fact, in average, we always need less than 5 steps to converge. Coalitional deviations without consideration always take more steps to converge than deviations with consideration. This is surely due to the fact that coalitional deviations are rarer when the consideration assumption applies. As for convergence, the number of steps before reaching convergence is higher for Nash deviations than for partition deviations which

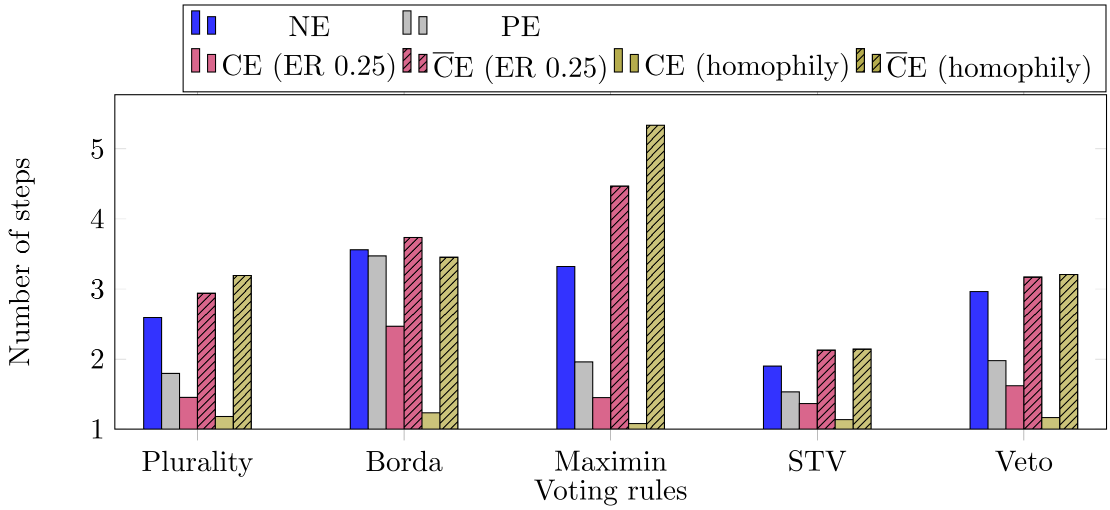

Figure 2.9: Average number of steps for convergence when $n = 10$ and $m = 5$ under impartial culture

is also higher than for considerate deviations.

2.6.3 Quality of the equilibria

Let us now explore the quality of the reachable equilibria in itself. With the same parameters for our experiments but with in addition the results of the truthful profile ("No manip."), we investigate the Condorcet efficiency (see Section 1.3.2), the Veto-SE efficiency (see Section 2.4.2), and the Borda closeness (see Section 1.3.2) of the equilibria that are reached by the dynamics. Concerning Condorcet efficiency, we restrict ourselves to instances where a Condorcet winner exists. The results for the Condorcet efficiency are presented in Figure 2.10.

In general, the Condorcet efficiency of the equilibria is good, especially for Maximin and STV. A very good point for the considerate equilibrium is that the Condorcet efficiency of its reachable equilibria is almost always better than without the consideration assumption (except for Plurality and STV with a small gap). The gap is particularly important for the Borda rule. In particular, for Borda and Veto, the considerate equilibrium is better than the partition equilibrium which is also better than the Nash equilibrium. For Maximin and STV, the gap is not very significant. Moreover, surprisingly, Condorcet efficiency is better without consideration for Plurality, which would imply that the absence of altruism in local improvements leads to better outcomes for the whole society at the end, under Plurality. However, the gap is not sufficiently large to conclude. Furthermore, except for Borda and Maximin without consideration, the iterative voting process enables to improve the Condorcet efficiency of the outcome, compared to the truthful outcome, in accordance with the conclusions of Grandi et al. [2013]. This improvement is particularly significant for Plurality where Condorcet efficiency is increased by around 10% in the reachable equilibria, compared to the truthful profile. This may due to the fact that the agents give more information about their preferences during

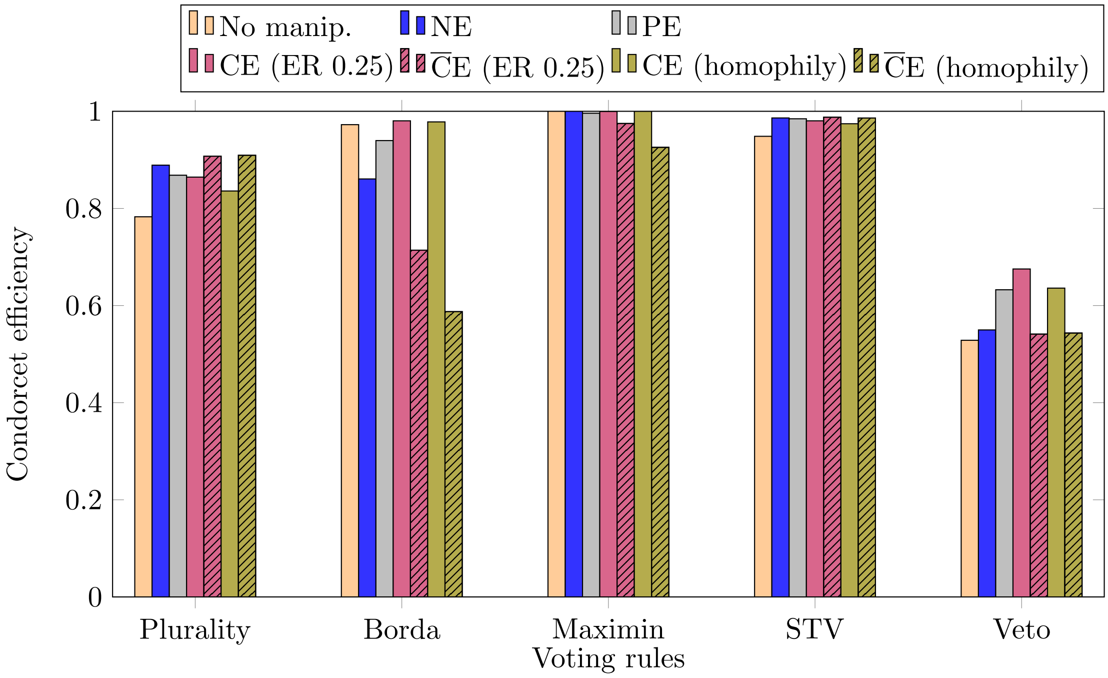

Figure 2.10: Condorcet efficiency of reachable equilibria when $n = 10$ and $m = 5$ under impartial culture

their deviations than in the Plurality ballot. The same explanation holds for Veto, for which the iterative version always does better regarding Condorcet efficiency than the truthful profile, although with smaller gaps.

The results for Veto-SE efficiency are presented in Figure 2.11. The conclusions are almost the same as for the Condorcet efficiency, especially concerning the fact that reachable considerate equilibria provide better results than coalitional equilibria without the consideration assumption (except for Plurality too). The gap is also more important for the Borda rule. Nevertheless, except for Borda, the gaps are not significant.

Let us now talk about the Borda closeness of the reachable equilibria. The results are presented in Figure 2.12.

In this case, the considerate equilibria provide the best results, the only exception is naturally the truthful outcome of the Borda rule. Like for the other criteria, the gap of improvement compared to the version without consideration is more important for Borda. But for any voting rule, the considerate equilibria are clearly better than the coalitional equilibria without consideration. Similarly to the other criteria, the reachable considerate equilibrium achieves a better Borda closeness than the partition equilibrium which is also better than the Nash equilibrium. Like for the other quality measures, Borda closeness is improved with the iterative version compared to the truthful state for Plurality. However, for the other voting rules, the truthful outcome is always better regarding Borda closeness than the reachable Nash equilibria and the reachable coalitional equilibria without consideration.

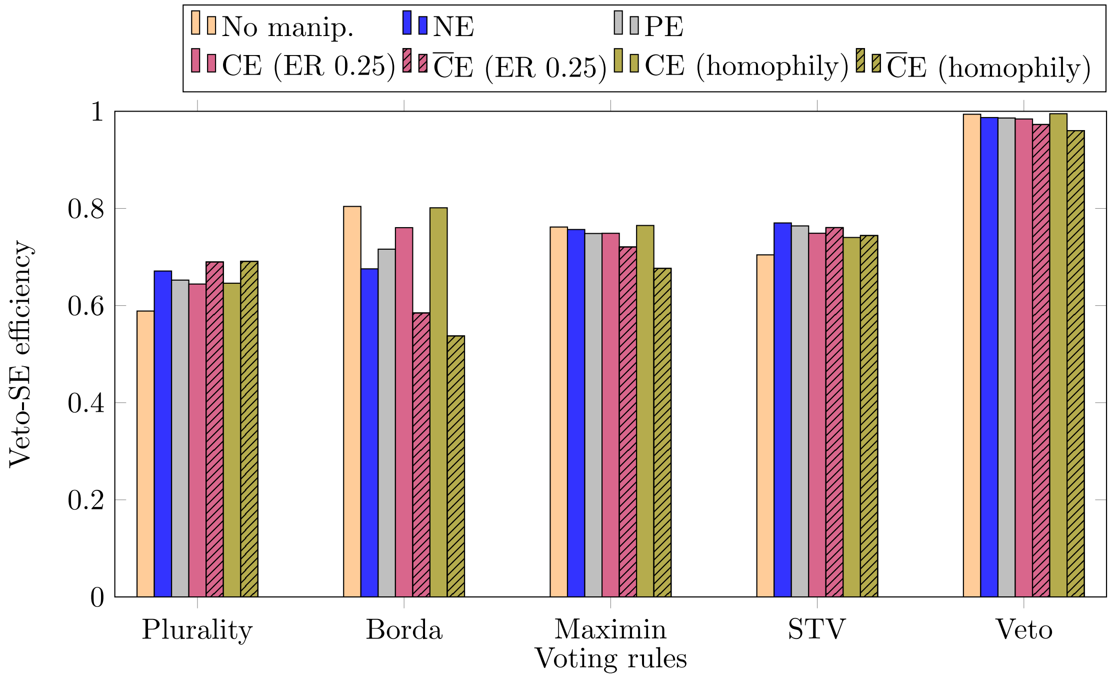

Figure 2.11: Veto-SE efficiency of reachable equilibria when $n = 10$ and $m = 5$ under impartial culture

2.7 Concluding remarks

We have proposed to explore voting games from a strategic and social point of view. The considerate equilibrium captures a stable outcome regarding every set of coalitions arising from a social network. Indeed, connected agents in the social network can collaborate and establish common strategies for manipulating the outcome of the election.

Table 2.1 summarizes the existence and convergence results for the different voting rules that we have studied. They are classified by solution concepts, which allows to spot the gap between existence/non-existence and convergence/non-convergence since the different equilibria are related (if there exists a considerate equilibrium for any social network, then there exists a partition equilibrium for any partition over the agents, and then there exists a Nash equilibrium, as stated in (2.2)).

We can remark that convergence under best responses rarely occurs, even for classical best responses defined by unilateral deviations (even if they are refined as in Propositions 2.20 and 2.21), and for best responses restricted to unanimous direct deviations for Plurality and Veto (Propositions 2.17 and 2.19). A possible extension would be to analyze the convergence to a considerate equilibrium for some restricted manipulation moves, as studied for Nash equilibria [Reijngoud and Endriss, 2012, Grandi et al., 2013, Obraztsova et al., 2015]. Another possible way to achieve convergence would be to focus on specific classes of graphs or coalition families, provided that it matches with a realistic social structure.

Despite these negative theoretical results regarding convergence, our experimental

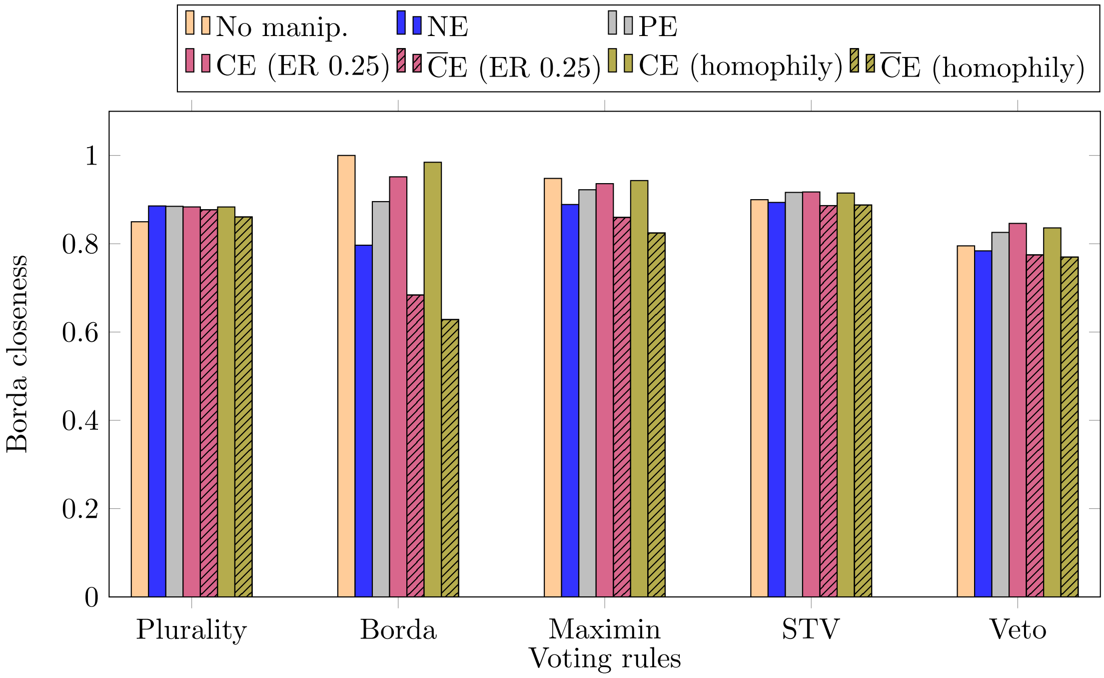

Figure 2.12: Average of Borda closeness of reachable equilibria when $n = 10$ and $m = 5$ under impartial culture

studies show that the cycles in the considerate dynamics are very rare (Figure 2.8). Moreover, when convergence occurs, the process ends very quickly after only few steps.

Beyond convergence, we were able to prove the existence of a considerate equilibrium for a significant number of voting rules, namely Plurality, Veto, Plurality with run-off, STV and Maximin. These results are encouraging because the notion of considerate equilibrium covers a large spectrum of families of coalitions. Moreover, we were able to prove a stronger result for the Veto rule, since a strong equilibrium is always guaranteed to exist. One can check in polynomial time whether a given candidate can be elected in a strong equilibrium under Veto, property that we call *Veto-SE*. Therefore, it is possible to derive, from the Veto-SE property, a quality measure for the outcome of a voting rule, that can be applied for any domain of preferences, and that is complementary to the Condorcet consistency/efficiency.

As a balance for the existence of a considerate equilibrium, the assumption of consideration within this equilibrium — the fact that no coalition harms its neighbors — is rather strong. Actually, without the consideration assumption, it is not possible to generalize the existence of such an equilibrium to every class of graphs. As an example, take the complete graph, for which the associated solution concept without consideration would correspond to a super strong equilibrium (see Example 2.8). This counterexample holds for every voting rule for which we have proved the existence of a considerate equilibrium, showing the importance of the consideration assumption.

The consideration assumption prevents many deviations to occur. This notably induces a huge number of considerate equilibria for a given instance (Figures 2.3 and

		PSRs				Run-off rules		Pairwise comparisons rules	
		Plurality	Veto	k-approval	Borda	STV	PwRO	Copeland	Maximin
strong	Existence	✗	✔	✗	✗	✗	✗	✗	✗
	Convergence	✗	✗	✗	✗	✗	✗	✗	✗
considerate	Existence	✔	✔	✗	✗	✔	✔	✗	✔
	Convergence	✗	✗	✗	✗	✗	✗	✗	✗
partition	Existence	✔	✔	✗	✗	✔	✔	✗	✔
	Convergence	✗	✗	✗	✗	✗	✗	✗	✗
Nash	Existence	✔	✔	✗	✗	✔	✔	✗	✔
	Convergence	✔[1]	✔[2]	✗[2]	✗[2]	✗	✗	✗	✗

[1] [Meir et al., 2010]
[2] [Lev and Rosenschein, 2012, Reyhani and Wilson, 2012]

Table 2.1: Global results on the existence of different solution concepts and convergence of the associated dynamics in voting games under different voting rules

2.5). This large number of equilibria is the main drawback of the considerate equilibrium: since for some instances almost all states are stable regarding considerate deviations, this solution concept is not sufficient to filter good voting profiles or to predict the plausible outcomes of the voting game. However, in the iterative voting perspective where the voters deviate by rounds, the reachable considerate equilibria are qualitatively good. In fact, they are clearly better regarding the Condorcet efficiency, the Veto-SE efficiency or the Borda closeness criteria, than other coalitional equilibria, such as equilibria without the consideration assumption. Therefore, although taking an arbitrary considerate equilibrium is not particularly relevant, a considerate equilibrium that is reached within an iterative voting process is really meaningful.

Moreover, the consideration assumption is relevant if we assume that the agents are not fully selfish and that they care about their relatives. It notably allows to integrate a social dimension into the voting game. If agents are connected in the social network, then they can be reluctant to act in a way that harm their partners.

If one wants to escape from the consideration assumption, then an option is to restrict to specific classes of graphs, or to specific families of coalitions. For example, with the partition equilibrium, our existence results hold without the consideration assumption. Partition equilibria are a compromise: they are guaranteed to exist as soon as a considerate equilibrium exists for any social network. Furthermore, they are slightly less numerous than considerate equilibria. However, the reachable partition equilibria have also a lower quality than considerate equilibria.

Another option is to relax the consideration assumption. For instance, a coalition could have consideration for a given neighbor that does not belong to the coalition if this agent has a sufficient number of neighbors within the coalition. Alternatively, a coalition could manipulate if its deviation does not harm more than a certain quota of its neighbors, for instance if no more than the half of the neighbors of the coalition disapproves the deviation.

Collaboration in strategic voting is typically an interaction that can be modeled by a social network. Some agents, driven by a common concern, collaborate in order to find a way to improve the outcome of an election that will impact all of them. One could also think of other settings, such as resource allocation, in which the agents can collaborate only for their own interest, via a transaction that is similar to a trade.

Chapter 3

Swap Dynamics in House Allocation

Abstract

We examine a resource allocation problem with indivisible goods where each agent is to be assigned exactly one object. In the housing market setting, each agent has preferences over objects and, starting from an initial endowment, agents may exchange their objects so as to improve their allocation. Assuming, like in the classical housing market framework, that any agent is able to communicate and thus to trade with any other agent is a strong assumption. We aim at relaxing this hypothesis by considering that the agents are embedded in a social network, which models their ability to trade their object with other agents.

We propose to study the possible allocations emerging from a sequence of simple swaps of objects between pairs of neighbors in the network. This model raises natural questions regarding (*i*) the reachability of a given full allocation, (*ii*) the ability of an agent to obtain a given object, and her final guarantee regarding the objects that she could obtain, and (*iii*) the search of Pareto-efficient allocations. Although this study is more oriented towards the analysis of the distributed process of swaps among the agents, a centralized perspective in the spirit of *the TTC algorithm* is also adopted.

We investigate the complexity of these problems by providing, according to the structure of the social network, polynomial and intractable cases. These questions are also investigated through parameterized complexity, focusing on budget constraints such as the number of exchanges an agent may be involved in or the total duration of the process of swaps.

Résumé

On examine dans ce chapitre un problème d'allocation de ressources avec des biens indivisibles dans lequel exactement un objet doit être affecté à chaque agent. A partir d'une allocation initiale, les agents peuvent améliorer leur situation en échangeant leurs objets. Faire l'hypothèse, comme dans le cadre classique de ce problème, que tout

This chapter is an extension of [Gourvès, Lesca, and Wilczynski, 2017] and [Saffidine and Wilczynski, 2018].

agent est capable de communiquer et donc d'échanger avec n'importe quel autre, est une hypothèse forte. On se propose de relâcher cette hypothèse en considérant que les liens entre les agents, conditionnant leur capacité à échanger des objets, sont modélisés par un réseau social.

On s'intéresse aux allocations possibles résultant d'une séquence de simple trocs entre voisins dans le réseau social. Ce modèle soulève notamment des questions naturelles concernant (i) l'atteignabilité d'une certaine allocation, (ii) la possibilité pour un agent d'obtenir un certain objet, ainsi que le niveau de garantie qu'il peut obtenir au final, et (iii) la recherche d'allocations optimales au sens de Pareto. Bien que cette étude soit plus largement orientée vers l'analyse du processus distribué des trocs entre des voisins dans le réseau, une approche centralisée dans l'idée de l'algorithme TTC est également adoptée.

La complexité de ces problèmes est abordée en fournissant, en fonction de la structure du réseau social, des cas polynomiaux ou difficiles. Ces questions sont également étudiées sous l'angle de la complexité paramétrée, en se concentrant sur des contraintes de budget impliquant le nombre d'échanges qu'un agent peut effectuer ou la durée totale du processus d'échanges.

Contents

3.1 Introduction

Reallocating indivisible items when each agent is initially endowed with exactly one resource is an important concern in Economics [Abdulkadiroğlu and Sönmez, 1999] and Computer science [Aziz and De Keijzer, 2012, Damamme et al., 2015]. This framework is known as *housing market* [Shapley and Scarf, 1974]. Many real-life situations can be modeled by such a setting, for example the reallocation of tasks among employees, the reassignment of time slots in schedules, or the problem of kidney exchanges.

3.1.1 Restricting the trades to neighbors of the social network

From an initial allocation, a typical way to reassign objects among agents comes from direct trades among the agents. In the housing market setting, it is implicitly assumed that all the agents have the capacity to perform direct deals. In fact, in the *top trading cycle* algorithm [Shapley and Scarf, 1974], any group of agents may form a trading cycle, provided that they obtain better objects after the deals.

However, this assumption is unrealistic in large scale instances, where some agents may be unable to communicate. Indeed, some logistics difficulties, such as communication and geographical distance, may prevent some trades to occur. Moreover, other simple reasons due to human relationships also matter: if the agents do not know each other or do not have affinities at all, they will not be willing to meet or make contact in order to exchange their items, as illustrated in the following example.

Example 3.1 *Let us take an example about students and dormitories that highlights the idea of "house" allocation in the spirit of Shapley and Scarf [1974]. Five new students, Anna, Brahim, Clara, David and Elena, arrive in a residence for students and are randomly assigned a room by the authority of the residence. However, they have different preferences over these rooms, due for instance to a specific location (street/garden side, ground floor, etc.). Here are summarized the preferences of the five students about the five rooms,* r_1, r_2, r_3, r_4 *and* r_5, *as well as their initial endowment (framed room).*

$$\begin{array}{rl}
\textit{Anna:} & r_2 \succ r_5 \succ \boxed{r_1} \succ r_4 \succ r_3 \\
\textit{Brahim:} & r_1 \succ \boxed{r_2} \succ r_3 \succ r_5 \succ r_4 \\
\textit{Clara:} & r_2 \succ r_4 \succ \boxed{r_3} \succ r_5 \succ r_1 \\
\textit{David:} & \boxed{r_4} \succ r_1 \succ r_2 \succ r_3 \succ r_5 \\
\textit{Elena:} & r_1 \succ \boxed{r_5} \succ r_3 \succ r_4 \succ r_2
\end{array}$$

The residence authority informs each of them that it is possible, if they are not satisfied, to exchange their room in the next few weeks but only with other new entrants, because the other students of the residence are already settled into their rooms for a few years. All the exchanges are not likely to occur among the new entrants. Indeed, Anna and Brahim hold grudges against each other because of a love story ending badly, whereas Clara, David and Elena do not know each other. Globally, the possibility of trades between the agents according to their relationships can be represented by the following undirected graph.

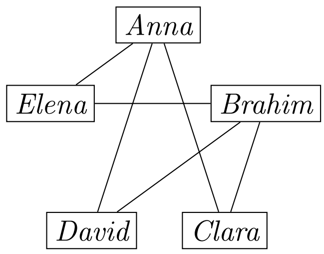

Actually, by only focusing on the preferences of the students (a student does not want to exchange her room against a room she prefers less), there are two possible exchanges: between Anna and Brahim and between Anna and Elena. However, as underlined by the graph, the communication between Anna and Brahim is not possible, even if they both prefer the room of the other one, excluding the possibility of an exchange between them. Finally, Anna and Elena decide to exchange their rooms because they both prefer the room of the other one and know sufficiently each other to have talked about their preferences.

Such a situation where trades are conditioned by social relationships among agents could appear quite marginal, but they can actually restrict the possibilities of exchange in real scenarios. Nevertheless, one could cite more "rational" examples, related to the idea of online platforms for exchanges, burgeoning nowadays, where geographical distance plays an important role.

Example 3.2 *Take an online forum on a website dedicated to car enthusiasts. Users have the possibility to exchange their cars. Since cars are cumbersome resources, this online platform only offers to the users the possibility to make contact with other passionate people. Every user is invited to enter into the data base her own car (or what she wants to trade) and her preferences over what she is looking for. Moreover, she is asked to input in the system her geographical constraints: her actual location as well as her willingness to move from this location ("radius of how many kilometers?"). Consequently, in order to filter the interesting proposals of trading for a given user, only proposals coming from individuals located within the radius of kilometers that she has entered are presented to her. Actually, in such a case, the geographical constraints define the possible trades, and they can be naturally represented by a graph over the agents.*

In a nutshell, restricting the set of direct exchanges to the ones which are actually possible seems realistic and relevant, especially in large scale instances. Indeed, natural obstacles may inhibit the agents in the trades. The ability of agents to exchange resources can be easily modeled as a social network. Moreover, the assumption of rational trades, where every agent needs to obtain a strictly better object (assumption already present in the top trading cycle setting), makes sense as we can see in Examples 3.1 and 3.2. Actually, lack of trust or uncertainty about the future exchanges in the process may lead agents to adopt a greedy behavior so as to be immediately better off in their new acquisition.

This chapter deals with a variant of housing market, where the agents are embedded in a social network which determines their ability to exchange their objects. Each participant is initially endowed with a single object, and she has strict ordinal preferences over objects. The agents may exchange their items under two conditions: they find the trade mutually profitable, and they are neighbors in the social network. We define such exchanges as *swaps*. Though sophisticated exchanges involving multiple agents have been analyzed [Sandholm, 1998, Dunne and Chevaleyre, 2008, Damamme et al., 2015], we focus on simple trades between pairs of neighbors. The exchanges are made without payments or monetary compensations.

3.1.2 Related work on trades and social network

Many works have studied how to reallocate indivisible items within a set of agents (see Section 1.4.2.b). One of the closest is the seminal work of Shapley and Scarf [1974] adopting a centralized approach in housing market for guiding the exchanges to an acceptable allocation for the agents which is, among other properties, Pareto-efficient (TTC algorithm).

In a distributed perspective, many studies focus on designing realistic conditions for trades and aim at analyzing the dynamics of exchanges, for instance with respect to its convergence properties. Rational exchanges where agents are necessarily better off after a trade is a common assumption. A model very close to ours is the *swap deals* model [Damamme et al., 2015]. This work considers deals within groups of agents of a given size in a housing market setting. To the best of our knowledge, this is the first work analyzing the housing market framework within a distributed perspective. Damamme et al. [2015] notably prove that when the deals are restricted to exchanges between two agents and the preferences are single-peaked, the dynamics of exchanges is guaranteed to converge to a Pareto-efficient allocation. However, in their model, there is no restriction on the possible exchanges according to some relations over the agents. When restricted to exchanges between two agents, the exchanges defined by Damamme et al. [2015] correspond to our definition of swaps where the social network is a complete graph.

Combining distributed process of exchanges and graphs, the notion of *negotiation topology graph* [Chevaleyre et al., 2007c], where the exchanges are restricted to agents belonging to the same clique of the social network, is noteworthy. Although this work does not deal with housing market but with a general resource reallocation setting where agents can have more than one item, this study is completely connected to ours, with a graph defining the possible trades. The main differences are that Chevaleyre et al. [2007c] do not restrict to swaps between neighbors of the network, as in our model, and study a model with cardinal preferences and monetary side payments. Moreover, their goal is to improve a specific fairness notion defined according to the graph.

3.1.3 Contributions and organization

We consider a housing market framework with exchanges between neighbors in a social network. Our main question concerns the analysis of the distributed process of exchanges: starting from the initial endowment, which allocation of the objects can emerge? Indeed, some solutions are ruled out because of the agents' preferences over the objects: we assume that no one is interested in exchanging her current object with something that she considers as worse. In addition, the network limits the access of certain participants to each other. Consequently, it is particularly challenging to understand how the combination of these two natural ingredients influence the outcome of a dynamics in which connected agents agree on mutually profitable swaps of objects.

One of the main issues is the Reachable Object (RO) problem: Given a target agent A and a target object x, is there a sequence of exchanges ensuring A is allocated x? Globally, the question of the emergence of a specific allocation is also interesting, leading to the Reachable Assignment (RA) problem: Is there a sequence of exchanges leading to a given allocation π? Moreover, we introduce Guaranteed Level of Satisfaction (GLS), a problem related to RO but more realistic. GLS asks whether an agent can be guaranteed to be eventually allocated an item at least as good as the input target item, regardless of the exchanges other agents perform. While RO takes an optimistic perspective by asking the existence of an appropriate sequence of exchanges, GLS adopts a more pessimistic point of view by considering the outcome of any possible sequence of exchanges, and not only focusing on "lucky" configurations.

Globally, these three decision problems naturally arise when we analyze the distributed process of exchanges, but they can also model concrete issues. Let us take the context of Example 3.2 with an online exchange platform where users input in the system which item they hold as well as their preference. A user may request a target object to the centralized system which could then suggest a series of intermediate exchanges to bring it to her. Even in such a context, restricted rational exchanges are relevant: geographical constraints can still prevent two agents to trade, moreover the guarantee of getting a better object is essential as otherwise an agent could be left worse off than she started, should an intermediate agent exit the system during the process.

These problems are very appealing but, as we will prove, they turn out to be computationally hard, sometimes even for simple social networks. Consequently, we attempt to mitigate these negative results by also looking at more realistic constrained settings. We draw inspiration from the fact that an agent may not be willing to perform a large number of swaps or to wait a long time before getting a desired object, and introduce natural budget constraints: the number of exchanges agents may make and the total duration of the process of exchanges. By considering such additional constraints, we perform a refined complexity analysis.

Moreover, in another direction, it appears interesting to consider the quality of an allocation resulting from a sequence of profitable exchanges. In the context of ordinal preferences, Pareto-efficiency appears as the minimal requirement for an allocation to be socially acceptable. Pareto-efficiency has been widely studied in the context of house allocation [Abraham et al., 2005] and housing market [Aziz et al., 2016a]. As far as we know, the computation of Pareto-efficient allocations has not been investigated when the possible allocations are constrained by a social network. Therefore, we also consider, within a centralized perspective similar to the idea of the top trading algorithm, whether it is possible for a central authority to guide the exchanges of the agents in order to reach a Pareto-efficient allocation.

We prove that all the decision problems that we consider, namely RO, RA and GLS, are intractable when the graph describing the social network is not restricted. However, we identify some specific classes of graphs, such as the trees for RA, that enable the tractability of the problems. In order to refine the complexity of these problems and obtain more positive results, we also adopt a parameterized complexity perspective. When parameterizing the problems by the maximal number of swaps per agent, we show intractability even for highly structured graphs. However, when constraining the duration of the process, we obtain more promising results: RO and RA are tractable for a class of graph that is meaningful for representing real social networks, namely the bounded degree graphs. Indeed, while online social networks with many hubs do not exhibit a bounded degree, this assumption is relevant in scenarios where the social network models possibility of collaboration or close relationship. In fact, the number of agents with who a given agent can actually interact cannot be too important. In general, the problems are tractable when the length of the sequence of swaps does not depend on the input size. This is also a realistic scenario because the time that an agent is willing to wait before getting a target object is independent from the input size (the patience of an agent has no reason to increase with the number of agents).

Within a centralized approach, a central coordinator tries to direct the agents to exchanges in order to reach a Pareto-efficient allocation among all possible outcomes. We show that such a coordinating strategy is computationally hard to design for a central authority in general, but we also provide tractable cases for simple classes of graphs.

The chapter is organized as follows. We firstly introduce formally the model of object reallocation along a social network in Section 3.2, as well as the four problems under consideration, and the budget parameters that we will use to refine our analysis. We then devote one section to each of our problems, Reachable Object, Reachable Assignment, Guaranteed Level of Satisfaction, and Pareto Reachability (Section 3.3–3.6). In each section dealing with a decision problem, we first treat the general case, before examining budget constraints. Finally, we conclude and highlight some avenues for future work (Section 4.8).

3.2 Swap dynamics model

Let us consider a resource allocation problem, with a set of agents $N = \{1, \ldots, n\}$ and a set of indivisible items $O = \{o_1, \ldots, o_r\}$, where the number of resources is equal to the number of agents, i.e., $r = n$. Each agent is initially endowed with an object, via an initial allocation π^0 assigning exactly one object per agent, i.e., $\pi^0 : N \to O$. The agents have strict ordinal preferences over the objects, represented by a linear order $\succ_i$ over O, for a given agent $i \in N$. Any allocation π of objects to agents assigns exactly one item per agent and is such that $\pi(i) \neq \pi(j)$, for all agents i and j. We sometimes write allocations as n-tuples where i^{th} coordinate refers to the object assigned to i. As described in Section 1.4.2.b, this setting corresponds to the housing market framework. The agents are embedded in a social network, represented by an undirected graph $G = (N, E)$, where the edges capture the possibility of trade between two agents. An instance of the swap dynamics model is then a tuple $\langle N, O, \succ, G, \pi^0 \rangle$.

3.2.1 Rational deals conditioned by a social network

Agents can trade their objects so as to obtain better objects, but not all exchanges are plausible. The exchange possibilities depend on the social network and on the preferences of the agents. We only admit *swaps*, rational trades between neighbors. Formally, a swap in an allocation π is an exchange between two agents i and j such that $\{i, j\} \in E$ and the exchange is rational, i.e., $\pi(i) \succ_j \pi(j)$ and $\pi(j) \succ_i \pi(i)$.

A sequence of swaps is a sequence of allocations $(\pi^0, \ldots, \pi^t)$ such that a swap is performed between every two consecutive allocations $\pi^{t'}$ and $\pi^{t'+1}$ for any $0 \leq t' < t$. An allocation π is *reachable* if there is a sequence of swaps leading to it starting from the initial allocation, i.e., there exists a sequence $(\pi^0, \ldots, \pi^t)$ such that $\pi^t = \pi$. Let us denote by *RAll* the set of all reachable allocations. An object $o \in O$ is reachable for agent i if there is a sequence of swaps $(\pi^0, \ldots, \pi^t)$ where $\pi^t(i) = o$. An allocation π is *stable* if no swap is possible from π.

Swap dynamics refers to a distributed process where agents may rationally exchange their objects when they are neighbors in the network, until a stable allocation is reached.

Example 3.3 *Consider an instance with four agents, where $N = \{1, 2, 3, 4\}$ and $O = \{o_1, o_2, o_3, o_4\}$. The social network, the preferences of the agents and the initial allocation are given as follows, where the framed objects represent the initial object of each agent.*

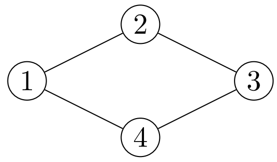

$$
\begin{aligned}
1 &: \; o_2 \succ o_3 \succ \boxed{o_1} \succ o_4 \\
2 &: \; o_3 \succ o_1 \succ \boxed{o_2} \succ o_4 \\
3 &: \; o_4 \succ o_1 \succ o_2 \succ \boxed{o_3} \\
4 &: \; o_1 \succ o_2 \succ \boxed{o_4} \succ o_3
\end{aligned}
$$

Initially, only the swaps between agents 1 and 2, and 2 and 3 are possible. Indeed, the rational exchange between agents 1 and 3 is not possible because the agents are not connected in the graph. Moreover, the exchange between connected agents 1 and 4 is not possible because it is not rational for 1. The sequence of swaps between the following pairs of agents, $\{1, 2\}$, $\{2, 3\}$, and $\{3, 4\}$, gives rise to a reachable allocation where every agent gets her best object. This allocation is stable: no further swap can be performed.

Observe that, by rationality of the exchanges, the rank of the object owned by agent i never increases within $\succ_i$, during any sequence of swaps for any agent i.

Observation 3.1 *For any sequence of swaps $(\pi^0, \ldots, \pi^t)$ and any steps t' and t'' such that $0 \leq t' < t'' \leq t$, it must hold that $\pi^{t''}(i) \succ_i \pi^{t'}(i)$ or $\pi^{t''}(i) = \pi^{t'}(i)$.*

This implies that, for any social network, an object cannot pass twice by the same agent, in the sense that once an agent exchange a given object, she cannot get back it.

Observation 3.2 *For any sequence of swaps $(\pi^0, \ldots, \pi^t)$, any agent i and any steps t_0, t_1 and t_2 such that $0 \leq t_0 < t_1 < t_2 \leq t$, if $\pi^{t_0}(i) \neq \pi^{t_1}(i)$, then $\pi^{t_2}(i) \neq \pi^{t_0}(i)$.*

Another implication of the rationality assumption in the swaps is that every agent can make at most $n - 1$ swaps, leading to a trivial quadratic bound for the length of any sequence of swaps.

Observation 3.3 *The length of any sequence of swaps is bounded by n^2.*

Let us consider that time is discretized according to the moments where swaps are made. The *makespan* of a sequence of swaps is the minimum time that elapses from the beginning of the sequence to the end, when we allow parallel swaps which simultaneously occur. This notion can be formalized as follows. Let $s = (\pi^0, \ldots, \pi^t)$ be a sequence of swaps. A parallel decomposition of s is a tuple of integers $\ell = (\ell_0, \ell_1, \ldots, \ell_p)$ of length $|\ell|= p$, such that $0 = \ell_0 < \ell_1 < \cdots < \ell_p = t$, and for all $0 \leq i < p$, all the swaps occurring between allocation π^{ℓ_i} and allocation $\pi^{\ell_{i+1}}$ do not involve the same agents. In other words, all the swaps occurring between π^{ℓ_i} and $\pi^{\ell_{i+1}}$ can be performed simultaneously. The makespan is the length p of the minimal parallel decomposition. In Example 3.3, the makespan of the sequence of swaps is equal to three, which is also the total number of exchanges, because no swap can be performed in parallel.

3.2.2 Decision problems and parameters

In order to analyze the distributed process of swap dynamics, we are interested in the allocations obtained from π^0 by sequences of swaps. Natural questions in this context are: *what are the objects that a given agent can obtain?* more globally, *what are the assignments that we can reach?* or more pessimistically, *what are the level of satisfaction that each agent is ensured to obtain?* These questions are transcribed into the following decision problems.

Reachable Object (RO):

Instance:	Swap dynamics instance $\langle N, O, \succ, G, \pi^0 \rangle$, target agent $A \in N$, target object $x \in O$.
Question:	Is there a sequence of swaps (π^0,..., π^t) such that $\pi^t(A) = x$?

Reachable Assignment (RA):

Instance:	Swap dynamics instance $\langle N, O, \succ, G, \pi^0 \rangle$, target allocation π.
Question:	Is there a sequence of swaps (π^0,..., π^t) such that $\pi^t = \pi$?

Guaranteed Level of Satisfaction (GLS):

Instance:	Swap dynamics instance $\langle N, O, \succ, G, \pi^0 \rangle$, target agent $A \in N$, target object $x \in O$.
Question:	Is it the case that *for all* sequences of swaps $(\pi^0, \ldots, \pi^t)$ where π^t is stable, either $\pi^t(A) = x$ or $\pi^t(A) \succ_A x$?

While Reachable Object asks whether a given object can be reached by a sequence of swaps, Guaranteed Level of Satisfaction adopts a more pessimistic perspective by considering a quasi "dual" problem: beyond the possible advantageous configurations, it asks the level of satisfaction that a given agent is guaranteed to obtain, whatever the sequence of swaps that is chosen. This is illustrated in the following example.

Example 3.4 *Let us consider an instance with four agents, where $N = \{1, 2, 3, 4\}$ and $O = \{o_1, o_2, o_3, o_4\}$. The social network, the preferences of the agents and the initial allocation (framed objects) are defined as follows.*

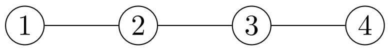

$$
\begin{aligned}
1 &: o_4 \succ o_2 \succ \boxed{o_1} \succ o_3\\
2 &: o_1 \succ o_3 \succ o_4 \succ \boxed{o_2}\\
3 &: o_2 \succ o_4 \succ \boxed{o_3} \succ o_1\\
4 &: o_2 \succ o_3 \succ \boxed{o_4} \succ o_1
\end{aligned}
$$

Observe that object o_4 is reachable for agent 1, with the sequence of swaps involving the following pairs of agents: $\{3,4\}$, $\{2,3\}$ and $\{1,2\}$. Therefore, optimistically, agent 1 could hope to obtain object o_4. However, there is no guarantee that she will obtain it in any case. Indeed, if agents 2 and 3 firstly swap their objects, allocation (o_1, o_3, o_2, o_4) is reached where no further swap can occur: the exchange between 1 and 2 is not rational as well as the exchange between 3 and 4. Therefore, in this case, we reach an allocation that is stable but where agent 1 does not obtain an object preferred or equal to o_4. Hence, this instance is a yes-*instance of* Reachable Object *with target agent 1 and target object o_4, whereas this is a* no-*instance of* Guaranteed Level of Satisfaction *with target agent 1 and target object o_4.*

When asking whether an agent can obtain a given object, a large number of swaps may not be realistic. Indeed, the agents may not be willing to perform many swaps or wait for a long time before getting their target objects. Consequently, we study three variants of our decision problems, namely RO, RA and GLS, referred to as {RO/RA/GLS }-{max/sum/makespan}, where the number of swaps in a solution sequence is limited. In each variant, this quantity is measured differently, leading to different complexity-theoretic characterizations of the problem.

- *max*: Every agent is involved in no more than k swaps.
- *sum*: The total length of the sequence is no more than k.
- *makespan*: The makespan of the sequence is no more than k.

Observe that the *sum* parameter provides an upper bound for the makespan by giving the duration of the process in case no parallel swaps take place.

The three parameters under consideration restrict the sequence of swaps. They are illustrated in the following example.

Example 3.5 *Let us consider an instance with four agents where $N = \{1, 2, 3, 4\}$ and $O = \{o_1, o_2, o_3, o_4\}$. The preferences, the social network and the initial allocation (framed objects) are defined as follows.*

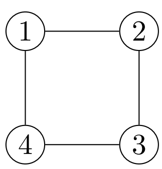

$$
\begin{aligned}
1 &: o_4 \succ o_2 \succ \boxed{o_1} \succ o_3\\
2 &: o_4 \succ o_1 \succ o_3 \succ \boxed{o_2}\\
3 &: o_1 \succ o_4 \succ \boxed{o_3} \succ o_2\\
4 &: o_2 \succ o_3 \succ \boxed{o_4} \succ o_1
\end{aligned}
$$

From the initial allocation, take the sequence of swaps between the following pairs of agents: $\{1,2\}, \{3,4\}, \{2,3\}$. This sequence has a parameter max equal to 2 because agents 2 and 3 perform the maximal number of swaps per agent in the sequence, which *is equal to two. Concerning the length of the sequence, the parameter sum is equal to 3 because there are 3 swaps in total in the sequence. However, the parameter makespan is equal to 2 because the two first swaps between the agents $\{1, 2\}$ and $\{3, 4\}$ can be performed in parallel since they do not involve the same agents.*

Motivated by practical considerations on real-life networks, we also examine the impact of some restrictions on the graph underlying the social network. For instance, in a scenario with physical item exchanges, geographical as well as social aspects will impact the maximum number of neighbors an agent may have. In other words, the social network in such scenarios has a small degree.

3.2.3 Pareto optimization problem

The swaps dynamics always converges to a stable allocation. However, the following example shows that we can reach a rather bad assignment if the agents exchange their objects in an uncoordinated way.

Example 3.6 *Consider an instance with n agents, where $N = \{1, \ldots, n\}$ and $O = \{o_1, \ldots, o_n\}$. The social network, the preferences and the initial allocation are described below ($[\ldots]$ stands for an arbitrary ranking over the remaining objects).*

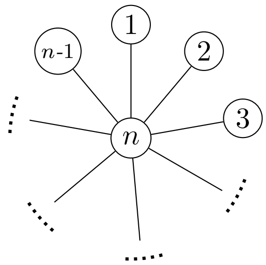

$$\begin{array}{rl} 1: & o_n \succ [\ldots] \succ \boxed{o_1} \\ 2: & o_1 \succ [\ldots] \succ \boxed{o_2} \\ \vdots & \vdots \\ n-2: & o_{n-3} \succ [\ldots] \succ \boxed{o_{n-2}} \\ n-1: & o_{n-2} \succ [\ldots] \succ o_n \succ \boxed{o_{n-1}} \\ n: & o_{n-1} \succ o_{n-2} \succ [\ldots] \succ o_1 \succ \boxed{o_n} \end{array}$$

Consider allocation π^1 resulting from the swap between agents $n-1$ and n. The center agent n obtains her most preferred object, thus no further swap is possible and π^1 is stable. However, if we consider the sequence of swaps performed by the pairs of agents $\{n, 1\}$, $\{n, 2\}$, $\ldots$, $\{n, n-2\}$, $\{n, n-1\}$, then we reach an allocation where every agent obtains her most preferred object.

Recall that an allocation π is *Pareto-efficient* if there is no allocation π' such that for all $i \in N$, $\pi'(i) \succ_i \pi(i)$ or $\pi'(i) = \pi(i)$, and there exists $j \in N$ such that $\pi'(j) \succ_j \pi(j)$. By focusing on allocations emerging from swap dynamics, there is no need to consider allocations that are not reachable. In consequence, we restrict the definition of Pareto-efficiency within the set *RAll* of reachable allocations, with a notion of efficiency that we call *RAll-efficiency*.

Definition 3.1 (RAll-efficiency) *An allocation π is RAll-efficient if $\pi \in RAll$ and there is no reachable allocation $\pi' \in RAll$ such that for all $i \in N$, $\pi'(i) \succ_i \pi(i)$ or $\pi'(i) = \pi(i)$, and for at least one agent $j \in N$, $\pi'(j) \succ_j \pi(j)$.*

Clearly, a RAll-efficient allocation is stable. Otherwise, at least two agents can be better off by swapping their objects, leading to a new reachable allocation dominating the previous one. Note that a Pareto-efficient allocation that is reachable is RAll-efficient, but a RAll-efficient allocation may not be Pareto-efficient.

As informally noticed in Example 3.6, swap dynamics can reach an allocation that is not RAll-efficient. This observation leads us to also consider the swap dynamics via a centralized approach, in addition of a distributed analysis through the decision problems presented in the previous subsection. The swaps, though constrained by the social network and the mutual benefit for the two involved agents, are guided by a coordinator, in order to reach a RAll-efficient allocation. This goal is transcribed into the following optimization problem.

PARETO REACHABILITY:	
Instance:	Swap dynamics instance $\langle N, O, \succ, G, \pi^0 \rangle$
Problem:	Find a RAll-efficient allocation

Note that PARETO REACHABILITY is not a decision problem because a RAll-efficient allocation is always guaranteed to exist. Actually, in the PARETO REACHABILITY problem, we seek a sequence of swaps that leads to one of the best possible outcomes, in the Pareto meaning. The goal of the PARETO REACHABILITY problem is described in the following example.

Example 3.7 *Let us consider an instance with four agents where $N = \{1, 2, 3, 4\}$ and $O = \{o_1, o_2, o_3, o_4\}$. The social network, the preferences of the agents and the initial allocation are defined as follows.*

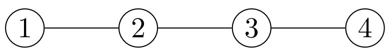

$1: \; o_4 \succ o_3 \succ o_2 \succ \boxed{o_1}$

$2: \; o_3 \succ o_1 \succ o_4 \succ \boxed{o_2}$

$3: \; o_1 \succ o_2 \succ o_4 \succ \boxed{o_3}$

$4: \; o_2 \succ o_1 \succ o_3 \succ \boxed{o_4}$

In this instance, there are only three maximal sequences of swaps, that is to say sequences which lead to stable allocations. We denote the three maximal sequences of swaps by s_1, s_2 and s_3. The sequences of neighbors involved in the swaps of s_1, s_2 and s_3 are $(\{3,4\}, \{2,3\}, \{1,2\})$, $(\{2,3\})$, and $(\{1,2\}, \{2,3\})$, respectively. Sequence s_1 leads to allocation $\pi_1 := (o_4, o_1, o_2, o_3)$, sequence s_2 to allocation $\pi_2 := (o_1, o_3, o_2, o_4)$ and sequence s_3 to allocation $\pi_3 := (o_2, o_3, o_1, o_4)$. By definition, they are all reachable allocations. Among the reachable allocations, only allocations π_1 and π_3 are Pareto-efficient. Indeed, allocation π_2 is Pareto-dominated by π_3: agent 1 prefers object o_2 to object o_1 and agent 3 prefers object o_1 to object o_2. Moreover, all the other reachable allocations are not stable. Consequently, π_1 and π_3 are RAll-efficient, whereas π_2 is not. Therefore, the reachable allocations π_1 and π_3 are possible solutions to the PARETO REACHABILITY *problem. Observe that these two solutions are Pareto-dominated by allocation (o_4, o_3, o_1, o_2). However, since this allocation is not reachable by a sequence of swaps, this is not a plausible outcome and we do not take it into account in our definition of the* PARETO REACHABILITY *problem.*

3.3 Reachable Object

In this section we focus on Reachable Object, the decision problem asking whether a given target agent A can obtain a given target object x via a sequence of swaps.

3.3.1 Reachable Object with no budget consideration

In this subsection, Reachable Object is analyzed in a context where no restriction is imposed on the sequence of swaps. We prove that the problem is NP-complete even when the network is a tree. However, for some further restrictions on the graph, the problem becomes polynomial.

Theorem 3.1 Reachable Object *is* NP*-complete even when the network is a tree.*

Proof: Given a sequence of swaps $(\pi^0, \dots, \pi^t)$, it is easy to determine whether the swaps are rational and lead to give object x to agent A, for some step t' such that $0 \leq t' \leq t$. Combined with the fact that any sequence of swaps has a length bounded by n^2 (Observation 3.3), this implies the membership to NP.

We use a reduction from the 2P1N-SAT problem, known to be NP-complete (Theorem 1.3). In 2P1N-SAT, we are given a set $X = \{x_1, \dots, x_v\}$ of variables, and a collection $\mathcal{C} = \{C_1, \dots, C_s\}$ of clauses over X such that each positive (respectively, negative) literal occurs exactly twice (respectively, once) in $\mathcal{C}$. The question is whether $\mathcal{C}$ is satisfiable by a truth assignment of the variables in X.

Let x_ℓ^i (respectively, $\overline{x}_\ell^i$) denote the positive (respectively, negative) literal x_ℓ if present in clause C_i. Index p_j^ℓ (respectively, n^ℓ) refers to the clause containing the j^{th} occurrence (with $j \in \{1,2\}$) of x_ℓ (respectively, the occurrence of $\overline{x}_\ell$).

From an instance $\mathcal{I} = \langle \mathcal{C}, X \rangle$ of 2P1N-SAT, we construct a swap dynamics instance $\langle N, O, \succ, G, \pi^0 \rangle$ for Reachable Object as follows. Each literal x_ℓ^i (respectively, $\overline{x}_\ell^i$) is associated with an agent Y_ℓ^i (respectively, $\overline{Y}_\ell^i$) who is initially endowed with object y_ℓ^i (respectively, $\overline{y}_\ell^i$). Every clause C_i is associated with an agent K_i initially endowed with object k_i. We add an agent Z initially endowed with object z, leading to $|N| = |O| = s + 3v + 1$. The graph $G = (N, E)$ is constructed as described in Figure 3.1.

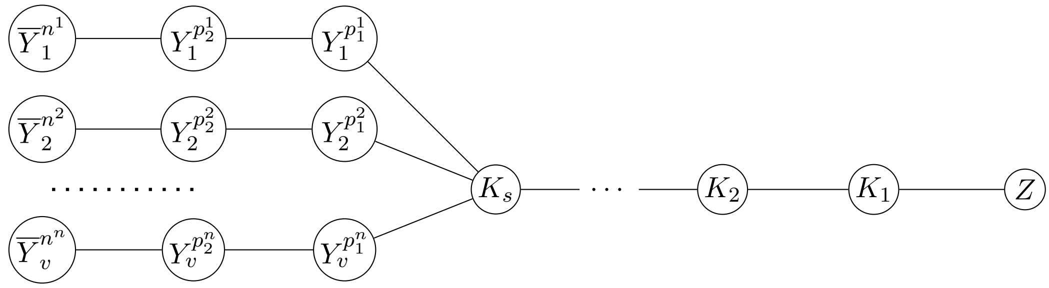

Figure 3.1: Construction of social network G

The preference profile $\succ$ is defined below. Notation $\{\ell_i\}$ stands for the set of the objects associated with the literals of clause C_i, and $[\dots]$ for all the remaining objects, that are not presented in the preference order; they are both ranked in arbitrary order.

$$\begin{aligned}
Z:&\quad \{\ell_1\} \succ \boxed{z} \succ [\dots]\\
K_i:&\quad \{\ell_{i+1}\} \succ z \succ \{\ell_i\} \succ k_1 \succ \{\ell_{i-1}\} \succ \dots \succ k_{i-1} \succ \{\ell_1\} \succ \boxed{k_i} \succ [\dots]\\
K_s:&\quad z \succ \{\ell_s\} \succ k_1 \succ \{\ell_{m-1}\} \succ \dots \succ k_{s-1} \succ \{\ell_1\} \succ \boxed{k_s}\\
Y_\ell^{p_1^\ell}:&\quad k_{s-n^\ell+1} \succ k_{s-p_2^\ell+1} \succ y_\ell^{p_2^\ell} \succ k_{s-p_1^\ell+1} \succ \overline{y}_\ell^{n^\ell} \succ \boxed{y_\ell^{p_1^\ell}} \succ [\dots]\\
Y_\ell^{p_2^\ell}:&\quad k_{s-p_1^\ell+1} \succ y_\ell^{p_1^\ell} \succ \overline{y}_\ell^{n^\ell} \succ \boxed{y_\ell^{p_2^\ell}} \succ [\dots]\\
\overline{Y}_\ell^{n^\ell}:&\quad y_\ell^{p_2^\ell} \succ \boxed{\overline{y}_\ell^{n^\ell}} \succ [\dots]
\end{aligned}$$

We claim that $\mathcal{C}$ is satisfiable in $\mathcal{I}$ if and only if object z is reachable for agent K_s in the constructed instance of Reachable Object.

Observe, on one hand, that in the dipath τ from Z to K_s, every agent prefers z to her initial object, and only accepts exchanging z with her successor in τ for an object corresponding to a literal that satisfies the clause associated with the successor. Therefore, the only way to move z from Z to K_s is to give to each agent K_i an object associated with a literal that satisfies clause C_i, by increasing order of the clause indices. Thus, from $i = 1$ to s, an object of $\{\ell_i\}$ moves to K_i.

On the other hand, we ensure, in each branch of the graph corresponding to a variable x_ℓ, i.e., in each path $[Y_\ell^{p_1^\ell}, Y_\ell^{p_2^\ell}, \overline{Y}_\ell^{n^\ell}]$, that if an object $y_\ell^{p_j^\ell}$ for $j \in \{1,2\}$ (respectively, $\overline{y}_\ell^{n^\ell}$) moves out of the branch to be exchanged with agent K_s, then object $\overline{y}_\ell^{n^\ell}$ (respectively, an object $y_\ell^{p_j^\ell}$ for $j \in \{1,2\}$) cannot move out thereafter. Indeed, to move out of the branch, object $y_\ell^{p_j^\ell}$ for $j \in \{1,2\}$ (respectively, $\overline{y}_\ell^{n^\ell}$) must pass by agent $Y_\ell^{p_1^\ell}$ and then be exchanged with K_s against object $k_{s-p_j^\ell+1}$ (respectively, object $k_{s-n^\ell+1}$). This is due to the fact that, by construction of the preferences, the only objects coming from outside the branch that agent $Y_\ell^{p_1^\ell}$ accepts are clause-objects, and to the remark of the previous paragraph implying that, each time an object from $\{\ell_i\}$ comes to the clause branch, a clause-object k_{s-i+1} must move to a variable branch. Therefore, if $\overline{y}_\ell^{n^\ell}$ is the first object to move out of the branch, then agent $Y_\ell^{p_1^\ell}$ must obtain her most preferred object $k_{s-n^\ell+1}$, after having swapped $\overline{y}_\ell^{n^\ell}$ with agent K_s. So, no other object in this branch can move out afterwards. Otherwise, if $y_\ell^{p_1^\ell}$ or $y_\ell^{p_2^\ell}$ moves first, then agent $Y_\ell^{p_1^\ell}$ has received from K_s an object that she prefers to $\overline{y}_\ell^{n^\ell}$, so object $\overline{y}_\ell^{n^\ell}$ is blocked and must stay within the branch.

Suppose that there exists a truth assignment ϕ of the variables which satisfies all the clauses of $\mathcal{C}$. Then, it suffices to choose, for each clause C_i, one literal ℓ_i that makes C_i true with assignment ϕ, and move the corresponding object to agent K_i, by increasing order of the clause indices. This sequence of swaps is possible. Remark that for each K_j on the path between K_s and K_i, i.e., such that $j > i$, K_j must receive object ℓ_i and, later, object k_i and later again, object z. This passage order is guaranteed by construction of the preferences: $z \succ_{K_j} k_i \succ_{K_j} \{\ell_i\} \succ_{K_j} k_j$. In the branch of the graph corresponding to the variable associated with ℓ_i, if some other objects have previously moved out of the branch, then their corresponding literal has necessarily the same truth value as ℓ_i, since the chosen variable-objects are associated with true literals in ϕ, and belongs to a clause C_j where $j < i$. Thus, by construction, we can perform the swaps within the branch to move the object associated with ℓ_i to agent K_s.

Now, suppose that z is reachable for K_s and denote by π a reachable allocation where $\pi(K_s) = z$. Consider an assignment ϕ of the variables where we set to true every variable corresponding to an object possessed by the agents $Z, \ldots, K_{s-1}$ in π. As previously observed, in order to move object z from agent Z to agent K_s, each K_i has to obtain an object $o \in \{\ell_i\}$ corresponding to one of the literals of C_i. Thus, once K_s obtains z, object o is owned by agent K_{i-1} (or Z if $i = 1$). By construction, agents $Z, \ldots, K_{s-1}$ cannot obtain a better object thereafter. It suffices to prove now that there are no two literals ℓ_1 and ℓ_2 in ϕ such that $\ell_1 = \overline{\ell_2}$. It is guaranteed by construction of each branch associated with a variable. □

The network in the previous proof is a particular tree, with at most one vertex of degree greater than 2, which corresponds to a spider (Definition 1.11). A spider generalizes a star, but while Reachable Object is NP-complete when the social network is a spider, the problem becomes polynomial in a star.

Proposition 3.2 *When $G = (N, E)$ is a star, there exists a polynomial algorithm for* Reachable Object.

Proof: The problem asks whether agent A can get object x. The network consists, without loss of generality, of a center denoted by n, and $n-1$ leaves denoted by $1, \ldots, n-1$, as depicted in Example 3.6. A swap always involves the center and a leaf. Once a leaf i has exchanged her initial object o_i, she is not involved in a subsequent swap because otherwise the center-agent would get back an object that she has previously owned (Observation 3.2). Thus, any sequence of swaps reduces to an ordered list (without repetition) of leaves, indicating with which agents the center exchanges her object.

Let us first focus on the case where agent A is the center. The problem reduces to the search of a dipath in a directed graph $G_D = (N, E')$ where $(i, j) \in E'$ with $i \in N$ and $j \in N \setminus \{n\}$ if and only if center-agent n and leaf-agent j can rationally trade when center-agent n owns object o_i, the initial object of i, and j still owns her initial object o_j. There is a dipath from n to j in the directed graph G_D if and only if the center can get o_j, the initial object of agent j. A linear algorithm solves this path problem and the construction of G_D can be done in polynomial time.

In case A is a leaf, the problem reduces to the previous one: the center gets object x and then, A and the center swap their objects. Return "yes" if these two steps are feasible. □

Now we suppose that the network is a path. Without loss of generality, we may assume that $N = \{1, \ldots, n\}$, $E = \{\{i, i+1\} : 1 \leq i < n\}$ and that the initial object of each agent i is denoted by o_i. Observation 3.2 implies that when the network is a path, once an object "moves" in a given direction in the path, then it cannot "move" in the opposite direction.

Let us define as *canonical sequence of exchanges* $\kappa(j, i)$ the sequence of exchanges which assigns object o_j to agent i by directly moving o_j from j to i along the path. This is the sequence of exchanges between the following pairs of agents if $j < i$: $\{j, j+1\}, \{j+1, j+2\}, \ldots, \{i-1, i\}$. This sequence transforms π^0 into an assignment π where $\pi(\ell) = o_\ell$ if $\ell < j$ or $\ell > i$, $\pi(\ell) = o_{\ell+1}$ if $j \leq \ell < i$, and $\pi(i) = o_j$. This sequence is said to be a sequence of swaps if all its exchanges are rational.

Proposition 3.3 *When $G = (N, E)$ is a path, if object o_j is reachable for agent n, then $\kappa(j, n)$ is the minimal sequence of swaps leading to give o_j to agent n.*

Proof: Assume for a contradiction that there exists $i \in \{j+1, \ldots, n\}$ such that the exchange between $i-1$ and i is not rational, i.e., $o_i \succ_i o_j$ or $o_j \succ_{i-1} o_i$, whereas o_j is actually reachable for agent n.

Assume first that $o_i \succ_i o_j$ holds. Note that if there exists a sequence of swaps leading to assign o_j to n, then o_j must be assigned once to agent i since there is a unique path between j and n. Because $o_i \succ_i o_j$, agent i will never accept o_j, contradiction.

Assume now that $o_j \succ_{i-1} o_i$ holds. If there exists a sequence of swaps leading to assign o_j to n, then o_j must be exchanged once with o_i, because o_j must reach n and there is no agent after n to receive o_i. This swap cannot be performed between $i-1$ and i since $o_j \succ_{i-1} o_i$, and thus occurs between some agents with indices lower than i (otherwise object o_i would pass twice by agent i, contradicting Observation 3.2). Consequently, agents $i-1$ and i must have performed an earlier swap in order to make o_i move to the agents with lower indices than i. After this swap, the object currently owned by $i-1$ must move to the direction of n before o_j, but there is no agent after n to receive this object and o_j cannot overtake it, contradiction. □

By Proposition 3.3, for testing that an object is reachable for a leaf of the path, it suffices to verify that the associated canonical sequence of exchanges is a sequence of swaps.

Corollary 3.4 *When $G = (N, E)$ is a path and agent A a leaf,* Reachable Object *is solvable in polynomial time.*

Other solvable cases of Reachable Object in a path can be listed. The main one is when the distance in the path between the agent and the owner of the target object is a constant, as formulated in the next proposition (this corresponds to the parameterized complexity class XP, see Definition 1.33).

Proposition 3.5 *Reachable Object, parameterized by the distance between the target agent A and the owner of the target object x in path G, is in* XP.

Proof: Let k be the distance in $G = (N, E)$ between the target agent A, say i, and the owner of target object x, say j such that $i < j$. We assume, without loss of generality, that j is a leaf of the path, i.e., $j = n$. Indeed, by Observation 3.2, the sequence of swaps leading to give o_j to i cannot involve objects o_ℓ such that $\ell > j$.

Let $(\pi^0, \ldots, \pi^t)$ be the minimal sequence of swaps leading to give object o_j to agent i, i.e., $\pi^t(i) = o_j$. The principle of our algorithm is to guess the k objects $\{o_{i_1}, \ldots, o_{i_k}\}$ that are assigned to the agents $\{i+1, \ldots, j\}$ in π^t. There are $\mathcal{O}(n^k)$ such partial allocations. For each guess with the objects of $\{o_{i_1}, \ldots, o_{i_k}\}$, a complete allocation π' over N is constructed such that $\pi'(i) = o_j$ and $\pi'(i+\ell) = o_{i_\ell}$ for each $\ell \in \{1, \ldots, k\}$. We aim at reproducing in π' the final allocation of a minimal sequence of swaps leading to give o_j to i. Therefore, provided there is no unnecessary swap, there is no agent $i' < i$ that receives in π' an object $o_{j'}$ such that $j' < i'$. We complete the rest of the allocation π' according to the following rule, consequence of Observation 3.2 in a path and the fact that we do not add unnecessary swaps: for two agents i' and j' such that $i' < j' < i$, the object received by agent i' must have a smaller index than the object received by agent j'. This rationality rule implies that the completion of π' is unique. Thus, by minimality of the sequence of swaps leading to π^t, π' must correspond to allocation π^t if the guess about the objects of $\{o_{i_1}, \ldots, o_{i_k}\}$ is correct. Then, it suffices to test whether allocation π' is reachable by a rational sequence of swaps, a problem that is solved in polynomial time in the next section (Proposition 3.16). Clearly, if there exists a minimal sequence of swaps $(\pi^0, \ldots, \pi^t)$ such that $\pi^t(i) = o_j$ then, by construction of π', for a certain guess we build $\pi' := \pi^t$ and thus when we test the reachability of π', we get a positive answer. □

Despite its apparent simplicity, Reachable Object in a path is a challenging open problem when no restriction on the agent's location is made.

3.3.2 Maximum number of swaps per agent

Let us now impose some restrictions on the sequence of swaps. We first consider that the agents are not willing to perform an important number of swaps in the whole swap process. Consequently, we study the problem RO-max, asking whether there exists a sequence of swaps leading to give target object x to target agent A, where every agent cannot make more than k swaps in the sequence. Surprisingly, in this context, Reachable Object remains computationally hard even for a very small maximum number of swaps.

Theorem 3.6 *For fixed $k \geq 2$, RO-max is NP-complete, even when the degree Δ_G of the graph is equal to 4.*

Proof: Membership in NP is straightforward, as it is a special case of the unconstrained RO problem, known to be in NP (Theorem 3.1).

For hardness, we fix $k = 2$ and reduce from (3, B2)-SAT, a variant of the Satisfiability problem known to be NP-complete (Theorem 1.4). The (3,B2)-SAT problem is the restriction of sat to instances where each clause contains three literals and each variable occurs exactly twice as a positive literal and twice as a negative literal.

From an instance of (3, B2)-SAT with a set $X = \{x_1, \ldots, x_v\}$ of v variables and a set $\mathcal{C} = \{C_1, \ldots, C_s\}$ of s clauses, we construct an instance of Reachable Object as follows. We create a literal-agent Y_j^ℓ (respectively, $\overline{Y_j^\ell}$) for each ℓ^{th} (with $\ell \in \{1, 2\}$) occurrence of literal x_j (respectively, $\overline{x_j}$), and a variable-agent Y_j for each variable x_j. Two clause-agents K_i and K_i' are created for each $i \in [s-1]$. Three other agents Y_0, K_0' and K_s are added. Each agent initially owns an object denoted by the lower-case version of her name, e.g., agent K_i gets object k_i.

In the network, we have the paths $[Y_{j-1}, Y_j^1, Y_j^2, Y_j]$ and $[Y_{j-1}, \overline{Y_j^1}, \overline{Y_j^2}, Y_j]$ for each $j \in [v]$, and the edge $\{K_i, K_i'\}$ for each $i \in [s-1]$. If the ℓ^{th} occurrence of literal x_j (respectively, $\overline{x_j}$) appears in clause C_i, then we have the path $[K_{i-1}', Y_j^\ell, K_i]$ (respectively, $[K_{i-1}', \overline{Y_j^\ell}, K_i]$). We connect K_s and Y_v. See Figure 3.2 for an example of the graph construction.

The preferences of the agents are given below. Notation $\{\ell_i\}$ stands for the literal-objects of clause C_i ranked in arbitrary order and, like in the proof of Theorem 3.1, index

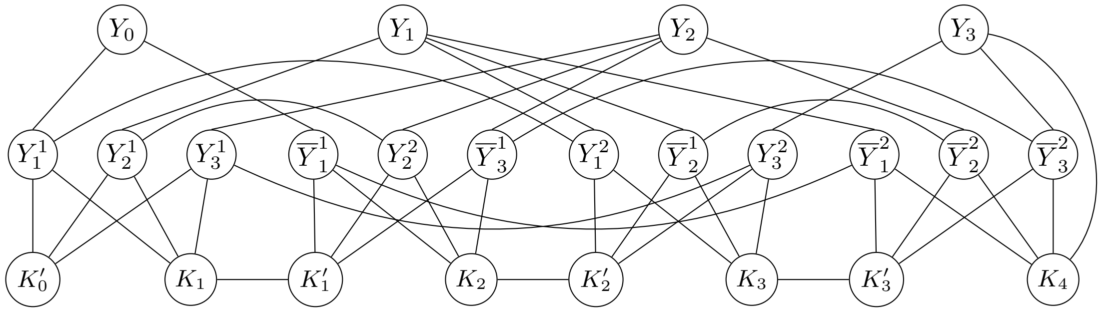

Figure 3.2: Graph construction for an instance of (3, B2)-SAT with four clauses where $C_1 = (x_1 \vee x_2 \vee x_3)$, $C_2 = (\overline{x_1} \vee x_2 \vee \overline{x_3})$, $C_3 = (x_1 \vee \overline{x_2} \vee x_3)$, and $C_4 = (\overline{x_1} \vee \overline{x_2} \vee \overline{x_3})$.

p_ℓ^j (respectively, n_ℓ^j) refers to the clause containing the ℓ^{th} occurrence (with $\ell \in \{1, 2\}$) of x_j (respectively, $\overline{x}_j$). The objects that are not mentioned in the preferences are ranked in arbitrary order after the initial endowment (notation $[\dots]$).

$$
\begin{array}{ll}
K_0': & \{\ell_1\} \succ \boxed{k_0'} \succ [\dots] \\
K_i: & k_i' \succ k_0' \succ \boxed{k_i} \succ [\dots] \\
K_i': & \{\ell_{i+1}\} \succ k_0' \succ \boxed{k_i'} \succ [\dots] \\
K_s: & y_0 \succ k_0' \succ \boxed{k_s} \succ [\dots] \\
Y_0: & \overline{y}_1^1 \succ y_1^1 \succ \boxed{y_0} \succ [\dots] \\
Y_v: & k_0' \succ y_0 \succ \boxed{y_v} \succ [\dots]
\end{array}
\qquad
\begin{array}{ll}
Y_j^1: & k_{p_1^j} \succ k_0' \succ y_j^2 \succ y_0 \succ \boxed{y_j^1} \succ [\dots] \\
Y_j^2: & k_{p_2^j} \succ k_0' \succ y_j \succ y_0 \succ \boxed{y_j^2} \succ [\dots] \\
\overline{Y}_j^1: & k_{n_1^j} \succ k_0' \succ \overline{y}_j^2 \succ y_0 \succ \boxed{\overline{y}_j^1} \succ [\dots] \\
\overline{Y}_j^2: & k_{n_2^j} \succ k_0' \succ y_j \succ y_0 \succ \boxed{\overline{y}_j^2} \succ [\dots] \\
Y_j: & \overline{y}_{j+1}^1 \succ y_{j+1}^1 \succ y_0 \succ \boxed{y_j} \succ [\dots]
\end{array}
$$

We claim that all the clauses in $\mathcal{C}$ are satisfiable if and only if object k_0' reaches agent Y_v in the constructed instance of RO-max where k is fixed to 2.

Note that the only way for Y_v to get hold of k_0' is by swapping y_0 with K_s. Indeed, by construction of the preferences, Y_v can only obtain object k_0' via a swap with agent K_s, and K_s accepts to exchange k_0' only against object y_0. Therefore, object k_0' must reach agent K_s while object y_0 must reach agent Y_v. Object k_0' can only reach K_s via clause-agents and literal-agents, while y_0 can only reach Y_v via variable-agents and literal-agents. Agents perform at most two swaps, so no literal-agent can be involved in the move of both y_0 and k_0'.

Suppose that truth assignment ϕ satisfies all the clauses of $\mathcal{C}$. Let T_i be a literal-agent of clause C_i related to a true literal in ϕ. Since all the clauses are satisfiable, there exists such an agent T_i for each clause C_i and thus, object k_0' can reach K_s via the path $[K_0', T_1, K_1, K_1', T_2, \dots, T_{s-1}, K_{s-1}, K_{s-1}', T_s, K_s]$. For variable x_j, let Z_j^1 and Z_j^2 be the literal-agents associated with the literals of x_j that are false in ϕ. By definition, these agents do not belong to $\bigcup_{i \in [s]} T_i$. Therefore, it suffices for y_0 to reach Y_v via the path $[Y_0, Z_1^1, Z_1^2, Y_1, \dots, Y_{v-1}, Z_v^1, Z_v^2, Y_v]$.

Suppose now that object k_0' is reachable for agent Y_v. By construction, the path of k_0' to K_s goes through exactly one literal-agent per clause, while the path of y_0 to Y_v goes through exactly two literal-agents associated with the same literal for each variable. Thus, the truth assignment of variables that sets to true the literals related to literal-agents in the path of object k_0', satisfies all the clauses.

If $k > 2$, we adapt the reduction via a delay gadget added to each agent. □

One could think that the problem is much easier when the structure of the network is restricted to trees. Yet, it is possible to prove that RO-max on trees is W[SAT]-hard (where W[SAT] is a "high" parameterized complexity class in the W-hierarchy, as we can see in Figure 1.4), which notably rules out, under standard complexity assumptions, the existence of an FPT algorithm.

Theorem 3.7 RO-*max is* W[SAT]-*hard in a tree.*

Sketch of proof: We reduce from the MONOTONE WEIGHTED SATISFIABILITY problem, known to be W[SAT]-complete (Theorem 1.8). MONOTONE WEIGHTED SATISFIABILITY asks whether a given monotone propositional formula φ over a set of variables $X = \{x_1, \ldots, x_v\}$ is satisfiable with a truth assignment over X where at most k variables are set to true. A formula φ is said to be monotone if all the variables in X only occur through positive literals in φ. Recall that the syntax tree of a propositional formula φ is a tree whose internal nodes, i.e., nodes of degree greater than one, are labeled by conjunctives of the formula and whose leaves represent propositional variables. Each node in the syntax tree corresponds to a subformula whose syntax tree is the subtree rooted at that node. Indeed, the children of a node corresponding to a connective of φ are the immediate subformulas on which the connective applies in φ. We suppose, without loss of generality, that all the relations in φ are binary (because operators $\wedge$ and $\vee$ satisfy $\varphi_1 \wedge \varphi_2 \wedge \varphi_3 \equiv (\varphi_1 \wedge \varphi_2) \wedge \varphi_3$ and $\varphi_1 \vee \varphi_2 \vee \varphi_3 \equiv (\varphi_1 \vee \varphi_2) \vee \varphi_3$), and that the root of the syntax tree of φ is a conjunction. The number of occurrences of variable x_i in φ is denoted by n_i.

An instance of RO-max in a tree is constructed by building a graph based on the syntax tree T_φ of formula φ. The goal is to make objects associated with variables in X reach agents symbolizing their occurrences in the formula. We construct a swap dynamics instance $\langle N, O, \succ, G, \pi^0 \rangle$ for RO-max where by convention we denote all the agents with upper-case letters and all the objects with lower-case letters. The initial allocation π^0 assigns to each agent the object denoted by the lower-case version of its name, i.e., an agent $T \in N$ initially owns object $t \in M$. For understanding the details of the construction, referring to the example provided in Figure 3.3 can be useful.

For each binary relation of formula φ, we create an agent R_i^0 and an object r_i^0, which refer to the i^{th} binary relation of the formula when considering a fixed order induced by the syntax tree of φ. The set of all these relation-agents is denoted by R.

For every variable x_j and every $\ell \in \{1, \ldots, n_j\}$, we create an agent $Y_{j,\ell}^0$ and an object $y_{j,\ell}^0$, which represent the ℓ^{th} occurrence of variable x_j in formula φ. The set of all these literal-agents is denoted by Y. A variable-agent Z_j^φ and her initial object z_j^φ are created for every variable x_j; they refer to the global variable x_j in the formula.

Let us denote by $\Psi_{i,1}$ and $\Psi_{i,2}$ (respectively, $\psi_{i,1}$ and $\psi_{i,2}$) the two agents (respectively, the two objects) related to the two members of the binary relation represented by agent $R_i^0 \in R$. The agents $\Psi_{i,1}$ and $\Psi_{i,2}$ can be either members of R, or members of Y. Each agent R_i^0 is connected in the social network G to the agents $\Psi_{i,1}$ and $\Psi_{i,2}$, like in the associated syntax tree T_φ.

For each relation-agent $R_i^0 \in R$, $v+1$ copies are created: the agents R_i^t and the associated objects r_i^t, for every $t \in \{1, \ldots, v+1\}$. For each literal-agent $Y_{j,\ell}^0$, v copies are created: the agents $Y^t_{j,\ell}$ and the associated objects $y^t_{j,\ell}$ for every $t \in \{1, \dots, v\}$. Copies $Y^t_{j,\ell}$ and $y^t_{j,\ell}$ are associated with variable x_t for $t \in [v]$. All the copies R^t_i of agent R^0_i are connected to R^0_i in the social network, and all the copies $Y^t_{j,\ell}$ of agent $Y^0_{j,\ell}$ are connected to $Y^0_{j,\ell}$. Here is a part of the preferences of R^0_i and $Y^0_{j,\ell}$ which will be further detailed in the description of the next gadgets as well as the preferences of their copies:

$$\begin{array}{rl} R^0_i: & \ldots \succ r^{v+1}_i \succ \ldots \succ r^v_i \succ \ldots \succ r^1_i \succ \ldots \succ \boxed{r^0_i} \\ Y^0_{j,\ell}: & \ldots \succ y^v_{j,\ell} \succ \ldots \succ y^1_{j,\ell} \succ \ldots \succ \boxed{y^0_{j,\ell}} \end{array}$$

We add three agents C, C' and D with their associated initial object c, c' and d. In the social network, agent C' is connected to agent C who is herself connected to agent D. Moreover, agent C is connected to all agents Z^φ_j for $j \in \{1, \dots, v\}$, and to agent R^0_1, corresponding to the connective at the root of the syntax tree (which is a conjunction by assumption). These agents have the following preferences, where notation [...] represents the rest of the objects ranked in arbitrary order, and t_1 will be explained later:

$$\begin{array}{rl} Z^\varphi_1: & c \succ \boxed{z^\varphi_1} \succ [\ldots] \\ Z^\varphi_j: & r^{j-1}_1 \succ \boxed{z^\varphi_j} \succ [\ldots] \\ C: & d \succ t_1 \succ c' \succ r^{n-1}_1 \succ z^\varphi_v \succ \ldots \succ r^1_1 \succ z^\varphi_2 \succ r^0_1 \succ z^\varphi_1 \succ \boxed{c} \succ [\ldots] \\ C': & r^{v-1}_1 \succ \boxed{c'} \succ [\ldots] \\ D: & t_1 \succ \boxed{d} \succ [\ldots] \end{array}$$

The occurrence of variable x_j in the global formula φ is represented by agent Z^φ_j and object z^φ_j. We present here a gadget allowing to "duplicate" this object each time it (or a previous duplicated object) passes by a relation-agent R^0_i: this object is duplicated into two objects associated with the two members of the relation in φ represented by R^0_i, then each copy associated with each subformula can move to the part of the network representing this subformula. For each agent R^0_i and each variable x_j, two agents $Z^{\psi_{i,1}}_j$ and $Z^{\psi_{i,2}}_j$ and two objects $z^{\psi_{i,1}}_j$ and $z^{\psi_{i,2}}_j$ are created. Agent R^0_i is connected in the social network G to agent $Z^{\psi_{i,1}}_j$ who is herself connected to agent $Z^{\psi_{i,2}}_j$. The agents involved in this gadget have the following preferences (for $t \in [v]$):

$$
\begin{array}{rl}
R_1^0: & \ldots \succ r_1^j \succ \psi_{1,2}^{j-1} \succ z_j^{\psi_{1,2}} \succ \psi_{1,1}^{j-1} \succ z_j^{\psi_{1,1}} \succ z_j^{\varphi} \succ r_1^{j-1} \succ \ldots \\
R_i^0: & \ldots \succ r_i^j \succ \psi_{i,2}^{j-1} \succ z_j^{\psi_{i,2}} \succ \psi_{i,1}^{j-1} \succ z_j^{\psi_{i,1}} \succ z_j^{r_i^0} \succ r_i^{j-1} \succ \ldots \\
R_i^t: & \psi_{i,2}^{t-1} \succ \ldots \succ \psi_{i,2}^{1} \succ r_i^{t-1} \succ \ldots \succ r_i^1 \succ \boxed{r_i^t} \succ [\ldots] \\
Y_{j,\ell}^0: & \ldots \succ y_{j,\ell}^n \succ \ldots \succ y_{j,\ell}^j \succ \ldots \succ z_j^{y_{j,\ell}^0} \succ y_{j,\ell}^{j-1} \succ \ldots \succ y_{j,\ell}^1 \succ z_1^{y_{j,\ell}^0} \succ \boxed{y_{j,\ell}^0} \succ [\ldots] \\
Y_{j,\ell}^t: & z_{t-1}^{y_{j,\ell}^0} \succ \ldots \succ z_1^{y_{j,\ell}^0} \succ y_{j,\ell}^{t-1} \succ \ldots \succ y_{j,\ell}^1 \succ \boxed{y_{j,\ell}^t} \succ [\ldots] \\
Z_j^{\psi_{1,1}}: & \psi_{1,1}^{j-1} \succ z_j^{\psi_{1,2}} \succ z_j^{\varphi} \succ \boxed{z_j^{\psi_{1,1}}} \\
Z_j^{\psi_{1,2}}: & z_j^{\varphi} \succ \boxed{z_j^{\psi_{1,2}}} \succ [\ldots] \\
Z_j^{\psi_{i,1}}: & \psi_{i,1}^{j-1} \succ z_j^{\psi_{i,2}} \succ z_j^{r_i^0} \succ \boxed{z_j^{\psi_{i,1}}} \succ [\ldots] \\
Z_j^{\psi_{i,2}}: & z_j^{r_i^0} \succ \boxed{z_j^{\psi_{i,2}}} \succ [\ldots]
\end{array}
$$

We are now able to design a gadget to check the validation, i.e., the satisfaction, of a given subformula rooted at a node U of the syntax tree. This node U can either corresponds in G to a literal-agent in Y or to a relation-agent in R; this corresponding agent is denoted by $N(U)$ and her associated object $M(U)$. Conversely, for an agent $B \in R \cup Y$, the associated node in the syntax tree of φ is given by $T_\varphi(B)$. The parent of a node U in T_φ is denoted by $Pre(U)$. For the sake of simplicity, for any agent $B \in (R \cup Y) \setminus \{R_1^0\}$, we denote $M(Pre(T_\varphi(B)))$ (respectively, $N(Pre(T_\varphi(B))))$ by $p(b)$ (respectively, $P(b)$), which represents the object (respectively, the agent) associated with the predecessor of the node corresponding to B in T_φ (which is necessarily a relation-agent). If node U corresponds to an agent in Y and more precisely to an agent $Y_{j,\ell}^0$, then two agents $T_{j,\ell}$ and $T_{j,\ell}^1$ are created, with their associated initial object $t_{j,\ell}$ and $t_{j,\ell}^1$. In the social network G, agent $Y_{j,\ell}^0$ is connected to agent $T_{j,\ell}^1$, who is herself connected to agent $T_{j,\ell}$. Literal-agent $Y_{j,\ell}^0$ is validated when she receives the variable-object $z_t^{y_{j,\ell}^0}$ corresponding to her variable, i.e., $z_j^{y_{j,\ell}^0}$. In such a case, the exchange with agent $T_{j,\ell}^1$ is allowed, in order to keep in memory that $Y_{j,\ell}^0$ is validated. The preferences of the agents involved in this gadget are the following for every $j \in [v]$ and every $\ell \in [n_j]$:

$$
\begin{array}{rl}
Y_{j,\ell}^0: & p(y_{j,t})^{v+1} \succ p(y_{j,t})^{v} \succ t_{j,\ell} \succ y_{j,\ell}^v \succ \ldots \succ y_{j,\ell}^j \succ t_{j,\ell}^1 \succ z_j^{y_{j,\ell}^0} \succ y_{j,\ell}^{j-1} \succ \ldots \succ \boxed{y_{j,\ell}^0} \succ [\ldots] \\
Y_{j,\ell}^j: & t_{j,\ell}^1 \succ \ldots \succ \boxed{y_{j,\ell}^j} \succ [\ldots] \\
T_{j,\ell}^1: & y_{j,\ell}^v \succ t_{j,\ell} \succ z_j^{y_{j,\ell}^0} \succ \boxed{t_{j,\ell}^1} \succ [\ldots] \\
T_{j,\ell}: & z_j^{y_{j,\ell}^0} \succ \boxed{t_{j,\ell}} \succ [\ldots]
\end{array}
$$

Now consider the case where node U corresponds to an agent in R, i.e., to an agent R_i^0. From the set R of agents related to the binary relations of formula φ, we distinguish the agents associated with conjunctions, whose set is denoted by $\mathscr{A}$, and the agents associated with disjunctions, whose set is denoted by $\mathscr{O}$. If agent $R_i^0 \in \mathscr{O}$, one agent T_i and her associated object t_i are created. Agents R_i^0 and T_i are connected in the social network G. By simplicity, we denote by $T_{\psi_{i,j}}$ and $t_{\psi_{i,j}}$ the validation-agent and the validation-object associated with $\Psi_{i,j}$. Relation-agent R_i^0 is validated if at least one of the agents associated with the subformulas, i.e., agents $\Psi_{i,1}$ and $\Psi_{i,2}$, is validated, i.e., R_i^0 receives from them object $t_{\psi_{i,1}}$ or $t_{\psi_{i,2}}$, respectively. The agents involved in this gadget have the following preferences:

$$\begin{aligned} R_i^0 :&\quad p(r_i)^{v+1} \succ p(r_i)^v \succ t_i \succ t_{\psi_{i,1}} \succ t_{\psi_{i,2}} \succ r_i^{v+1} \succ r_i^v \succ \dots \\ R_i^{v+1} :&\quad r_i^v \succ \boxed{r_i^{v+1}} \\ T_i :&\quad t_{\psi_{i,1}} \succ t_{\psi_{i,2}} \succ \boxed{t_i} \succ [\dots] \end{aligned}$$

Note that object r_i^{v+1} is not necessary for the gadget when $R_i^0 \in \mathscr{O}$ but since it is needed when $R_i^0 \in \mathscr{A}$, we also introduce it for disjunctions in order to be more general and avoid to check whether the predecessor relation is a conjunction.

Otherwise, that is if $R_i^0 \in \mathscr{A}$, two agents T_i^1 and T_i are created with their associated objects t_i^1 and t_i. Agent R_i^0 is connected to agent T_i^1 who is herself connected to T_i in the social network G. Relation-agent R_i^0 is validated if both of the agents associated with the subformulas, i.e., agents $\Psi_{i,1}$ and $\Psi_{i,2}$, are validated, i.e., R_i^0 receives from them objects $t_{\psi_{i,1}}$ and $t_{\psi_{i,2}}$, respectively. The preferences of the agents involved in this gadget are the following:

$$\begin{aligned} R_1^0 :&\quad c' \succ t_1 \succ t_{\psi_{1,2}} \succ r_1^{v+1} \succ t_1^1 \succ t_{\psi_{1,1}} \succ r_1^v \succ \dots \\ R_i^0 :&\quad p(r_i)^{v+1} \succ p(r_i)^v \succ t_i \succ t_{\psi_{i,2}} \succ r_i^{v+1} \succ t_i^1 \succ t_{\psi_{i,1}} \succ r_i^v \succ \dots \\ R_i^{v+1} :&\quad t_i^1 \succ \boxed{r_i^{v+1}} \succ [\dots] \\ T_i :&\quad t_{\psi_{i,1}} \succ \boxed{t_i} \succ [\dots] \\ T_i^1 :&\quad t_{\psi_{i,2}} \succ t_i \succ t_{\psi_{i,1}} \succ \boxed{t_i^1} \succ [\dots] \end{aligned}$$

We claim that formula φ is satisfiable with weight k if and only if object d is reachable by agent C in instance $\langle N, O, \succ, G, \pi^0 \rangle$ of RO-max where the maximum number of exchanges per agent is $6k + 7$. The global idea is, given a truth assignment ϕ of weight k which satisfies φ, for each variable x_j assigned to true in ϕ (by increasing order of the variable index), the associated object z_j^φ must move from agent Z_j^φ. This object is duplicated at each binary relation-agent R_i^0 into literal-objects $z_j^{\psi_{i,1}}$ and $z_j^{\psi_{i,2}}$ corresponding to the subformulas of R_i^0, until the duplicated literal-objects reach literal-agents. A literal-agent $Y_{j,\ell}^0$ is then validated when she receives a literal-object corresponding to her variable, namely object $z_j^{y_{j,\ell}^0}$. By induction, a relation-agent R_i^0 corresponding to a conjunction is validated when both agents $\Psi_{i,1}$ and $\Psi_{i,2}$ associated with its subformulas are validated, and a relation-agent R_i^0 corresponding to a disjunction is validated when at least one of the agents $\Psi_{i,1}$ and $\Psi_{i,2}$ is validated. Since truth assignment ϕ satisfies formula φ, it follows that binary-agent R_1^0, corresponding to the binary relation at the root of the syntax tree of φ, is eventually validated. This enables the exchanges among the agents C, C' and D in order to give object d to agent C. Refer to Section A of the appendix to see the complete proof of the equivalence. □

The previous hardness result highlights the computational challenge of RO-max on trees, and enables to stress the relevance of the next positive results when both the graph structure and the parameter are constrained.

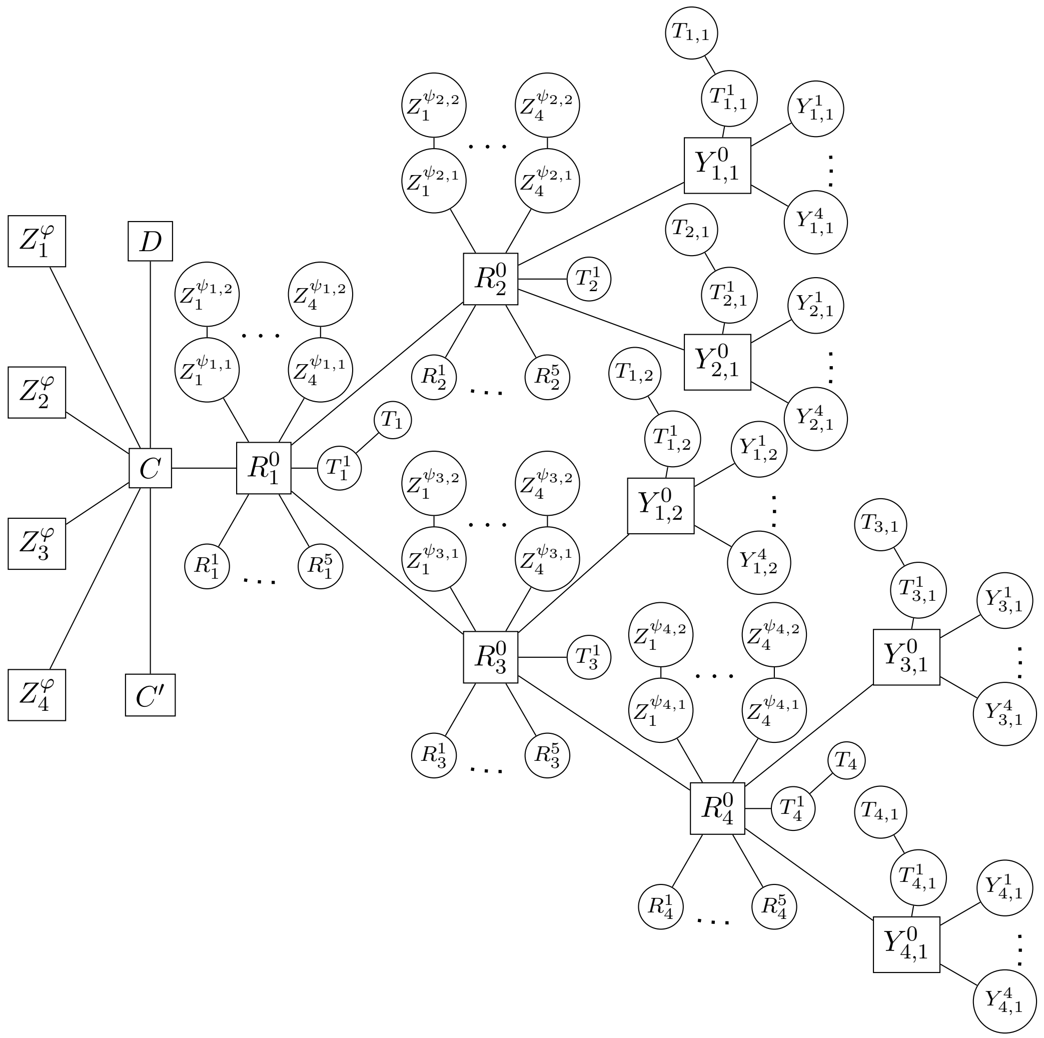

Figure 3.3: Example of the construction of the social network for a monotone propositional formula φ such that $\varphi \equiv (x_1 \vee x_2) \wedge (x_1 \vee (x_3 \wedge x_4))$.

Proposition 3.8 RO-*max is solvable in polynomial time in a tree if* $k = 2$.

Proof: There exists a unique path τ_x in G between the owner of object x, say X, and agent A. All the agents in τ_x, except X and A, perform their two swaps to move x towards A. However, in order to get x, A can perform another swap so as to obtain an object y preferred to x by her neighbor in τ_x. It suffices to check the reachability to A of any such y, when agent A is only allowed to make the swap to get y, provided τ_x and τ_y do not cross. □

By restricting even more the structure of the graph, we can derive a tractability result when parameter k is a constant, which corresponds to class XP (see Definition 1.33). This assumption is relevant since the maximum number of swaps that an agent is able to perform has no reason to increase with the number of agents involved in the process.

Proposition 3.9 RO-*max is in* XP *if the network is a path.*

Proof: Like in Section 3.3.1 for the case of the path, we may assume, without loss of generality, that $E = \{\{j, j+1\} : 1 \leq j < n\}$. Like in the proof of Proposition 3.5, we consider, without loss of generality, that we aim at determining the reachability of object o_n, initially owned by n, for agent i. Note that each agent j such that $i < j < n$ necessarily performs two swaps to transfer o_n towards agent i. Since $G = (N, E)$ is a path, only two directions are allowed for the objects: to the "left" (decreasing agent indices) or to the "right" (increasing indices). By rationality of the swaps, once an object moves in some direction, it cannot come back the other direction (Observation 3.2). Thus, two objects o_j and o_ℓ respectively possessed by agents j and ℓ for $j < \ell$, which both move in the same direction, must be finally possessed by some agents j' and ℓ' such that $j' < \ell'$. Therefore, it suffices to know which objects move to the left or to the right to know the objects which are finally assigned to the agents $i+1, \ldots, n$, and so to deduce the associated sequence of swaps and determine whether it satisfies the rationality condition. Since at most k swaps are allowed per agent, then only $\lfloor k/2 \rfloor$ objects can pass by each agent. This implies that at most $\lfloor k/2 \rfloor - 1$ objects can pass by agent i, including objects in O_r, defined as the set of objects possessed by agents in $\{1, \ldots, i-1\}$ which move to her right, and objects in O_ℓ, defined as the set of objects possessed by agents in $\{i+1, \ldots, n-1\}$ which move to her left. Thus, knowing the objects in O_r and O_ℓ which pass by i boils down to know the final allocation for which we must test the reachability, and this test can be done in polynomial time (Proposition 3.16). There are $\mathcal{O}(n^k)$ different such allocations, hence the problem is solvable in $\mathcal{O}(n^k)$. □

In a nutshell, the tractable cases of RO-max that we have identified are limited to very specific configurations. It would be interesting to determine an upper bound for the complexity of RO-max in trees. Although the W[SAT]-hardness of RO-max in trees (Theorem 3.7) does not prevent the membership of the problem to XP, we conjecture that it is not the case.

3.3.3 Length of the sequence of swaps

Two parameters are used to bound the length of the sequence of swaps: the total number of exchanges and the makespan. These two parameters enable to define two parameterized versions for Reachable Object: RO-sum and RO-makespan. Contrary to the previous parameter *max*, these parameters on the length of the sequence lead to circumscribe the problems into parameterized complexity classes that are not so high in the hierarchy. For instance, they allow to obtain tractability results when the parameters are bounded by a constant for any social network. Moreover, for bounded degree graphs, relevant in the context of a social network modeling the possibility of collaboration, we obtain fixed-parameter tractability.

Nevertheless, in general, the problems are still computationally hard, even when the social network is a tree.

Theorem 3.10 RO-*sum and* RO-*makespan are* W[1]-*hard even for trees.*

Sketch of proof: We perform a reduction from Clique, problem known to be W[1]-complete (Theorem 1.6). Clique is the problem of deciding whether there exists a clique of size k in an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ such that $\mathcal{V} = \{1, ..., s\}$ and $|\mathcal{E}| = p$. Assume

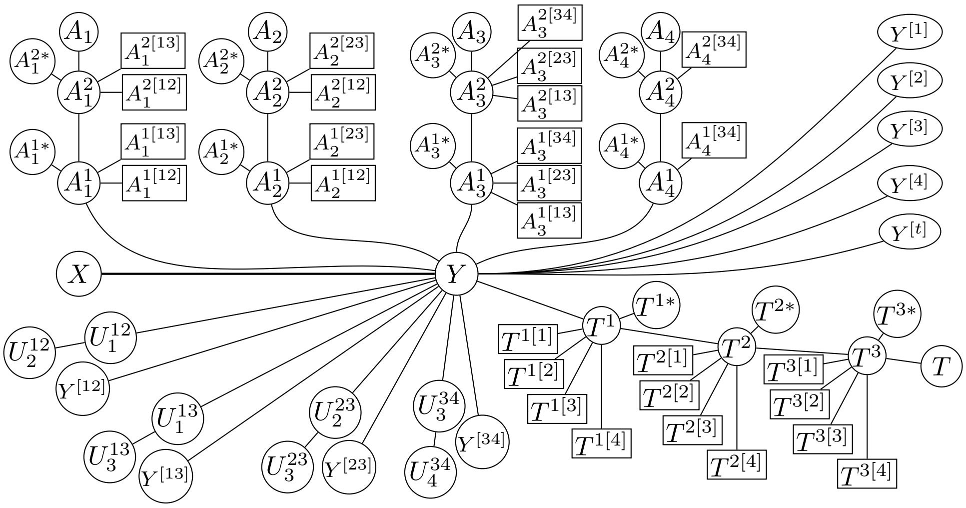

Figure 3.4: Graph construction for an instance of Clique with vertices $\mathcal{V} = \{1, 2, 3, 4\}$ and $k = 3$. The edges are: $\{1, 2\}$, $\{1, 3\}$, $\{2, 3\}$ and $\{3, 4\}$.

that each edge in $\mathcal{E}$ is written (v, w) such that $v < w$, and consider the lexicographical order over $\mathcal{E}$. Let us denote by e_i^1 and e_i^2 the first and second vertex of the i^{th} edge. Recall that the degree of a vertex v in $\mathcal{G}$ is denoted by $\delta_{\mathcal{G}}(v)$. For the sake of simplicity, we omit the reference to $\mathcal{G}$ and simply write $\delta(v)$ for the degree of v in $\mathcal{G}$ and denote by $\delta^d(v)$ the d^{th} edge incident to v, for $1 \leq d \leq \delta(v)$. We construct an instance $\mathcal{I}'$ of RO (see Figure 3.4 for an example) by creating:

- two connected agents X and Y, and two vertex-agents U_v^{vw} and U_w^{vw} for each edge $(v, w) \in \mathcal{E}$, connected via a path $[Y, U_v^{vw}, U_w^{vw}]$.
- agents T and T^ℓ, for $1 \leq \ell \leq k$, representing the k vertices of the clique that we must choose. They are connected via a path to Y: $[Y, T^1, \ldots, T^k, T]$.
- agents A_v and A_v^ℓ, for $v \in \mathcal{V}$ and $1 \leq \ell < k$, representing the choice of the $k-1$ edges of the clique that are incident to v if v belongs to the clique. They are connected via a path to Y for each v: $[Y, A_v^1, \ldots, A_v^{k-1}, A_v]$.
- agent $T^{\ell*}$ connected to T^ℓ, for $1 \leq \ell \leq k$, and agent $A_v^{\ell*}$ connected to A_v^ℓ, for $1 \leq \ell < k$ and $v \in \mathcal{V}$. They are used to "validate" their associated agent by giving to her their initial object once they own an expected object.
- auxiliary agents used to facilitate the passage of some objects: if an agent B has a connected auxiliary agent $B^{[z]}$, then the swap with $B^{[z]}$ must precede a swap for getting an object associated with z. The auxiliary agents we use are agents $Y^{[vw]}$ corresponding to edge (v, w) and connected to Y, agents $Y^{[v]}$ corresponding to vertex v and connected to Y, agent $Y^{[t]}$ corresponding to object t and connected to Y, agents $A_v^{\ell[\delta^d(v)]}$ corresponding to edge $\delta^d(v)$, for $1 \leq d \leq \delta(v)$, and connected to agent A_v^ℓ, for $v \in \mathcal{V}$ and $1 \leq \ell < k$, and agents $T^{\ell[v]}$ corresponding to vertex v and connected to agent T^ℓ for $1 \leq \ell \leq k$.

The initial object of an agent is denoted by the lower-case version of her name, e.g., agent $Y^{[v]}$ gets object $y^{[v]}$. The preferences of the agents are as follows (objects in brackets may not exist for all indices).

$$
\begin{aligned}
T &: \quad t^{k*} \succ \boxed{t} \succ [\dots] \\
A_v &: \quad a_v^{k-1*} \succ \boxed{a_v} \succ [\dots] \\
X &: \quad t \succ \boxed{x} \succ [\dots] \\
U_w^{vw} &: \quad y^{[vw]} \succ \boxed{u_w^{vw}} \succ [\dots] \\
U_v^{vw} &: \quad a_v^{1[vw]} \succ u_w^{vw} \succ y^{[vw]} \succ \boxed{u_v^{vw}} \succ [\dots] \\
A_v^{\ell*} &: \quad u_v^{\delta^{\delta(v)}(v)} \succ \dots \succ u_v^{\delta^1(v)} \succ \boxed{a_v^{\ell*}} \succ [\dots] \\
A_v^{\ell[\delta^d(v)]} &: \quad (a_v^{\ell+1[\delta^{\delta(v)}(v)]}) \succ \dots \succ (a_v^{\ell+1[\delta^1(v)]}) \succ a_v^\ell \succ \boxed{a_v^{\ell[\delta^d(v)]}} \succ [\dots] \\
T^{\ell*} &: \quad a_1 \succ a_2 \succ \dots \succ a_s \succ \boxed{t^{\ell*}} \succ [\dots] \\
T^{\ell[v]} &: \quad (t^{\ell+1[v-1]}) \succ \dots \succ (t^{\ell+1[1]}) \succ t^\ell \succ \boxed{t^{\ell[v]}} \succ [\dots] \\
Y^{[v]} &: \quad (t^{1[v-1]}) \succ \dots \succ t^{1[1]} \succ a_{e_p^2}^{1[e_p]} \succ \dots \succ a_{e_1^2}^{1[e_1]} \succ \boxed{y^{[v]}} \succ [\dots] \\
Y^{[vw]} &: \quad a_{e_p^2}^{1[e_p]} \succ \dots \succ a_{e_1^2}^{1[e_1]} \succ y \succ \boxed{y^{[vw]}} \succ [\dots] \\
A_v^\ell &: \quad y^{[v]} \succ (a_v^{\ell-1*}) \succ a_v \succ a_v^{\ell*} \succ u_v^{\delta^{\delta(v)}(v)} \succ a_v^{\ell[\delta^{\delta(v)}(v)]} \succ \dots \succ (a_v^{\ell+1[\delta^2(v)]}) \succ \\
&\quad\;\; u_v^{\delta^2(v)} \succ a_v^{\ell[\delta^2(v)]} \succ (a_v^{\ell+1[\delta^1(v)]}) \succ u_v^{\delta^1(v)} \succ a_v^{\ell[\delta^1(v)]} \succ \boxed{a_v^\ell} \succ [\dots] \\
T^\ell &: \quad y^{[t]} \succ (t^{\ell-1*}) \succ t \succ t^{\ell*} \succ a_s \succ t^{\ell[s]} \succ \dots \succ (t^{\ell+1[2]}) \succ a_2 \succ t^{\ell[2]} \succ \\
&\quad\;\; (t^{\ell+1[1]}) \succ a_1 \succ t^{\ell[1]} \succ \boxed{t^\ell} \succ [\dots] \\
Y &: \quad x \succ t \succ y^{[t]} \succ t^{1[s]} \succ a_s \succ y^{[s]} \succ \dots \succ t^{1[1]} \succ a_1 \succ y^{[1]} \succ a_{e_p^2}^{1[e_p]} \succ u_{e_p^2}^{e_p} \succ \\
&\quad\;\; a_{e_p^1}^{1[e_p]} \succ u_{e_p^1}^{e_p} \succ y^{[e_p]} \succ \dots \succ a_{e_1^2}^{1[e_1]} \succ u_{e_1^2}^{e_1} \succ a_{e_1^1}^{1[e_1]} \succ u_{e_1^1}^{e_1} \succ y^{[e_1]} \succ \boxed{y} \succ [\dots]
\end{aligned}
$$

We claim that there exists a clique of size k in graph $\mathcal{G}$ if and only if object x can reach agent Y within a total of k^3+4k^2+k+2 swaps or a makespan of $5k(k-1)/2+3k+4$ in instance $\mathcal{I}'$ of Reachable Object. An agent A_v^ℓ (or T^ℓ) is said to be "validated" if she obtains at a moment object $a_v^{\ell*}$ (or $t^{\ell*}$). In order to make object x reach agent Y, all the k agents T^ℓ and all the $k-1$ agents A_v^ℓ of k branches A_v need to be validated. The associated clique in graph $\mathcal{G}$ is given by the vertices v for which all the $k-1$ agents A_v^ℓ have been validated. All the agents A_v^ℓ, for $1 \leq \ell < k$, are validated if we can bring in the branch $k-1$ objects u_v^{vw} (or u_v^{wv}, following the order) representing an edge incident to v. Observe that the given budget allows bringing in the branches only $k(k-1)$ objects u_v^{vw} and the construction forces to choose u_w^{vw} if u_v^{vw} has been chosen. Refer to Section B of the appendix to see the complete proof of the equivalence. □

This W[1]-hardness for RO parameterized by the length of the sequence rules out the existence of FPT algorithms, even in trees, under standard complexity assumptions. However, the following W[1] membership result, based on a reduction to a model checking problem of existential first-order formulas (see Section 1.5.2.b), shows that the problem is not so hard. It is notably in XP for any graph, thus tractable when the parameter is a constant.

Theorem 3.11 RO-*sum is in* W[1].

Proof: An instance $\mathcal{I}$ of RO with a swap dynamics instance $\langle N, O, \succ, G, \pi^0 \rangle$, agent A and object x, and k as a total number of swaps, is transformed into an instance $\mathcal{I}' = \langle \mathcal{A}, \varphi \rangle$ of MC(Σ_1), known to be W[1]-complete (Theorem 1.9). Let us consider the following vocabulary $\tau_{RO} := \{E, SUCC, ALLOC, TARG_A, TARG_x\}$. Structure $\mathcal{A}$ is a τ_{RO}-structure over the universe $N \cup O$ where:

- $E^{\mathcal{A}}$ is the edge relation over N^2 in graph G, i.e., $E^{\mathcal{A}}(i,j) \Leftrightarrow \{i,j\} \in E, \ \forall i,j \in N$
- $SUCC^{\mathcal{A}}$ is a ternary relation over $N \times O^2$ representing the preference relation $\succ$, i.e., $SUCC^{\mathcal{A}}(i,a,b) \Leftrightarrow a \succ_i b, \ \forall i \in N, \ \forall a,b \in O$
- $ALLOC^{\mathcal{A}}$ is a binary relation over $N \times O$ representing the initial allocation π^0, i.e., $ALLOC^{\mathcal{A}}(i,o) \Leftrightarrow \pi^0(i) = o, \ \forall i \in N, \ \forall o \in O$
- $TARG_A^{\mathcal{A}}$ is an unary relation over N representing agent A, i.e., $TARG_A^{\mathcal{A}}(i) \Leftrightarrow i = A, \ \forall i \in N$
- $TARG_x^{\mathcal{A}}$ is an unary relation over O representing object x, i.e., $TARG_x^{\mathcal{A}}(o) \Leftrightarrow o = x, \ \forall o \in O$

The Σ_1-formula φ is defined as $\varphi = \exists x_0 \exists b_0 \exists x_1 \exists y_1 \exists a_1 \exists b_1 \dots \exists x_k \exists y_k \exists a_k \exists b_k \Big(ALLOC(x_0, b_0) \wedge \bigvee_{0 \leq k' \leq k} \psi^{k'} \Big)$ with

$$\psi^{k'} \equiv TARG_A(x_{k'}) \wedge TARG_x(b_{k'}) \wedge \bigwedge_{i=1}^{k'} \Big(E(x_i, y_i) \wedge SUCC(x_i, b_i, a_i) \wedge SUCC(y_i, a_i, b_i) \wedge \mathrm{obj}_i(x_i, a_i) \wedge \mathrm{obj}_i(y_i, b_i) \Big)$$

where for all i, $\mathrm{obj}_i(q,s)$ stands for $\Big(ALLOC(q,s) \wedge \bigwedge_{j=1}^{i-1} x_j \neq q \wedge y_j \neq q \Big) \vee \bigvee_{j=1}^{i-1} \Big[\Big(\bigwedge_{p=j+1}^{i-1} x_p \neq q \wedge y_p \neq q \Big) \wedge \Big(\Big(x_j = q \wedge \mathrm{obj}_j(y_j, s) \Big) \vee \Big(y_j = q \wedge \mathrm{obj}_j(x_j, s) \Big) \Big) \Big]$.

We claim that object x is reachable for agent A by a sequence of swaps of at most k swaps in instance $\mathcal{I}$ if and only if $\mathcal{A}$ is a model of formula φ in instance $\mathcal{I}'$. The global idea of φ is the following: there exists a sequence of at most k swaps involving, at each step $i \leq k'$ (for $k' \leq k$), a pair of connected agents $\{x_i, y_i\}$ owning objects a_i and b_i, respectively, who rationally exchange their objects, such that agent A finally owns object x at step k'.

Let us first prove by induction over i that formula $\mathrm{obj}_i(q,s)$ is true if and only if object s is owned by agent q before the i^{th} swap. The base case is trivial: formula $\mathrm{obj}_1(q,s)$ is true if and only if $ALLOC(q,s)$ is true, i.e., agent q initially owns object s in π^0. Consider step i, agent q and object s, and suppose that $\mathrm{obj}_j(.,.)$ is correct for all steps $j = 1, \dots, i-1$. If agent q has not performed a swap before step i, then the second member of the disjunction in the formula is false, but the first one is true if and only if s is the object initially owned by q. Otherwise, if agent q has made at least one swap before step i, then the object actually owned by q is the object that was possessed by the agent with who she made her last swap, just before this swap. This is expressed by the second member of the disjunction stating that there exists a step $j < i$ such that agent q was involved in the swap at step j but not in the subsequent swaps until step i, and the agent with who q has exchanged at step j was possessing object s, fact translated by $\mathrm{obj}_j(., s)$. Therefore, $\mathrm{obj}_i(q, s)$ is true if and only if $\mathrm{obj}_j(q', s)$ is true where q' denotes the agent with who q has made her last swap before step i. This is true by induction assumption. Hence, formula $\mathrm{obj}_i(q, s)$ is true if and only if agent q currently owns object s before the swap at step i.

Now, observe that formula $\psi^{k'}$ is true if and only if the sequence of exchanges between the agents (x_i, y_i) exchanging the objects (a_i, b_i), for $i \in \{1, \dots, k'\}$, is a sequence of swaps leading to give object x to agent A. Indeed, $\psi^{k'}$ repertories the two conditions for a swap. The two agents x_i and y_i involved in the i^{th} swap must be connected in the social network and the exchange must be rational: x_i must prefer object b_i, which must be the object of agent y_i just before swap i, to object a_i, and vice versa. Moreover, $\psi^{k'}$ imposes that one of the agent involved in the last swap must be agent A and obtain in this last swap object x. Thus, $\psi^{k'}$ expresses the reachability of object x for agent A after exactly k' swaps.

Hence, φ is true if and only if object x is reachable for agent A after k' swaps, for $k' \leq k$, or x is initially owned by A. □

The same idea and a slightly different first-order formula for model checking work for RO-makespan.

Proposition 3.12 RO-*makespan is in* W[1].

Proof: We reduce to MC(Σ_1) but face a new difficulty. We cannot quantify over all potential exchanges within makespan k as it would lead to a formula of size $\Omega(n)$. The crux of this proof is to observe that not all exchanges are relevant to decide the problem. Assume we process independent swaps in parallel for up to k time steps. Looking at it from the end, the only relevant swap in the last step k involves agent A, so we quantify over a single swap and ignore all concurrent ones. In the one-before-last, only swaps involving A or A's partner at step k may be relevant. So considering two swaps happening at step $k-1$ and ignoring all other concurrent ones suffices. All in all, we need to quantify over no more than 2^{k+1} exchanges. The rest of the proof is similar to that of Theorem 3.11.

An instance $\mathcal{I}$ of RO with a swap dynamics instance $\langle N, O, \succ, G, \pi^0\rangle$, agent A and object x, and k as the makespan of the sequence of swaps, is transformed into an instance $\mathcal{I}' = \langle \mathcal{A}, \varphi\rangle$ of MC(Σ_1). Structure $\mathcal{A}$ is a τ_{RO}-structure over the universe $N \cup O$, which is defined as in the proof of Theorem 3.11.

The Σ_1-formula φ is defined as $\varphi = \exists x_0^1 \exists b_0^1 \exists x_1^1 \exists y_1^1 \exists a_1^1 \exists b_1^1 \dots \exists x_1^{L_1} \exists y_1^{L_1} \exists a_1^{L_1} \exists b_1^{L_1} \dots \exists x_k^1 \exists y_k^1 \exists a_k^1 \exists b_k^1 \dots \exists x_k^{L_k} \exists y_k^{L_k} \exists a_k^{L_k} \exists b_k^{L_k} \Big(ALLOC(x_0^1, b_0^1) \wedge \bigvee_{0 \leq k' \leq k} \bigvee_{\ell_0 \leq L_0} \bigvee_{\ell_1 \leq L_1} \dots \bigvee_{\ell_{k'} \leq L_{k'}} (\psi^{k'}_{\ell_0,\dots,\ell_{k'}} \wedge \chi^{k'}_{\ell_0,\dots,\ell_{k'}})\Big)$ with $L_i = 2^{k-i}$ for all $i \in [k]$, $L_0 = 1$, and

$$\psi^{k'}_{\ell_0,\dots,\ell_{k'}} \equiv \bigvee_{j=0}^{\ell_{k'}} \Big(TARG_A(x^j_{k'}) \wedge TARG_x(b^j_{k'})\Big) \wedge \bigwedge_{i=1}^{k'} \bigwedge_{j=1}^{\ell_i} \Big(E(x_i^j, y_i^j) \wedge SUCC(x_i^j, b_i^j, a_i^j) \wedge SUCC(y_i^j, a_i^j, b_i^j) \wedge \mathrm{obj}_{ij}(x_i^j, a_i^j) \wedge \mathrm{obj}_{ij}(y_i^j, b_i^j)\Big)$$

$$\chi^{k'}_{\ell_0,\ldots,\ell_{k'}} \equiv \bigwedge_{i=1}^{k'} \bigwedge_{j=1}^{\ell_i} \bigwedge_{p=j+1}^{\ell_i} \Big((x_i^j \neq x_i^p) \wedge (y_i^j \neq y_i^p) \wedge (x_i^j \neq y_i^p) \wedge (y_i^j \neq x_i^p) \wedge (a_i^j \neq a_i^p) \wedge (b_i^j \neq b_i^p) \wedge (a_i^j \neq b_i^p) \wedge (b_i^j \neq a_i^p) \Big)$$

where, for all i, j, $\text{obj}_{ij}(q, s)$ stands for $\Big(ALLOC(q, s) \wedge \bigwedge_{p=1}^{i-1} \bigwedge_{r=1}^{\ell_p} \left(x_p^r \neq q \wedge y_p^r \neq q\right) \Big) \vee \bigvee_{p=1}^{i-1} \Big[\Big(\bigwedge_{h=p+1}^{i-1} \bigwedge_{t=1}^{\ell_h} (x_h^t \neq q \wedge y_h^t \neq q) \Big) \wedge \bigvee_{r=1}^{\ell_p} \Big(\Big(x_p^r = q \wedge \text{obj}_{pr}(y_p^r, s) \Big) \vee \Big(y_p^r = q \wedge \text{obj}_{pr}(x_p^r, s) \Big) \Big) \Big]$.

We claim that object x is reachable for agent A by a sequence of swaps of makespan at most k in instance $\mathcal{I}$ if and only if $\mathcal{A}$ is a model of formula φ in instance $\mathcal{I}'$. The global idea of φ is the following: there exists a sequence of swaps with makespan at most k involving, at each step $i \leq k'$ (for $k' \leq k$), a group of parallel swaps between the pairs of connected agents $\{x_i^j, y_i^j\}$ owning objects a_i^j and b_i^j, respectively, who rationally exchange their objects, such that agent A finally owns object x within the group of parallel swaps at step k'.

Formula $\psi^{k'}_{\ell_0,\ldots,\ell_{k'}}$ says, similarly as its definition in the proof of Theorem 3.11, that any exchange must be a swap and that the last swap must give object x to agent A.

Formula $\text{obj}_{ij}(q, s)$, like in the previous proposition, is true if and only if agent q gets object s just before the j^{th} swap of the i^{th} group of parallel swaps in the sequence. We omit the proof because it is similar to the proof in Theorem 3.11.

Formula $\chi^{k'}_{\ell_0,\ldots,\ell_{k'}}$ says that any group of parallel swaps must involve different pairs of agents and different pairs of objects. Consequently, φ is true if and only if x is reachable for A by a sequence of swaps which can be decomposable into a sequence of at most k sets of parallel swaps. We now verify that the size of the formula is in function of k only, by proving that the i^{th} set of parallel swaps has at most $L_i = 2^{k-i}$ parallel swaps. Observe first that at the k^{th} set, the last one, at most one swap is required: the swap between an agent Y and A, that makes A getting object x. All the other parallel swaps are useless for the reachability of x for A. Therefore, there is at most $2^{k-k} = 1$ useful swaps at the k^{th} set. Consider the i^{th} set and assume that at the $(i+1)^{\text{th}}$ set at most 2^{k-i-1} parallel swaps are useful. All the useful parallel swaps at i^{th} set can only involve agents that are involved in all the j^{th} sets for $j \in \{i+1, \ldots, k\}$. There is at most $2\times$ {the maximum number of useful parallel swaps at the $(i+1)^{\text{th}}$ set} such agents, and thus at most 2^{k-i} parallel swaps at the i^{th} set. □

Combining Theorem 3.10 with Theorem 3.11 and Proposition 3.12 leads to the following corollary.

Corollary 3.13 RO-*sum and* RO-*makespan are* W[1]-*complete.*

The previous result shows that RO is not "so hard" considering the length of the sequence as a parameter, because it is circumscribed to the first level of the W-hierarchy (see Figure 1.4). For instance, with such parameters, Reachable Object is in XP and thus is solvable in polynomial time when the length of the sequence is a constant.

Furthermore, for some natural classes of graphs, the problem is even fixed-parameter tractable with respect to these parameters.

Proposition 3.14 RO-*sum/makespan is* FPT *on bounded degree graphs.*

Proof: The proof follows the idea developed for Proposition 3.12, by considering the sequence of exchanges from the end. We enumerate all the possible sequences of swaps by only taking into account relevant swaps.

At the last step of a sequence of swaps leading to give x to A, the only relevant exchange involves agent A and a neighbor Y of A, otherwise the swap is not useful for the reachability of x to A. So the number of relevant possible swaps at the last step is equal to the degree of A. Therefore, there are at most Δ_G, i.e., the degree of the graph, relevant possible swaps at the last step for both RO-sum and RO-makespan. Following the same idea, the one-before-last step must involve one of the agents making the swap at the last step, otherwise the exchange is not useful. Therefore, there are less than $2\Delta_G$ possible swaps for RO-sum, by counting as a possible swap an exchange between A and a neighbor of A, or between Y and a neighbor of Y. However, for RO-makespan, either one or two parallel swaps can be performed at this step. In the latter case, the two swaps must be within different pairs of agents with A belonging to one, and Y to the other one, and obviously the other neighbor with A and with Y cannot be the same agent. Therefore, with a rough upper bound, we have at most $2\Delta_G + \Delta_G^2$ possible swaps for RO-makespan at the one-before-last step. Globally, one can draw a tree for enumerating all the possible sequences of swaps following the same principle at each step of the sequence. The same argument applies at any of the $(\text{last}-t)^{\text{th}}$ step, where we need to test at most $(t+1).\Delta_G$ swaps for RO-sum and at most $\sum_{j=1}^{2^t} \binom{2^t \Delta_G}{j}$ different groups of parallel swaps for RO-makespan. Since the global number of steps is bounded by k, we have to test in total at most $\Delta_G^k.k!$ sequences of swaps for RO-sum and at most $\prod_{t=1}^{k} \sum_{j=1}^{2^t} \binom{2^t \Delta_G}{j}$ sequences of swaps for RO-makespan, of maximal authorized length. Then, it suffices to verify after each step if one sequence leads to an allocation assigning object x to agent A. □

In a nutshell, Reachable Object parameterized by some budget constraints on the sequence of swaps remains computationally hard in general. However, the parameters *sum* and *makespan* restricting the length of the sequence enable to determine some interesting tractable cases with respect to these parameters.

3.4 Reachable Assignment

In this section, we address the decision problem Reachable Assignment: given a swap dynamics instance, is a target allocation π reachable?

Like for Reachable Object, we analyze the complexity of Reachable Assignment with respect to classical and parameterized complexity, that is where no budget constraint is imposed on the solution, i.e., on the sequence of swaps, or when the sequence of swaps is limited regarding the number of swaps that can be performed by the agents.

3.4.1 Reachable Assignment with no budget consideration

First of all, let us investigate the unconstrained version of Reachable Assignment where no budget consideration is imposed on the solution size. We prove that Reach-

ABLE ASSIGNMENT is computationally hard in general but polynomial when the network is a tree.

Theorem 3.15 REACHABLE ASSIGNMENT *is* NP*-complete.*

Proof: One can easily verify that the problem is in NP: given a sequence of swaps $(\pi^0, \ldots, \pi^t)$, whose length is bounded by n^2 (Observation 3.3), it is easy to verify whether the sequence is rational and there exists a step $t' \leq t$ such that $\pi^{t'} = \pi$.

For hardness, we propose a reduction from REACHABLE OBJECT, that we have proved to be NP-complete in Theorem 3.1. Take an instance $\mathcal{I} = \langle N, O, \succ, G, \pi^0, A, x \rangle$ of RO. The problem asks whether target agent $A \in N$ can reach target object $x \in O$. We construct an instance of REACHABLE ASSIGNMENT $\mathcal{I}' = (N \cup N' \cup \{Y\}, O \cup O' \cup \{y\}, \succ'$ $, \pi'^0, E')$ where each element of N (respectively, O) has a copy in N' (respectively, O'), and one more agent Y with her associated initial object y is added. The copies of agent $i \in N$ and $o \in O$ are denoted by i' and o', respectively. The edge set E' is a superset of E where each agent $i \in N$ is connected to her copy $i' \in N'$ and G' mimics graph G on N by replacing agent A by agent Y, i.e., $\{\{i', j'\} : \{i, j\} \in E \text{ and } i, j \neq A\} \subseteq E'$ and $\{\{Y, i\} : i \in \mathcal{N}_G(A)\} \subseteq E'$. See Figure 3.5 for an example of the construction.

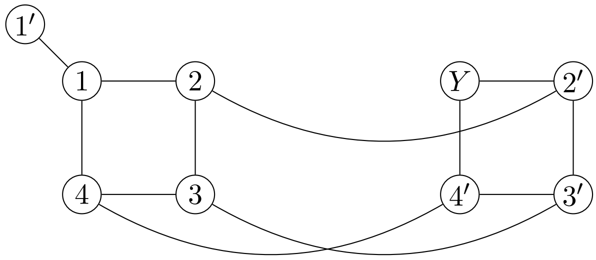

Figure 3.5: Construction of G' from initial graph $G = (N, E)$ where $N = \{1, 2, 3, 4\}$ and $E = \{\{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 1\}\}$, with a target agent $A = 1$.

Let us denote by o_i the initial endowment of each agent $i \in N \setminus \{A\}$ in $\mathcal{I}$, and by o'_i its copy in O'. Similarly, let us denote by a the initial endowment of agent A in $\mathcal{I}$ and by a' its copy. The initial assignment π'^0 in $\mathcal{I}'$ is such that $\pi'^0(i) = \pi^0(i)$ when $i \in N$, $\pi'^0(i') = o'_i$ when $i' \in N' \setminus \{A'\}$, $\pi'^0(A') = a'$, and $\pi'^0(Y) = y$. For every $i \in N$, $\succ'_i$ consists of o'_i on top, followed by $\succ_i$, and the remaining objects are put on the last positions. Agent A' only prefers target object x to her initial endowment a' and all the other remaining objects are ranked in the last positions of $\succ'_{A'}$. For any other agent $i' \in N'$, the only objects that are preferred to the initial endowment o'_i are objects of $(O \setminus \{x\}) \cup \{y\}$ that are preferred or equal to o_i in the preferences $\succ_i$ of her related agent i in N. All the remaining objects are ranked in the last positions of $\succ'_i$. More precisely, the objects preferred by agent i' to o'_i are ranked according to the reversed preferences of agent i from her top object to her initial object o_i where object x (which is one of the $\{o_j\}_{j \in N}$) is replaced by object y. For instance, if the preferences of any agent $i \neq A$ are the following: $o_1 \succ_i o_2 \succ_i \ldots \succ_i o_q \succ_i x \succ_i o_{q+1} \succ_i \ldots \succ_i o_{q+\ell} \succ_i \boxed{o_i} \succ_i [\ldots]$, then the preferences of i' are: $o_i \succ'_{i'} o_{q+\ell} \succ'_{i'} \ldots \succ'_{i'} o_{q+1} \succ'_{i'} y \succ'_{i'} o_q \succ'_{i'} \ldots \succ'_{i'} o_2 \succ'_{i'} o_1 \succ'_{i'}$ $\boxed{o'_i} \succ'_{i'} [\ldots]$. The construction of the preferences of agent Y follows the same principle but based on the preferences $\succ_A$ of agent A from target object x to the initial object a. For instance, if the preferences of agent A are the following: $o_1 \succ_A o_2 \succ_A \ldots \succ_A o_q \succ_A x \succ_A o_{q+1} \succ_A \ldots \succ_A o_{q+\ell} \succ_A \boxed{a} \succ_A [\ldots]$, then the preferences of agent Y are: $a \succ'_Y o_{q+\ell} \succ'_Y \ldots \succ'_Y o_{q+1} \succ'_Y \boxed{y} \succ'_Y [\ldots]$.

We claim that x is reachable for agent A in $\mathcal{I}$ if and only if every agent gets her most preferred object in $\mathcal{I}'$.

Suppose that x is reachable for agent A in $\mathcal{I}$. By construction, it is also the case in $\mathcal{I}'$. Once x has reached agent A in $\mathcal{I}'$, every agent $i \in N$ exchanges her object with her copy $i' \in N'$. Thus, each $i \in N$ possesses her most preferred object o'_i, as well as agents A' who gets her top object x, and any agent i' who is the copy of an agent i not involved in the first sequence of swaps. Then, is performed the reverse sequence of swaps s' within N' of the first sequence s leading to give x to A within N, where each agent $i \in N$ involved in s is replaced by her copy $i' \in N'$ in s', except for A who is replaced by Y. By construction of the preferences, we reach an allocation where every agent in N' obtains the initial object of the agent for who she is the copy and Y obtains object the initial object a of agent A. Therefore, this final allocation gives to every agent her most preferred object.

Suppose that every agent can get her most preferred object in $\mathcal{I}'$. The only possibility for agent A' to obtain her most preferred object x is via agent A. Observe that x cannot reach agent A through agents in N' because no agent among them prefers x to her initial object. Moreover, if an agent $i \in N$ makes a swap with her copy $i' \in N'$, then i necessarily obtains her best object o'_i because this is the only object from O' that she prefers to her initial object. Therefore, she cannot be involved in a subsequent swap and cannot play a role for making object x reaching agent A. Therefore, object x must reach agent A via a sequence of swaps that only involves agents in N. Since $G'[N] = G$ and the restriction of $\succ'_i$ to objects in O is equal to $\succ_i$ for every agent $i \in N$, object x is also reachable for agent A in $\mathcal{I}$. □

In the previous reduction, we have constructed a graph that contains cycles. However, when $G = (N, E)$ is a tree, a polynomial algorithm (Algorithm 3.1) solves REACHABLE ASSIGNMENT.

The global idea of the algorithm is that every object must move along a unique dipath in order to reach its owner in π from its owner in π^0. Therefore, it suffices to verify that every dipath intersects another one of the opposite direction within a swap for the involved agents. In the pseudocode, list L stores the first arc of each dipath and `pop`(P) outputs the first arc of dipath P and deletes this arc from P.

Let us illustrate Algorithm 3.1 with an example.

Example 3.8 *Consider an instance with five agents, where $N = \{1, 2, 3, 4, 5\}$ and $O = \{o_1, o_2, o_3, o_4, o_5\}$. The social network, the preferences and the initial allocation are as follows.*

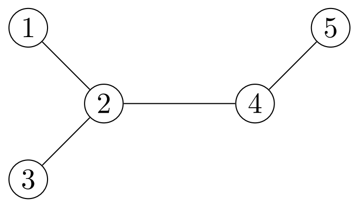

$$
\begin{array}{ll}
1: & o_4 \succ \boxed{o_1} \succ o_2 \succ o_5 \succ o_3 \\
2: & o_5 \succ o_3 \succ o_1 \succ o_4 \succ \boxed{o_2} \\
3: & o_1 \succ \boxed{o_3} \succ o_5 \succ o_2 \succ o_4 \\
4: & o_3 \succ o_5 \succ o_2 \succ \boxed{o_4} \succ o_1 \\
5: & o_2 \succ \boxed{o_5} \succ o_1 \succ o_4 \succ o_3
\end{array}
$$

Algorithm 3.1:

Input: Swap dynamics model $\langle N, O, \succ, G, \pi^0 \rangle$, assignment π

Output: Whether π is reachable from π^0

1 $L \leftarrow \emptyset$; $\pi' \leftarrow \pi^0$;
2 **foreach** $x \in O$ **do**
3     $P_x \leftarrow$ unique dipath in G from the owner of x in π^0 to the owner of x in π;
4     $L \leftarrow L \cup \{\texttt{pop}(P_x)\}$;
5 **while** $L \neq \emptyset$ **do**
6   **if** $\exists i, j \in N, i \neq j$ *such that* $(i, j) \in L$ *and* $(j, i) \in L$ **then**
7   **if** $\pi'(i) \succ_i \pi'(j)$ *or* $\pi'(j) \succ_j \pi'(i)$ **then**
8   **return** false;
9   Update π' with exchange between i and j;
10         $L \leftarrow L \setminus \{(i, j), (j, i)\}$;
11         $L \leftarrow L \cup \{\texttt{pop}(P_{\pi'(i)})\} \cup \{\texttt{pop}(P_{\pi'(j)})\}$;
12   **else return** false;
13 **return** true;

The question is whether the agents can reach, by a sequence of swaps, allocation $\pi = (o_4, o_5, o_1, o_3, o_2)$, where every agent obtains her best object.

The dipaths P_x computed at line 3 of Algorithm 3.1 are:

$$P_{o_1} = ((1,2),(2,3)) \quad P_{o_3} = ((3,2),(2,4)) \quad P_{o_5} = ((5,4),(4,2))$$
$$P_{o_2} = ((2,4),(4,5)) \quad P_{o_4} = ((4,2),(2,1))$$

The following table illustrates the differents steps of the while loop (lines 5–12):

L	*Swap*	π'
$\{(1,2),(2,4),(3,2),(4,2),(5,4)\}$	$2 \leftrightarrow 4$	$(o_1, o_4, o_3, o_2, o_5)$
$\{(1,2),(4,5),(3,2),(2,1),(5,4)\}$	$1 \leftrightarrow 2$	$(o_4, o_1, o_3, o_2, o_5)$
$\{(2,3),(4,5),(3,2),\emptyset,(5,4)\}$	$2 \leftrightarrow 3$	$(o_4, o_3, o_1, o_2, o_5)$
$\{\emptyset,(4,5),(2,4),\emptyset,(5,4)\}$	$4 \leftrightarrow 5$	$(o_4, o_3, o_1, o_5, o_2)$
$\{\emptyset,\emptyset,(2,4),\emptyset,(4,2)\}$	$2 \leftrightarrow 4$	$(o_4, o_5, o_1, o_3, o_2)$
$\emptyset$	-	-

At each step, list L stores the first arc of each P_{o_i}. At step 1, only one exchange is possible: between agents 2 and 4. This exchange being rational, π' is updated by performing the swap. Arcs $(2,4)$ and $(4,2)$ are then removed from L. Arcs $(4,5)$ and $(2,1)$, which are respectively the new first arcs of P_{o_2} and P_{o_4}, are inserted in L. The algorithm stops when L is empty, implying that π is reached.

Proposition 3.16 *Algorithm 3.1 solves* Reachable Assignment *in polynomial time when $G = (N, E)$ is a tree.*

Proof: The principle of Algorithm 3.1 is that each object x must follow a unique directed path P_x in graph G from its current owner in initial allocation π^0 to its target owner in allocation π. Actually, this is the case for any sequence of swaps leading to π from π^0 because there is a unique path between two nodes in a tree and the objects cannot pass twice by the same agent (Observation 3.2). Algorithm 3.1 performs a swap between two connected agents i and j whose respective objects must follow edge $\{i, j\}$ in the two opposite directions, as soon as the exchange is rational. In Algorithm 3.1, these swaps are made in any arbitrary order. Let us check whether this matters.

Assume that Algorithm 3.1 returns *false* whereas π is actually reachable. Denote by s_1 the sequence of swaps leading to reach π from π^0 and by s_2 the sequence of swaps that has been performed so far from π^0 in Algorithm 3.1 when it returns false. This sequence of swaps is rational by construction of the algorithm. Sequences s_1 and s_2 are necessarily different. Let us consider the first agent i that makes a different swap in s_1 and in s_2, and denote by π_1 and π_2 the allocation reached by s_1 and s_2, respectively, just before this different swap involving i. Note that globally π_1 and π_2 may not be the same because some agents can have made a swap in s_1 but not yet in s_2, and vice versa. However, this is clear that $\pi_1(i) = \pi_2(i)$, otherwise it would not be the first different swap of i. Let us denote by o_i the object owned by i at this moment in both allocations, i.e., $o_i := \pi_1(i) = \pi_2(i)$.

Suppose that this swap involving i which is different between s_1 and s_2 is made with different agents, say j_1 and j_2, respectively. It follows that object o_i follows arc (i, j_1) in s_1 whereas it follows arc (i, j_2) in s_2 with $j_1 \neq j_2$. This is impossible because it would imply two different paths between i and the target owner of object o_i in π, and G is a tree.

Suppose now that the swap involving i which is different between s_1 and s_2 is made with the same agent j but against another object, say o_1 and o_2, respectively, with $o_1 \neq o_2$. This implies that agent j has made a previous swap in one of the sequence that was not made in the other (she cannot have performed a different swap in the two sequences, otherwise j would be the first agent for which s_1 and s_2 differ, instead of i). Say, without loss of generality, that this previous swap of j has been made in s_1 but not in s_2. Agent j has made this swap in s_1 with an agent i' different from i otherwise, by rationality of the swaps, j could not exchange directly after with i again. The object possessed by j before her swap with i' was necessarily object o_2, by assumption that i is the first agent for which s_1 and s_2 differ. Therefore, in s_1, object o_2 follows arc (j, i') whereas in s_2 it follows arc (j, i) with $i \neq i'$. Like in the previous case, this implies two different paths for o_2 reaching her target agent in π, contradicting the fact that G is a tree. □

Contrasting with Reachable Object which is hard even for trees (Theorem 3.1), Reachable Assignment is tractable for such a class of graphs.

3.4.2 Reachable Assignment under budget constraints

While Theorem 3.15 has shown that Reachable Assignment is NP-complete in general, Proposition 3.16 has established tractability on trees. We now focus on whether introducing budget constraints can let us derive efficient algorithms on networks more general than trees. We first show that Reachable Assignment parameterized by the number of swaps per agent, called RA-max, remains hard on general graphs.

Proposition 3.17 RA-*max is* NP-*complete even for fixed* $k \geq 3$ *and graphs of degree* Δ_G *equal to 5.*

Proof: We consider the reduction given in the proof of Theorem 3.15, based on an instance of Reachable Object constructed in the proof of Theorem 3.6, where the swaps per agent are limited to two.

Observe that the construction of the instance of Reachable Assignment only requires an additional swap with a copy-agent for the agents of the original instance of Reachable Object. Moreover, the copy-agents introduced in the reduction make the reverse swaps of the agents of the original instance, in addition to make one swap with them. Therefore, the proof holds for a maximum number of swaps per agent limited to three.

Moreover, the graph constructed in the proof of Theorem 3.6 for RO has a degree equal to four and, in the reduction from RO to RA, we only add one more edge to the agents of the original instance while the copy agents form a copy of the graph of the original instance. Therefore, the proof holds for a graph of degree 5. □

Similarly as Reachable Object, let us denote by RA-sum and RA-makespan the parameterized versions of Reachable Assignment with respect to the parameters *sum* and *max*.

So far, our results regarding the complexity of RO-sum and RO-makespan are the same. However this remark does not hold for Reachable Assignment: we prove that RA-sum is fixed-parameter tractable whereas RA-makespan is W[1]-hard. Let us observe the reduction from Reachable Object given in the proof of Theorem 3.15. An instance of RA is constructed where two copies of the graph of the RO instance are connected via edges connecting each agent to her copy. During the reduction, at one point, each agent must exchange with her copy agent. This is done in n total swaps where n is the number of agents in the RO instance, therefore the total number of swaps cannot be expressed in function of k. However, all these swaps can be performed in parallel since they do not involve the same agents, leading to a makespan equal to 1 for this subsequence of swaps. Therefore, by using an instance of RO for which the length of the sequence is bounded (like in Theorem 3.10), we get that RA-makespan is hard since the makespan of the sequence can be expressed only according to the parameter k. But this is not the case of RA-sum for which this reduction does not hold.

Proposition 3.18 RA-*makespan is* W[1]-*hard.*

Proof: The reduction is the same as given in the proof of Theorem 3.15 but based on an instance of RO-makespan, like the one constructed in the proof of Theorem 3.10. □

Contrasting with the hardness of RA-makespan, RA-sum is fixed-parameter tractable.

Proposition 3.19 RA-*sum is* FPT.

Proof: We can assume, without loss of generality, that the target object of each agent is different from her initial object, i.e., $\pi(i) \neq \pi^0(i)$ for every agent i. Otherwise, we can remove from the instance all the agents i such that $\pi(i) = \pi^0(i)$, without affecting the reachability of the other objects to other agents. It follows that each agent must perform at least one swap to obtain her target object, and thus it must hold that $k \geq \lceil n/2 \rceil$, otherwise we have a trivial *no*-instance. Therefore, $n = \mathcal{O}(k)$ and an exponential function depending on n can be expressed only depending on k. □

An interesting open question would be to analyze the complexity of RA-makespan when the network is a bounded degree graph: is the problem easy to solve like RA-sum and RO-makespan on this type of graphs, or computationally hard like RA-makespan in a general graph?

3.5 Guaranteed Level of Satisfaction

We dedicate this section to the Guaranteed Level of Satisfaction problem, very close to the dual problem of Reachable Object, which enables to provide minimal guarantees for the agents concerning their situation at the end of the swap dynamics.

3.5.1 Relation between Reachable Object and Guaranteed Level of Satisfaction

Reachable Object (RO) asks whether an agent A can obtain an object x by a sequence of swaps. We have proved that RO is NP-complete even for trees (Theorem 3.1). Guaranteed Level of Satisfaction (GLS) asks whether agent A is guaranteed to obtain object x or an object preferred to x in any stable reachable allocation. GLS appears even more natural than RO since it offers guarantees for the agent and does not only focus on lucky configurations. It is close to the complement of RO, because co-GLS is formulated as follows: is an object y reachable for agent A in a stable allocation, such that $x \succ_A y$? Thus, the study of RO also contributes to the understanding of GLS.

Proposition 3.20 *Any instance $\langle N, O, \succ, G = (N, E), A, x\rangle$ of* co-Reachable Object *is linearly reducible to an equivalent instance $\langle N \cup \{Y\}, O \cup \{y\}, \succ', \pi'^0, G' = (N \cup \{Y\}, E \cup \{\{Y, A\}\}), A, y\rangle$ of* GLS.

Proof: We reduce from the co-Reachable Object problem, asking whether object x is unreachable for agent A, which is co-NP-complete by NP-completeness of RO (Theorem 3.1).

Let $\mathcal{I} = \langle N, O, \succ, G = (N, E), A, x\rangle$ be an instance of co-RO. An instance $\mathcal{I}' = \langle N \cup \{Y\}, O \cup \{y\}, \succ', \pi'^0, G' = (N \cup \{Y\}, E \cup \{\{Y, A\}\}), A, y\rangle$ of GLS is constructed by adding an agent Y and an object y. The initial allocation π'^0 is the same as π^0 for all agents in N and assigns y to Y. The output social network has the same structure as the input one, with one more edge $\{Y, A\}$. Denote by a the object of agent A in π^0. If A prefers x to a, then denote by P_a the set containing a and the objects that are preferred to a and less preferred than x in $\succ_A$. Otherwise, P_a contains a and the objects that A prefers to a. Denote by P_x the set of objects that A prefers to x. The ranking $\succ'_A$ is constructed from $\succ_A$, by moving the objects in P_x to the end of $\succ'_A$, and by putting y

at the top of $\succ'_A$. The agents in $N \setminus \{A\}$ keep the same preferences as in $\succ$ but rank y last, and Y only prefers the objects of P_a to object y.

We claim that x is not reachable for A in $\mathcal{I}$ if and only if A obtains y or an object preferred to y in any reachable stable allocation in $\mathcal{I}'$, i.e., if and only if there is no stable reachable allocation where A gets z such that $y \succ_A z$.

Suppose that x is not reachable for A in $\mathcal{I}$. At a stable allocation in $\mathcal{I}$, either (i) agent A gets an object from P_x without having owned x during the sequence, but it is not possible in $\mathcal{I}'$ because $a \succ'_A P_x$, and then agent A can only get objects from P_a in $\mathcal{I}'$, or (ii) agent A gets an object from P_a, and it will be also the case in $\mathcal{I}'$. In both cases, agent A gets an object in P_a and thus can exchange with Y in order to obtain y. Hence, agent A eventually gets an object preferred or equal to y.

Suppose now that x is reachable for A in $\mathcal{I}$. It follows that x is also reachable for A in $\mathcal{I}'$ since $G'[N]$ and G are the same and A does not use the objects of P_x to obtain x in $\mathcal{I}$. However, by obtaining x, agent A cannot exchange any more with Y because of the construction of the preferences $\succ'$ and the output social network (agent Y is only connected to A). Thus, agent A cannot obtain an object preferred or equal to y in this sequence of swaps.

Observe that in this reduction, GLS asks whether agent A is guaranteed to obtain her best object. The generalization to an object ranked at k^{th} position is straightforward by inserting $k+1$ dummy agents and objects, not accessible for A. □

Since RO is NP-complete even on trees (Theorem 3.1) and the previous reduction adds only one agent and possibly one swap to an instance of RO, GLS is co-NP-hard on trees. Observe that GLS is in co-NP: after guessing a sequence of swaps, whose length is polynomial (Observation 3.3), which leads to a stable allocation, one can directly check whether this allocation is a certificate for a *no*-instance of GLS, by observing whether the object obtained by A is less preferred than object x.

Corollary 3.21 GLS *is* co-NP*-complete even for trees.*

We refine the complexity of GLS using natural parameters: the number of swaps per agent and the length of the sequence. Although RO and GLS are close to be dual problems, they are indeed not complementary because GLS focuses on stable allocations. This is even more visible when GLS is parameterized by the length of the sequence of swaps. A k-bound on the sequence of swaps introduces a dependency on k on the notion of stability: we focus on allocations reachable after exactly k swaps or stable allocations reachable after less than k swaps. Stability is not necessary in RO because for an assignment solution π where agent A gets object x, all the stable allocations reachable from π assign to A an object preferred to x or x itself.

3.5.2 Guaranteed Level of Satisfaction under budget constraints

Let us define by GLS-max, GLS-sum and GLS-makespan the parameterized versions of the Guaranteed Level of Satisfaction problem with respect to the parameters on the maximum number of swaps per agent, on the total number of swaps in the sequence, and on the makespan of the sequence of swaps, respectively.

From Proposition 3.20 and its proof, if RO-max is computationally hard then GLS-max too, in an instance of GLS where the target agent A has one more neighbor and performs one more swap comparing to the instance of RO. Consequently, Theorem 3.6 implies the hardness of GLS-max for $k \geq 3$ on graphs of degree 5. Moreover, since the structure of tree is kept by the linear reduction from RO to GLS (Proposition 3.20), Theorem 3.7 implies the hardness of GLS-max in trees.

Corollary 3.22 *For $k \geq 3$,* GLS*-max is* co-NP*-complete, even on graphs of degree Δ_G equal to 5.*

Corollary 3.23 GLS*-max is* co-W[SAT]*-hard on trees.*

Globally, GLS, like RO, remains difficult in very simple graphs even when the number of swaps per agent is limited. Alternatively, the parameters on the length of the sequence enable to determine upper bounds for the complexity of GLS, even if the problem in general remains hard.

Combining the proofs of Theorem 3.10 and Proposition 3.20 leads to hardness for GLS parameterized by the length of the sequence of swaps.

Corollary 3.24 GLS*-sum and* GLS*-makespan are* co-W[1] *hard even for trees.*

In terms of complexity upper bounds, a reasoning similar to that of Theorem 3.11 and Proposition 3.12 can be applied to GLS. While we could show W[1] membership for RO-sum and RO-makespan, leading to a tight characterization, here we are only able to obtain co-A[2] membership (see Figure 1.4 for an idea of the location of class A[2] in the A-hierarchy). Actually, we reduce GLS to the model-checking problem on more sophisticated first-order formulas by using Σ_2 formulas (see Section 1.5.2.b). Still, the results confirm that GLS is tractable when the parameter is a constant.

Proposition 3.25 GLS*-sum and* GLS*-makespan are in* co-A[2]*.*

Proof: We reduce co-GLS to MC(Σ_2), known to be A[2]-complete (Theorem 1.10), by following an approach similar to that of Theorem 3.11 and Proposition 3.12.

An instance $\mathcal{I}$ of co-GLS-sum with a swap dynamics model $\langle N, O, \succ, G, \pi^0\rangle$, agent A, object x, and at most k swaps in total, is transformed into an instance $\mathcal{I}' = \langle \mathcal{A}, \varphi\rangle$ of MC(Σ_2). The co-GLS problem asks whether there exists a reachable allocation, stable under the condition of at most k exchanges, where agent A obtains an object less preferred than x. Structure $\mathcal{A}$ is a τ_{RO}-structure over the universe $N \cup O$, which is defined as in the proof of Theorem 3.11.

The Σ_2-formula φ is defined as $\varphi = \exists c \exists z \exists x_0 \exists x_1 \exists y_1 \exists a_1 \exists b_1 \ldots \exists x_k \exists y_k \exists a_k \exists b_k$ $\forall x \forall y \forall a \forall b \left(\psi^k \vee \bigvee_{0 \leq k' < k} \psi^{k'} \wedge \chi^{k'}\right)$ with

$$\psi^{k'} \equiv \bigwedge_{i=1}^{k'} \Big(E(x_i, y_i) \wedge SUCC(x_i, b_i, a_i) \wedge SUCC(y_i, a_i, b_i) \wedge \text{obj}_i(x_i, a_i) \wedge \text{obj}_i(y_i, b_i)\Big)$$
$$\wedge TARG_A(x_{k'}) \wedge \text{obj}_{k'+1}(x_{k'}, c) \wedge TARG_x(z) \wedge SUCC(x_{k'}, z, c)$$

$$\chi^{k'} \equiv \Big(E(x,y) \wedge \mathrm{obj}_{k'+1}(x,a) \wedge \mathrm{obj}_{k'+1}(y,b)\Big) \rightarrow \Big(SUCC(x,a,b) \vee SUCC(y,b,a)\Big)$$

where for all i, $\mathrm{obj}_i(.,.)$ is defined as in Theorem 3.11.

We claim that there exists a reachable allocation, stable under the condition of at most k swaps where agent A obtains an object less preferred than x in $\mathcal{I}$ if and only if $\mathcal{A}$ is a model of formula φ in $\mathcal{I}'$. We omit the details of this proof because they are similar to those of Theorem 3.11.

Formula φ expresses the reachability of an object c for agent A, either in a stable allocation or within at most k swaps, such that A prefers x to c. More precisely, formula $\psi^{k'}$ says that, after a sequence of k' swaps, agent A obtains an object c that she prefers less than object x. Formula $\chi^{k'}$ is related to the variables with universal quantifiers and says that any two connected agents must not mutually prefer the object of the other one, after the k' $^{\text{th}}$ swap. This formula forces the allocation reached after k' swaps to be stable. It is defined in φ only for $k' < k$ because since the problem is parameterized by k, an allocation reached after k swaps does not need to be stable. Hence, co-GLS-sum is correctly translated.

The reasoning for co-GLS-makespan is the same, based on the formula of the proof in Proposition 3.12. □

Note that the previous result does not enable us to provide a completeness result for GLS. However, by referring to the parameterized hierarchies (see Figure 1.4), we know that the exact parameterized complexity class for GLS-sum and GLS-makespan is somewhere between co-W[1], co-W[2] and co-A[2].

3.6 Reachable Pareto-efficient allocations

The question of this section is how to coordinate the swaps in order to reach a Pareto-efficient allocation within *RAll*, i.e., a RAll-efficient allocation. Note that a RAll-efficient allocation is stable, otherwise at least two agents will benefit from a possible swap. However, the reverse does not hold, as we can see in Example 3.6 with allocation π^1 that is stable but Pareto dominated by another reachable allocation.

First of all, we show that it is computationally hard to find a sequence of swaps leading to a RAll-efficient allocation.

Proposition 3.26 PARETO REACHABILITY *is not solvable in polynomial time unless* P =NP[1].

Proof: We base ourselves on the reduction provided in the proof of Theorem 3.15, which constructs an instance $\mathcal{I}'$ of REACHABLE ASSIGNMENT from an instance $\mathcal{I}$ of REACHABLE OBJECT. Instance $\mathcal{I}$ is a *yes*-instance if and only if every agent gets her most preferred object in $\mathcal{I}'$. This allocation must be the unique RAll-efficient allocation in case it is reachable. Therefore, if an algorithm computing a RAll-efficient assignment could exist, then it would be used to recognize a *yes*-instance of REACHABLE OBJECT. □

[1]We do not state that PARETO REACHABILITY is NP-hard since this is not a decision problem. Actually, we do not know for which precise complexity class PARETO REACHABILITY is hard, but some complexity classes like PPAD could be adapted since a solution for PARETO REACHABILITY always exists.

This negative result does not prevent the existence of a polynomial algorithm for constructing a RAll-efficient allocation in a specific class of instances. The remainder of this section is devoted to the resolution of PARETO REACHABILITY in a path and in a star.

A classical algorithm for achieving Pareto-efficiency is the *Serial Dictatorship* mechanism [Abdulkadiroğlu and Sönmez, 1998]. It ranks the agents in an arbitrary manner, and assigns them in turns their most favorite object within the set of unassigned objects, until each object is assigned. When the social network is a path, one can use this idea to compute a RAll-efficient allocation.

Algorithm 3.2: Serial dictatorship in a path

Input: Swap dynamics instance $\langle N, O, \succ, G, \pi^0 \rangle$, agent i
Output: a RAll-efficient allocation π
if $i = 1$ **then**
  return π^0;
$o_j \leftarrow$ best reachable object for i;
$j \leftarrow$ owner of o_j in π^0;
$\pi \leftarrow$ apply $\kappa(j, i)$ on π^0;
return *Algorithm 3.2* ($\langle N, O, \succ, G, \pi \rangle$, $i - 1$);

The parameter i of Algorithm 3.2 designates the dictator who chooses her best reachable object, and $\kappa(j, i)$ is the canonical sequence (see Proposition 3.3).

Proposition 3.27 *When $G = (N, E)$ is a path, Algorithm 3.2 with $i = n$ solves* PARETO REACHABILITY *in polynomial time.*

Proof: The algorithm starts with a leaf i of the path, modifies the current allocation in such a way that the leaf-agent obtains her best object (Corollary 3.4 is used), and continues on the subpath from $i - 1$ to 1. Let us denote by π the allocation given by Algorithm 3.2. The proof is by induction on i: when the decision for agent i is made, the allocation π restricted to the agents $\{i+1, \ldots, n\}$ is RAll-efficient. The base case where $i = n$ follows from the definition of the algorithm: no previous swap has been done and we choose the sequence of swaps that gives to n her best possible object. Assume that the partial allocation of π over the agents $\{i+1, \ldots, n\}$ is RAll-efficient, and let us consider the step of the algorithm where the serial dictator is i. At step i, object $o_j := \pi(i)$ cannot be currently owned by an agent ℓ such that $\ell > i$, otherwise agent ℓ is still able to perform rational swaps where she obtains an object that she prefers to o_j, contradicting the fact that π restricted to the agents $\{i + 1, \ldots, n\}$ is RAll-efficient. Therefore, o_j is actually owned by an agent j such that $j \leq i$. Suppose that there exists a reachable allocation π' where $\pi'(\ell) = \pi(\ell)$ for any agent ℓ such that $\ell > i$, but $\pi'(i) \succ_i \pi(i)$. This implies that we have previously performed in the algorithm swaps that prevent us to make $\pi'(i)$ reachable for agent i at step i. However, the only swaps that we have performed are swaps of canonical sequences. And the canonical sequence is the minimal and only possible way to move each object o_ℓ to agent ℓ for $\ell > i$ in a path (Proposition 3.3). So, this contradicts the assumption that $\pi'(\ell) = \pi(\ell)$ for any agent ℓ such that $\ell > i$.

Therefore, by induction assumption, the allocation π restricted to the agents $\{i, \ldots, n\}$ is RAll-efficient, i.e., Pareto-efficient among all reachable allocations. □

Now we study the case of a star. Assume that the network consists of a center denoted by n, and $n-1$ leaves denoted by $1, \ldots, n-1$ (see Example 3.6 for an illustration), and that each agent $i \in N$ is initially endowed with object o_i. We can suppose, without loss of generality, that for any $i \in [n-1]$, $o_i \succ_n o_{i+1}$. For instance, if $o_n \succ_n o_j$ for some j then the center will never exchange her object with j, so j keeps her object in any sequence of swaps.

Proposition 3.28 *When $G = (N, E)$ is a star, there exists a linear time algorithm for* Pareto Reachability.

Proof: Let us consider the following algorithm: for $i = n-1$ down to 1, exchange the objects of n and i if it is rational. As already mentioned, a leaf who has exchanged her initial object is not involved in a subsequent swap. The algorithm considers the objects by increasing order of preference of the center agent.

Suppose by contradiction that the allocation π returned by the algorithm is Pareto dominated by another reachable allocation π'. Let us denote by s and s' the sequences of exchanges leading respectively to π and π'. The key observation is that the pairs of agents involved in any feasible sequence of swaps are composed of the center agent n and a leaf-agent i where the index of the leaves progressively decreases. In the first step for which s and s' differ, the center agent n swaps her object with ℓ in s and ℓ' in s', with $\ell \neq \ell'$. By construction of our algorithm, in which we perform the swaps with each leaf-agent by decreasing order of the indices if this is rational, it must hold that $\ell' < \ell$. But this implies, by rationality of the swaps, that agents ℓ and n cannot later swap their objects in s' since, by exchanging with agent ℓ', agent n has obtained an object that she prefers to object o_ℓ. Therefore $\pi(\ell) \succ_\ell \pi'(\ell)$, contradicting the fact that π is Pareto-dominated by π'. □

As an open problem, it appears interesting to see if Pareto Reachability is polynomial time solvable in a spider graph, class of graphs for which RO is hard (Theorem 3.1), by a combination of the techniques used to solve the cases of paths and stars.

3.7 Concluding remarks

We have investigated some natural problems arising when a group of agents exchange their object along a social network. Our results show that, beyond the agents' preferences, the social network as a collaboration tool can widely influence and constrain the possible allocations. In particular, we have proved that deciding whether an agent can obtain a given object (Reachable Object) or be guaranteed to obtain a given level of satisfaction (Guaranteed Level of Satisfaction) is computationally difficult, even if the social network is a tree. Nevertheless, an efficient algorithm can determine if a complete allocation (Reachable Assignment) is reachable in a tree. Concerning simple graph structures like paths, we were able to decide if a leaf-agent can acquire a given object. This result can be extended to the case where the distance between a

		No constraint	Budget constraints		
			max	sum	makespan
RO	General graph	NP-c (Thm. 3.1)	NP-c for $k \geq 2$ (Thm. 3.6)	W[1]-c (Cor. 3.13)	W[1]-c (Cor. 3.13)
	Tree	NP-c (Thm. 3.1)	W[SAT]-hard (Thm. 3.7)	W[1]-c (Cor. 3.13)	W[1]-c (Cor. 3.13)
	Bounded degree	NP-c (Thm. 3.6)	NP-c for $k \geq 2$	FPT	FPT
RA	General graph	NP-c (Thm. 3.15)	NP-c for $k \geq 3$ (Prop. 3.17)	FPT (Prop. 3.19)	W[1]-hard (Prop. 3.18)
	Tree	P (Prop. 3.16)	P	P	P
	Bounded degree	NP-c (Prop. 3.17)	NP-c for $k \geq 3$ (Prop. 3.17)	FPT (Prop. 3.19)	?
GLS	General graph	co-NP-c (Cor. 3.21)	co-NP-c for $k \geq 3$ (Cor. 3.22)	co-W[1]-hard (Cor. 3.24) / co-A[2] (Prop. 3.25)	co-W[1]-hard (Cor. 3.24) / co-A[2] (Prop. 3.25)
	Tree	co-NP-c (Cor. 3.21)	co-W[SAT]-hard (Cor. 3.23)	co-W[1]-hard / co-A[2]	co-W[1]-hard / co-A[2]
	Bounded degree	co-NP-c (Cor. 3.22)	co-NP-c for $k \geq 3$	co-A[2] (Prop. 3.25)	co-A[2] (Prop. 3.25)

Table 3.1: Complexity results of RO, RA and GLS under classical and parameterized complexity

non-leaf-agent and the original location of the object is bounded by a constant. We left open the question whether Reachable Object can be efficiently solved in a path, without restriction on the agent's location.

Despite the fact that in general all these decision problems are computationally hard, the parameterized approach allows us to escape this difficulty under realistic assumptions. We consider natural parameters constraining the number of swaps per agent or the duration of the sequence. Assuming that they remain small is reasonable in practice as the patience of the agents, i.e., their willingness to wait before obtaining a target object, typically does not increase with the instance size. In the case of few swaps *per agent*, RO, GLS and RA remain hard even on bounded degree graphs. So, this parameterization, although natural, does not help us to grasp the problems. However, considering the *length* of the sequence, although the problems are intractable (even for trees for the case of RO and GLS), this hardness is circumscribed to not "so hard" parameterized complexity classes, with for example a complexity upper bound to the first level of the W-hierarchy for RO. This leads for instance to the possibility of handling the problems when the parameters do not depend on the instance size, a very natural assumption. Furthermore, unlike the first parameter, the length of the sequence enables to obtain fixed-parameter tractability for some problems, such as RA with respect to the total number of swaps in the sequence, and RO on bounded degree graphs, graphs which typically model social networks representing a possibility of collaboration among agents. See Table 3.1 for an overview of the complexity results.

Note that the case of RA parameterized by the makespan of the sequence of swaps is particular. In general this problem is computationally hard whereas RA parameterized by the total number of swaps in the sequence is fixed-parameter tractable. The difference of complexity between these two parameterizations contrasts with the cases of RO and GLS for which the parameters on the length of the sequence lead to the same complexity classes. This suggests that the possibility of making parallel swaps makes the problem harder, when all the agents need to reach a particular object together. Moreover, we left open the question of an upper bound for the complexity of RA with respect to the makespan, as well as the restricted case where the network is a bounded degree graph. We suspect that this could be computationally hard since the tricks used for solving RO parameterized by the makespan in bounded degree graphs and RA parameterized by the total number of swaps cannot apply in such a case.

The social quality of reachable allocations has also been studied through the search of Pareto-efficient allocations. This problem is shown difficult in general networks. On the positive side, polynomial algorithms have been presented for paths and stars. It seems challenging to settle the complexity of computing a Pareto-efficient allocation when the network is a tree.

As future works, many additional aspects of the model deserve attention. For example, we have not investigated the impact of a strategic behavior of the agents on the process of swaps. Reasoning strategically can drive an agent to refuse an immediate profitable deal, for instance in a look-ahead search. Studying the social welfare of the possible outcomes would also be interesting, even if classical measures such as the egalitarian or the utilitarian social welfare are undefined in our model without utilities. Like in the price of anarchy/stability, how bad a stable outcome can be, compared to an allocation that is reached after swaps guided by a central authority?

Another future direction is to allow more than two agents to exchange their objects along the network and see which allocations emerge. This extension would be very interesting because it could enable to generalize the conditions of the top trading cycle algorithm where any cycle of exchanges is allowed provided they are rational: in such a context the top trading cycle can be applied when the social network is a complete graph. In such a perspective, one could think that the groups of agents that are enable to exchange their object could be modeled for instance as cliques of the network such as in the work of Chevaleyre et al. [2007c], in an idea similar to Chapter 2.

The parameterized approach allows progress in the understanding of the problems and leads to significant and realistic positive results. So far, we have considered restrictions on the network as well as on the solution size. A natural extension is to investigate the influence of a third dimension: constraints on the preference profile, e.g., single-peaked or single-crossing domains. Furthermore, assuming the full knowledge of the preferences and the network is not relevant in all the contexts. Relaxing this assumption could be a challenging future work.

We have seen in this chapter that, when the social network models the ability of agents to collaborate in a "win-win" perspective, the structure of the graph can deeply affect the complexity of the problems under study, as well as their outcome. Beyond collaboration, the social network can model what the agents are able to observe about the situation of other agents. Typically, in a context of resource allocation, the links of the network could tell the partial allocation that the agents are able to observe. The model of swap dynamics can also be viewed under this point of view: the swap is possible between two agents if they both know the object assigned to the other and they prefer it to their current endowment (note that it is impossible to trade an object for another whose owner is not known). In resource allocation, by considering that the network models a partial vision of the agents, one could think that the graph topology would significantly impact the perceived *fairness* of the allocation.

Part II

The Social Network as an Informative Tool

Introductory comments

So far, we have examined two problems of social choice where the possibility of collaboration among agents is modeled by a social network, represented by an undirected graph over the agents. In this part, we focus on another type of interaction, more related to the communication and the visibility of the agents with one another. We assume that the agents only have information about a limited part of the agents, even if they know that other agents exist. This visibility relation among the agents can also be modeled by a social network. However, this time, there is no reason to assume by default that the visibility relation is symmetric. Indeed, one agent can communicate some information about her situation to another agent but this agent is not obliged to disclose her own situation.

Consequently, we represent in this part the social network $G = (N, E)$ as a directed graph, where the binary relation over N that is induced by E does not need to be symmetric. Nevertheless, for some specific results, it makes sense to impose some restrictions on the graph, for instance that G is undirected to express the fact that the binary relation represented by E is symmetric.

We investigate the two problems that were already studied in the first part, namely house allocation and strategic voting. However, since the social network has a different interpretation, we study these problems under different views: house allocation is examined through perceived fairness whereas we deal with uncertainty in strategic voting. More precisely, in Chapter 4, given an allocation of items to agents, agents can only be envious of their successors in the graph. The vision of the agents is then partial and given by the social network, allowing to define a local notion of the envy-freeness criterion. In Chapter 5, voters devise strategies for manipulation in iterative voting, based on a partial information given by the social network: they can only observe the votes of their successors in the graph.

Chapter 4

Envy-Freeness in House Allocation

Abstract

We study the fair division problem consisting in allocating exactly one item per agent, which is called *house allocation*, so as to avoid, or minimize, envy. We focus on a particular setting where the agents can only be envious of their direct successors in the social network, modeled as a directed graph over the agents. The agents only perceive from the allocation the objects assigned to their successors in the graph, or more generally their successors are the only agents they can envy because of some specific relationships. This defines a notion of *local envy* according to the social network. The existence of a *locally envy-free* allocation is investigated, as well as the minimization of local envy in case a locally envy-free allocation does not exist. We also examine a variant of the problem where the agents can be located on the network by a central authority, in such a way to avoid local envy. These problems turn out to be computationally hard even on very simple graph structures, but we identify several tractable cases. We further provide efficient algorithms and experimental insights.

Résumé

Ce chapitre est dédié à un problème de partage équitable consistant à affecter exactement une ressource par agent, de manière à éliminer ou minimiser l'envie entre les agents. On s'intéresse à un cadre particulier où les agents ne peuvent envier que leurs successeurs directs dans le réseau social, modélisé par un graphe orienté. En effet, les agents ne perçoivent de l'allocation que les objects affectés à leurs successeurs dans le graphe, ou plus généralement leurs sucesseurs sont les seuls qu'ils peuvent envier en raison de relations sociales particulières. Ceci permet de définir une notion d'*envie locale* dépendant du réseau social. On se propose d'étudier les conditions d'existence d'une allocation localement sans envie, ainsi que la minimisation de l'envie locale, dans une perspective d'optimisation, lorsqu'il n'existe pas d'allocation localement sans envie. Dans une variante du problème, les agents eux-mêmes peuvent être affectés à des noeuds du graphe

This chapter is an extension of [Beynier, Chevaleyre, Gourvès, Lesca, Maudet, and Wilczynski, 2018].

par l'autorité centrale, de manière à éviter l'envie locale. Tous ces problèmes s'avèrent être difficiles d'un point de vue calculatoire, même avec des structures de graphes très simples, mais on relève néanmoins certains cas pouvant être résolus efficacement. Des algorithmes efficaces et des résultats expérimentaux sont également fournis.

Contents

4.1 Introduction

Fairly allocating resources to agents is a fundamental problem in Economics and Computer science, and has been the subject of intense investigations [Young, 1995, Brams and Taylor, 1996, Moulin, 2004, Bouveret et al., 2016, Chevaleyre et al., 2017]. Recently, several papers have explored the consequences of assuming in such settings an underlying network connecting agents [Chevaleyre et al., 2007c, Abebe et al., 2017, Bei et al., 2017]. The most intuitive interpretation is that agents have limited information regarding the overall allocation. Two agents can perceive each other if they are directly connected in the graph.

4.1.1 Envy-freeness and social network

A fairness measure, very sensitive to the information available to agents, is the notion of *envy* [Tinbergen, 1946, Foley, 1967, Varian, 1974]. Envy occurs when an agent prefers the share of some other agents over her own. Accounting for a network topology boils down to replace "other agents" by "neighbors". The notion of envy can thus naturally be extended to account for the limited visibility of the agents. Intuitively, an allocation is *locally envy-free* if none of the agents envies her neighbors in the social network. This notion has been referred as *graph*, *social*, or *local envy-freeness*, and has been introduced for cake-cutting problems [Cato, 2010, Abebe et al., 2017, Bei et al., 2017], as well as for allocation problems with indivisible resources [Chevaleyre et al., 2017, Aziz et al., 2018, Bredereck et al., 2018, Flammini et al., 2018].

We are concerned with the allocation of indivisible goods within a group of agents. Our study is focused on *house allocation*, the specific setting where each agent should receive exactly one object. The agents are embedded in a social network, which captures the possibility of envy among them. Since the relation of envy may not be symmetric, the social network is represented by a directed graph over the agents. Nevertheless, we sometimes restrict ourselves to the symmetric case of undirected graphs. The classical notion of envy-freeness, which corresponds to the case where the network is complete, is not a very exciting notion in house allocation. Indeed, for an allocation to be envy-free, each agent must get her top object, implying that the only positive instances are trivial and very rare. However, when an agent can only be envious of a subset of the other agents, given by the links of the social network, she may not need to get her top-resource to be envy-free. The connections between the agents are then crucial issues in order to compute a locally envy-free allocation, as we can observe in the following scenario.

Example 4.1 *Suppose there is a team of workers taking their shifts in sequence, to which their employer must assign different jobs. The workers have preferences regarding the jobs. Concretely, there are three jobs, "chop the tree", "mow the lawn", and "trim the hedge", and three gardeners, Alice, Bruno and Carlos, with the following preferences.*

$$\begin{aligned} \textit{Alice:} \quad & \textit{chop} \succ \textit{mow} \succ \textit{trim} \\ \textit{Bruno:} \quad & \textit{mow} \succ \textit{chop} \succ \textit{trim} \\ \textit{Carlos:} \quad & \textit{chop} \succ \textit{trim} \succ \textit{mow} \end{aligned}$$

Clearly, there is no envy-free allocation in such a context because Alice and Carlos both prefer the job "chop the tree". However, the shifts are contiguous and the employees *work at the same place. Therefore, they only have the opportunity to see the job allocated to some other workers, as one ends and the other one begins her shift, inducing a symmetric visibility relation among the workers. Actually the three workers take shifts in the following order: Alice, then Bruno and finally Carlos. This can be modeled as a path topology, as depicted on the graph below, because of the symmetry of the visibility relation and the chronological order of the shifts.*

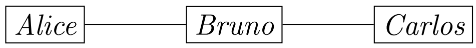

Since Alice and Carlos never meet during the shifts, they may not know each other or at least may not feel envy towards the other, because they are not able to observe the object of envy. This implies that this may not be useful to look for an allocation that is envy-free according to these two agents. By allocating the job "chop the tree" to Alice, "mow the lawn" to Bruno, and "trim the hedge" to Carlos, we get an envy-free allocation if we disregard the fact that agent Alice and Carlos may be envious of each other. Therefore, restricting to locally envy-free allocations, where the possibility of envy is given by a graph, appears as a relevant and realistic way to allocate resources in a fairly manner.

One could object that, in Example 4.1, Carlos may still be envious of Alice, because he knows that she must have received the task that Bruno did not get, i.e., "chop the tree". This is a valid point, to which we provide two counter-arguments. First, as a technical response, note that in general, agents would not know exactly who gets the items they do not see. Thus, although agents may know that they must be envious of *some* agents, they cannot identify which one, which makes a significant difference in the case of envy. In a similar idea, the links of the network can be interpreted as membership in the same social class, and thus agents may only envy someone who is similar to themselves. For instance, one may want to have a better car than the people she went to school with but be indifferent if a random person has a nicer car. Our second point is more fundamental and concerns the model and the motivation of this work. Clearly, the existence of a network may be due to an underlying notion of proximity (either geographical, or temporal as in our example) in the problem. However, another interpretation of the meaning of links must be emphasized: links may represent envy the central authority is concerned with. In other words, although there may theoretically be envy among all agents, the central authority may have reasons to only focus on some of these envy links. For instance, one may wish to avoid envy among members of the same team in an organization, because they actually work together on a daily basis (in that case links may capture team relationships). Under this interpretation, an undirected network of degree $n - 2$ could for instance model a situation where agents team-up in pairs and conduct a task together, sharing their resources. In a similar vein, one may focus on avoiding envy among "similar" agents, because they may be legitimate to complain if they are not treated equally despite similar competences.

Our previous example shows how modeling the possibility of envy by a network can be meaningful. By describing a situation where tasks should be assigned to workers with the same grade, envy is defined as symmetric. However, this is not a general rule and many natural real-life examples highlight the fact that envy may not be symmetric. Contexts where a hierarchical relation exists among the agents, like in Example 4.2, are typical examples.

Example 4.2 *Let us consider a scenario where the director general of a company aims at rewarding her employees at the end of the year. The company is composed of 7 employees, including one deputy director, one project manager, one marketing manager, and four other employees either in the marketing branch or in the project branch. They are all tied by the following hierarchy.*

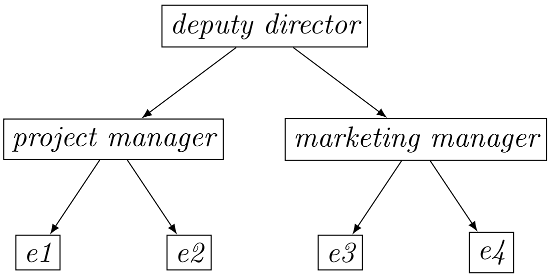

This hierarchy among the employees of the company can tell us the possible relations of envy. If an agent estimates that the reward obtained by one of the employees she manages is better than hers, then she can be envious. However an employee has less reasons to be envious of the reward of her manager. In this example, envy is not symmetric and the network can be represented by the directed graph describing the hierarchy among the employees. Observe that the two managers are not envious of each other because they belong to different branches and do not consider that their two fields can be compared. The typical networks of such hierarchical examples are directed acyclic graphs like in this example, or can be for instance transitive graphs (the binary relation represented by the graph is transitive, see Definition 1.3) if employees may be envious of all the employees that are lower than them in the hierarchy.

Generally, there is no reason to assume that envy has to be symmetric and thus we model the possibility of envy via a directed graph. We nevertheless sometimes assume a symmetric context for envy and then use an undirected graph as a special case.

4.1.2 Related work

Our work is connected to a number of other recent contributions. Among the works dealing with local envy-freeness that we have already cited, one could mention in particular the work of Bredereck et al. [2018], where the authors also investigate the existence of a locally envy-free allocation. The main differences are about the model of fair allocation in itself: they consider cardinal preferences whereas we deal with ordinal ones, and they study general resource allocation problems whereas we focus on the house allocation setting.

Other works combine fair allocation and graphs although they do not assume that the agents are embedded in the graph. For instance, the allocation of a graph has recently been studied [Bouveret et al., 2017]. In this context, the nodes of the graph represent indivisible resources to allocate and edges formalize connectivity constraints between the resources. Some computational aspects of allocating agents on a line (or undirected path) are also discussed by Aziz et al. [2017]: in that case the line concerns the items (e.g., slots) to be allocated, and induces a domain restriction (stronger than single-peakedness). In these cases the model is really different, since the graph is capturing dependencies between the resources (like spatial dependencies for pieces of land).

Several ways for a central authority to control fair division have been discussed by Aziz et al. [2016b]: the structure of the allocation problem can be changed by adding or removing items to improve fairness. Interestingly, our model introduces a new type of control action: locating agents on a graph. Finally, because envy-freeness is difficult to achieve in general (with indivisible items) [de Keijzer et al., 2009], different notions of degree of envy have been studied (see e.g., Nguyen and Rothe [2014], Caragiannis et al. [2009], Lipton et al. [2004], and a recap of these degrees in Section 1.4.1).

4.1.3 Contributions and organization

In this chapter, we study a local notion of envy-freeness for the specific problem of house allocation. The relaxation of the envy-freeness criterion is particularly relevant in house allocation because the standard requirement of envy-freeness is too restrictive in such a context. Moreover, whereas the literature on envy-freeness conditioned by graphs deals with cardinal preferences, we focus on the specific case of ordinal preferences, which implies for instance particular measures for optimizing the non-envy of the society.

A formal definition of the model, together with the definition of the main problems that we address, are provided in Section 4.2. Section 4.3 is dedicated to the problem of deciding whether a central planner can allocate the objects, such that no agent will envy a successor in the network. In this problem, called Dec-LEF, the central planner has a complete knowledge of the social network and the agents' rankings of the objects. We identify intractable and polynomial cases of this decision problem, with respect to the number of successors of each agent, that is the out-degree of the nodes in the graph representing the social network. The problem is also investigated through another relevant parameter that is the size of a (minimal) *vertex cover* in the graph.

Section 4.4 is dedicated to optimization problems with two different perspectives: maximizing the number of locally envy-free agents, and maximizing the degree of non-envy of the society. We provide approximation algorithms for both approaches.

A variant of Dec-LEF called Dec-Location-LEF is studied in Section 4.5. This problem asks if one can decide both the placement of the agents and the object allocation so as to satisfy local envy-freeness. A natural interpretation for this problem relies on a network which models connections related to working hours of employees in a team, like in Example 4.1. In this idea, the manager of the team could decide to organize the time schedule of her employees in the same time she assigns to them tasks. The problem is shown to be NP-complete, and a special case is resolved in polynomial time. Another variant of the problem is examined in Section 4.6, where the reachability of a locally envy-free allocation, by means of swap dynamics as defined in Chapter 3, is investigated. Before concluding, we report some experimental results in Section 4.7.

4.2 Problems related to local envy-freeness

The model of resource allocation that we study is a house allocation problem, as described in Section 1.4.2.a. House allocation refers to a resource allocation problem where, given a set of agents $N = \{1, \dots, n\}$ and a set of indivisible resources $O = \{o_1, \dots, o_r\}$, the number of resources is equal to the number of agents, i.e., $r = n$, and each agent $i \in N$ should receive exactly one resource $o \in O$. An allocation of objects to agents is then an allocation $\pi : N \to O$ where $\pi(i) \neq \pi(j)$. The links of the social network, represented by a directed graph $G = (N, E)$ over the agents, provide for each agent $i \in N$ the set of agents that she may envy. More precisely, agent i can envy agent j only if there exists an arc from i to j in E. An instance of a house allocation problem is thus described by a tuple $\mathcal{I} = \langle N, O, \succ, G = (N, E)\rangle$, where we assume that the agents have strict preferences over the items. When the social network G is dense, it may be easier to describe it through its complementary graph $\overline{G}$ which is the unique graph defined on the same vertex set and such that there is an arc between two vertices if and only if this arc does not exist in G.

We are now able to define the notion of *local envy-freeness*.

Definition 4.1 (Local envy-freeness (LEF)) *An allocation π is locally envy-free (LEF) if no ordered pair of agents $(i, j) \in E$ satisfies $\pi(j) \succ_i \pi(i)$.*

Classical envy-freeness (Definition 1.30) refers to the case where the social network G is a complete graph. For a given allocation, an agent is said to be locally envy-free (LEF) if she prefers her object to the object(s) of her successor(s) in the social network.

Several notions of degrees of envy have been studied [Chevaleyre et al., 2017, Nguyen and Rothe, 2014, Caragiannis et al., 2009, Lipton et al., 2004]. In our context, we shall study the number of envious agents, and a measure capturing some simple notion of intensity of envy, in terms of the difference of ranks between items. These two notions would correspond to $e^{sum,max,bool}$ and $e^{sum,sum,raw}$, up to normalization, under the classification of Chevaleyre et al. [2017] (see Section 1.4.1).

Definition 4.2 (Degrees of (non)-envy) *Given an allocation π, the degree of envy of agent i towards an agent j such that $(i, j) \in E$ is*

$$e(\pi, i, j) = \frac{1}{n-1} \max\{0, \mathfrak{r}_i(\pi(i)) - \mathfrak{r}_i(\pi(j))\}$$

where $\mathfrak{r}_i(o)$ is the rank of object o in i's preferences. The average degree of envy (respectively of non-envy) is $\mathscr{E}(\pi) = \frac{1}{|E|}\sum_{(i,j)\in E} e(\pi, i, j)$ (respectively is $\mathcal{N}\mathscr{E}(\pi) = 1 - \mathscr{E}(\pi)$).

Note that for a given allocation π, an agent i envies a successor j if and only if $e(\pi, i, j) > 0$.

We mainly address four problems: Dec-LEF, Max-LEF, Max-NE and Dec-Location-LEF. The first one is a decision problem regarding the existence of an LEF allocation over a given social network.

Dec-LEF:	
Instance:	Instance $\langle N, O, \succ, G\rangle$
Question:	Is there an LEF allocation π?

The second and the third ones are optimization problems in which an allocation that is as close as possible to local envy-freeness is sought, using the aforementioned criteria.

Max-LEF:	
Instance:	Instance $\langle N, O, \succ, G\rangle$
Problem:	Find an allocation that maximizes the number of LEF agents

MAX-NE:

Instance: Instance $\langle N, O, \succ, G\rangle$

Problem: Find an allocation that maximizes the average degree of non-envy $\mathcal{NE}(\pi)$

In DEC-LOCATION-LEF, one has to place the agents on the network in addition to the allocation. This placement makes sense if we consider Example 4.1 where the agents take shifts.

DEC-LOCATION-LEF:

Instance: Instance $\langle N, O, \succ\rangle$, graph $G = (V, E)$

Question: Are there two bijections $\pi : N \to O$ and $\mathcal{L} : N \to V$ (π and $\mathcal{L}$ determine the allocation of the objects and the location of the agents on the network, respectively) such that $\pi(i) \succ_i \pi(j)$ for every arc $(\mathcal{L}(i), \mathcal{L}(j)) \in E$?

Example 4.3 *As a warm-up, let us consider an instance with five agents, where* $N = \{1, 2, 3, 4, 5\}$ *and* $O = \{o_1, o_2, o_3, o_4, o_5\}$*. The visibility relation among the agents is supposed to be symmetric and to be based on a chronological order, like in Example 4.1. More precisely, the agents are connected to each other via a path. The network and the preferences are given as follows.*

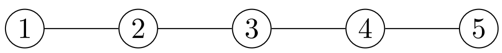

$1: \; o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5$

$2: \; o_3 \succ o_1 \succ o_2 \succ o_4 \succ o_5$

$3: \; o_1 \succ o_2 \succ o_4 \succ o_3 \succ o_5$

$4: \; o_2 \succ o_1 \succ o_4 \succ o_3 \succ o_5$

$5: \; o_3 \succ o_5 \succ o_2 \succ o_1 \succ o_4$

Is there an LEF allocation of goods to agents? If not, what is the minimum number of envious agents? Finally, is it possible to find an LEF allocation by relocating agents on this path?

Let us try to construct an allocation π *that is LEF. Observe that the agents* 3 *and* 4*, who are neighbors, both rank objects* o_1 *and* o_2 *as their first two preferred objects and rank the remaining objects in the last positions of their preference ranking following the same order. This implies that they cannot obtain one of the remaining objects in an LEF allocation, i.e., an object within* $\{o_3, o_4, o_5\}$*. Indeed, if only one agent between* 3 *and* 4 *obtains an object in this subset, then she will be envious of the other agent. Otherwise, if they both get an object from this subset, since their preferences over these objects are the same, one of them will necessarily envy the other. Therefore, we have to assign object* o_1 *to agent* 3 *in* π*, as well as* o_2 *to agent* 4*, because they both prefer this item to the other one, respectively. Consequently, agent* 2*, neighbor of agent* 3*, must obtain an object preferred to* o_1*, which is assigned to agent* 3*. She only prefers object* o_3 *to* o_1*, so we have to assign* o_3 *to agent* 2 *in* π*. Agent* 5*, neighbor of agent* 4*, must get an object preferred to* o_2*, which is assigned to agent* 4*. The only possible objects are* o_3 *and* o_5*, but* o_3 *is already assigned to agent* 2*, thus we assign object* o_5 *to agent* 5 *in* π*. Finally, there only remains object* o_4 *and agent* 1*. Agent* 1 *prefers* o_3*, the object assigned to her* *neighbor (agent 2), to o_4. Therefore, by assigning o_4 to agent 1 in π, we get that agent 1 is envious of agent 2. Thus, there is no LEF allocation in this instance, implying that this is a no-instance of* DEC-LEF.

Observe that allocation π is almost LEF since only agent 1 is envious. Therefore, there exists an allocation with only one envious agent. Because there is no LEF allocation, this is the minimum number of envious agents that we can obtain in any allocation.

Finally, remark that, in allocation π, the only envious agent 1 gets object o_4, and the only object that agent 1 likes less than o_4 is object o_5. Object o_5 is owned by agent 5 who is located at a leaf of the path and who on the opposite prefers o_5 to o_4. Therefore, by considering a new location of the agents which is the same as the current graph but with agent 1 as a leaf of the path who is only connected to agent 5 (i.e., the new path $[2, 3, 4, 5, 1]$), allocation $\pi = (o_4, o_3, o_1, o_2, o_5)$ is LEF. Hence, this is a yes-instance of DEC-LOCATION-LEF.

4.3 Existence of a locally envy-free allocation

This section is devoted to DEC-LEF. Our main findings settle the computational status of DEC-LEF with respect to the degree of the nodes in the social network, as well as the size of a vertex cover.

Trivially, every allocation is LEF if the social network is empty. In a complete graph, the only configuration that enables the existence of an LEF allocation is when the most preferred object of each agent is different. Under such a condition, there is a unique LEF that assigns to each agent her best object.

Our first result shows that DEC-LEF is computationally difficult, even if the social network is very sparse and undirected. This is somewhat surprising as such a network offers very little possibility for an agent to be envious.

Theorem 4.1 DEC-LEF *is* NP*-complete, even if G is a matching, i.e., an undirected regular graph of degree 1.*

Proof: The reduction is from 3-SAT, known to be NP-complete (Theorem 1.2). In 3-SAT, we are given a set of clauses $\mathcal{C} = \{C_1, \cdots, C_s\}$ defined over a set of variables $X = \{x_1, \cdots, x_v\}$, where each clause contains exactly three literals. The question is whether there exists a truth assignment which satisfies all the clauses.

Take an instance $\mathcal{I} = \langle \mathcal{C}, X \rangle$ of 3-SAT and create an instance $\mathcal{J}$ of DEC-LEF as follows.

The set of objects is $O = \{u_i^j : i \in [v],\ j \in [s]\} \cup \{\overline{u}_i^j : i \in [v],\ j \in [s]\} \cup \{q_j : j \in [s]\} \cup \{t_{ij} : i \in [v],\ j \in [s]\} \cup \{h_\ell : \ell \in [s.(v-1)]\}$. Here, u_i^j and $\overline{u}_i^j$ correspond to the unnegated and negated literals of x_i in clause C_j, respectively, q_j corresponds to clause C_j, and objects t_{ij} and h_ℓ are gadgets. Thus, $|O|= 4sv$.

The set of agents N is built as follows. For each $(i, j) \in [v] \times [s]$, create a pair of variable-agents X_{ij} and X'_{ij} which are linked in the social network. For each $j \in [s]$, create a pair of clause-agents K_j and K'_j which are linked in the network. For each $\ell \in [s.(v-1)]$, create a pair of garbage-agents L_ℓ and L'_ℓ which are linked in the network. Thus, the network consists of a perfect matching with $4sv$ agents.

Each clause C_j is associated with the pair of clause-agents (K_j, K'_j), q_j and three objects corresponding to its literals. For example, $C_2 = x_1 \vee x_4 \vee \overline{x}_5$ is associated with objects q_2, u_1^2, u_4^2, and $\overline{u}_5^2$. The preferences of the clause-agents are:

$$\begin{aligned} K_j: &\quad q_j \succ \text{(the 3 objects related to the literals of } C_j) \succ \text{rest}_j \\ K'_j: &\quad \text{(the 3 objects related to the literals of } C_j) \succ q_j \succ \text{rest}_j \end{aligned}$$

where "rest_j" means the remaining objects. Both "rest_j" and the three objects corresponding to the literals of c_j are arbitrarily ordered, but in the same way for K_j and K'_j. Each variable x_i is associated with the s pairs of variable-agents (X_{ij}, X'_{ij}), for $j \in [s]$. The preferences of these variable-agents are defined as follows (note that $j+1 = 1$ when $j = s$):

$$\begin{aligned} X_{ij}: &\quad u_i^j \succ t_{ij} \succ \overline{u}_i^j \succ t_{ij+1} \succ \text{rest}_i^j \\ X'_{ij}: &\quad t_{ij} \succ u_i^j \succ t_{ij+1} \succ \overline{u}_i^j \succ \text{rest}_i^j \end{aligned}$$

where "rest_i^j" means the remaining objects arbitrarily ordered, but in the same way for X_{ij} and X'_{ij}. The preferences of the garbage-agents (L_ℓ, L'_ℓ), for $\ell \in [s.(v-1)]$, are:

$$\begin{aligned} L_\ell: &\quad h_\ell \succ U \succ \text{rest}_\ell \\ L'_\ell: &\quad U \succ h_\ell \succ \text{rest}_\ell \end{aligned}$$

where $U = \{u_i^j, \overline{u}_i^j : i \in [v],\ j \in [s]\}$, "$\text{rest}_\ell$" is the set of remaining objects, and both U and "rest_ℓ" are arbitrarily ordered in the same way for L_ℓ and L'_ℓ.

We claim that there is an LEF allocation in instance $\mathcal{J}$ if and only if there is a truth assignment of the variables in X satisfying all the clauses in instance $\mathcal{I}$.

Take a truth assignment ϕ of the variables in X which satisfies all the clauses in $\mathcal{C}$. One can allocate objects to each pair of variable-agent (X_{ij}, X'_{ij}) in such a way that it is LEF. More precisely, if x_i is set to true in ϕ, then agent X_{ij} gets object $\overline{u}_i^j$ and agent X'_{ij} gets object $t_{i\,j+1}$ (where $t_{i\,m+1} := t_{i1}$). Otherwise, i.e., if x_i is set to false in ϕ, X_{ij} gets u_i^j and X'_{ij} gets t_{ij}. One can allocate objects to each clause-agent pair (K_j, K'_j) in such a way that it is LEF. Indeed, clause C_j is satisfied thanks to one of its literals. Therefore, it suffices to give to agent K_j object q_j and to agent K'_j an unallocated object corresponding to a literal that makes clause C_j true in ϕ. Finally, allocate objects to each garbage-agent pair (L_ℓ, L'_ℓ) in such a way that it is LEF by assigning object h_ℓ to agent L_ℓ and any unallocated object of U to agent L'_ℓ.

Suppose an LEF allocation exists for $\mathcal{J}$. Consider a variable x_i. By construction of the preferences of the variable-agent pair (X_{i1}, X'_{i1}), we observe that there is absence of envy in only two cases: X_{i1} gets u_i^1 and X'_{i1} gets t_{i1}, otherwise X_{i1} gets $\overline{u}_i^1$ and X'_{i1} gets t_{i2}. If we are in the first case, then there is absence of envy between X_{is} and X'_{is} only if X_{is} gets u_i^s and X'_{is} gets t_{is} because t_{i1} is already allocated, and so on; the X_{ij}'s get all the u_i^j's (i is fixed but $1 \le j \le s$). If we are in the second case, then there is absence of envy between X_{i2} and X'_{i2} only if X_{i2} gets $\overline{u}_i^2$ and X'_{i2} gets t_{i3} because t_{i2} is already allocated, and so on; the X_{ij}'s get all the $\overline{u}_i^j$'s (i is fixed but $1 \le j \le s$). Thus, set x_i to $false$ (respectively, x_i to $true$) if every X_{ij} gets u_i^j (respectively, X_{ij} gets $\overline{u}_i^j$).

Consider any clause C_j. By construction of the preferences of the clause-agent pair (K_j, K'_j), we observe that there is absence of envy in only three cases: K_j gets q_j and

K'_j gets one of the three objects associated with the literals of C_j. Since the allocation is LEF, there is some i^* such that K'_j gets either $u^j_{i^*}$ or $\overline{u}^j_{i^*}$, and this object is not allocated to a variable-agent. Thus, C_j is satisfied by the above truth assignment. To conclude, all the clauses are satisfied. □

The strength of this result lies on the fact that the network structure is extremely simple. As a consequence, it can easily be used as a building block to show hardness of a large variety of graphs. Despite this negative result, we can establish that the problem is solvable in polynomial time when the social network is a directed graph G which does not contain directed cycles.

Proposition 4.2 *An LEF allocation is guaranteed to exist if G is a directed acyclic graph (DAG) and finding one can be done in polynomial time.*

Proof: A DAG has at least one source, i.e., a vertex with in-degree 0. If a source of a DAG is deleted, then we still get a (possibly empty) DAG. The algorithm computing an LEF allocation works as follows: while the social network is non-empty, find a source agent i, allocate i her most preferred object $o_i \in O$, remove o_i from O, and delete i. This algorithm guarantees by construction to find a locally envy-free allocation. Indeed, any agent can only envy her successors in the graph but she always chooses an object before them. □

Observe that the previous algorithm also returns a Pareto-efficient allocation and mimics a well known procedure for allocating goods called *serial dictatorship* [Satterthwaite and Sonnenschein, 1981] (see Section 1.4.2.a).

Note that DAGs actually characterize exactly those digraphs which guarantee the existence of an LEF allocation for any preference profile. Indeed, if a directed cycle (dicycle) exists, then it suffices to set the preferences of all the agents within the dicycle to be exactly the same to get a *no*-instance of Dec-LEF.

One could think that if the dicycles of the graph are actually cliques, like in a transitive graph (see Definition 1.3), recognizing a *yes*-instance of Dec-LEF could be easy like in a complete graph. However, even when the graph is transitive, the problem is hard, by Theorem 4.1, because the matching is a transitive graph.

We will now investigate more in details the impact of the degree of the network on the existence of an LEF allocation.

4.3.1 Local envy-freeness and degree of nodes

First of all, note that some objects cannot be assigned to certain agents for the allocation to be LEF. For example, the best object of an agent cannot be assigned to one of her successors. More generally, for a given agent, all her successors in the network must obtain an object that she likes less than the object assigned to her, leading to the following observations.

Observation 4.1 *In any LEF allocation, an agent with k successors must get an object ranked among her $n - k$ top objects.*

Observation 4.2 *In any LEF allocation, the best object for an agent is either assigned to herself or to one of her successors in* $\overline{G}$.

These observations imply that an agent having $n-1$ successors in the network must receive her best object in any LEF allocation.

Surprisingly, Dec-LEF is computationally hard even when the graph is undirected and the degree of each agent in the graph is equal to 1, as stated by Theorem 4.1, with the case of the matching. This hardness result can be extended to any undirected graph of constant degree greater or equal to 1.

Proposition 4.3 Dec-LEF *is* NP*-complete on a path, on a cycle, and generally on undirected regular graphs of degree* k *for* $k \geq 1$ *constant.*

Proof: The idea is to introduce additional dummy agents connected to agents from the original matching instance in the proof of Theorem 4.1, and to make sure that each dummy agent D_i will obtain her associated dummy resource d_i. For this purpose, each dummy agent ranks first "her" dummy resource d_i, followed by a copy of the ranking (minus d_i) of an (arbitrary) neighbor, while all her neighbors rank d_i after the initial objects of the instance in their ranking. Consequently, in any LEF allocation, each dummy agent must receive her associated dummy resource, and the rest of the allocation must satisfy the conditions of the existence of an LEF allocation on a matching.

More precisely, for obtaining a path from a matching with n agents, we introduce $\lceil \frac{n}{2} \rceil - 1$ dummy agents in order to connect each edge of the matching along a path, and for a cycle we introduce one more dummy agent in order to close the path by connecting its two leaves. For obtaining a general undirected regular graph of degree k (k constant), we add $k-1$ dummy agents connected within a clique to the two vertices of each edge of the matching in the original instance. □

Given this result, one may suspect Dec-LEF to be hard on any graph structure that is denser than the matching. However, our next result shows that if the social network is dense enough, then Dec-LEF is polynomial.

Theorem 4.4 Dec-LEF *in graphs of minimum out-degree* $n-2$ *is solvable in polynomial time.*

Proof: Note that $\overline{G}$ is a graph where every node has at most one successor. Let us define as $\phi : N \to N$, the mapping such that $\phi(i)$ is the only successor of agent i in $\overline{G}$, if the out-degree of i in G is equal to $n-2$, and $\phi(i) = i$ if the out-degree of i in G is equal to $n-1$, i.e., i has no successor in $\overline{G}$.

We reduce the problem to 2-SAT which is solvable in linear time [Aspvall et al., 1979]. Let us consider boolean variables x_{ij} for $1 \leq i, j \leq n$, such that x_{ij} is true if and only if object o_j is assigned to i. Denote by o_i^ℓ the index j of object o_j at position ℓ in the preference order of agent i. Consider the following formula φ:

$$\bigwedge_{i\in[n]} (x_{io_i^1} \vee x_{io_i^2}) \wedge \bigwedge_{\substack{1\leq i<\ell\leq n \\ 1\leq j\leq n}} (\neg x_{ij} \vee \neg x_{\ell j}) \wedge \bigwedge_{i\in[n]} (x_{io_i^1} \vee x_{\phi(i)o_i^1})$$

The first part of formula φ expresses that each agent must obtain an object within her top 2, as noted in Observation 4.1. By combination with the second part of φ, we get that the solution must be an assignment: each agent must obtain her first or second choice but not both since every object is owned by at most one agent and $|N|= |O|$. Observation 4.2 implies that the best object for agent i must be assigned either to agent i or $\phi(i)$. This condition is given by the last part of the formula. Hence, formula φ exactly translates the constraints of an LEF allocation. □

Interestingly, the status of Dec-LEF changes between social networks of out-degree at least $n-2$ and those of out-degree $n-3$. Actually, the problem is hard even for undirected graphs of degree $n-3$.

Theorem 4.5 Dec-LEF *is* NP*-complete even if G is a regular undirected graph of degree $n-3$.*

Proof: The reduction is from (3,B2)-SAT, known to be NP-complete (Theorem 1.4). This problem is a restriction of 3-SAT where each literal appears exactly twice in the clauses, and therefore, each variable appears four times in total. From an instance $\mathcal{I} = \langle \mathcal{C}, X \rangle$ of (3,B2)-SAT, where $\mathcal{C} = \{C_1, \dots, C_s\}$ is a set of s clauses and $X = \{x_1, \dots, x_v\}$ is a set of v variables, we create an instance $\mathcal{J} = \langle N, O, \succ, G \rangle$ of Dec-LEF where G is an undirected graph of degree $|N|-3$.

Instead of describing G, we describe its complement $\overline{G}$. Note that $\overline{G}$ is a regular graph of degree 2. Hence, $\overline{G}$ is a collection of cycles. For each variable x_i, we introduce dummy variable-objects q_i^1 and q_i^2 and literal-objects $u_i^1, u_i^2, \overline{u}_i^1$ and $\overline{u}_i^2$ corresponding to its first and second occurrence as an unnegated and negated literal, respectively, as well as a cycle in $\overline{G}$ containing literal-agents $X_i^1, \overline{X}_i^1, X_i^2$ and $\overline{X}_i^2$, connected in this order. The preferences of the literal-agents are as follows, for $i \in [v]$ ($[\dots]$ stands for an arbitrary order over the remaining objects):

$$X_i^1 : \; q_i^1 \succ q_i^2 \succ u_i^1 \succ [\dots] \qquad \overline{X}_i^1 : \; q_i^1 \succ q_i^2 \succ \overline{u}_i^1 \succ [\dots]$$
$$\overline{X}_i^2 : \; q_i^2 \succ q_i^1 \succ \overline{u}_i^2 \succ [\dots] \qquad X_i^2 : \; q_i^2 \succ q_i^1 \succ u_i^2 \succ [\dots]$$

Note that only the three top objects are represented since no object ranked below can lead to an LEF allocation (see Observation 4.1). Note also that in any LEF allocation, either q_i^1 and q_i^2 are allocated to agents X_i^1 and X_i^2, or q_i^1 and q_i^2 are allocated to agents $\overline{X}_i^1$ and $\overline{X}_i^2$. Indeed, if q_i^1 (respectively, q_i^2) is allocated to an agent Y not in $X_i := \{X_i^1, X_i^2, \overline{X}_i^1, \overline{X}_i^2\}$, then since Y is the neighbor in G of each of the agents in X_i, the agents X_i^1 and $\overline{X}_i^1$ (respectively, X_i^2 and $\overline{X}_i^2$) will envy Y. Moreover, if q_i^1 is owned by X_i^1 and q_i^2 by $\overline{X}_i^2$ (respectively, by $\overline{X}_i^1$ and X_i^2), then agent $\overline{X}_i^1$ (respectively, X_i^1) will be envious of $\overline{X}_i^2$ (respectively, X_i^2), her only neighbor in G within X_i, because she cannot have q_i^1 and this is the only object that she prefers to q_i^2. The case where q_i^1 and q_i^2 are allocated to agents X_i^1 and X_i^2 can be interpreted in $\mathcal{I}$ as setting x_i to true, and the case where q_i^1 and q_i^2 are allocated to agents $\overline{X}_i^1$ and $\overline{X}_i^2$ as setting x_i to false.

For each clause C_j we introduce dummy clause-objects d_j^1 and d_j^2, as well as a cycle in $\overline{G}$ containing clause-agents K_j^1, K_j^2, K_j^3. The preferences of clause-agent K_j^h, for $j \in [s]$ and $h \in [3]$, are:

$$K_j^h: \quad d_j^1 \succ d_j^2 \succ \ell(j,h) \succ [\ldots]$$

where $\ell(j,h)$ is the literal-object corresponding to the h^{th} literal of C_j. Note that an allocation is LEF if d_j^1, d_j^2 and one literal-object corresponding to a literal of C_j are assigned to K_j^1, K_j^2, K_j^3. This can be interpreted in $\mathcal{I}$ as the requirement for at least one literal of C_j to be true.

The reduction is almost complete but it remains to describe gadgets collecting all unassigned objects. Indeed, so far we have introduced $4v+3s$ agents and $6v+2s$ objects. It remains to construct garbage collectors for the $2v-s$ remaining objects. Note that no dummy object (neither variable nor clause) may be part of the remaining objects since they must be assigned to literal-agents or clause-agents in any LEF allocation. Let $\mathcal{L} = \{u_i^j, \overline{u}_i^j : i \in [v],\ j \in [2]\}$ denote the set of literal-objects, where literal-objects are ordered arbitrarily, and let $\mathcal{L}(\ell)$ denote the ℓ^{th} element of $\mathcal{L}$.

Let us now describe a gadget collecting a single object of $\mathcal{L}$. For each $\ell \in [4v]$, we introduce dummy-objects t_ℓ^1 and t_ℓ^2 and a cycle in $\overline{G}$ containing gadget-agents L_ℓ^1, L_ℓ^2 and L_ℓ^3. Furthermore, for each $\ell \in [4v-1]$, we introduce gadget-object g_ℓ. Globally, in this gadget, we introduce $12v$ new agents and $12v-1$ new objects. Preferences are as follows (we assume that the notations g_0 and g_{4v} actually refer to the objects g_1 and g_{4v-1}, respectively).

$$\begin{aligned} L_\ell^1: &\quad t_\ell^1 \succ t_\ell^2 \succ g_{\ell-1} \succ [\ldots] \\ L_\ell^2: &\quad t_\ell^1 \succ t_\ell^2 \succ \mathcal{L}(\ell) \succ [\ldots] \\ L_\ell^3: &\quad t_\ell^1 \succ t_\ell^2 \succ g_\ell \succ [\ldots] \end{aligned}$$

Note that in any LEF allocation, objects t_ℓ^1 and t_ℓ^2 are allocated to agents belonging to $\{L_\ell^1, L_\ell^2, L_\ell^3\}$, and the remaining unassigned agent receives either $g_{\ell-1}, g_\ell$ or $\mathcal{L}(\ell)$. Since no more than $4v-1$ agents can receive a gadget-object, at least one literal-object is assigned to agent L_ℓ^2 for some $\ell \in [4v]$. Moreover, all gadget-objects must be assigned to gadget-agents since no other agent has a gadget-object in her top three objects. Therefore, in every LEF allocation, exactly one literal-object is allocated to an agent belonging to the gadget.

We use exactly $2v-s$ copies of this gadget in order to collect all the remaining literal-objects of the first part of the construction, and thus obtaining as many agents as objects in the whole reduction.

Now let us show that one can allocate objects without envy in the gadget. Let $\mathcal{L}(\ell)$ be the literal-object assigned in the gadget. This object must be assigned to L_ℓ^2. Assign t_ℓ^1 and t_ℓ^2 to agents L_ℓ^1 and L_ℓ^3, respectively. For any $\ell' \neq \ell$, assign $t_{\ell'}^1$ to agent $L_{\ell'}^2$. For any $\ell' > \ell$, object $g_{\ell'-1}$ is assigned to $L_{\ell'}^1$ and object $t_{\ell'}^2$ is assigned to $L_{\ell'}^3$. Finally, for any $\ell' < \ell$, object $g_{\ell'}$ is assigned to $L_{\ell'}^3$ and object $t_{\ell'}^2$ is assigned to $L_{\ell'}^1$.

We claim that $\mathcal{C}$ is satisfiable in instance $\mathcal{I}$ if and only if $\mathcal{J}$ has an LEF allocation.

Suppose first that there exists a truth assignment ϕ of the variables in X such that all the clauses in $\mathcal{C}$ are satisfiable. For each variable x_i, if x_i is true (respectively, false) in ϕ then we assign the objects q_i^1 and q_i^2 to the agents X_i^1 and X_i^2 (respectively, $\overline{X}_i^1$ and $\overline{X}_i^2$), and the literal-objects associated with the negative (respectively, positive) occurrence, namely $\overline{u}_i^1$ and $\overline{u}_i^2$ (respectively, u_i^1 and u_i^2), to the agents $\overline{X}_i^1$ and $\overline{X}_i^2$ (respectively, X_i^1 and X_i^2). This set of agents envies nobody because for each agent in the cycle, she either obtains her most preferred object or all the objects preferred to her assigned object are owned by neighbors in the non-envy graph. Now, only literal-objects associated with literals true in ϕ are available. Since all the clauses are satisfiable by ϕ, there exists at least one available literal-object per clause and we assign it to one of the clause-agents K_j^h for each clause C_j. The other clause-agents receive the dummy-objects d_j^1 and d_j^2. By construction of the preferences, no clause-agent can be envious. Finally, it suffices to assign the remaining literal-objects to garbage-agents, as previously described in the construction of the gadgets, in such a way that no garbage-agent can be envious. Therefore, we finally obtain an LEF allocation.

Suppose now that there exists an LEF allocation. As previously remarked, in any LEF allocation, the objects q_i^1 and q_i^2 must be assigned either to agents X_i^1 and X_i^2, or to $\overline{X}_i^1$ and $\overline{X}_i^2$. In the first case, the negative literal-objects associated with x_i must be assigned to agents $\overline{X}_i^1$ and $\overline{X}_i^2$, and thus only the positive literal-objects associated with x_i will be available for other agents. Otherwise, in the second case, the situation is reversed, and thus only the negative literal-objects associated with x_i will be available for other agents. Moreover, we know that in any LEF allocation, exactly one literal-object associated with clause C_j must be assigned to an agent K_j^h. Consequently, to be LEF, the allocation must have one available literal-object per clause and all these objects cannot correspond to opposite literals. Hence, if we set to true every variable x_i such that q_i^1 and q_i^2 are possessed by X_i^1 and X_i^2 in the allocation, and to false every variable x_i such that q_i^1 and q_i^2 are possessed by $\overline{X}_i^1$ and $\overline{X}_i^2$, then we obtain a truth assignment of the variables such that all the clauses are satisfiable. □

As for the generalization of Theorem 4.1 to Proposition 4.3, we can extend Theorem 4.5 to more general classes of graphs. It suffices to add, in the graph of the previous proof, dummy agents who are connected to three other agents. They have a dummy resource on top of their ranking, followed by the whole ranking of one of her neighbors. Each agent of the original instance ranks last the dummy resources.

Corollary 4.6 Dec-LEF *is* NP*-complete on undirected graphs of minimum degree* $n-k$ *for* $k \geq 3$ *constant.*

Related to the question of the degree of the nodes, it appears interesting to determine how the computational hardness of Dec-LEF evolves on cluster graphs (see Section 1.2.2.a). The cluster graphs are relevant in the context of a social network, because they may represent several groups of agents that do not have interconnections. One can think for instance to group of families or different sport teams. In fact, the problem is computationally hard when the cluster graph is composed of $n/2$ cliques because this is the case of the matching (Theorem 4.1). This hardness is extended to any cluster graph composed of n/k cliques (for $k \geq 2$ constant) according to the construction in the proof of Proposition 4.3. Note that the case of n clusters is trivial since it is the empty graph. Moreover, the problem is solvable in polynomial time when there is only one clique in the cluster graph (the easy case of the complete graph). A natural question is then the complexity of Dec-LEF when the cluster graph is only composed of two cliques. The next theorem shows that even in this case, the problem is NP-complete.

Theorem 4.7 DEC-LEF *is* NP*-complete even when the social network is restricted to two cliques of equal size.*

Proof: The reduction is from (3,B2)-SAT, known to be NP-complete (Theorem 1.4), which is the restriction of 3-SAT to instances where each literal appears exactly twice in the clauses. We consider an instance $\mathcal{I} = \langle \mathcal{C}, X \rangle$ of (3,B2)-SAT where the set of clauses is $\mathcal{C} = \{C_1, \dots, C_s\}$ and the set of variables is $X = \{x_1, \dots, x_v\}$. Let us denote by s_j and $\overline{s}_j$ the number of positive and negative literals of clause C_j, respectively.

We construct an instance $\mathcal{J}$ of DEC-LEF with the following objects in O:

- variable-objects y_i^ℓ and $\overline{y}_i^\ell$, for $i \in [v]$ and $\ell \in \{1, 2\}$, respectively associated with the ℓ^{th} occurrence of x_i and $\overline{x}_i$,
- literal-objects u_{jh} and $\overline{u}_{jh'}$, for $j \in [s]$, $h \in [s_j]$ and $h' \in [\overline{s}_j]$, respectively associated with the h^{th} positive and the h'^{th} negative literal of clause C_j,
- clause-object q_j for $j \in [s]$,
- object z_j for every clause C_j containing at least one negated and one unnegated literal.

The social network G contains two cliques of equal size that we denote by K and $\overline{K}$. The set of agents is composed of the following ordered pairs of agents where the first agent belongs to K and the second to $\overline{K}$: variable-agents $(Y_i^\ell, \overline{Y}_i^\ell)$, for $i \in [v]$ and $\ell \in \{1, 2\}$, clause-agents $(Q_j, \overline{Q}_j)$ for $j \in [s]$. The preferences of the agents are given as follows ("rest$_A$" refers to an arbitrary order over the remaining objects that is the same for all the agents associated with a given A, and $u(x_i^\ell)$, (respectively, $\overline{u}(\overline{x}_i^\ell)$), stands for the literal-object corresponding to the ℓ^{th} occurrence of x_i, (respectively, $\overline{x}_i$)):

K

$$\begin{aligned}
Y_i^1 :&\quad y_i^1 \succ y_i^2 \succ u(x_i^1) \succ u(x_i^2) \succ \overline{y}_i^1 \succ \overline{y}_i^2 \succ \text{rest}_{Y_i} \\
Y_i^2 :&\quad y_i^2 \succ y_i^1 \succ u(x_i^1) \succ u(x_i^2) \succ \overline{y}_i^2 \succ \overline{y}_i^1 \succ \text{rest}_{Y_i} \\
Q_j :&\quad u_{j1} \succ \dots \succ u_{js_j} \succ q_j \succ u_{j+1\ 1} \succ \dots \succ u_{j+1\ s_{j+1}} \succ q_{j+1} \succ u_{j+2\ 1} \succ \dots \succ u_{j+2\ s_{j+2}} \succ \\
&\quad q_{j+2} \succ \dots \succ u_{j-1\ 1} \succ \dots \succ u_{j-1\ s_{j-1}} \succ q_{j-1} \succ \text{rest}_Q
\end{aligned}$$

$\overline{K}$

$$\begin{aligned}
\overline{Y}_i^1 :&\quad y_i^1 \succ y_i^2 \succ \overline{u}(\overline{x}_i^1) \succ \overline{u}(\overline{x}_i^2) \succ \overline{y}_i^1 \succ \overline{y}_i^2 \succ \text{rest}_{\overline{Y}_i} \\
\overline{Y}_i^2 :&\quad y_i^2 \succ y_i^1 \succ \overline{u}(\overline{x}_i^1) \succ \overline{u}(\overline{x}_i^2) \succ \overline{y}_i^2 \succ \overline{y}_i^1 \succ \text{rest}_{\overline{Y}_i} \\
\overline{Q}_j :&\quad \overline{u}_{j1} \succ \dots \succ \overline{u}_{j\overline{s}_j} \succ q_j \succ \overline{u}_{j+1\ 1} \succ \dots \succ \overline{u}_{j+1\ \overline{s}_{j+1}} \succ q_{j+1} \succ \overline{u}_{j+2\ 1} \succ \dots \succ \overline{u}_{j+2\ \overline{s}_{j+2}} \succ \\
&\quad q_{j+2} \succ \dots \succ \overline{u}_{j-1\ 1} \succ \dots \succ \overline{u}_{j-1\ \overline{s}_{j-1}} \succ q_{j-1} \succ \text{rest}_{\overline{Q}}
\end{aligned}$$

Garbage-agents are added for each clause C_j according to the number of positive and negative literals in C_j ("$\succ_A \backslash O'$" in the preferences stands for the preference ranking of agent A without considering the objects in O'):

1. if $s_j = 0$, agents Z_j^1 and Z_j^2 belonging to K, with the preferences:

K

$$\begin{aligned}
Z_j^1 :&\quad \overline{u}_{j1} \succ \overline{u}_{j2} \succ \overline{u}_{j3} \succ \text{rest}_j \\
Z_j^2 :&\quad \overline{u}_{j3} \succ \overline{u}_{j2} \succ \overline{u}_{j1} \succ \text{rest}_j
\end{aligned}$$

2. if $s_j = 1$, agents Z_j^1 and Z_j^2 belonging to K and agent $\overline{Z}_j$ belonging to $\overline{K}$, with the preferences:

$$\begin{array}{llll} & K & & \overline{K} \\ Z_j^1: & \overline{u}_{j1} \succ z_j \succ \overline{u}_{j2} \succ \text{rest}_j & \overline{Z}_j: & u_{j1} \succ z_j \succ [\succ_{\overline{Q}_j} \setminus \{u_{j1}, z_j\} \\ Z_j^2: & \overline{u}_{j2} \succ z_j \succ \overline{u}_{j1} \succ \text{rest}_j & & \end{array}$$

3. if $s_j = 2$, agent Z_j belonging to K and agents $\overline{Z}_j^1$ and $\overline{Z}_j^2$ belonging to $\overline{K}$, with the preferences:

$$\begin{array}{llll} & K & & \overline{K} \\ Z_j: & \overline{u}_{j1} \succ z_j \succ [\succ_{Q_j} \setminus \{\overline{u}_{j1}, z_j\} & \overline{Z}_j^1: & u_{j1} \succ z_j \succ u_{j2} \succ \text{rest}_j \\ & & \overline{Z}_j^2: & u_{j2} \succ z_j \succ u_{j1} \succ \text{rest}_j \end{array}$$

4. if $s_j = 3$, agents $\overline{Z}_j^1$ and $\overline{Z}_j^2$ belonging to $\overline{K}$, with the preferences:

$$\begin{array}{ll} & \overline{K} \\ \overline{Z}_j^1: & u_{j1} \succ u_{j2} \succ u_{j3} \succ \text{rest}_j \\ \overline{Z}_j^2: & u_{j3} \succ u_{j2} \succ u_{j1} \succ \text{rest}_j \end{array}$$

If at this point, K and $\overline{K}$ have not the same size, then we add to the clique with the minimum size as many dummy agents and dummy objects as the difference between the two sizes. The dummy agents all share the same preferences as an arbitrary agent in the clique, except that her associated dummy object is her most preferred object, and all the agents of the clique rank the dummy objects in the last positions of their preference ranking. By construction, in an LEF allocation, each dummy agent should get her associated dummy resource and no other agent of the clique envies her.

We claim that there exists an LEF allocation in $\mathcal{J}$ if and only if $\mathcal{C}$ is satisfiable in $\mathcal{I}$.

Suppose that there exists a truth assignment ϕ of the variables in X such that $\mathcal{C}$ is satisfiable. For each variable x_i, if x_i is true in ϕ, then assign object y_i^ℓ to agent Y_i^ℓ and object $\overline{y}_i^\ell$ to agent $\overline{Y}_i^\ell$ for each $\ell \in \{1, 2\}$. Otherwise, assign object $\overline{y}_i^\ell$ to agent Y_i^ℓ and object y_i^ℓ to agent $\overline{Y}_i^\ell$ for each $\ell \in \{1, 2\}$. Observe that there is no envy within this set of agents because agents Y_i^1 and Y_i^2 are not connected to the agents $\overline{Y}_i^1$ and $\overline{Y}_i^2$ in G. Now, let us choose exactly one literal that is true for each clause C_j. If this literal is positive, then assign the associated positive literal-object, say u_{jh}, to agent Q_j and object q_j to agent $\overline{Q}_j$. Otherwise, assign the associated negative literal-object, say $\overline{u}_{jh'}$, to agent $\overline{Q}_j$ and object q_j to agent Q_j. Clearly, there is no envy between Q_j and $\overline{Q}_j$. Moreover, there is no envy between the variable-agents and the clause-agents of the same clique. More precisely, the only possible case of envy would be when a literal-object associated with literal x_i is chosen to be allocated to an agent Q_j. But, this means that x_i is present in clause C_j and set to true in ϕ, and thus by construction we have assigned to Y_i^1 and Y_i^2 their best object, so they cannot envy Q_j (the same reasoning applies in clique $\overline{K}$). At this point, it remains to assign two literal-objects associated with each clause C_j. If these two objects both correspond to positive (respectively, negative) literals, then assign them to the agents $\overline{Z}_j^1$ and $\overline{Z}_j^2$ (respectively, Z_j^1 and Z_j^2) in such a way that they do not envy each other (this is always possible by construction of the preferences), and if $s_j \geq 1$ and $\overline{s}_j \geq 1$, also assign z_j to agent Z_j (respectively, $\overline{Z}_j$). Otherwise, if it remains one positive and one negative literal-objects, and $s_j < \overline{s}_j$ (respectively, $s_j > \overline{s}_j$), then assign the positive one to agent $\overline{Z}_j$ (or to $\overline{Z}_j^\ell$ if this literal is the ℓ^{th} positive literal of the clause) and assign the negative one to agent Z_j (or to Z_j^ℓ if this literal is the ℓ^{th}

negative literal of the clause). By construction, there is no envy among the garbage-agents. Moreover, there is no envy between the garbage-agents and the clause-agents because if they are in the same clique, then they own literal-objects associated with opposite literals. Therefore, to summarize, the constructed allocation is LEF.

Suppose now that there exists an LEF allocation π. Then, this allocation must assign literal-objects or garbage-objects z_j to the garbage-agents. Moreover, by construction of the preferences, π must assign to each clause-agent Q_j (respectively, $\overline{Q}_j$) either a literal-object associated with a positive (respectively, negative) literal of clause C_j, or q_j. They are the only agents to accept objects q_j in an LEF allocation. Therefore, each q_j must be assigned either to agent Q_j or to agent $\overline{Q}_j$. Consequently, the agent Q_j or $\overline{Q}_j$ who does not obtain q_j must get a literal-object associated with respectively a positive or a negative literal of clause C_j. This simulates the fact that each clause must be satisfied. Concerning the variable-agents Y_j^ℓ and $\overline{Y_j^\ell}$, they must own in any LEF allocation either variable-object y_j^ℓ or variable-object $\overline{y}_j^\ell$. Moreover, by construction, there will be envy if Y_i^1 gets $\overline{y}_i^1$ and Y_i^2 gets y_i^2, or vice-versa, and the same holds for agents $\overline{Y}_i^1$ and $\overline{Y}_i^2$. Therefore, the only possibility is that the agents Y_i^1 and Y_i^2 get the objects y_i^1 and y_i^2 together or the objects $\overline{y}_i^1$ and $\overline{y}_i^2$ together. Therefore, this simulates a truth assignment of the variables. Finally, there is envy between a variable-agent Y_i^ℓ (respectively, $\overline{Y}_i^\ell$) and a clause-agent Q_j (respectively, $\overline{Q}_i$) only if Q_j (respectively, $\overline{Q}_j$) owns a literal-object associated with a positive (respectively, negative) literal x_i (respectively, $\overline{x}_i$) and Y_i^ℓ (respectively, $\overline{Y}_i^\ell$) owns the variable-object $\overline{y}_i^\ell$. Thus, to be LEF, an allocation must assign object y_i^ℓ to agent Y_i^ℓ (respectively, $\overline{Y}_i^\ell$) if a positive literal associated with x_i (respectively, $\overline{x}_i$) is owned by an agent Q_j. Hence, globally, if we consider a truth assignment of the variables in X that sets to true every variable x_i such that Y_i^ℓ gets y_i^ℓ in π and sets to false every variable x_i such that Y_i^ℓ gets $\overline{y}_i^\ell$ in π, then all the clauses in $\mathcal{C}$ are satisfiable. □

By adding clusters of dummy agents having their associated dummy resource on top of their preference ranking, we can generalize the previous negative result to any cluster graph with $k \geq 2$ (k constant) clusters.

Corollary 4.8 DEC-LEF *is* NP*-complete in any cluster graph with* $k \geq 2$ *clusters or* n/k *(*$k \geq 2$*) clusters for* k *constant.*

It follows that the easy case of the complete graph for deciding whether there exists a locally envy-free allocation cannot be generalized to a graph composed of disjoint cliques.

4.3.2 Local envy-freeness and vertex cover

So far the complexity of DEC-LEF has been investigated through the degree of its nodes, but other parameters can be taken into account. Let us show how the size of a (smallest) *vertex cover* can help. For the convenience of the definition of a vertex cover, more adapted to an undirected context, we restrict in this subsection to the case where social network G is an undirected graph.

Let us first state that DEC-LEF is solvable in polynomial time if G admits a vertex cover of constant size.

Theorem 4.9 *If the social network G admits a vertex cover of size k, then* Dec-LEF *can be answered in $\mathcal{O}(n^{2k+3})$. In other words,* Dec-LEF *is in* XP *with respect to the parameter k equal to the size of the smallest vertex cover of G.*

Proof: Let us find a vertex cover C of the social network. See Kleinberg and Tardos [2006] for a $\mathcal{O}(2^k n)$ algorithm which decides and builds a vertex cover of size k in a graph with n vertices. Then, use brute force to assign k objects of O to C (the time complexity is $\mathcal{O}(n^{2k})$). For each partial allocation π without envy within C, let $O_{-\pi}$ be the set of unassigned objects (if no such partial allocation exists, then we can immediately conclude that no LEF allocation exists).

Now, let us consider the rest of the agents, that is the set $I := N \setminus C$. By definition, I is an independent set of graph G. Build a bipartite undirected graph $G' := (I \cup O_{-\pi}, E')$ with an edge connecting agent $i \in I$ to object $o \in O_{-\pi}$ if assigning o to i does not create envy. Actually, since I is an independent set, there is no envy among the agents of I. Therefore, an edge connecting $i \in I$ to $o \in O$ simply means that agent i prefers o to the objects currently assigned to her neighbors in C. Hence, there is an LEF allocation for the entire network if and only if bipartite graph G' admits a perfect matching (which can be verified in $\mathcal{O}(n^3)$). □

Consequently, the method described in the previous proof is efficient when k is small. For instance, Dec-LEF is polynomial if the social network is a star because the central node of a star is a vertex cover.

One could expect that Dec-LEF also belongs to FPT for the same parameter since the problem of finding a vertex cover of size k is FPT [Downey and Fellows, 1995b]. However, the following theorem shows that there is no hope that Dec-LEF can be fixed-parameter tractable, with respect to the size of a vertex cover, under standard complexity assumptions.

Theorem 4.10 Dec-LEF *parameterized by the size of a vertex cover is* W[1]*-hard.*

Proof: We present a parameterized reduction from Multicolored Independent Set, known to be W[1]-complete (Theorem 1.7). An instance $\mathcal{I}$ of Multicolored Independent Set consists of an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, an integer k, and a partition $(\mathcal{V}_1, \ldots, \mathcal{V}_k)$ of $\mathcal{V}$. The task is to decide whether there is an independent set of size k in $\mathcal{G}$ containing exactly one vertex from each set $\mathcal{V}_i$.

We construct an instance $\mathcal{J}$ of Dec-LEF as follows. For each vertex v in $\mathcal{V}$, we introduce object o_v. Let O_i denote the set of objects $\{o_v : v \in \mathcal{V}_i\}$, and let $O_i^{\uparrow}$ denote an arbitrary order over the objects of O_i. For each edge $e = \{v, v'\}$ in $\mathcal{E}$, we introduce two agents X_e^v and $X_e^{v'}$, and two edge-objects o_e and o'_e. Let $O_{\mathcal{E}}$ denote the set of edge-objects, and let $O_{\mathcal{E}}^{\uparrow}$ denote an arbitrary ranking over the objects of $O_{\mathcal{E}}$. For each integer $1 \leq i \leq k$, we introduce agent K_i. The agents of $\{K_i\}_{1 \leq i \leq k}$ form a clique in the social network G. Furthermore, for each vertex $v \in \mathcal{V}_i$ and for each edge $e = \{v, v'\}$ in $\mathcal{E}$, agent X_e^v is connected to agent K_i in G. Finally, for each integer $1 \leq j \leq |\mathcal{V}| - k$, we introduce agent D_j. The preferences are the following:

$$
\begin{aligned}
K_i: &\quad O_i^{\uparrow} \succ O_1^{\uparrow} \succ \ldots \succ O_{i-1}^{\uparrow} \succ O_{i-1}^{\uparrow} \succ \ldots O_k^{\uparrow} \succ O_{\mathcal{E}}^{\uparrow} \\
X_e^v: &\quad o_e \succ o_v \succ o'_e \succ [\ldots]
\end{aligned}
$$

Since agent D_j is isolated in G, her preferences may be arbitrary. It is easy to check that $\{K_i\}_{1\leq i\leq k}$ forms a vertex cover in network G.

We show that $\mathcal{G}$ has an independent set of size k containing one vertex in each set $\mathcal{V}_i$ in $\mathcal{I}$ if and only if an LEF allocation exists in $\mathcal{J}$.

Assume first that $\{v_1, \ldots, v_k\}$ is an independent set in $\mathcal{G}$, where $v_i \in \mathcal{V}_i$ for each $i \in [k]$. We construct an LEF allocation as follows. For each $i \leq k$, assign o_{v_i} to K_i, and for each edge $e = (v_i, v')$ in $\mathcal{E}$, assign o_e to $X_e^{v_i}$. For each agent X_e^v such that v is not selected in the independent set, assign o_e to X_e^v if it is still available, and otherwise assign o'_e to X_e^v. Finally, assign the remaining objects arbitrarily. We claim that this allocation is locally envy-free. Indeed, each agent K_i receives an object of O_i. Furthermore, for each vertex v in $\mathcal{V}_i$ and for each edge e in $\mathcal{E}$, agent X_e^v has a single neighbor who is K_i. If K_i receives o_{v_i} and $v = v_i$ then $X_e^{v_i}$ receives o_e, and otherwise X_e^v receives o'_e.

Assume now that an LEF allocation exists. We claim that each agent K_i should receive an object of O_i. By contradiction, assume that agent K_i receives object $o \notin O_i$. Note that for any $j \neq i$, K_i and K_j are neighbors. Hence, for any object o', if $o \notin O_j$ and $o' \notin O_i \cup O_j$ then $o \succ_{V_i} o'$ if and only if $o' \succ_{V_j} o$ holds. This implies that if $o \notin O_j$ then an object of O_j must be assigned to K_j to avoid envy between agents K_i and K_j. Therefore, if $o \in O_{\mathcal{E}}$ then agent K_i envies agent K_j, a contradiction. On the other hand, if $o \in O_j$ for some $j \neq i$ then either K_i envies K_j or K_j envies K_i, since $o \succ_{V_i} o'$ if and only if $o \succ_{V_j} o'$ holds because $o' \in O_j$, a contradiction. Let o_{v_i} denote the object assigned to K_i. We claim that $\{v_1, \ldots, v_k\}$ forms an independent set in $\mathcal{G}$. By contradiction assume that edge e connects v_i and v_j in $\mathcal{G}$. This implies by construction that $X_e^{v_i}$ and $X_e^{v_j}$ are neighbors of K_i and K_j in G, respectively. On one hand, if $X_e^{v_i}$ does not receive o_e then she envies K_i. On the other hand, if $X_e^{v_j}$ does not receive o_e then she envies K_j. Therefore, o_e must be assigned to both $X_e^{v_i}$ and $X_e^{v_j}$, leading to a contradiction since o_e cannot be assigned twice. □

Let us conclude the section dedicated to DEC-LEF with a table summarizing our results. Our findings for DEC-LEF with respect to the out-degree of the nodes and the size of the vertex cover (for the case where the network is an undirected graph) are summarized in Table 4.1.

out-degree δ_G^+	$\delta_G^+ \leq k$ ($k \geq 1$ fixed)	NP-c	Th. 4.1
	$\delta_G^+ \geq n - k$ ($k \geq 3$ fixed)	NP-c	Th. 4.5
	$\delta_G^+ \geq n - 2$	P	Th. 4.4
number of clusters c in a cluster graph	$c = n/k$ ($k \geq 2$ fixed)	NP-c	Th. 4.1
	$c = k$ ($k \geq 2$ fixed)	NP-c	Th. 4.7
	$c = 1$ or $c = n$	P	
parameter k on the vertex cover size in G undirected		XP	Th. 4.9
		W[1]-hard	Th. 4.10

Table 4.1: Complexity results of DEC-LEF

4.4 Maximization of the local non-envy

In light of Section 4.3, we know that both MAX-LEF and MAX-NE, the two optimization problems that we consider for minimizing local envy, are computationally hard even on very simple graph structures. We present in this section approximation algorithms for MAX-LEF and MAX-NE.

4.4.1 Maximizing the number of non-envious agents

This subsection is dedicated to MAX-LEF, whose goal is to find an allocation maximizing the number of LEF agents (see Section 4.2). A general method is proposed in Algorithm 4.1 for the case of undirected graphs. For a maximization problem, an algorithm is ρ-approximate, with $\rho \in [0, 1]$, if it outputs a solution whose value is at least ρ-times the optimal value, for any instance (see Section 1.5.1.c).

Algorithm 4.1: Maximization of the number of LEF agents when G is undirected

Input: Instance $\langle N, O, \succ, G\rangle$
Output: Allocation π

1 $I \leftarrow$ independent set of the network; `# found in any opportune way`
2 **foreach** $i \in I$ **do**
3     $\pi(i) \leftarrow i$'s most preferred object within O;
4     $O \leftarrow O \setminus \{\pi(i)\}$;
5 Complete π for agents in $N \setminus I$; `# in any opportune way`
6 **return** π

The strategy used in Algorithm 4.1 is to find an independent set $I \subseteq V$ in the network, and then to assign by turns to each agent in I her best object in O among the remaining items. This algorithm simulates a serial dictatorship mechanism on I (see Section 1.4.2.a). Then, the allocation is completed arbitrarily.

By construction, every member of I is LEF, and a direct upper bound on the optimal number of LEF agents is $|N| = n$, leading to the following observation.

Observation 4.3 *Algorithm 4.1 is $\frac{|I|}{n}$-approximate for* MAX-LEF.

Let us specify a way to construct independent set I (Step 1 of Algorithm 4.1) in order to get an explicit approximation ratio.

Set I is initially empty. While the set of agents N is non-empty, an agent i is selected from N and added to I. In the same time, agent i and her neighbors are deleted from N.

Since a node has at most Δ_G neighbors (for Δ_G being the degree of G, see Section 1.2.2.a), I is an independent set of size at least $n/(\Delta_G + 1)$. Using Observation 4.3, we get a ratio of $(\Delta_G + 1)^{-1}$ for Algorithm 4.1.

Observation 4.4 *The construction of I in Algorithm 4.1 (Step 1) can be done so that a polynomial time $(\Delta_G + 1)^{-1}$-approximation is produced.*

The $(\Delta_G + 1)^{-1}$-approximation algorithm is long known for the MAXIMUM INDEPENDENT SET problem (that is, find an independent set of maximum cardinality), see for example Paschos [1992]. The following lemma shows that MAX-LEF shares exactly the same inapproximability results as MAXIMUM INDEPENDENT SET.

Lemma 4.11 *Any ρ-approximate algorithm for* MAX-LEF *is also a ρ-approximate algorithm for* MAXIMUM INDEPENDENT SET.

Proof: Suppose we have a ρ-approximation algorithm for MAX-LEF. Let us consider an instance of MAX-LEF where all the agents have identical preferences, i.e., $\succ_1 = \succ_2 = \ldots \succ_n$. Our ρ-approximation algorithm computes for this instance an allocation π and a set J of non-envious agents. Because preferences are identical, a pair of connected agents cannot be locally envy-free, whatever the allocation is. In this setting, the set J is thus necessarily an independent set of G. Hence, our algorithm is also a ρ-approximation algorithm for MAXIMUM INDEPENDENT SET. □

MAXIMUM INDEPENDENT SET in general is Poly-APX-complete [Bazgan et al., 2004], meaning that it belongs to the hardest problems that can be approximated within a polynomial factor. Lemma 4.11 implies that MAX-LEF is Poly-APX-hard. Thus, Algorithm 4.1 is asymptotically optimal.

Interestingly, there are graph classes where the size of an independent set can be expressed as a fraction of n. Therefore, this fraction corresponds to the approximation ratio of Algorithm 4.1.

Proposition 4.12 *A polynomial time* 0.5*-approximate algorithm for* MAX-LEF *exists if social network G is an undirected bipartite graph.*

Proof: Suppose the social network is a bipartite graph $G = (N_1 \cup N_2, E)$. By definition of a bipartite graph, both N_1 and N_2 are independent sets. If $|N_1| \geq |N_2|$, then run Algorithm 4.1 with $I := N_1$, otherwise run Algorithm 4.1 with $I := N_2$. Since $2|I| \geq |N_1| + |N_2| = |N| = n$, a polynomial time 0.5-approximation is reached. □

Proposition 4.12 can be easily extended to undirected *k-partite graphs* (whose vertex set can be partitioned into k different independent sets), leading to a polynomial time k^{-1}-approximation algorithm.

Note that if the social network admits a vertex cover V' of size k, then Algorithm 4.1 with $I := N \setminus V'$ provides a $(1 - k/n)$-approximate solution to MAX-LEF.

4.4.2 Optimizing the degree of (non)-envy

Instead of simply counting the number of non-envious agents, we will now focus on a more subtle criterion, measuring the degree of envy among agents. This leads to the MAX-NE optimization problem (defined in Section 4.2) which consists in minimizing the *average degree of envy* $\mathscr{E}(\pi)$ (or equivalently maximizing the *average degree of non-envy* $\mathcal{N}\mathscr{E}(\pi) = 1 - \mathscr{E}(\pi)$). In this subsection, there is no need to restrict anymore to undirected graphs and thus we consider G as a directed graph.

Before describing our approximation algorithm for MAX-NE, we first state the following lemma.

Lemma 4.13 *Let $\mathcal{U}_n$ be the uniform distribution over all allocations of n objects to n agents (with exactly one object per agent). Then we have $\mathbb{E}_{\pi\sim\mathcal{U}_n}[\mathcal{N}\mathscr{E}(\pi)] = \frac{5}{6} - o(1)$.*

Proof: In the following, the notations $x, x' \sim O$ means x and x' are two random objects drawn uniformly and independently from O. Also, $v \sim [n]$ means v is a random integer drawn uniformly from $[n]$. By linearity of expectation, we get:

$$\begin{aligned}
\mathbb{E}_{\pi\sim\mathcal{U}_n}[\mathscr{E}(\pi)] &= \frac{1}{|E|}\sum_{(i,j)\in E}\mathbb{E}_{\pi\sim\mathcal{U}_n}[e(\pi,i,j)] \\
&= \frac{1}{|E|\,(n-1)}\sum_{(i,j)\in E}\mathbb{E}_{x,x'\sim O}\left[\max\left\{0, \mathscr{v}_i(x) - \mathscr{v}_i\left(x'\right)\right\} : x \neq x'\right] \\
&= \frac{1}{|E|\,(n-1)}\sum_{(i,j)\in E}\mathbb{E}_{v,v'\sim[n]}\left[\max\left\{0, v - v'\right\} : v \neq v'\right]
\end{aligned}$$

By the law of total expectation, we have:

$$\begin{aligned}
\mathbb{E}_{v,v'\sim[n]}\left[\max\left\{0, v - v'\right\} : v \neq v'\right] &= \mathbb{E}_{v,v'\sim[n]}\left[v - v' : v > v'\right]\cdot\mathbb{P}\left[v > v' : v \neq v'\right] \\
&\quad +0\cdot\mathbb{P}\left[v < v' : v \neq v'\right] \\
&= \frac{1}{2}\mathbb{E}_{v,v'\sim[n]}\left[v - v' : v > v'\right] \\
&= \frac{1}{n(n-1)}\sum_{k\in[n-1]}k(n-k) \\
&= \frac{n+1}{6}
\end{aligned}$$

So $\mathbb{E}_{\pi\sim\mathcal{U}_n}[\mathscr{E}(\pi)] = \frac{n+1}{6(n-1)}$ and $\mathbb{E}_{\pi\sim\mathcal{U}_n}[\mathcal{N}\mathscr{E}(\pi)] = \frac{5n-7}{6(n-1)} = \frac{5}{6} - \frac{2}{6(n-1)} = \frac{5}{6} - o(1)$. □

This tells us that with high probability, random allocations of objects yield high degrees of non-envy. To get a deterministic algorithm based on this idea, we apply a standard derandomization technique. In Algorithm 4.2, at each step i, agent i receives one of the remaining unallocated objects. This object is chosen so as to minimize the conditional expectation of $\mathscr{E}$ (line 7): given the partial current allocation π, we choose among all possible remaining objects in O', the object x^* which, once assigned to i, minimizes the expected average degree of envy $\mathscr{E}$. We will show below that this conditional expectation can be computed efficiently.

Proposition 4.14 *Algorithm 4.2 is a polynomial-time $\left(\frac{5}{6} - o(1)\right)$-approximation algorithm for* Max-NE.

Proof: First, by standard arguments of the derandomization method (similar to e.g., page 132 of Vazirani [2001]) together with Lemma 4.13, this algorithm outputs an allocation π such that $\mathcal{N}\mathscr{E}(\pi) \geq \frac{5}{6} - o(1)$. By design we have $\mathcal{N}\mathscr{E}(\pi) \leq 1$, so the approximation ratio holds. To show that the algorithm runs in polynomial time, we need to bound the computation time of v_x (line 7) for a given agent i, a given object $x \in O'$, and the current allocation π^x within the two nested loops of the algorithm where

Algorithm 4.2: Minimization of the average degree of envy

$\pi \leftarrow \{\}$; # empty allocation
$O' \leftarrow O$; # set of available objects
$S \leftarrow \emptyset$; # set of agents already assigned in π
foreach *agent* $i \in [n]$ **do**
    foreach *object* $x \in O'$ **do**
        $\pi^x \leftarrow \pi;\ \pi^x(i) \leftarrow x$;
        $v_x \leftarrow \mathbb{E}_{\pi' \sim \mathcal{U}_n}\left[\mathscr{E}(\pi') : \pi'_{|S\cup\{i\}} = \pi^x_{|S\cup\{i\}}\right]$;
    $x^* \leftarrow \arg\min_{x \in O'} v_x$;
    $\pi(i) \leftarrow x^*$;
    $O' \leftarrow O' \setminus \{x^*\}$;
    $S \leftarrow S \cup \{i\}$;
return π

$\pi^x(i) = x$. If π is a partial allocation, define $P(\pi, \ell)$ as the set of goods that agent ℓ can own without violating π. For example, if π is a complete allocation, $P(\pi, \ell) = \pi(\ell)$ and if $\pi = \{\}$, then $P(\pi, \ell) = O$. First note that v_x can be expressed as the following sum of conditional expectations: $\frac{1}{|E|}\sum_{(\ell,h)\in E} \mathbb{E}_{\pi' \sim \mathcal{U}_n}\left[e(\pi', \ell, h) : \pi'_{|S\cup\{i\}} = \pi^x_{|S\cup\{i\}}\right]$. Next, note that for any $\ell, h \in N$ the expectation $\mathbb{E}_{\pi' \sim \mathcal{U}_n}[e(\pi', \ell, h) : \pi'_{|S\cup\{i\}} = \pi^x_{|S\cup\{i\}}]$ is equal to $\frac{1}{|Z_{\ell,h}|.(n-1)}\sum_{(a,b)\in Z_{\ell,h}} \max\{0, \boldsymbol{r}_\ell(a) - \boldsymbol{r}_\ell(b)\}$ where $Z_{\ell,h} = \{(a, b) \in P(\pi^x, \ell) \times P(\pi^x, h) : a \neq b\}$. The computation of v_x can thus be done in $O(n^4)$. □

It follows that, despite the computational hardness of Dec-LEF, it is possible to construct efficiently an allocation whose average degree of envy is guaranteed to be close to the minimum.

4.5 Location and allocation

This section is dedicated to Dec-Location-LEF, the decision problem asking whether there exists a way of assigning both the agents on the directed graph and the resources to the agents such that we obtain an LEF allocation (see Section 4.2).

In this problem, the central authority has more freedom than in the simple allocation problem Dec-LEF since she can also choose the location of the agents on the graph. Consequently, there should be more positive instances for this problem than for Dec-LEF. The following example illustrates this intuition.

Example 4.4 *Let us consider an instance with three agents, where* $N = \{1, 2, 3\}$ *and* $O = \{o_1, o_2, o_3\}$*. The social network and the preferences are as follows.*

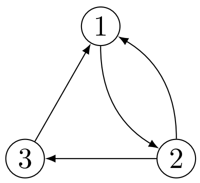

$$\begin{aligned} 1: &\quad o_1 \succ o_2 \succ o_3 \\ 2: &\quad o_1 \succ o_2 \succ o_3 \\ 3: &\quad o_3 \succ o_2 \succ o_1 \end{aligned}$$

Clearly, there is no locally envy-free allocation in this instance because agent 1 *and* 2 *can both be envious of the other and there have exactly the same preferences. So, this is a no-instance of* Dec-LEF*. However, if we only have the structure of the network, we can locate the agents as follows.*

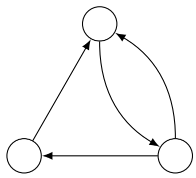

$\longrightarrow$

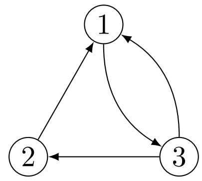

Then, it suffices to assign object o_2 *to agent* 1*,* o_1 *to agent* 2 *and* o_3 *to agent* 3 *to obtain a locally envy-free allocation, leading to a yes-instance of* Dec-Location-LEF*.*

However, since the central authority has to make more choices in this decision problem, it is difficult to imagine that Dec-Location-LEF could be easier than Dec-LEF. Indeed, the following theorem shows that this problem is computationally challenging, even for the restriction to undirected graphs.

Theorem 4.15 Dec-Location-LEF *is* NP*-complete even when the network is undirected.*

Proof: The reduction is from Independent Set, known to be NP-complete (Theorem 1.5). Given an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where $\mathcal{V} = \{v_1, \cdots, v_s\}$ and a positive integer $k \leq |\mathcal{V}|$, the problems asks whether there is an independent set $I \subseteq \mathcal{V}$ of size k.

From an instance $\mathcal{I}$ of Independent Set, let us construct an instance $\mathcal{J}$ of Dec-Location-LEF as follows. The set of objects is $O = Q \cup T$, where $Q = \{q_1, \cdots, q_{s-k}\}$ and $T = \{t_1, \cdots, t_k\}$. The set of agents is $N = \{X_1, \ldots, X_{s-k}\} \cup \{L_1, \ldots, L_k\}$. Let Q_{-i} denote the set $Q \setminus \{q_i\}$, and let $Q^{\uparrow}_{-i}$, $Q^{\uparrow}$ and $T^{\uparrow}$ denote partial orders over Q_{-i}, Q and T, respectively, where objects are ranked by increasing order of indices. Preferences are as follows for $i \in [s-k]$ and $j \in [k]$.

$$\begin{aligned} X_i: &\quad q_i \succ Q^{\uparrow}_{-i} \succ T^{\uparrow} \\ L_j: &\quad T^{\uparrow} \succ Q^{\uparrow} \end{aligned}$$

Finally, the social network is $G = \mathcal{G} = (\mathcal{V}, \mathcal{E})$.

We claim that $\mathcal{J}$ is a yes-instance of Dec-Location-LEF if and only if $\mathcal{G}$ contains an independent set of size k in $\mathcal{I}$.

Assume that I is an independent set of size k in $\mathcal{G}$. We can assume without loss of generality that $I = \{v_1, \ldots, v_k\}$. We construct an allocation π of the objects to the agents and an allocation $\mathcal{L}$ of the agents to the vertices of G as follows. If $v_i \in I$ then $\mathcal{L}(L_i) = v_i$ and $\pi(L_i) = t_i$. Otherwise, agents are placed arbitrarily on G and receive their best item, i.e, $\pi(X_i) = q_i$. It is easy to check that π is LEF with respect to $\mathcal{L}$ since no two vertices $\mathcal{L}(L_i)$ and $\mathcal{L}(L_j)$ are neighbors in G.

Assume now that there exists an allocation $\mathcal{L}$ of the agents to the vertices of G and an allocation π of the objects to the agents such that π is LEF with respect to $\mathcal{L}$. If $\mathcal{L}(L_i)$ and $\mathcal{L}(L_j)$ are neighbors in G then either $\pi(L_i) \succ_{L_j} \pi(L_j)$ or $\pi(L_j) \succ_{L_i} \pi(L_i)$ holds since L_i and L_j have the same preferences, leading to a contradiction with π LEF. Hence, $\{\mathcal{L}(L_1), \dots, \mathcal{L}(L_k)\}$ forms an independent set of size k in $G = \mathcal{G}$. □

Interestingly, the above reduction also holds when π is fixed, i.e., the allocation of objects to agents is imposed by the problem.

Nevertheless, as for Dec-LEF, we provide a polynomial result for Dec-Location-LEF in very dense graphs. However, unlike Dec-LEF, we restrict ourselves to the case of undirected graphs.

Theorem 4.16 Dec-Location-LEF *is solvable in polynomial time in undirected graphs of minimum degree $n - 2$.*

Proof: Observation 4.2 implies that two agents having the same top object must be neighbors in $\overline{G}$, otherwise one of them will be envious. Therefore, one can focus on $\mathbb{L}_{\succ}$, defined as the set of location functions such that each pair of agents having the same top object are neighbors in $\overline{G}$ (or equivalently, not neighbors in G).

If an instance contains three (or more) agents with the same top object then it must be a *no*-instance since each vertex in $\overline{G}$ has degree at most 1. The following lemma shows that the location functions of $\mathbb{L}_{\succ}$ are all equivalent for the search of an LEF allocation.

Lemma 4.17 *If π is an LEF allocation for some location function $\mathcal{L}$, and π cannot be improved by a swap between two agents without violating the LEF condition, then π is also LEF for any location function of $\mathbb{L}_{\succ}$.*

Proof: First of all, $\mathcal{L}$ must belong to $\mathbb{L}_{\succ}$ for π to be LEF. Let $\mathcal{L}'$ be a function of $\mathbb{L}_{\succ}$. It is easy to check that any pair of agents having the same top object have the same set of neighbors in G for both $\mathcal{L}$ and $\mathcal{L}'$. Therefore, if π is LEF for these agents under $\mathcal{L}$, then π is also LEF for these agents under $\mathcal{L}'$.

Let i be an agent who solely ranks some object o at the first position in her preferences. On one hand, if $\mathcal{L}(i)$ is a vertex of degree $n - 1$ then Observation 4.2 implies that she must receive o. On the other hand, if $\mathcal{L}(i)$ is a vertex of degree $n - 2$ and j is the unique neighbor of i in $\overline{G}$ then Observation 4.2 implies that o is assigned either to i or to j. But j must also be the unique agent to have some object o' ranked first in her preferences, where $o \neq o'$. Therefore, either agent i or j must receive o'. Since, by assumption, π cannot be improved by swapping objects o and o', o must be assigned to i and o' must be assigned to j. In all, agent i must receive her top object in π. Therefore, π is also LEF for agent i. □

In order to solve Dec-Location-LEF, one can compute a function $\mathcal{L}$ of $\mathbb{L}_{\succ}$ by assigning the agents having the same top object to vertices connected in $\overline{G}$, and by assigning the other agents arbitrarily. Once $\mathcal{L}$ is fixed, one can use the algorithm presented in Theorem 4.4 to compute an LEF allocation if such an allocation exists. If an

LEF allocation is returned then the algorithm returns $\mathcal{L}$ and π. Otherwise, we know by Lemma 4.17 that no function in $\mathbb{L}_{\succ}$ can lead to an LEF allocation, and the algorithm returns *false*. This algorithm clearly runs in polynomial time. □

Intuitively, Dec-Location-LEF seems harder than Dec-LEF. Actually, our results show that, in general, Dec-Location-LEF is NP-complete as well as Dec-LEF. But the tractable case that we have identified for Dec-Location-LEF is slightly weaker for than for Dec-LEF. An interesting question would be to further investigate the links between the two problems.

4.6 Reaching a locally envy-free allocation

So far, we have studied the problem related to local envy-freeness within a centralized perspective where the goal is to allocate the objects to the agents in such a way that the resulting allocation is locally envy-free. Another possibility is to decide whether the agents are able to coordinate themselves in order to reach an LEF allocation. Indeed, the intuition is that if two agents are connected in the social network and are both envious of the object owned by the other, then a win-win strategy for them is to exchange their objects.

In this section, we consider a housing market setting where the agents are embedded in a social network represented as an undirected graph, exactly as in Chapter 3 (the interaction relation must be symmetric for the cooperation of the agents via a trade). In this context, the agents are initially endowed with an object – π^0 denotes this initial allocation – and they can perform rational swaps with a neighbor. Recall that a swap between two connected agents in the social network is rational if both agents prefer the object owned by the other. An allocation is then reachable if there exists a sequence of swaps leading to this allocation. In addition to the housing market framework of Chapter 3, we assume that the social network also defines the possibility of envy.

This leads to consider a new decision problem, called Reachable LEF, relying on both swap dynamics as defined in Chapter 3 and local envy-freeness.

Reachable LEF:	
Instance:	Instance $\langle N, O, \succ, G, \pi^0 \rangle$
Question:	Does there exist a reachable allocation that is LEF?

Note that a reachable LEF allocation π must be stable, in the sense that no swap can be performed from π. Otherwise, if two neighbors can make a swap, by definition they both prefer the object assigned to the other agent and π is not LEF.

Under the condition that the allocation must be reachable, the problem of the LEF existence (Reachable LEF) becomes easier in some classes of graphs than the existence question in general (Dec-LEF). This is for instance the case of the matchings, for which Dec-LEF is NP-complete (Theorem 4.1) but Reachable LEF is solvable in polynomial time: we only need to check a unique reachable allocation that is stable, which is the allocation obtained where all the agents linked by an edge in the matching make a swap if possible. Nevertheless, we prove that in general Reachable LEF is intractable. Actually, for a given instance, if the problem of deciding the reachability of an object

(Reachable Object in Chapter 3) is hard, then there is no hope that Reachable LEF is easy to solve.

Theorem 4.18 Reachable LEF *is* NP*-complete even when the social network is a tree.*

Proof: If a sequence of swaps until a given allocation π is given, then it is easy to check whether the sequence of swaps is valid and π is LEF, proving the membership to NP.

Let us perform a reduction from the Reachable Object problem, defined in Section 3.2.2, and proved to be NP-complete in Theorem 3.1 even when the social network is a tree. Reachable Object asks whether there exists a reachable allocation assigning a specific target object x to a specific target agent A.

We consider an instance $\mathcal{I} = \langle N, O, \succ, G, \pi^0 \rangle$ of Reachable Object where the target agent is A and the target object is x. Let us construct an instance $\mathcal{J} = \langle N \cup N', O \cup O', \succ', G' = (N \cup N', E \cup E'), \pi'^0 \rangle$ of Reachable LEF where each agent $i \in N \setminus \{A\}$ is connected to a new agent $i' \in N'$ with initial endowment $o_{i'} \in O'$, and agent A is connected to a new agent A' with initial endowment $a' \in O'$. The preferences $\succ'_i$ of agent $i \in N \setminus \{A\}$ are exactly the same as in $\succ_i$, except that object $o_{i'}$ is the top object of $\succ'_i$ and all the other objects in O' are ranked at the end of $\succ'_i$. Concerning agent A, ranking $\succ'_A$ is exactly the same as in $\succ$, except that all the objects of O', including object a', are ranked at the end of $\succ'_A$. The preferences of every agent $i' \in N' \setminus \{A'\}$, initially endowed with $o_{i'}$, are such that all the objects in O are arbitrarily ranked before $o_{i'}$ in $\succ'_{i'}$ and all the objects in $O' \setminus \{o_{i'}\}$ are arbitrarily ranked after $o_{i'}$ in $\succ'_{i'}$. Concerning agent A', initially endowed with object a', the construction of $\succ'_{A'}$ is the same as for the other agents in N', except that only the objects in $O \setminus \{x\}$ are ranked before a' in $\succ'_{A'}$.

We claim that there exists a reachable LEF allocation in $\mathcal{J}$ if and only if object x is reachable for agent A in $\mathcal{I}$. Observe that if G is a tree, then G' is still a tree.

Suppose first that there exists a reachable allocation where target agent A gets target object x in instance $\mathcal{I}$. Then, this is also the case in instance $\mathcal{J}$ by the same sequence of swaps that involves only agents in N. Observe that in this allocation, there is no envy between agent A and agent A' because A' prefers her initial object a' to x and A prefers x to a'. Let all the agents i in $N \setminus \{A\}$ swap with their copy agent $i' \in N' \setminus \{A'\}$. By construction of the preferences, these swaps are rational. At this point, all the agents in $N \setminus \{A\}$ obtain their most preferred object so they do not envy anyone. The agents in $N' \setminus \{A'\}$ do not envy their only neighbor who actually gets their previous object by rationality of the swaps. Finally, agent A does not envy any agent in N because she prefers x to any object in O'. Hence, the resulting allocation is LEF.

Suppose now that there exists a reachable allocation in $\mathcal{J}$ that is LEF. First observe that no rational swap will occur between agent A and her neighbor A' because A prefers her initial object to the initial object a' of A' and agent A' cannot obtain another object since A is her only neighbor. Therefore, all along the sequence of swaps A' will stay with her initial object a'. Consequently, the only possibility for agent A' to not be envious of her neighbor A is that A gets an object in $O' \cup \{x\}$ which gathers the objects that A' prefers less than her initial object a'. However, agent A cannot obtain an object in O' via rational swaps because she prefers her initial object to any object in O', so the only solution is that agent A obtains object x. Once an agent $i \in N \setminus \{A\}$ has exchanged with her copy agent $i' \in N'$, she cannot perform further swaps because object $o_{i'}$ is her most preferred object. It follows that no agent in N' is involved in the sequence leading to get object x to agent A. Hence, in order to reach an LEF allocation in $\mathcal{J}$, a sequence of swaps within agents in N leading to give object x to agent A is needed, implying that x is reachable for agent A in $\mathcal{I}$. □

This result rules out the possibility that Reachable LEF would be solvable in polynomial time for many graphs, since it is already hard for trees. However, Reachable LEF is tractable for certain simple classes of graphs, such as the stars (see Definition 1.10), which are a special case of trees, for which Reachable Object is also easy.

Proposition 4.19 Reachable LEF *is solvable in polynomial time when the social network is a star.*

Proof: We present an algorithm that is very close to the algorithm used to solve Reachable Object in a star, in the proof of Proposition 3.2. Observe that the center agent of a star must be assigned her best object, say o_c, since she has $n-1$ neighbors. Therefore, we need to find a sequence of swaps such that the center-agent eventually gets o_c. However, to be LEF, any leaf-agent must be assigned an object that she prefers to o_c, and by rationality of the swaps such a leaf-agent can deviate at most once with the center-agent. Therefore, during her only swap with the center-agent, a leaf-agent must obtain an object that she prefers to o_c. Consequently, we construct a directed graph $G_D = (N, E')$ where there is an arc $(i, j) \in E'$ if and only if the center-agent can exchange with leaf agent j, when she owns the initial object of agent i, within a rational swap that gives to j an object that she prefers to o_c. In other words, there is an arc $(i, j) \in E'$ if and only if agent j prefers the initial object of i to her own object and the center-agent prefers the initial object of j to the initial object of i. We then search for a directed path in this directed graph from the center-agent to the initial owner of o_c. □

It seems that Reachable LEF is easier than Dec-LEF. At least, we have not identified classes of graphs for which Reachable LEF is harder than Dec-LEF but the reverse holds (take the example of a graph which is a matching). However, in general, the tractability of Reachable LEF is limited by the Reachable Object problem (Theorem 4.18) which is hard for many classes of graphs (see Chapter 3).

4.7 Experiments

In order to better understand the impact of the structure of the graph on local envy-freeness, we run some experiments where we investigate the influence of different characteristics of the network. In particular, we observe the impact of the degree of the nodes in regular undirected graphs, the density in random graphs, and some specific classes of graphs.

4.7.1 Impact of the degree of the nodes

In this subsection, we generate random instances of our decision and optimization problems, and we solve these instances exactly using mixed integer linear program formulations. We build on the ones proposed by Dickerson et al. [2014], which address envy-freeness and the minimization of maximum pairwise envy among any two agents [Lipton et al., 2004] (see Section 1.4.1), in a context of additive utilities with several goods per agent. To fit our setting, we adapt it so as to account for graph constraints and allocations of exactly one item per agent. We further design three variants, two where the objective functions correspond to MAX-LEF and MAX-NE, and another one where the locations of agents on the graph are treated as decision variables, to address the more challenging problem DEC-LOCATION-LEF.

For these experiments, we generate random undirected graphs that are regular with degree k, with 8 vertices, for k ranging from 1 to 7. The preferences of the agents are randomly drawn from impartial culture. The results, averaged over 1000 runs, are presented in Table 4.2 for the likelihood of finding an LEF allocation ("LEF"), the minimum number of envious agents ("max-LEF"), the maximum average degree of non-envy ("max-NE"), the minimum of maximum pairwise envy ("MMPE"), and the likelihood of finding both an allocation of the agents to the nodes of the graph and an allocation of objects to the agents in such a way that it is LEF ("Loc-LEF").

Degree	1	2	3	4	5	6	7
LEF	1	0.72	0.22	0.05	0.02	<0.01	<0.01
max-LEF	0	0.28	0.93	1.52	1.95	2.44	2.78
max-NE	1	0.99	0.99	0.99	0.98	0.98	0.98
MMPE	0	0.28	0.83	1.19	1.42	1.69	1.91
Loc-LEF	1	1	1	0.92	0.49	0.07	<0.01

Table 4.2: Likelihood of LEF, number of LEF agents, $\mathcal{NE}$, minmax pairwise envy and likelihood of LEF with location, in regular undirected graphs of different degrees with 8 agents and impartial culture

A natural question is how the likelihood to find an LEF allocation evolves as the degree of the graph augments. It must clearly decrease (in the extreme case of a complete graph, recall that all agents must have a different preferred item, which occurs with a probability as low as $n!/n^n$). The question is how this drop will occur. Our experiments show that this decrease is sharp, and from a degree equal to half of the agents, it actually becomes highly unlikely to find an LEF allocation. On the other hand, for graphs of small degrees, it is often the case that an LEF allocation can be found, and, as expected, it becomes even more so as the number of agents and items augments. Further experiments on a higher number of agents confirm this. As a rule of thumb, this means for instance that from 20 agents, it is almost always possible to find an LEF allocation on a grid-like network.

The ability to allocate agents on the network gives the central authority some extra-power when it comes to find an LEF. However, note that this power heavily depends on the structure of the graph (for instance, it is useless when the graph is complete, as all the different ways to label the graph with agents are isomorphic). Table 4.2 shows that this power can be significant: the likelihood to find an LEF remains above 90% until degree 4, while it was as low as 5.5% in the basic problem.

We also report in Table 4.2 results regarding the measures we optimize (as well as the "classical" minimization of maximum pairwise envy (MMPE) of Lipton et al. [2004], which in our context can be interpreted as minimizing the maximum number of agents envied by any agent). Note in particular that even with a complete graph, it is on average possible to allocate items so as to make envious only about a third of the agents, and that no agent envies more than two other agents in our instance with 8 agents.

4.7.2 Influence of the density in random networks

We do not restrict anymore to undirected graphs and consider Erdös-Rényi random directed graphs, with different densities (see Section 1.2.2.b). This type of graphs is notably known for reproducing in expectation the small-world phenomenon, property that usually appears in real social networks.

In this subsection, we are only concerned with the frequency of positive instances of Dec-LEF, that is how often a locally envy-free allocation exists, and how many when they exist. We also look at the number of such allocations that are Pareto-efficient. The preferences are generated either from impartial culture (IC), or are single-peaked uniform (SP-U), i.e., uniformly drawn from the urn containing all single-peaked profiles, or single-peaked uniform peak (SP-UP), i.e., single-peaked generated uniformly according to the top object of each agent (see Section 1.2.1.b). As in the first subsection, we run 1000 instances with 8 agents. The results about frequency of existence of LEF allocations are given in Figure 4.1 for different densities of the graphs.

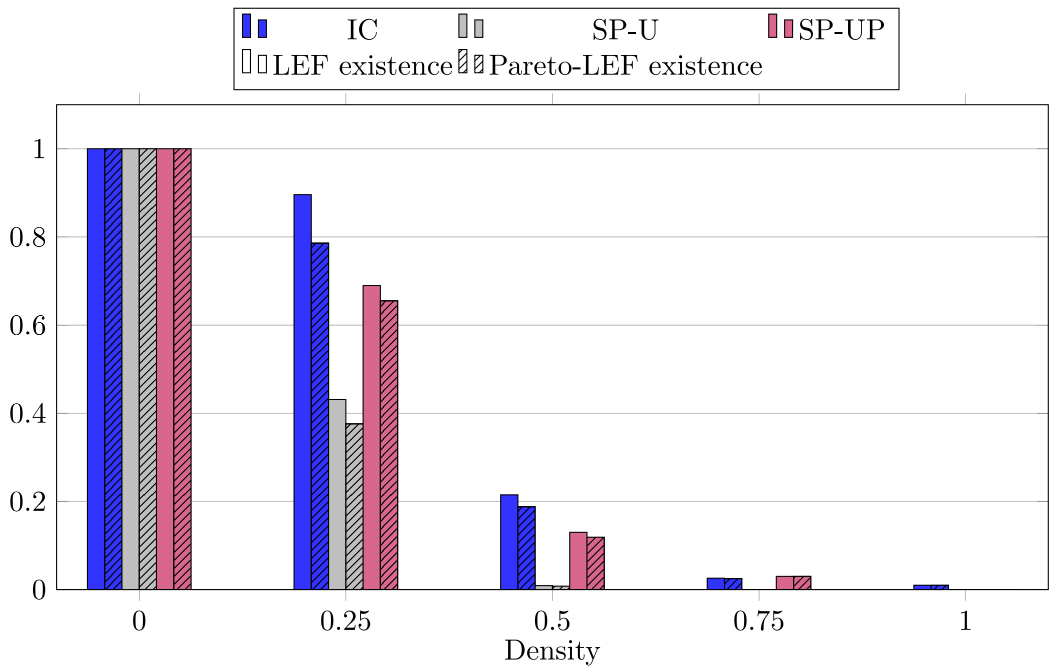

Figure 4.1: Proportion of instances for which there exists an LEF allocation in Erdös-Rényi graphs of different densities with 8 agents

The experiments confirm the intuition that the likelihood of finding an LEF allocation decreases when the density increases. However, we did not expect that the decrease would be as fast. Indeed, when the density is equal to 0 (empty graph), the case is trivial since an LEF is always guaranteed to exist, leading to a likelihood equal to 1 for the existence of an LEF allocation. But when the density increases to 0.25 (which still corresponds to a sparse graph), the frequency of finding an LEF allocation decreases by more than the half for SP-U profiles, by 25% for SP-UP profiles and by 10% for uniformly generated profiles. Moreover, the gap between a density of 0.25 and a density of 0.5 is very important as well, since the likelihood of finding an LEF allocation does not exceed 20% for any type of preferences, and is almost null for SP-U profiles.

In general, the likelihood of finding a locally envy-free allocation is significantly higher in profiles generated uniformly (IC). This can be explained by the fact that the preferences of the agents are less correlated and thus globally the agents do not desire the same objects. In the same idea, the likelihood of finding an LEF allocation is higher in SP-UP profiles than in SP-U profiles because the preferences of the agents are also more diverse. Indeed, by considering for instance the two extreme points of the single-peaked axis, say o_1 and o_n, the probability of drawing the unique single-peaked preference order with o_1 or o_n as a peak (i.e., as the top object) is equal to $\frac{1}{2^{n-1}}$ in the SP-U distribution whereas it is equal to $\frac{1}{n}$ in the SP-UP distribution.

It is noteworthy that the frequency of existence of an LEF allocation that is also Pareto-optimal is very close to the frequency of existence of a simple LEF allocation. This means that there are only a few instances where the existence of an LEF allocation does not guarantee the existence of an allocation that is both LEF and Pareto-efficient.

Let us now examine the number of LEF allocations when they exist. Figure 4.3 displays the number of LEF allocations in instances where there exists at least one such allocation, as well as the number of LEF allocations that are Pareto-optimal when such allocations exist. The numbers are given in average without counting the negative instances for existence in order to have a clearer idea and not being noised by the numerous instances with no such allocations. The instances are the same as those considered for testing the likelihood of existence. Recall that the total number of possible house allocations for instances with 8 agents is equal to 40320.

	IC		SP-U		SP-UP	
Density	Nb LEF	Nb Pareto-LEF	Nb LEF	Nb Pareto-LEF	Nb LEF	Nb Pareto-LEF
0	40320	163.74	40320	2001.31	40320	395.38
0.25	66.82	8.92	80.45	16.20	90.54	10.15
0.5	3.16	1.88	2.11	1.625	3.86	1.92
0.75	1.08	1.07	0	0	1.25	1.17
1	1	1	0	0	0	0

Table 4.3: Number of LEF allocations and Pareto-LEF allocations in positive instances of Dec-LEF for Erdös-Rényi graphs of different densities with 8 agents

The same observations as for the likelihood of existence hold: the gap is terribly high between empty graphs and density 0.25, and between density 0.25 and density 0.5 (recall that in an empty graph all the allocations are LEF). However, although there are more instances with LEF allocations under impartial culture, the number of LEF allocations in average is not the largest. Indeed, the largest number of LEF allocations in average is found for SP-UP profiles and then (at least for density 0.25) for SP profiles. This may be explained by the fact that there are less instances with existence of LEF allocations in single-peaked profiles and thus when they exist in a given instance, they may exhibit an opportune configuration where the LEF allocations are very numerous. On the contrary, under impartial culture there can be many configurations where only a few LEF allocations can be found. We finally observe that in general the number of LEF allocations that are also Pareto-optimal is very small.

4.7.3 Specific classes of graphs

We conduct the same type of experiments as in the previous subsection, but this time on specific classes of graphs, namely the directed acyclic graphs (DAGs), the Barabási-Albert random graphs, and graphs with homophily. The Barabási-Albert graphs are typical scale-free graphs, property that is usually found in real networks (see Section 1.2.2.b). Note that they are not generated for a given density. The graphs with homophily are generated according to the preferences of the agents, as described in Section 1.2.2.b. We also consider graphs complementary to graphs with homophily (we denote them as $\overline{\text{homophily}}$).

The results are presented in Figure 4.2 and Table 4.4.

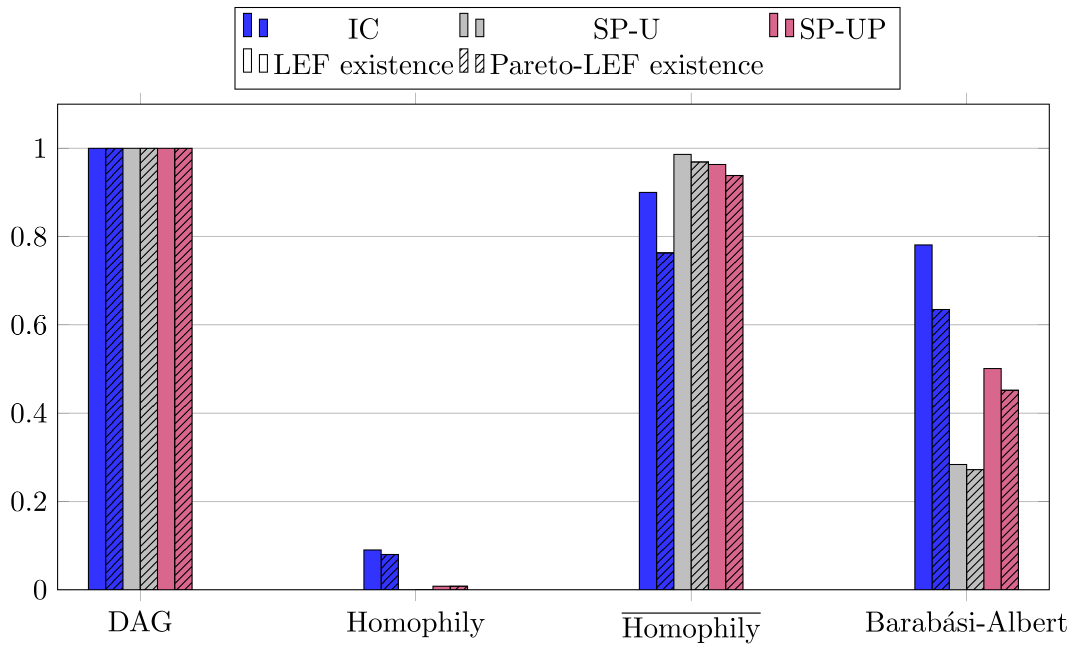

Figure 4.2: Proportion of instances for which there exists an LEF allocation for different classes of graphs with 8 agents

As expected by Proposition 4.2, there always exists an LEF allocation in directed acyclic graphs. However, the likelihood of finding an LEF allocation in a graph with homophily is extremely low. Indeed, closer the preferences of two agents, more likely they are to be connected in the network with homophily. Therefore, it appears natural

	IC		SP-U		SP-UP	
Graph	Nb LEF	Nb Pareto-LEF	Nb LEF	Nb Pareto-LEF	Nb LEF	Nb Pareto-LEF
DAG	4837.45	40.845	5144.162	311.209	5804.745	102.371
Homophily	2.27	1.375	0	0	1	1
$\overline{\text{Homophily}}$	208.23	16.97	6396.40	663.15	1037.72	123.07
Barabasi-Albert	7131.24	47.91	18006.51	963.60	11744.64	236.52

Table 4.4: Number of LEF allocations and Pareto-LEF allocations in positive instances of Dec-LEF for different classes of graphs with 8 agents

that this is difficult to find an LEF allocation in such instances. On the contrary, when this is the complementary graph, i.e., the non-envy graph $\overline{G}$, that respects homophily, the likelihood of finding an LEF allocation is very high, due to the same reasons. Following the same idea, the likelihood of existence of an LEF allocation in such graphs is even higher for single-peaked profiles. Concerning the Barabási-Albert graphs, the results are more intermediate and illustrate the general behavior of local envy-freeness regarding preferences: the likelihood of finding an LEF allocation is important under impartial culture (around 80%), low for SP-U profiles (around 25%) and medium for SP-UP profiles (around 50%). As for Erdös-Rényi graphs, the frequency of existence of an LEF allocation that is also Pareto-efficient is not far from the frequency for a simple LEF allocation.

Regarding the number of LEF allocations, in accordance with the likelihood of existence which is very weak for networks with homophily, the number of locally envy-free allocations is also extremely low. For the other types of graphs, the number of LEF allocations is rather high, especially for Barabási-Albert graphs. As for Erdös-Rényi graphs, the number of LEF allocations is less important under impartial culture. However, except for the graph with homophily, the number of LEF allocations that are also Pareto-efficient is largely higher than in Erdös-Rényi graphs.

4.8 Concluding remarks

We have studied different aspects of local envy-freeness in house allocation settings. First of all, deciding whether a locally envy-free allocation exists in a given instance is computationally hard even for very simple and sparse graphs. Nevertheless, we were able to provide some tractable cases according to the network topology. See Table 4.1 for the details of the complexity results and polynomial cases, with respect to some parameters of the graph. Interestingly, this problem is solvable in polynomial-time in directed acyclic graphs and graphs of out-degree at least $n-2$. These cases are very interesting because they rely on meaningful envy-graphs. Indeed, the DAGs can model some hierarchical situations (see Example 4.2). Moreover, the graphs with out-degree at least $n-2$ refer to the case where the non-envy graph is of out-degree at most 2, and thus includes the case where the non-envy graph is composed of couples of agents within which there is no reason for envy to exist.

We have also investigated an optimization perspective, and have tried to minimize the number of envious agents or the average envy with respect to specific degrees of envy. We have provided for both cases approximation algorithms with rather good approximation ratios.

In a third direction, we have studied the power of the central authority by assuming that, given the structure of the social network, she can assign in addition to the items to the agents, the agents to the locations on the graph. This problem can be understood as assigning tasks to workers as well as time slots (see Example 4.1). Although hard in general, this problem is solvable in polynomial time for the interesting case of graphs of degree at least $n-2$.

Furthermore, in relation with a process of swap dynamics in house allocation (see Chapter 3), we have asked the question of the reachability of an LEF allocation by a sequence of swaps. It turns out that the problem is computationally difficult even for simple graphs, and tractable cases can be found only for basic graphs.

Finally, the experiments confirm the intuition that the likelihood of finding a locally envy-free allocations is higher in sparse graphs. But they also highlight the fact that for some graphs close to real social networks, such as non-envy graphs with neighbors having similar preferences or scale-free networks, the probability of existence of an LEF allocation is rather high as well as the number of such allocations.

There are several interesting future directions to explore. First of all, one could examine further natural constraints on Dec-LEF. For instance, given a partial allocation of the objects, can a full LEF allocation be found? Or, given some forbidden object-agent pairs, can an LEF allocation be found?

Another relevant direction is to assume domain restrictions for preferences. There is a long tradition in social choice to consider domain restrictions on the preferences of the agents to obtain positive results. This would be natural to study our setting under such assumptions. For example, we can fix the number of different rankings. To take a concrete question, how difficult Dec-LEF and Dec-Location-LEF are when there are only two categories of agents: those with ranking $\succ_1$ on the objects and those with ranking $\succ_2$? More generally, can well-known domain restrictions, such as single-peakedness, be useful? Since the relevance of this domain restriction in the context of house allocation has recently been emphasized [Bade, 2017, Damamme et al., 2015], this might be an interesting road to pursue.

In this chapter, we have investigated a model where the fairness of a social choice solution is evaluated locally according to the perception of the agents, that is conditioned by the social network. The information provided by the social network enables the agents to be able to compare their own situation with the situation of other agents. This information can also be useful for other settings, such as voting, in a strategic perspective where agents could design strategies for manipulation whose accuracy would depend on the social network.

Chapter 5

Uncertainty in Iterative Voting

Abstract

This chapter deals with strategic voting under incomplete information. We propose a descriptive model, inspired by political elections, where the information about the vote intentions of the electorate comes from public opinion polls and a social network, modeled as a directed graph over the voters. The social network represents a visibility relation, in the sense that the voters can only observe the votes of their successors in the graph. The voters are assumed to be confident in the poll and they update the communicated results given by the poll with the information they get from their relatives in the social network. We consider an iterative voting model based on this behavior and study the associated "poll-confident" dynamics. Two configurations are investigated: an election with a single initial poll and one where several polls are performed and communicated all along the period. We analyze the convergence of the poll-confident dynamics for both configurations and the quality of the outcomes, with respect to the structure of the social network. In this context, we also ask the question of manipulation by the polling institute.

Résumé

On s'intéresse dans ce chapitre à un problème de vote stratégique avec information incomplète. Un modèle descriptif, inspiré des élections politiques, est proposé, dans lequel l'information portant sur les intentions de vote de l'électorat provient des sondages d'opinion et d'un réseau social, modélisé par un graphe orienté sur les agents. Le réseau social représente une relation de visibilité, dans le sens où les agents peuvent observer l'intention de vote de leurs successeurs dans le graphe. Les votants sont supposés avoir confiance dans le sondage et se basent sur les résultats communiqués, qu'ils mettent à jour avec ce qu'ils peuvent observer des intentions de vote de leurs proches dans le réseau social. On considère un modèle de vote itératif basé sur ce comportement et on se propose d'étudier la dynamique de déviations associée. Pour cela, deux configurations

This chapter is an extension of [Wilczynski, 2019].

sont analysées : une élection avec un unique sondage initial, et une élection avec plusieurs sondages ponctuant la période électorale. Les conditions de convergence de la dynamique étudiée, ainsi que la qualité de son issue, sont examinées pour les deux configurations, en fonction de la structure du réseau social. Dans un tel contexte, on se propose également d'évaluer la possibilité de manipulation de la part de l'institut de sondage.

Contents

5.1 Introduction

Strategic voting occurs in many scenarios. Modeling manipulation in voting via a strategic game makes sense in order to capture the plausible outcomes of an election where voters may vote tactically. In voting games, the players are traditionally assumed to know all the others' votes. However, this assumption appears highly unrealistic, especially for political elections where the set of voters is actually very large. Therefore, it appears particularly relevant to relax this important assumption in order to design models closer to real life situations.

5.1.1 Uncertainty in voting, social network and public opinion polls

The most part of the literature dealing with uncertainty in voting (see Section 1.3.3.c) assumes that agents adopt a risk-averse behavior, in the sense that voters manipulate only under the guarantee that their deviation does not produce a worse outcome. Even if it appears as a rational behavior, it does not really capture the behavior of real voters who might actually manipulate even if they are not guaranteed to obtain a better outcome in any case. Indeed, if real electors were risk-averse, then they would never manipulate because uncertainty is too important in political elections, implying the possibility of a configuration where they would be worse off if they had manipulated. But, manipulation actually occurs in political elections [Blais, 2004, Alvarez et al., 2006, Daoust, 2015].

We argue that voters adopt an elementary, but not necessarily myopic, behavior regarding the information they have. Moreover, we reject the idea that real voters could think in terms of probabilities. Dealing with expected utility functions is cognitively hard for the agents [Tversky and Kahneman, 1974] and implies cardinal preferences that are difficult to elicit. Intuitively, we aim at modeling the following common strategic behavior in political elections.

Example 5.1 *Let us consider a context of a presidential election where the winner is determined via the Plurality rule. A preliminary opinion poll over the population announces the following results concerning the vote intentions of the electorate.*

Extreme-left (EL)	*Green (G)*	*Left (L)*	*Conservative (C)*	*Extreme-right (ER)*	*Other parties*
15%	*7%*	*20%*	*25%*	*18%*	*15%*

A certain voter, Alice, prefers the green candidate G. By observing the poll, Alice realizes that G will not win because his score is too far from the score of the candidate announced winner in the poll, namely the conservative candidate C. Consequently, Alice decides to change her vote intention and intends to vote for the left-wing candidate L, who is her second choice and whose score in the poll is not so far from the score of the conservative candidate. After several weeks, Alice notices that the most part of her relatives, who were initially supporters of the left-wing candidate, have changed their vote intention to the extreme-left candidate EL, who is also Alice's third preferred candidate. Finally, Alice decides to switch her vote intention from L to EL since she considers now that EL is more susceptible to win than L.

Inspired by political elections, we assume that voters have two sources of information: the *public opinion polls* and the *social networks*. The public opinion polls punctuate the election campaign in many countries, and represent an important part of the election process. Many studies aim at understanding how they affect voting [Brams, 1982, Forsythe et al., 1993, Fredén, 2017]. Concerning social networks, they occupy an increasingly important place in our lives. For a particular citizen, they constitute a tool to have an idea of the population opinions, even if this vision is biased. The social networks are a natural channel for acquiring information in context of uncertainty in voting [Chopra et al., 2004, Sina et al., 2015, Tsang and Larson, 2016].

In a context where the polls offer a large part of the information available to the voters, a natural question is: what happens if the polling institute does not say the truth about the scores of the candidates? In fact, can the polling institute manipulate the election? This question is related to the problem of *election control* [Faliszewski and Rothe, 2016], i.e., how an external agent can influence and manipulate the election. While usually the external agent manipulates by adding / deleting candidates or votes, or even links of a social network [Sina et al., 2015], we suppose that the polling institute may lie during the communication of the results of the poll.

5.1.2 Poll-confident model and related work

We study an iterative voting model where voters trust the results of an opinion poll and have all the same prior assumption about the vote distribution, which is given by the poll. The voters update this belief about the vote intentions of the society by observing the deviations of their relatives in the social network. The strategic behavior of voters is then conditioned by personal *pivotal thresholds* modeling their willingness to deviate and how strong is their belief, in a similar way as in the local-dominance approach [Meir et al., 2014, Meir, 2015] (see Section 1.3.3.c) or as in the non-myopic model [Obraztsova et al., 2016] (see Section 1.3.3.b). These parameters define a new best response dynamics, that we call *poll-confident*.

More precisely, our study generalizes the local-dominance or the non-myopic model, by assuming that the initial prior information possessed by the voters comes from the results of a public opinion poll, and that the current knowledge of some votes during the voting game is given by the connections of the social network. In consequence, these models correspond to ours when the social network is supposed to be complete. Nevertheless, the interpretation is quite different. On one hand, Meir et al. [2014] and Meir [2015] suppose in the local-dominance model that the voters consider the existence of some noise in the communicated scores. An uncertainty threshold associated with each voter delimits in consequence all the profiles that are estimated possible by a voter, according to a given distance from the actual voting profile. On the other hand, Obraztsova et al. [2016] do not take into account uncertainty in their model. Indeed, the voters are assumed to know all the current profile. The personal thresholds of the voters are not related to the information they have but to their strategic behavior, in a simple look-ahead perspective: they support the candidate that they consider the best among those whose score is reasonably close to the score of the winner. Our model combines the two different interpretations. In fact, the voters only know the current ballots of their relatives in the social network. For the rest of the electorate, for who they are not able to observe the vote intentions, they base their estimation on the scores provided by the poll institute. They believe in the poll results but up to a certain level, level abstractly represented by the personal pivotal thresholds, which also model the willingness to deviate in a non-myopic way. Actually, the voters consider as potential winners all the candidates that are close to win in the scores of the poll, with respect to their own pivotal threshold, by believing that the real scores can be slightly modified or that these candidates can win if a reasonable number of new voters decide to support them. When the social network is complete, the pivotal thresholds only model a personal strategic behavior, that is a willingness to deviate, in the spirit of Obraztsova et al. [2016].

Within the important literature dealing with uncertainty in voting, other works assume that the partial information possessed by the voters comes from a graph, which is a social network [Sina et al., 2015, Tsang and Larson, 2016, Tsang et al., 2018], or simply a *knowledge graph* [Chopra et al., 2004]. However, our work differs from these previous studies with respect to various parameters. First of all, Tsang and Larson [2016] and Tsang et al. [2018] consider that the voters have cardinal preferences and infer some probability distribution over the global repartition of the votes in the electorate from what they can observe in the network. Even though this approach highlights an existing bias of the voters according to their environment, it ignores the possibility of external information. However, external information is actually present in political elections, where voters can have some information about a prior distribution of the votes, via opinion polls for instance. We consequently assume the presence of an initial public opinion poll that is common knowledge for the voters. Our model is closer to Sina et al. [2015]'s, since they also consider ordinal preferences and an initial opinion poll, even if they suppose that the network is an undirected graph. Nevertheless, the voters are assumed to aggregate these two types of information very differently: they compute the new scores by adding the score of a candidate, that is observable in the network, with the score of this candidate in the poll. The voters deviate when they believe that they are exactly pivotal, with no notion of pivotal threshold. Our study is also similar on the election control question: whereas Sina et al. [2015] investigate the manipulation by an external agent controlling the network, we suppose that the polling institute can provide corrupted results. Without focusing on a specific voting rule or manipulation move, Chopra et al. [2004] shows a convergence property, when the knowledge graph is a directed acyclic graph, that is close to our Proposition 5.1, even if the model of deviations and the information aggregation are different. Globally, most of these works only focus on the Plurality rule.

5.1.3 Contributions and organization

We especially investigate two configurations for the poll-confident dynamics: a *local* dynamics with one initial poll, and a *global* dynamics where the results of several polls are announced along the election. We especially focus on the Plurality rule, where the interpretation of the ballot is simple, i.e., approving one candidate, and where the voters can easily deduce some potential winners from the results of the poll (see Endriss et al. [2016] for a study of manipulation regarding the information communicated to the voters).

Convergence results are provided, stating that the local dynamics is guaranteed to converge for some classes of graphs. However, in general, it is difficult to recognize convergent instances for both local and global poll-confident dynamics. Experiments are given for testing the practical convergence of the dynamics and the quality of their outcomes. Then, we study the computational complexity of manipulation by the polling institute, proving that it is hard even for simple graphs. Beyond the general hardness of the problem, we provide some simple heuristics which are efficient in practice.

We present in Section 5.2 our model and the poll-confident dynamics in their local and global version. Then we analyze their convergence in Section 5.3, and the experimental quality of their outcome in Section 5.4, before finally investigating the manipulation by the polling institute in Section 5.5.

5.2 Poll-confident dynamics

5.2.1 Strategic voters in a social and informative context

Let us consider a voting game $\langle N, M, \succ, \mathcal{F}_{\rhd} \rangle$ with pure strategies, as described in Section 1.3.3, where the agents are given by the set $N = \{1, \dots, n\}$ and the candidates are given by $M = \{a, b, c, \dots\}$. In particular, the preferences of the voters over the candidates are strict and modeled by linear orders. Voting rule $\mathcal{F}$ will mostly refer to the Plurality rule, even if some definitions and results hold for any voting rule based on scores over the candidates. Moreover, $\mathcal{F}_{\rhd}$ is the resolute voting rule where $\mathcal{F}$ is combined with the tie-breaking rule associated with $\rhd$, which is assumed to be the alphabetical order, i.e., $a \rhd b \rhd c \rhd \dots$. We study an iterative voting framework where the dynamics is associated with unilateral deviations and the manipulation moves are defined according to the information possessed by the voters.

Initially, an opinion poll is undertaken from the initial voting profile σ^0, where the voters are assumed to give their real preferences (the agents do not have enough information to manipulate yet), i.e., σ^0 is truthful. The result Δ^0 of this initial poll is communicated to the agents, via a vector of scores describing the Plurality score of each candidate, i.e., $\Delta^0 : M \to \mathbb{N}$. At this moment, the voters know the scores obtained by each candidate, but may not have the information on which voter has voted for which candidate. Let $(s_j)_{j \in M}$ be a vector of scores associated with voting rule $\mathcal{F}$. For the sake of simplicity, we sometimes write $\mathcal{F}(s)$ to designate the winner of a voting profile whose score vector is s.

Though the agents are aware of the initial poll, they do not know all the subsequent deviations of the agents. We suppose that the agents are embedded in a social network $G = (N, E)$, represented by a directed graph whose nodes are the agents. An arc $(i, j) \in E$ means that agent i can observe the ballot of agent j at any time. We denote by $\Gamma(i)$ the set of agents for which agent i can observe the current ballot, i.e., $\Gamma(i) = \{j \in N : (i, j) \in E\} \cup \{i\}$. For a voting profile σ, the score of candidate x that agent i is able to observe is denoted by $\mathrm{Sc}_i^{\sigma}(x)$, which corresponds to the score of x under voting rule $\mathcal{F}$ in the subprofile of σ restricted to i and the successors of i in G, i.e., $\mathrm{Sc}_i^{\sigma}(x) := \mathrm{Sc}^{\sigma_{\Gamma(i)}}(x)$. An instance of a *linked voting game*, where the informative context of agents is given by a graph, is then a tuple $\mathcal{I} = \langle N, M, \succ, \mathcal{F}_{\rhd}, G \rangle$.

5.2.2 Information aggregation and belief update

The voters have two sources of information about the current scores: the results of the initial poll and the current votes of their relatives, i.e., their successors in the network.

The question is how they aggregate these informations. We assume that the voters base their belief on the results of the poll, updated with the votes of their relatives.

In fact, each voter has her own belief about the score of each candidate. Let us denote by $\mathrm{BS}_i^t : M \to \mathbb{N}$ the *believed score function* of agent i at the t^{th} step of the game. Initially, $\mathrm{BS}_i^0 = \Delta^0$ for every voter i since the voters trust the poll. We may see BS_i^t and Δ^0 as m-tuples where j^{th} coordinate corresponds to the believed and the announced score of j^{th} candidate (using the alphabetical order), respectively. At step t of the game, the believed score function of agent i is updated only if a relative of i deviates.

Definition 5.1 (Score Belief Update) *The update of the believed score function of agent i at step $t+1$, after the deviation of an agent j from strategy σ_j to σ'_j at step t, is given by $BS_i^{t+1} := BS_i^t \otimes_G (j, \sigma_j, \sigma'_j)$. Operator $\otimes_G$ updates believed score function BS_i^t according to the differences between old ballot σ_j and new ballot σ'_j only if $j \in \Gamma(i)$.*

In particular, for Plurality, $\otimes_G$ computes the new score of each candidate z as follows, after the deviation of an agent j from candidate x to candidate y:

$$[BS_i^t \otimes_G (j, x, y)](z) = \begin{cases} BS_i^t(z) - 1 & \textit{if } z = x \textit{ and } j \in \Gamma(i) \\ BS_i^t(z) + 1 & \textit{if } z = y \textit{ and } j \in \Gamma(i) \\ BS_i^t(z) & \textit{otherwise} \end{cases}$$

Note that for any voter i such that $\Gamma(i) = \{i\}$ and i never deviates, $\mathrm{BS}_i^t = \Delta^0$ for any step t of the game. Similarly, for any voter i such that $\Gamma(i) = N$, $\mathrm{BS}_i^t = \mathrm{Sc}^t$ for any step t. If $\Gamma(i) = N$ for every voter i, then we get the classical iterative voting framework [Meir et al., 2010] (as depicted in Section 1.3.3.b), where the voters know the current vote of any other voter. This situation occurs when social network G is a complete graph.

At any step of the linked voting game, the voters have a belief on the scores of the current profile under voting rule $\mathcal{F}$. The voters then elaborate strategies for manipulation, based on this belief.

5.2.3 Manipulation moves

We are now able to define specific manipulation moves under Plurality for agent i at step t, regarding her belief and her strategic behavior. A *pivotal threshold*, represented by a fixed integer p_i, is associated with each voter i, and describes the strategic type of voter i under Plurality. We denote by $p = (p_1, \ldots, p_n)$ the vector of pivotal thresholds where p_i refers to the pivotal threshold of agent $i \in N$. When p_i is the same for all the voters, the pivotal thresholds are said to be *homogeneous*. The pivotal threshold of a voter represents her willingness to deviate to a new strategy on the basis of her belief about the scores. This can measure how much the voters are uncertain about their belief or represent the threshold from which they think that their vote can be pivotal and matter in the election, notably by influencing the other votes. This threshold defines for each voter i a set of potential winners at step t, denoted by PW_i^t, representing the candidates that are still able to win the election, according to voter i. The intuition is that a candidate is a potential winner if the number of additional votes which are necessary to make it the winner is lower than the pivotal threshold. At step t, agent i

believes that the winner ω_i^t is the candidate that maximizes the believed scores BS_i^t. If we do not consider the current vote of agent i, then let us denote by $\mathrm{BS}_i^{t,\setminus i}$ the believed score function of voter i at step t, and its associated winner $\omega_i^{t,\setminus i}$.

Definition 5.2 (Potential winner) *Candidate x is a potential winner for agent i at step t if and only if i believes that x has currently less than p_i points of difference with the winner, i.e., x belongs to PW_i^t if and only if $BS_i^{t,\setminus i}(\omega_i^{t,\setminus i}) - BS_i^{t,\setminus i}(x) + \mathbb{1}_{\{\omega_i^{t,\setminus i} \rhd x\}} \leq p_i$.*

Observe that the tie-breaking rule matters in the definition of a potential winner because, if x obtains the same score as the believed winner (in $\mathrm{BS}_i^{t,\setminus i}$) but x has a lower rank than $\omega_i^{t,\setminus i}$ in $\rhd$, then it is not sufficient to make it win. In such a case, one more point should be added to the score of x in order to make it the new believed winner. This is expressed by the last part of the left-hand side of the inequality, in the previous definition, where $\mathbb{1}_{\{\varphi\}}$ is equal to one if φ is true, and to zero otherwise.

The notion of potential winners is illustrated in the following example.

Example 5.2 *Consider an instance with 100 voters and six candidates, where $M = \{a, b, c, d, e, f\}$ and $\mathcal{F}$=Plurality. Take a voter i who, at a given step t of the iterative voting process, believes that the scores are: $BS_i^t = (20, 25, 19, 6, 12, 18)$. Agent i believes that the current winner at step t is candidate b with a score equal to 25. Assume that voter i currently votes for candidate e whose believed score is 12. Let us suppose that the pivotal threshold of i is equal to $p_i = 10$, that is she considers that any candidate with a gap of at most 10 points with the winner is a potential winner. The believed scores of the candidates according to voter i (minus her own vote) are graphically represented as follows.*

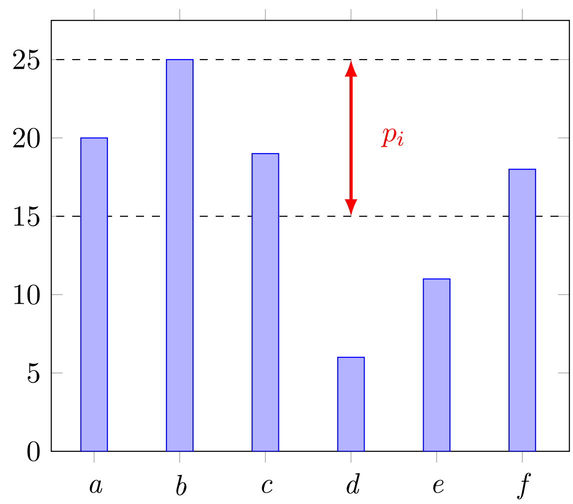

From the point of view of agent i, the potential winners are thus the candidates a, b, c and f. They are the candidates with a current believed score between the score of the believed winner and the score of the believed winner minus p_i.

Observe that, by definition, the believed winner when the current ballot of agent i is not taken into account, i.e., $\omega_i^{t,\setminus i}$, always belongs to the set of potential winners for agent i, as well as the current believed winner at step t for agent i, i.e., ω_i^t, if $p_i \geq 1$ (this may not be the same candidate if $\sigma_i^t = \omega_i^t$).

We are now able to define best response deviations for Plurality according to the set of potential winners.

A better response deviation according to the belief of agent i at step t must approve a candidate within PW_i^t. However, by assuming that the voters are rational, it makes sense to suppose that agent i directly changes for her preferred candidate within PW_i^t, leading to a concept of *believed best response*. We suppose that if there is a deviation for agent i, this is not to approve candidate ω_i^t, because agent i already believes that ω_i^t is winning, so there is no need to further support it.

Definition 5.3 (Believed best response (BBR$_p$)) *A deviation to a ballot $y \in M \setminus \{\sigma_i^t\}$ from ballot σ_i^t is a believed best response for agent i at step t if $y \in PW_i^t \setminus \{\omega_i^t\}$ and for all $x \in PW_i^t \setminus \{y\}$, it holds that $y \succ_i x$.*

Observe that the BBR$_p$ manipulation move is defined according to the strategic types of the voters, which are given by the vector p of pivotal thresholds, because the set of potential winners is defined with respect to pivotal threshold p_i. This manipulation move clearly generalizes the *direct best response* [Meir et al., 2010] (see Definition 1.27), which corresponds to the case where $p = (1, \ldots, 1)$ and the social network is complete. For arbitrary pivotal thresholds and a complete graph as a social network, a believed best response under Plurality corresponds to a *local-dominant strategy* [Meir et al., 2014] (if the initial voting profile σ^0 is truthful) (concept recalled in Section 1.3.3.c), where the distance metric between voting profiles is an ℓ_1-norm, and to an *NM-Plurality response* [Obraztsova et al., 2016] (recalled in Section 1.3.3.b).

Let us remark that the pivotal thresholds allow to model different strategic behaviors of voters. For instance, for a given voter i, if $p_i = 0$ then voter i never manipulates whatever her current ballot and the current winner. If $p_i = n$, then any candidate is a potential winner. In this case, by considering believed best response deviations, voter i is sincere and always gives her sincere ballot. If $p_i = 1$, then voter i deviates only when she is exactly pivotal, that is, according to her belief, she can make the outcome change.

In general, we define stable states with respect to the information given by the poll and a specific manipulation move μ.

Definition 5.4 *A strategy profile σ^t is a* (Δ^0, μ)-equilibrium *if and only if no agent $i \in N$ can perform a deviation of type μ from σ^t, according to her belief about σ^t, given by BS_i^t.*

Observe that the classical iterative voting, where the current votes are common knowledge for all the voters corresponds in our model to an instance where the social network is a complete graph. In such a context, a $(\Delta^0, BBR_{(1,\ldots,1)})$-equilibrium under Plurality is a *Nash equilibrium* (see Definition 1.23), as the essential condition for a Nash equilibrium is the complete knowledge about the current state [Aumann and Brandenburger, 2014].

5.2.4 Local and global dynamics

The information contained in the poll, updated by the observations given by the network, as well as a manipulation move μ and initial state σ^0, allow to define a new dynamics of the strategic game, that we call *poll-confident μ-dynamics*. We say that the poll-confident μ-dynamics converges if any sequence of deviations eventually leads to a (Δ^0, μ)-equilibrium, and we say that the dynamics can cycle if there exists a sequence of deviations where a strategy profile appears more than once. Note that in general, we do not consider a specific turn function τ for determining who deviates at a given step (see Section 1.3.3.b), unless we specify it as a part of instance $\mathcal{I}$ of the linked voting game, by mentioning $\mathcal{I}_\tau$.

A sequence of believed best responses deviations under Plurality is illustrated in the following example.

Example 5.3 *Let us consider an instance of the linked voting game with four voters and four candidates, where $N = \{1, 2, 3, 4\}$, $M = \{a, b, c, d\}$, and $a \rhd b \rhd c \rhd d$. The social network G and the preferences $\succ$ are as follows.*

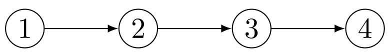

$$
\begin{aligned}
1:&\quad a \succ c \succ b \succ d\\
2:&\quad b \succ d \succ c \succ a\\
3:&\quad c \succ d \succ a \succ b\\
4:&\quad d \succ b \succ c \succ a
\end{aligned}
$$

The pivotal thresholds of the agents are such that $p_1 = 2$ and $p_2 = p_3 = p_4 = 1$. Consider the following sequence of believed best responses. At each step t, an arrow designates the deviation performed by voter i mentioned below the arrow, believing the score function $BS_i^{t,\setminus i}$ mentioned above the arrow. The real scores and winner are given in the brackets. The deviations are written in bold.

$$
\{(1,1,1,1) : a\} \xrightarrow[2]{(1,0,1,1)} \{(1,0,1,\mathbf{2}) : d\} \xrightarrow[3]{(1,1,0,1)} \{(1,0,0,\mathbf{3}) : d\} \xrightarrow[4]{(1,1,1,0)}
$$
$$
\{(1,\mathbf{1},0,2) : d\} \xrightarrow[3]{(1,2,0,0)} \{(\mathbf{2},1,0,1) : a\}
$$

From σ^0, agent 2 changes her vote to candidate d because she prefers d to the winner a. Agent 1 observes this move but has a pivotal threshold of 2: she still believes that her preferred candidate a can win so she does not deviate from a. Agent 3 did not observe any move, and thus believes in the initial scores, leading to her deviation to d. This belief also holds for agent 4 who deviates to b. Agent 3 observes this move while believing that agent 2 has not deviated and thus she moves to a. At this point, no agent can deviate given the information she gets, so profile (a, d, a, b) is a (Δ^0, BBR_p)-equilibrium.

5.2.4.a Global poll-confident dynamics

We define, analogously to the local poll-confident dynamics, a *global poll-confident dynamics* where several polls are communicated to the voters. More precisely, each time the local poll-confident dynamics converges, the scores of the candidates at the equilibrium are communicated via a poll. A sequence of deviations within global dynamics is a sequence of states $(\sigma^0, \sigma^{0,1}, \ldots, \sigma^{0,t_0}, \sigma^{1,1}, \ldots, \sigma^{1,t_1}, \ldots, \sigma^{t,1}, \ldots, \sigma^{t,t_t})$, where $\sigma^{i,j}$ is the strategy profile at i^{th} global step and j^{th} local step. For each $0 \leq k \leq t-1$, there is a poll Δ^{k+1} between state σ^{k,t_k} and state $\sigma^{k+1,1}$ which announces the scores of profile σ^{k,t_k} under $\mathcal{F}$. Initial poll Δ^0 is given between initial state σ^0 and state $\sigma^{0,1}$. For all $0 \leq k \leq t$, sequence $(\sigma^{k,1}, \ldots, \sigma^{k,t_k})$ represents the deviations of the local dynamics with initial poll Δ^k, which leads to the (Δ^k, μ)-equilibrium σ^{k,t_k}, for a given manipulation move μ. We assume that the voters do not keep in memory the history of the previous global steps. At each global step, the voters behave in the same way as in the initial one.

The global poll-confident dynamics can cycle if there exists a sequence of deviations $(\sigma^0, \sigma^{0,1}, \ldots, \sigma^{0,t_0}, \ldots, \sigma^{t,1}, \ldots, \sigma^{t,t_t})$ such that there are $k < k'$ where $\sigma^{k',t_{k'}} = \sigma^{k,t_k}$ or $\sigma^{k',t_{k'}} = \sigma^0$, or if the associated local poll-confident dynamics cycles. A stable state regarding the global dynamics is called a *global equilibrium.*

Observation 5.1 *For a homogeneous pivotal threshold $p = (1, \ldots, 1)$, a state σ is a global equilibrium, according to believed best response deviations, if and only if σ is a Nash equilibrium.*

A global equilibrium σ is either σ^0 or a (Δ^t, BBR_p)-equilibrium for some step t. In any case, the scores at σ are announced either by Δ^0 or by Δ^{t+1}, leading to complete information about σ.

5.3 Convergence to Poll Equilibria

In this section, we study the theoretical convergence properties of the poll-confident dynamics. Despite some positive results, in the general case, it is difficult to recognize instances for which the dynamics converges.

5.3.1 Local poll-confident dynamics

We investigate the convergence properties of the local poll-confident dynamics according to the topology of the social network.

5.3.1.a General convergence results

Let us first investigate the general convergence properties of the local poll-confident μ-dynamics for an arbitrary (deterministic) manipulation move μ and an arbitrary voting rule $\mathcal{F}$ based on scores.

The deviation of each agent depends on the information she has about the deviations of the other agents, which is given by social network G. If a voter does not see anybody deviating after her own deviation, then she has no reason to deviate again.

Observation 5.2 *Any agent i such that $\Gamma(i) = \{i\}$ deviates at most once for any voting rule $\mathcal{F}$, any manipulation move μ, and any initial profile σ^0.*

Observation 5.2 implies that the local poll-confident μ-dynamics converges for any voting rule $\mathcal{F}$, any manipulation move μ ad any initial profile σ^0 when E is empty. A condition on the graph can be further derived in order to ensure the convergence of the dynamics for a larger class of graphs.

Proposition 5.1 *If G is a directed acyclic graph (DAG), then the local poll-confident μ-dynamics converges within $\mathcal{O}(n^2)$ steps for any voting rule $\mathcal{F}$, any manipulation move μ, and any initial profile σ^0.*

Proof: Suppose that the local poll-confident dynamics cycles within the subset of agents $N' \subseteq N$. Observation 5.2 implies that the out-degree in G of each agent $i \in N'$ is different from zero. However, a DAG contains at least one sink, i.e., a vertex of null out-degree, and $G[N']$ is a DAG since it is a subgraph of a DAG. Contradiction.

Now let us count the deviations. A deviation of a sink-agent can push at most $n-1$ other agents to deviate. By iteratively removing a sink of G and counting the deviations caused by a sink-agent in the new graph, we get that the sequence of deviations is of size $\mathcal{O}(n^2)$. □

The previous convergence property holds for any voting rule and any type of manipulation move. Let us now focus on specific dynamics that are guaranteed to converge from any initial state when the network is complete. We can actually generalize the convergence result on a complete graph to any transitive graph (see Definition 1.3). Indeed, if a μ-dynamics is proved to converge in presence of complete information about the current state, then this dynamics also converges when the social network is a transitive graph, as stated in the following proposition.

Proposition 5.2 *Let μ be a manipulation move such that the local poll-confident μ-dynamics on a complete graph is guaranteed to converge from any initial profile within $\mathcal{O}(f_\mu(n,m))$ steps. If G is transitive, then the local poll-confident μ-dynamics converges within $\mathcal{O}(n \cdot f_\mu(n,m))$ steps to a (Δ^0, μ)-equilibrium from any initial profile.*

Proof: Let N' be a minimal subset of agents for which the local poll-confident μ-dynamics cycles in a given instance where G is a transitive graph. Remark that all the agents i in a clique K of G have the same believed scores BS_i^t at any step t: by transitivity, $\Gamma(i) = \Gamma(j)$ for all $i, j \in K$. So, by the convergence guarantee of the μ-dynamics when all the current votes are known from any initial profile, N' does not form a clique in G. By Proposition 5.1, there must be a directed cycle (dicycle) in $G[N']$. But a dicycle along agents $N'' \subseteq N$ in a transitive graph implies that N'' actually forms a clique in G. So, $G[N']$ includes a set of cliques. Since there is no visibility of the deviations between two disjoint connected components and N' is minimal, $G[N']$ is connected. The agents in N' can be partitioned into the groups $N'_1, \dots, N'_k$ following their level of knowledge: each $G[N'_i]$ for $i \in [k]$ is a set of disjoint cliques and any agent in N'_i can only observe some agents in N'_j for $j \leq i$. Each agent in a group N'_i deviates according to the deviations of agents in $\bigcup_{j \leq i} N'_j$, because she cannot observe other deviations. So, a cycle in the dynamics within N' implies a cycle in the dynamics within $\bigcup_{j \leq i} N_j$. By minimality of N', N' is a single clique, contradiction. □

Observe that if all the agents are linked according to their preferences in such a way that two agents who agree on their most preferred candidate are connected via two opposite arcs in G, then this graph with homophily (see Section 1.2.2.b) is transitive. Thus the local poll-confident μ-dynamics converges in such a graph, for any μ-dynamics guaranteed to converge in presence of complete information about the current profile.

5.3.1.b Believed best responses under Plurality

Let us now study, more specifically, the believed best responses under Plurality. Note that, by Proposition 5.1, the local poll-confident BBR$_p$-dynamics is guaranteed to converge if G is a directed acyclic graph, for any vector of pivotal thresholds even when the pivotal thresholds are heterogeneous.

Recall that, when the network is complete and $p = (1, \ldots, 1)$, the local poll-confident BBR$_p$-dynamics corresponds to a direct best response (Definition 1.27). By Theorem 3 of Meir et al. [2010], we get that the local poll-confident BBR$_p$-dynamics converges from any initial state within $\mathcal{O}(n \cdot m)$ steps, when the network is complete and $p = (1, \ldots, 1)$. Therefore, by Proposition 5.2, the following corollary holds.

Corollary 5.3 *The local poll-confident BBR$_p$-dynamics converges from any initial state within $\mathcal{O}(n^2 \cdot m)$ steps when the social network is a transitive graph and $p = (1, \ldots, 1)$.*

Let us now analyze the case of more general pivotal thresholds. Recall that the local poll-confident BBR$_p$-dynamics under Plurality, when G is a complete graph, corresponds to a local-dominant strategy [Meir et al., 2014] (see Section 1.3.3.c) if σ^0 is truthful, and to an NM-Plurality response [Obraztsova et al., 2016] (see Section 1.3.3.b). We deduce from the proofs of Theorem 5.3 of Meir et al. [2014] and Theorem 2 of Obraztsova et al. [2016] that the local poll-confident BBR$_p$-dynamics converges within at most $n \cdot (m-1)$ steps from the truthful voting profile, when the pivotal thresholds are homogeneous and the network is complete. However, these two proofs do not hold for an arbitrary initial profile, ruling out the possibility to use Proposition 5.2.

Nevertheless, a further restriction on turn function τ, in the spirit of Proposition 6 of Meir et al. [2014], enables us to derive a convergence property from any initial profile when the network is complete and the pivotal thresholds are homogeneous. We can consequently generalize this result to transitive graphs, thanks to Proposition 5.2.

For a BBR$_p$ deviation $a \xrightarrow{i} b$ occurring at step t, where voter i deviates from a to b, if $b \succ_i a$, then the associated deviation is called an *opportunity move* otherwise, if $a \succ_i b$, then it is called a *compromise move.*

Proposition 5.4 *The local poll-confident BBR$_p$-dynamics converges from any initial profile within $\mathcal{O}(n.m)$ steps, when the pivotal thresholds are homogeneous and the network is complete, in any instance $\mathcal{I}_\tau$ of the linked voting game where τ always chooses opportunity moves before compromise moves.*

Proof: Let us consider a sequence of BBR$_p$ deviations starting from any initial profile. Since the pivotal thresholds are homogeneous, let us denote by p_u the common pivotal threshold. Observe that there are at most $n(m-1)$ consecutive opportunity moves because each such a deviation implies that the rank, in the preferences of the manipulator, of the candidate approved in the new ballot is strictly better than in the previous ballot. Therefore, by definition of the turn function, there exists a state t_0 where no agent has incentive to perform an opportunity move. If no agent can deviate from σ^{t_0}, then we are done. Otherwise, there is some agent i that performs a compromise move $a \xrightarrow{i} b$ at step t_0. We will prove by a complete induction, for any step $t \geq t_0$ where an agent i performs a deviation $a \xrightarrow{i} b$, that:

(I) the score of the winner does not decrease,
(II) $\bigcup_{j \in N} PW_j^{t'} \subseteq \bigcup_{j \in N} PW_j^t$ for any step $t' > t$,
(III) $a \succ_i b$, i.e., this is a compromise move,
(IV) after this move no agent can vote for a.

First consider the state where $t = t_0$, with the BBR_p deviation $a \xrightarrow{i} b$. This is necessarily a compromise move by definition of t_0, thus $a \succ_i b$ and (III) holds. Since, by definition of a BBR_p, $b \in PW_i^t$, this is not possible that $a \in PW_i^t$ otherwise it must hold that $b \succ_i a$, which contradicts the assumption that there is no opportunity move at t_0. Consequently, $a \notin PW_i^t$, and in particular $a \neq \mathcal{F}(\sigma^{t_0})$, so the winner score cannot decrease, proving the base case of (I). Since the winner score has not decreased and agent i has only given one point to $b \in PW_i^t$, then $\bigcup_{j \in N} PW_j^{t+1} \subseteq \bigcup_{j \in N} PW_j^t$, proving the base case of (II). Observe that $PW_i^{t_0} = PW_i^{t_0+1}$ since agent i does not take into account her own ballot in the computation of the potential winners. Let us consider an agent $j \neq i$. If $a \notin PW_j^{t_0}$ then, since a loses one point after this deviation and the winner score does not decrease, it holds that $a \notin PW_j^{t_0+1}$. However, even if the thresholds are homogeneous, it is possible that $a \in PW_j^t$ whereas $a \notin PW_i^t$. Indeed, (1) i may not see a as a potential winner because she is currently voting for it, and/or (2) j can see more candidates as potential winners because her current ballot is the winner. Suppose first that the second case does not hold. Then, by definition of the potential winners, it holds that $\text{BS}_j^{t,\setminus j}(a) \geq \text{BS}_j^{t,\setminus j}(\omega_j^{t,\setminus j}) - p_u + \mathbb{1}_{\{\omega_j^{t,\setminus j} \rhd a\}}$, whereas $\text{BS}_i^{t,\setminus i}(a) < \text{BS}_i^{t,\setminus i}(\omega_i^{t,\setminus i}) - p_u + \mathbb{1}_{\{\omega_i^{t,\setminus i} \rhd a\}}$. However, the current ballot of i, as well as the current ballot of j at step t, are both different from the real winner $\mathcal{F}(\sigma^t)$, therefore $\omega_j^{t,\setminus j} = \omega_i^{t,\setminus i} = \mathcal{F}(\sigma^t)$. It follows that $\text{BS}_i^{t,\setminus i}(a) = \text{BS}_j^{t,\setminus j}(a) - 1 = \text{BS}_j^{t+1,\setminus j}(a)$. Because the score of $\omega_j^{t,\setminus j}$ has not decreased by the deviation of i, this implies that $a \notin PW_j^{t+1}$. Now suppose that (2) holds, that is $\sigma_j^t = \mathcal{F}(\sigma^t)$. Then, it is possible that, even after the deviation of i, a still belongs to the potential winners of j. However, it must hold that $\sigma_j^t \succ_j a$, otherwise agent j would have incentive to deviate within a compromise move at step t_0. Therefore, even if a still belongs to the potential winners of j, she has no incentive to deviate to it, at least while σ_i^t remains a potential winner.

Assume now that (I)-(IV) hold for any step from t_0 to some step $t > t_0$. We will prove that (I)-(IV) still hold for $t+1$. Suppose that there is a deviation $a' \xrightarrow{i'} b'$ at step t. Suppose that $b' = a$, that is some agent can vote for a that i has left at step t_0. From the proof of the base case, the only possibility is that the current ballot of i' is the winner of step t_0, which is not a potential winner anymore. But, if $\sigma_{i'}^t$ is not a potential winner, then this means that the score of the new winner has increased too much for σ_i^t remaining a potential winner. This is due to the fact that no agent currently voting for $\sigma_{i'}^t$ had incentive to move while it was winner, and to the assumption that the steps occur after t_0. However, if this is impossible for σ_j^t to belong to the potential winners, then it will be completely impossible for a who had no new support since the deviation of i at step t_0, which proves (IV). Suppose that i' performs an opportunity move. By induction assumption on (III) and the nature of the turn function, i' had no incentive to do it from t_0 to t. This means that a new candidate b', which has just entered into the potential winners of i', is preferred to her current ballot. But, no new candidate can enter into the potential winners of i' unless the winner score has decreased, which is impossible by induction assumption on (I). Therefore, (III) holds and the move of i' is a compromise move. This implies that $a' \notin PW_{i'}^t$, and thus the score of the winner cannot decrease after the move of i', proving (I). Agent i' cannot make the score of the winner decrease and b' must belong to $PW_{i'}^t$. This implies that (II) still holds. Hence, to summarize, the BBR$_p$-dynamics must converge.

Observe that there are at most $n \cdot (m-1)$ moves until t_0 is reached since only opportunity moves occur before t_0, and at most $n \cdot (m-1)$ moves after t_0, since only compromise moves can occur. Therefore, the whole process converges within $\mathcal{O}(n \cdot m)$ steps. □

Corollary 5.5 *The local poll-confident BBR$_p$-dynamics converges from any initial profile within $\mathcal{O}(n^2 \cdot m)$ steps, when the pivotal thresholds are homogeneous and the network is a transitive graph, in any instance $\mathcal{I}_\tau$ of the voting game where τ always chooses opportunity moves before compromise moves.*

The previous convergence result holds for a specific condition on the turn function. However, we conjecture that convergence is guaranteed even without this condition.

Despite the BBR$_p$-dynamics converges for specific graphs when the pivotal thresholds are homogeneous, the convergence is not guaranteed for any social network, even when $p = (1, \dots, 1)$, as we can observe in the following example.

Example 5.4 *Let us consider an instance with four voters and four candidates where $N = \{1,2,3,4\}$, $M = \{a,b,c,d\}$ and $a \rhd b \rhd c \rhd d$. The social network and the preference profile are as follows.*

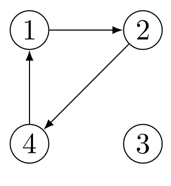

$1: \; c \succ d \succ b \succ a$

$2: \; b \succ c \succ d \succ a$

$3: \; a \succ c \succ b \succ d$

$4: \; d \succ b \succ c \succ a$

All the agents $i \in N$ have the same pivotal threshold $p_i = 1$. The following sequence of BBR$_p$-deviations, starting from the truthful profile, creates a cycle in the dynamics. Indeed, the states σ^1 to σ^7 form a cycle. Deviations are marked with arrows and bold letters. The believed scores $BS_i^{t,\setminus i}$ of each deviating agent i at step t are indicated above the arrow representing the manipulation move.

Steps	σ^0		σ^1		σ^2		σ^3		σ^4		σ^5		σ^6		σ^7
1 :	c	$\xrightarrow{(1,1,0,1)}$	$\mathbf{d}$		d	$\xrightarrow{(1,0,1,1)}$	$\mathbf{c}$		c		c	$\xrightarrow{(1,1,0,1)}$	$\mathbf{d}$		d
2 :	b		b	$\xrightarrow{(1,0,1,1)}$	$\mathbf{c}$		c		c	$\xrightarrow{(1,1,1,0)}$	$\mathbf{b}$		b		b
3 :	a		a		a		a		a		a		a		a
4 :	d		d		d		d	$\xrightarrow{(1,1,1,0)}$	$\mathbf{b}$		b		b	$\xrightarrow{(1,1,0,1)}$	$\mathbf{d}$
$\mathcal{F}(\sigma^t)$	a		d		d		c		c		b		b		d

It is easy to see that each manipulation move is a BBR_p deviation.

Moreover, note that when the pivotal thresholds are heterogeneous, there is no guarantee of convergence even for a complete graph as a social network. We can observe this phenomenon, for instance, in Theorem 3 of Obraztsova et al. [2016].

5.3.1.c Recognizing convergent instances

Globally, as Proposition 5.1 shows and Example 5.4 illustrates, the existence of a dicycle in the network is a necessary condition for the local dynamics to cycle. However, if all the subsets of agents that are connected by a dicycle in the graph actually form a clique in G, then the convergence is ensured under certain conditions, as Proposition 5.2 establishes.

These remarks are not sufficient for predicting the possibility of a cycle in the dynamics because the preferences of the agents also matter. For a general graph, we can even state that it is computationally hard to recognize instances for which the local poll-confident dynamics can cycle. This result holds even when all the agents have the same pivotal threshold equal to $p = (1, \dots, 1)$ and the voting rule is Plurality.

Theorem 5.6 *Deciding whether the local poll-confident BBR_p-dynamics can cycle is* NP*-hard, even for a homogeneous pivotal threshold $p = (1, \dots, 1)$.*

Proof: We perform a reduction from 2P1N-SAT, known to be NP-complete (Theorem 1.3). The 2P1N-SAT problem is a satisfiability problem defined on a set $\mathcal{C} = \{C_1, \dots, C_s\}$ of s clauses over a set $X = \{x_1, \dots, x_v\}$ of v variables. Recall that in 2P1N-SAT, each variable occurs twice as a positive literal and once as a negative literal. We assume that the clauses are indexed such that each first occurrence of a variable is a positive literal (see Section 1.5.1.a). Each clause $C_i \in \mathcal{C}$ contains s_i literals.

From an instance $\langle \mathcal{C}, X \rangle$ of 2P1N-SAT, we construct an instance $\mathcal{I} = \langle N, M, \succ, \mathcal{F}_{\rhd}, G \rangle$ of the linked voting game, where $\mathcal{F}$=Plurality. The set M of candidates includes candidates a, b, y, z and clause-candidates c_i for each $i \in [s]$. The tie-breaking $\rhd$ is defined according to the following linear order: $b \rhd y \rhd z \rhd c_s \rhd c_{s-1} \rhd \cdots \rhd c_1 \rhd a$. The set N of agents includes agents A, B, Y and Z, and literal-agents L^i_j for each j^{th} literal of clause C_i. We will usually refer to the literal-agents L^i_j, for $j \in [s_i]$, associated with clause C_i, as the agents of clause C_i. The preferences are as follows (all candidates not listed are ranked in arbitrary order within $[\dots]$, and the candidates in brackets may not exist for all the indices):

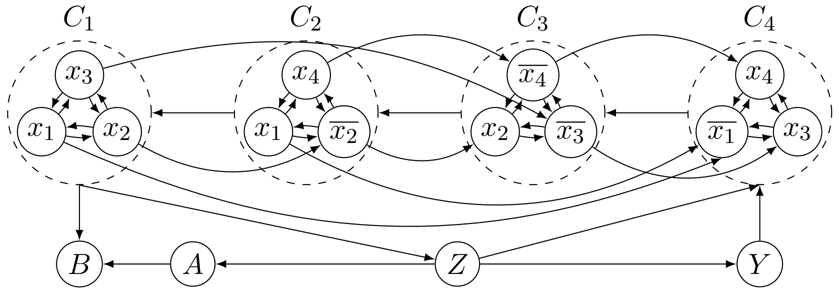

Figure 5.1: Construction of G for a 2P1N-SAT instance where $C_1 = (x_1 \vee x_2 \vee x_3)$, $C_2 = (x_1 \vee \overline{x_2} \vee x_4)$, $C_3 = (x_2 \vee \overline{x_3} \vee \overline{x_4})$ and $C_4 = (\overline{x_1} \vee x_3 \vee x_4)$. The literals (we keep for the figure the literal name) in a same circle are in a same clause; the arcs from or to the circles concern every agent inside.

$$L^i_j: \quad y \succ c_i \succ c_s \succ (c_{i-1}) \succ [\ldots]$$

$$A: \quad a \succ y \succ b \succ [\ldots] \qquad Y: \quad y \succ z \succ b \succ [\ldots]$$

$$B: \quad a \succ b \succ y \succ [\ldots] \qquad Z: \quad z \succ y \succ c_s \succ [\ldots]$$

We add $3v(s+2)+s$ dummy voters, where $3v+1$ among them rank c_i first for each candidate c_i, $3v$ rank b first and $3v$ rank z first. Every candidate, except a and b, obtains $3v+1$ votes, and b has $3v$ votes. The initial winner is y, thanks to $\rhd$. In the network G, all the agents of a same clause form a clique. There is an arc from each agent of clause C_i to each agent of clause C_{i-1}. There is an arc from each agent of clause C_1 to agents B and Z, and from agents Z and Y to the agents of clause C_s. Agent Z points to A and Y, and A points to agent B. There is an arc from agent L^i_j to agent L^ℓ_k if $i < \ell$ and the agents correspond to opposite literals. See Figure 5.1 for an illustration of the graph construction.

We claim that $\mathcal{C}$ is satisfiable if and only if the local poll-confident BBR$_p$-dynamics starting from the truthful profile where $p = (1, \ldots, 1)$ can cycle in instance $\mathcal{I}$. Indeed, we can prove that the only possible cycle within the dynamics involves the agents Y, Z and exactly one agent related to each clause. Moreover, if there are two agents corresponding to opposite literals among the deviating agents related to the clauses, then the cycle does not occur.

Suppose that $\mathcal{C}$ is satisfiable and consider a truth assignment ϕ of the variables. In instance $\mathcal{I}$, let agent B deviate to candidate b. Only agents A and the agents of clause C_1 can observe this deviation. Let A deviate now to candidate t and denote by σ' the resulting voting profile. The agents of clause C_1 do not see this deviation, they still believe they have incentive to deviate to ballot c_1. We choose a voter associated with a literal of clause C_1 that is true in ϕ and let her deviate to c_1, creating an incentive to manipulate for the agents of C_2. Following the same principle, we choose a deviating voter associated with a true literal in ϕ for each clause C_i, from C_1 until C_s. Afterwards, agent Z deviates to y because she still believes that y can win since she has observed A's deviation to y. Then, agent Y who can only observe the agents of C_s, deviates to z. By observing that Z gives one more point to y, the agent in C_1 who has previously deviated comes back to ballot y, leading progressively all the previous deviating agents within the clauses to come back to y. These new deviations are possible because a deviating agent in clause C_{j+1} can only observe the deviations of clause C_j, since there is no two opposite literals both true in ϕ. Now, agent Z who has observed the deviation of agent Y to ballot z, believes that z can win and thus deviates to it. Agent Y, who has only observed the deviation to y of an agent of clause C_s, deviates to candidate y, her first choice. We reach voting profile σ', leading to a cycle within the dynamics.

Now suppose that there is a cycle within the local poll-confident BBR$_p$-dynamics in instance $\mathcal{I}$. We claim that the only possible cycle within the dynamics involves the agents Y, Z and exactly one literal-agent per clause. Let us denote by Σ the agents involved in a cycle of the dynamics. The agents A, B and all the dummy voters do not belong to Σ because they have a null out-degree (and A only points to B). All the dummy voters are isolated, nobody can see their deviations, so we assume w.l.o.g. that they do not deviate. At the first step, only agent B has incentive to deviate. After this deviation to b, agent A and the agents of clause C_1 want to deviate. Only agent Z can see a deviation of A, but Z does not move because A deviates to y, which is the second best candidate of Z. She has incentive to deviate only if she observes another deviation, and the only possible one is from literal-agents of clause C_s. However, to make an agent of C_s deviate, there must have been a previous deviator within C_{s-1}, and so on along the clauses, leading to the previous deviation of at least one literal-agent in C_1. The agents of the clauses can progressively deviate from an initial deviation performed in C_1, because by observing the deviation of an agent in clause C_i to candidate c_i, an agent of clause C_{i+1} believes that voting c_{i+1} is a best response ($c_{i+1} \rhd c_i$).

We claim that no cycle can occur during this stage where only the agents of the clauses deviate. Indeed, as stated in Lemma 5.7, the literal-agents in clause C_i can only deviate to c_i or c_s. Moreover, once they deviate to c_s there is no possibility to deviate to candidate c_i since only the other literal-agents of C_i, where there is complete visibility, can deviate to c_i, and $c_s \rhd c_i$.

Lemma 5.7 *In the first stage where the literal-agents L_j^i in the clauses perform deviations and agents Y and Z have not deviated yet, each literal-agent L_j^i can only deviate to candidate c_i or candidate c_s.*

Proof: Initially, each literal-agent votes for y. So, once at least one literal-agent has deviated, no other literal-agent can deviate to y, because either she believes that y is not potential winner anymore, by having observed a deviation of another literal-agent, or she is currently voting for y. Suppose that at some step t and for some $\ell \in [s]$, at least one literal-agent of clause C_i has deviated, for each $i \in [\ell]$. Let us denote by $\overline{L_j^i}$ the literal-agent corresponding to the unique literal that is opposite to the literal associated with literal-agent L_j^i and that belongs to a clause i' such that $i' > i$. Moreover, $c(\overline{L_j^i})$ denotes the clause-candidate associated with this opposite literal-agent. A literal-agent L_j^1 of clause C_1 believes one of the following cases:

(1) Candidate b is winning with $3v + 1$ points. Consequently, L_j^1 deviates to c_1, which already has $3v + 1$ points.
(2) Candidate c_1 is winning. Then, L_j^1 has no incentive to move.
(3) Candidate c_s is winning. Therefore, L_j^1 deviates to c_1 if she can make it win or does not move otherwise.

(4) Candidate $c(\overline{L_j^1})$ is winning. This implies that agent $\overline{L_j^1}$ has voted for it. Since L_1^j could only observe this vote in favor of $c(\overline{L_j^1})$, it suffices to give one more vote to c_s to make it winner ($c_s \rhd c(\overline{L_j^1})$ if $c_s \neq c(\overline{L_j^1})$).

Therefore, agent L_j^1 can only deviate to candidate c_1 or c_s. Suppose now that the induction assumption holds for all the agents of C_{i-1} for some $i \in [\ell - 1]$ and let L_j^i be a literal-agent of C_i. By induction assumption, all the agents of C_{i-1} have not deviated, or have voted for c_{i-1}, or for c_s. Thus, agent L_j^i believes that the winner is within $\{c_{i-1}, c_i, c(\overline{L_j^i}), c_s\}$. We can detail each case:

(1) Candidate c_{i-1} is winning. Then, agent L_j^i can deviate to c_i or c_s if she is able to make win one of them. Otherwise she does not have incentive to move.
(2) Candidate c_s is winning. Then agent L_j^i can only try to deviate to c_i.
(3) Candidates c_i is winning. Consequently, agent L_j^i does not have incentive to move.
(4) Candidate $\overline{c_j^i}$ is winning. Then, agent L_j^i has only observed the vote of the opposite agent $\overline{L_j^i}$ who voted for $c(\overline{L_j^i})$. Thus, L_j^i can either vote for c_i, if another agent of clause C_i has already deviated to c_i, or to c_s because $c_s \rhd c(\overline{L_j^i})$ if $c_s \neq c(\overline{L_j^i})$.

This completes the proof. □

In order to summarize, for a deviation of Z to occur, it needs that exactly one agent in C_s deviates. The same holds for the deviator in C_s: she must have seen exactly one deviation from C_{s-1}, and so on. Consequently, at least one agent has deviated per clause. When Z finally deviates to candidate y, only the agents in C_1 can deviate in consequence, beginning a new phase of deviations within the clauses. The agents L_j^i who can come back to candidate y are those who have deviated alone in their clause C_i and for who the opposite agent $\overline{L_j^i}$ has not deviated, otherwise they do not believe that candidate y has enough points to win. These deviations still progress by increasing order of the clause indices. In order to get a cycle, since some agents of the clauses have come back to a strategy performed before the deviation of agent Z, Z must deviate again. Thus, these coming backs to old strategy y for agents of the clauses must affect each clause until C_s. Finally, if agent Y has not previously deviated, then we necessarily obtain convergence. Therefore, agent Y has previously deviated to candidate z as a believed best response. Consequently, agent Z comes back to candidate z and then Y to candidate y after having observed it, creating a cycle. Hence, by setting true the literals associated with the agents of the clauses who belong to the cycle, we obtain a truth assignment satisfying all the clauses of $\mathcal{C}$. □

Consequently, characterizing the instances for which the local dynamics can cycle, according to both the social network and the preferences of the agents, is a challenging issue.

5.3.2 Global poll-confident dynamics

The global poll-confident dynamics is not guaranteed to converge, even when the associated local poll-confident dynamics always converges. Actually, there exist examples where the dynamics can cycle even if the social network is empty. For a general graph, it is difficult to know whether the global poll-confident dynamics can cycle, even for BBR_p-dynamics and convergent local dynamics.

Let us first state the possibility of a cycle in the global poll-confident BBR_p-dynamics, even when the network is empty and $p = (1, \ldots, 1)$.

Proposition 5.8 *The global poll-confident BBR_p-dynamics may cycle even when G is empty and for a homogeneous pivotal threshold $p = (1, \ldots, 1)$.*

Proof: Let us consider an instance with four voters and four candidates, where $N = \{1, 2, 3, 4\}$, $M = \{a, b, c, d\}$ and $a \rhd b \rhd c \rhd d$. The social network is empty and the preferences of the agents are as follows.

$$\begin{aligned} 1:&\quad c \succ a \succ d \succ b \\ 2:&\quad d \succ a \succ b \succ c \\ 3:&\quad a \succ c \succ d \succ b \\ 4:&\quad d \succ b \succ c \succ a \end{aligned}$$

The pivotal thresholds are homogeneous and all equal to one, i.e., $p = (1, \ldots, 1)$. Here is a cycle of global poll-confident BBR_p-dynamics, occurring within two global steps (the polls are given in bold at the beginning of the lines, the deviations are in bold and $\text{BS}_i^{t,\setminus i}$ is mentioned above the arrow of a deviation made by agent i):

$$[\mathbf{1{,}0{,}1{,}2}]\ \{(1,0,1,2):d\} \xrightarrow[1]{(1,0,0,2)} \{(\mathbf{2},0,0,2):a\} \xrightarrow[3]{(0,0,1,2)} \{(1,0,\mathbf{1},2:d)\}$$

$$[\mathbf{1{,}0{,}1{,}2}]\ \{(1,0,1,2):d\} \xrightarrow[1]{(0,0,1,2)} \{(0,0,\mathbf{2},2):c\} \xrightarrow[3]{(1,0,0,2)} \{(\mathbf{1},0,1,2:d)\}$$

□

Interestingly, from the previous example, it appears that the most common cases of cyclicity in global poll-confident dynamics seem to reproduce the basic occurrences of cycles when voters deviate simultaneously.

However, by definition of the dynamics, the global poll-confident dynamics is equivalent to the local poll-confident dynamics when the graph is complete. In fact, in both cases, the voters have complete information about the vote of the other agents, thus the addition of new polls does not bring more information. Therefore, when the network is complete, the convergence results established for local dynamics still apply (consequence of Theorem 3 of Meir et al. [2010], and Proposition 5.4).

Corollary 5.9 *The global poll-confident BBR_p-dynamics converges from any initial profile within $\mathcal{O}(n \cdot m)$ steps when the network is complete and $p = (1, \ldots, 1)$.*

Corollary 5.10 *The global poll-confident BBR_p-dynamics converges from any initial profile within $\mathcal{O}(n \cdot m)$ steps, when the pivotal thresholds are homogeneous and the network is complete, for any instance $\mathcal{I}_\tau$ of the linked voting game where τ always chooses opportunity moves before compromise moves,*

For a general graph, we can prove that it is NP-hard to know whether the global dynamics can cycle, even for homogeneous pivotal thresholds $p = (1, \ldots, 1)$ and Plurality.

Theorem 5.11 *Deciding whether the global poll-confident BBR$_p$-dynamics can cycle is* NP*-hard even for homogeneous pivotal threshold* $p = (1, \dots, 1)$.

Proof: We perform a reduction from 2P1N-SAT, known to be NP-complete (Theorem 1.3). The 2P1N-SAT problem is a satisfiability problem defined on a set $\mathcal{C} = \{C_1, \dots, C_s\}$ of s clauses over a set $X = \{x_1, \dots, x_v\}$ of v variables. Recall that in 2P1N-SAT, each variable occurs twice as a positive literal and once as a negative literal. Each clause $C_i \in \mathcal{C}$ contains s_i literals. We assume that the clauses are indexed in such a way that each first occurrence of a variable is negative (see Section 1.5.1.a).

From an instance $\langle \mathcal{C}, X \rangle$ of 2P1N-SAT, let us construct an instance $\mathcal{I} = \langle N, M, \succ, \mathcal{F}_\rhd, G \rangle$ of the linked voting game, where $\mathcal{F}$=Plurality, as follows. The set M of candidates is $\{a_1, \dots, a_s, c_1, \dots, c_s, y, c\}$ where each c_i corresponds to a clause C_i for $i \in [s]$, and the tie-breaking is: $y \rhd a_s \rhd \cdots \rhd a_1 \rhd c_s \rhd \cdots \rhd c_1$. The set N of agents is $\{A \cup Y \cup \{L^i_j : i \in [s], j \in [s_i]\}\}$ where $A = \{A_1, \dots, A_s\}$ and $Y = \{Y_2, \dots, Y_s\}$ refer to the clauses, and L^i_j represents the j^{th} literal of clause C_i. The preferences are the following (the dots mean that the rest of the candidates are ranked arbitrarily):

$$\begin{array}{llll} A_1: & a_1 \succ c_1 \succ y \succ c \succ [\dots] & L^i_j: & c_i \succ y \succ a_i \succ c \succ [\dots] \\ A_i: & a_i \succ y \succ c \succ [\dots] & Y_i: & y \succ c_i \succ c \succ [\dots] \end{array}$$

Denote by q the maximum number of times a candidate is ranked first at this point, i.e., $q = \max\{\max_{i \in [s]}\{s_i\}, s\}$. Dummy voters are added in order to rank every candidate first exactly q times, leading to the election of candidate y under Plurality, thanks to $\rhd$. In the social network $G = (N, E)$, all the dummy voters are isolated. All the literal-agents L^i_j associated with a same clause C_i are connected. If an agent L^i_j $(i > 1)$ corresponds to a negative literal, then she points to all the agents of C_{i-1}, otherwise she points to all the agents of the clause containing its associated opposite literal except its opposite literal itself. The agents of clause C_1 are all linked to agent Y_s, and all the agents C^i_j are linked to Y_i if $i > 1$. Each agent Y_i points to Y_{i-1} for $i > 2$, and agent Y_2 to A_1. Agent A_i points to all L^{i-1}_j, and Y_i points to A_i for all $i > 1$. See Figure 5.2 for an illustration of the graph construction.

We claim that $\mathcal{C}$ is satisfiable if and only if the global poll-confident BBR$_p$-dynamics, starting from the truthful profile where $p = (1, \dots, 1)$, can cycle in instance $\mathcal{I}$.

Observe that initially only agent A_1 has incentive to deviate, since y is winner, and she deviates to c_1. Only Y_2 is able to observe this deviation and deviates herself to c_2. All agents Y_i deviate thereafter in the increasing order of their indices because agent Y_{i+1} can observe the deviation of Y_i $(i \in \{2, \dots, s\})$ and $c_{i+1} \rhd c_i$. No other agent has incentive to deviate because only agents L^i_j $(j \in [s_i])$ can also observe the deviation of Y_i, and the deviation to candidate c_i suits to them. Once agent Y_s has deviated, the agents of clause C_1 have incentive to deviate. By construction, an agent L^i_j has incentive to deviate if one of these two conditions holds: (1) she is associated with a negative literal and one agent in the previous clause has deviated, which is initially the case thanks to the deviation of someone in C_1 after the deviation of Y_s, or (2) she is associated with a positive literal and one agent in the clause of her opposite agent has deviated, and this deviation is not performed by the opposite agent herself: the agent observes all

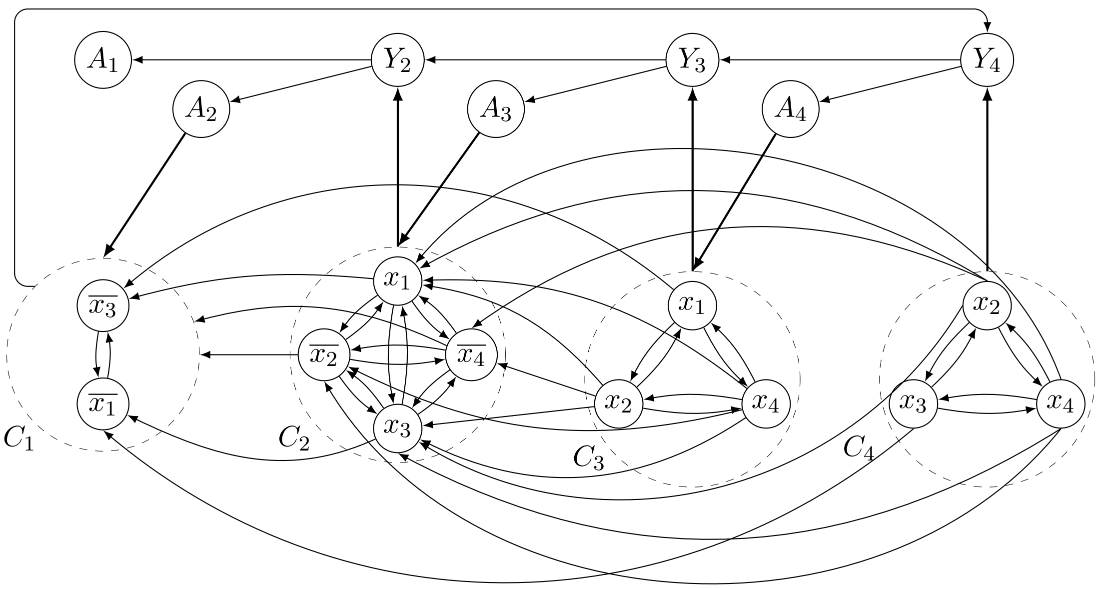

Figure 5.2: Construction of G for a 2P1N-SAT instance where $C_1 = (\overline{x_1} \vee \overline{x_3})$, $C_2 = (x_1 \vee \overline{x_2} \vee x_3 \vee \overline{x_4})$, $C_3 = (x_1 \vee x_2 \vee x_4)$ and $C_4 = (x_2 \vee x_3 \vee x_4)$. The literals (we keep for the figure the literal name) in a same circle are in a same clause; the arcs from or to the circles concern every agent inside.

the agents of this clause except her opposite agent. If an agent of clause C_i deviates, her best-response is not y but a_i, despite of her ranking of preferences, because she has previously observed the deviation of Y_i and thus does not believe that y could win. Once a voter deviates within clause C_i, the other agents of the clause do not incentive to do it because they are all connected and share the same preferences, but agent A_{i+1} has incentive to deviate to y because she still believes that it can win. This last deviation suits to agent Y_{i+1} who observes that. Remark that no dummy voter deviates because they all still believe that y is winner and y is in their two preferred candidates.

Suppose that $\mathcal{C}$ is satisfiable and let φ be a truth assignment of the variables satisfying $\mathcal{C}$. After the deviation of A_1 and all agents Y_i, we choose to make deviate for each clause, by increasing order of indices, an agent associated with a literal true in φ. Since there is no two opposite true literals in φ, every chosen agent at her turn has incentive to deviate since she is aware of the previous deviation, or the deviation of some agent in the clause of her opposite agent. Each chosen agent L_j^i deviates to a_i and then agent A_{i+1} deviates to y for all $i \in [s-1]$. There is a deviating agent in each clause because all the clauses are satisfied by φ. After all these deviations, we obtain a stable state (regarding the local dynamics) where every candidate has x points and thus y still wins. For the next global step, every agent who previously deviated, except the agents of Y, have incentive to deviate to their preferred candidate. If the agents L_j^i first deviate and then the agents of A, then the agents of Y have incentive to deviate, and by deviating in decreasing order of their indices, we obtain voting profile σ^0. Hence, we get a cycle.

Suppose now there is no truth assignment of the variables that satisfies $\mathcal{C}$. Consider a local dynamics from σ^0. As already mentioned, A_1 and then all the agents of Y must deviate. Thereafter, some agents of the clauses C_i deviate. Since $\mathcal{C}$ is unsatisfiable, all the clauses cannot be satisfied without setting true two opposite literals. It follows that there exists a clause C_i in instance $\mathcal{I}$ for which there is no deviating agent L_j^i because L_j^i cannot observe the deviation of another agent $L_{j'}^{i'}$. For the clauses for which there is a deviating agent L_j^i, the associated agent in A deviates. To summarize, the dynamics converges after the deviation of all agents of Y, of only some agents of A, and of one literal-agent for only some clauses. Therefore, the new winner is not y because in the best case, where only C_s is not satisfied, it obtains x points whereas there exists at least one candidate with $x + 1$ points, for instance a candidate associated with a clause for which there is no deviating agent. Consequently, at the next global step, all the dummy voters, except those who rank first the new winner, deviate to candidate c as a direct best response if y has less than x points or to candidate y if y has x points. It creates a too big gap within the scores to permit a further deviation in the global poll-confident dynamics, and thus the global poll-confident dynamics converges. □

The previous result implies that we cannot characterize efficiently the instances for which the global dynamics can cycle, when both the preferences of the agents and the social network are taken into account. In general, global dynamics appear less predictable than the local ones since the only case of convergence that we have identified is on complete graphs when the associated local dynamics is guaranteed to converge for such graphs.

5.4 Experimental analysis of the quality of the game

We present some experiments over 10,000 generated instances with 100 voters and 10 candidates, where the preferences are drawn from impartial culture. The social network is either a random directed Erdös and Rényi [1959]'s graph (see Section 1.2.2.b), or a random Barabási and Albert [1999]'s graph (see Section 1.2.2.b), or a graph generated following the protocol described in Section 1.2.2.b in order to incorporate homophily. The Erdös-Rényi graphs are generated with different densities, and enable to observe the impact of the number of links in the network. The Barabási-Albert graphs are realistic for representing real networks because for instance they are scale-free. The graphs with homophily seem to be realistic too because they model the fact that agents are more likely to be connected to agents with similar preferences.

We observe the frequency of convergence of the poll-confident BBR$_p$-dynamics for Plurality, as well as the quality of the equilibria that are reached by the dynamics, according to different measures, such as the Condorcet efficiency, the Borda closeness (see Section 1.3.2) or the Veto-SE efficiency (see Section 2.4.2). The results are given according to the density of the graphs (for Erdös-Rényi graphs) or the type of graph, and the pivotal thresholds. We examine homogeneous pivotal thresholds of value 1, 5 and 10 and heterogeneous ones uniformly distributed over the voters with values in [1..5] or [1..10].

5.4.1 Convergence in practice

Let us first analyze the experimental convergence of the poll-confident dynamics. During the experiments, we stop the iterative process only when an equilibrium is reached or a cycle is hit. We provide experimental results on the frequency of convergence of the poll-confident BBR_p-dynamics and the number of steps that are needed in average in order to reach convergence (for global dynamics we only mention the number of global steps).

5.4.1.a Impact of the density of the network

We present the experimental results of the frequency of convergence for the poll-confident BBR_p-dynamics with Erdös-Rényi graphs in Figure 5.3 and the number of steps before convergence in Table 5.1.

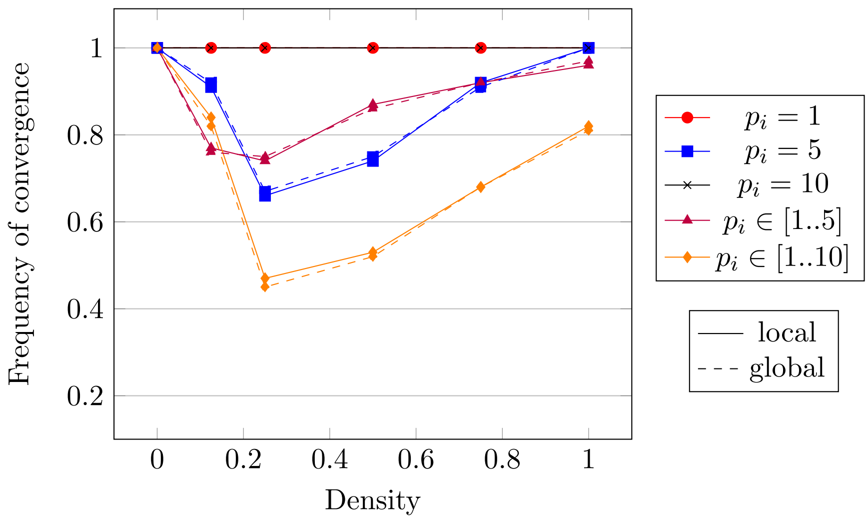

Figure 5.3: Frequency of convergence of the local and global BBR_p-dynamics for different pivotal thresholds and densities of network

The dynamics, both in their local and global versions, mostly converges, especially for sparse or dense graphs. When $p = (1, \dots, 1)$ (marked as $p_i = 1$ in the figure) or $p = (10, \dots, 10)$ (marked as $p_i = 10$ in the figure), they almost always converge, as opposed to the other thresholds for which we observe less convergent profiles, in particular when the density is around 0.25. This can be explained by the fact that when $p = (10, \dots, 10)$, most voters stay with their sincere ballot, and when $p = 1$, only a few candidates are potential winners. So, in both cases, the deviations are rather limited. This is not the case for heterogeneous thresholds and $p = (5, \dots, 5)$ (marked as $p_i = 5$ in the figure), because the agents can have very different best strategies. Moreover, the peak of non convergence for a density equal to 0.25 can be explained by the fact that the voters have enough information to deviate several times but this information is too partial to have a clear idea of the votes of the other agents, inducing an important bias in the deviations. Observe that in general the local and the global dynamics behave in the same way.

Our experimental results confirm the theoretical convergence results: when the graph is empty, the local dynamics always converges whatever the pivotal thresholds, and when $p = (1, \dots, 1)$, convergence is guaranteed for a complete graph. Moreover, the experimental results conform with our conjecture (sentences after Corollary 5.5) on the convergence of the local dynamics when the threshold is homogeneous and the graph is complete.

Density		$p_i = 1$	$p_i = 5$	$p_i = 10$	$p_i \in [1..5]$	$p_i \in [1..10]$
0	local	16.267	33.4263	8.5	29.1516	23.5802
	global	1.3987	2.9296	2.1460	2.3058	3.8208
0.25	local	15.6490	60.3313	9.2158	53.8849	66.6598
	global	1.4631	3.0609	2.0045	3.1577	4.5012
0.5	local	14.6245	61.1990	10.3078	49.2208	68.7952
	global	1.4729	3.5665	1.8741	3.7713	4.0875
0.75	local	14.5280	60.3270	10.6567	49.4782	64.8349
	global	1.4658	3.3960	1.7759	3.6381	3.6997
1	local	25.614	65.226	9.9492	64.4258	74.4068
	global	1.3662	1.9309	1.5823	1.9381	1.9990

Table 5.1: Number of steps before convergence in average for poll-confident BBR_p-dynamics for different pivotal thresholds and densities of network

Regarding the number of steps before convergence, globally no more than 75 local steps or 5 global steps are needed. This is not a huge number for an instance of 100 voters. The most important number of local steps is for $p = (5, \dots, 5)$ and $p_i \in [1..10]$, while the lowest number of local steps is for $p = (10, \dots, 10)$. Moreover, there are more local steps for $p_i \in [1..5]$ than for $p = (1, \dots, 1)$. These observations confirm our explanations for convergence. Indeed, very few steps for $p = (10, \dots, 10)$ show that only a small part of the agents has incentive to deviate. According to the definition of a believed best response deviation under Plurality, this means that for a large part of the voters, their preferred candidate is a potential winner. The number of local steps is slightly more important when $p = (1, \dots, 1)$ because the sincere ballot of the voters is much less likely to be a potential winner, so some voters must deviate to their preferred candidate within the set of potential winners. But this set is small, restricting the number of deviations. The set of potential winners should be larger for $p = (5, \dots, 5)$, $p_i \in [1..5]$ and $p_i \in [1..10]$, explaining that the number of local steps is larger.

In accordance with Observation 5.2, stating that the voters deviate at most once when the network is empty, the number of local steps is very limited (and clearly smaller than the number of agents) for such a network. Globally, the number of local steps is more important when the graph is complete. This appears natural because, in the poll-confident dynamics, the voters can deviate again, after a previous deviation, only if they have observed other deviations in the meantime. The voters can observe all the deviations when the network is complete, which is not the case for sparser graphs. However, for a density between 0.25 and 0.75, there is no configuration where the voters perform significantly more or less deviations.

5.4.1.b Realistic networks

So far, we have investigated the impact of the density of the network on the convergence of the dynamics. Let us now focus on other types of graphs, namely Barabási-Albert graphs and graphs with homophily, which are graphs with characteristics of real social networks.

The results are presented in Figure 5.4 (bars are displayed instead of curves because the results involve different types of graphs), and Table 5.2.

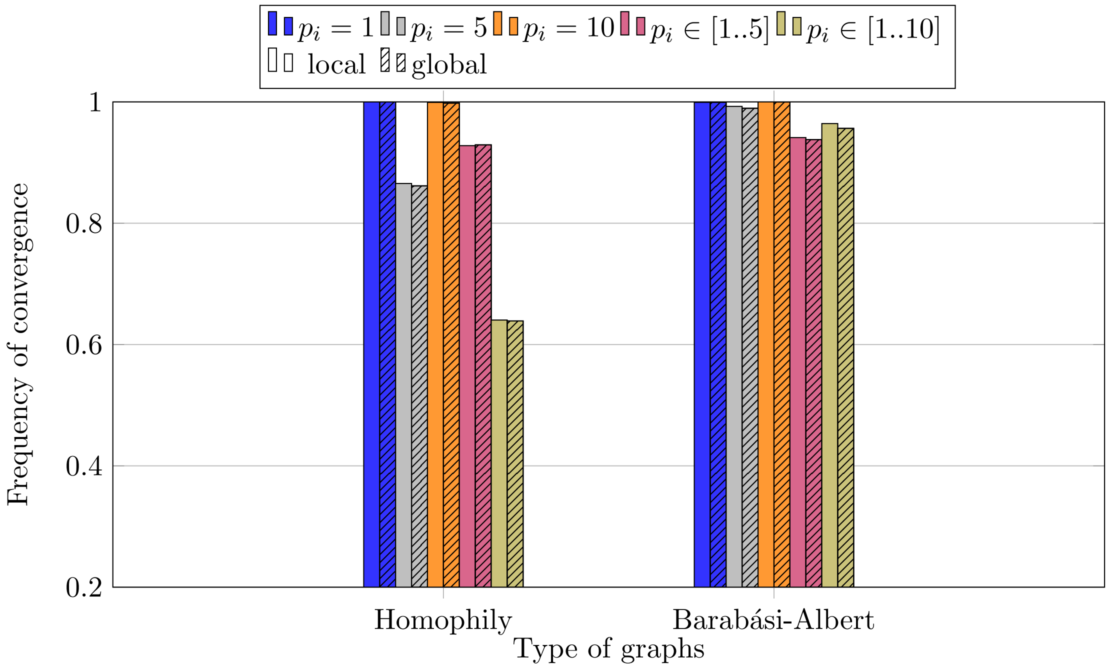

Figure 5.4: Frequency of convergence of the local and global BBR$_p$-dynamics for different pivotal thresholds and specific graphs

As for Erdös-Rényi graphs, the frequency of convergence with local and global dynamics is almost the same. The convergence rate is extremely high in Barabási-Albert graphs, where convergence is almost always ensured. In graphs with homophily, this is more conditioned by the pivotal thresholds. The behavior is the same as for Erdös-Rényi graphs. The lowest convergence rate is found for heterogeneous pivotal thresholds $p_i \in [1..10]$ which is around 60%. Then, the frequency of convergence is higher for homogeneous pivotal thresholds $p = (5, \ldots, 5)$, with a significant improvement since the convergence rate is above 80%, and even better are the heterogeneous pivotal thresholds $p_i \in [1..5]$ since the convergence rate is above 90%. Finally, for homogeneous thresholds $p = (1, \ldots, 1)$ or $p = (10, \ldots, 10)$, convergence is almost always ensured. The different behaviors of the dynamics with respect to the pivotal thresholds can be explained with the same arguments as in Erdös-Rényi graphs, related to the size of the set of potential winners.

One could think that the behavior of the dynamics in these graphs can be explained by their density. Indeed, Barabási-Albert graphs have a very low density (less than 0.1 in average), and by referring to Figure 5.3, the random graphs with such a density have a very good convergence rate. However, this explanation is not sufficient because the graphs with homophily that we have generated have a density around 0.5 but have a clear better convergence rate than the random graphs with density equal to 0.5 (see Figure 5.3). Actually, the behaviors of the dynamics in such graphs are closer to those of random graphs with density 0.75. This can be explained by the fact that the possibility of observing agents with similar preferences prevents basic cycles in the dynamics, such as the cycle in the example of the proof of Proposition 5.8, where two agents both deviate to the initial ballot of the other. In such a way, the connection to agents with similar preferences enables more coordination in the deviations, and thus improves the convergence rate of the dynamics.

Graph		$p_i = 1$	$p_i = 5$	$p_i = 10$	$p_i \in [1..5]$	$p_i \in [1..10]$
Homophily	local	14.4745	57.0317	9.5113	46.2714	66.2875
	global	1.5554	4.0458	1.8949	4.1642	4.5737
Barabási-Albert	local	15.8440	33.7376	8.4743	31.9387	25.2765
	global	1.3811	2.9536	2.1344	2.2792	4.0007

Table 5.2: Number of steps before convergence in average of poll-confident BBR_p-dynamics for different pivotal thresholds and specific graphs

Let us now analyze the number of steps before convergence. As for Erdös-Rényi graphs, generally, the number of steps does not exceed 70 local steps and 5 global steps. We do not detail the behavior with respect to the pivotal thresholds which is the same as for random graphs with different densities. As for convergence, Barabási-Albert graphs behave like very sparse graphs, which can be explained by their density which is very low. Following the same idea, the numbers of steps needed to reach convergence in graphs with homophily are close to those of random graphs with density in between density 0.25 and density 0.75 which are all more or less the same (see Table 5.1). However, it is noteworthy that the number of local steps in graphs with homophily is slightly smaller than in random graphs with these densities. This can be explained by the same argument as for convergence, by avoiding unnecessary deviations from agents who have already observed deviations that are convenient for them, performed by agents with similar preferences.

5.4.2 Quality of equilibria

We investigate the quality of the equilibria reached by the poll-confident BBR_p-dynamics under Plurality. The equilibria are evaluated through three criteria: the Condorcet efficiency, the Borda closeness and the Veto-SE efficiency.

5.4.2.a Impact of the density of the network

The results according to the Condorcet efficiency in Erdös-Rényi graphs are presented in Figure 5.5, where the frequency of electing the Condorcet winner (see Definition 1.21) is given with respect to different network densities and different pivotal thresholds for the agents. The preferences are drawn from impartial culture but with the restriction to the Condorcet domain (see Section 1.3.2), to be ensured that a Condorcet winner exists.

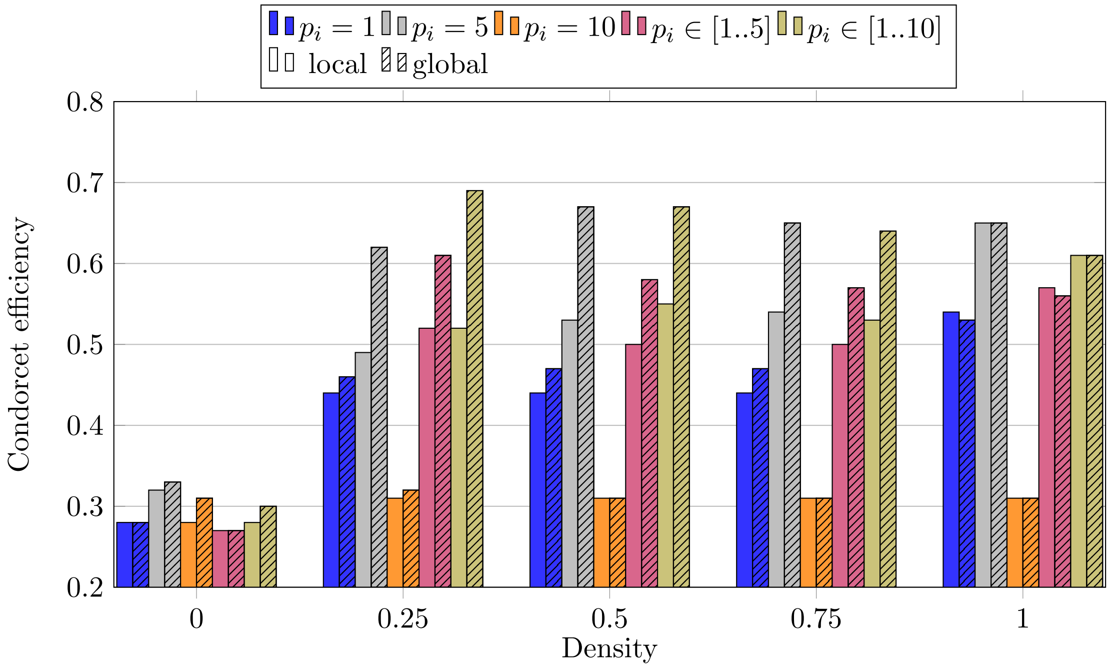

Figure 5.5: Condorcet efficiency of local and global BBR$_p$-dynamics for different pivotal thresholds and densities of network

Let us first recall that when the density is equal to 1, the local and global poll-confident dynamics corresponds to the case where the agents have complete information about the vote of the other agents.

The main observation is that the global dynamics always leads to better outcomes regarding the Condorcet efficiency. This appears natural since the voters regularly obtain more information about the current profile. Moreover, the frequency of election of the Condorcet winner seems to increase with the density of the network. This is clearly the case when $p = (1, \dots, 1)$, and for local dynamics for any threshold. The explanation is clear: the voters have more information and they can consequently perform more precise deviations. The positive point is that this precision does not only enable them to design better strategies for the next step, in a myopic way, but also strategies that benefit the whole electorate, since they lead to the election of the Condorcet winner more often.

However, for global dynamics and heterogeneous thresholds, it seems that the Condorcet efficiency of the equilibria is better when the density is around 0.25 and 0.5. This appears surprising but can be explained by the fact that this corresponds to the cases where the convergence of the dynamics is the lowest (see Figure 5.3). Yet, we analyze the quality of the equilibria only, by definition, when we reach an equilibrium. So, the results do not take into account the instances for which a cycle is reached. Therefore, it may be possible that we escape some critical instances for which it is difficult to reach a Condorcet winner.

The results when $p = (10, \ldots, 10)$ are relatively the same for the different network densities. This corroborates our previous explanation about the fact that when $p = (10, \ldots, 10)$, only a few voters deviate, starting from the truthful profile, because their preferred candidate is likely to be a potential winner.

Let us now analyze the Borda closeness of the equilibria. The results according to this criterion are presented in Figure 5.6.

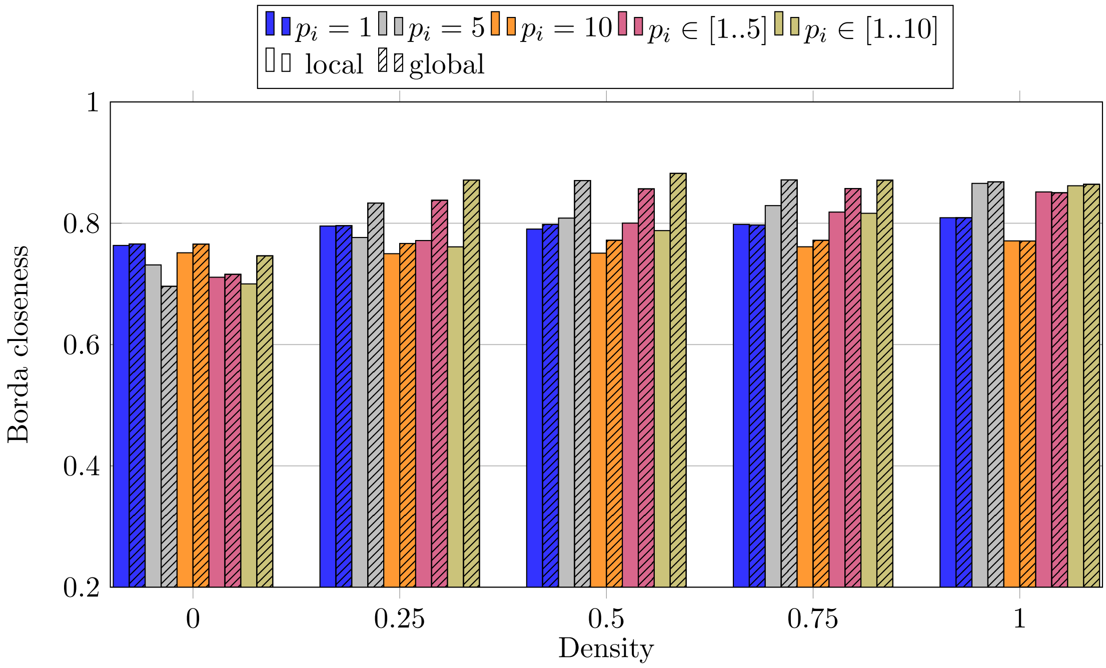

Figure 5.6: Average of Borda closeness for local and global BBR$_p$-dynamics for different pivotal thresholds and densities of network

The observations are similar to those provided for the analysis of the Condorcet efficiency. However, the differences regarding the quality of the outcomes on this criterion are less important. The main remark is that the global equilibria are better than the (Δ^0, BBR_p)-equilibria. The agents, thanks to the social network, have a local information all along the iterative process. Giving them a global information occasionally, via a public opinion poll, seems to improve the quality of the equilibria. This highlights the role of the information that the agents get in the quality of the equilibria. Following the same idea, the quality of the equilibria according to the Borda closeness increases with the density of the network. For instance, the Borda closeness is clearly poorer when the graph is empty.

The results concerning the Veto-SE efficiency are given in Figure 5.7.

The results on the Veto-SE criterion conform with the previous remarks for Condorcet efficiency and Borda closeness. In fact, global equilibria elect Veto-SE candidates more often and the frequency of electing a Veto-SE candidate increases with the density of

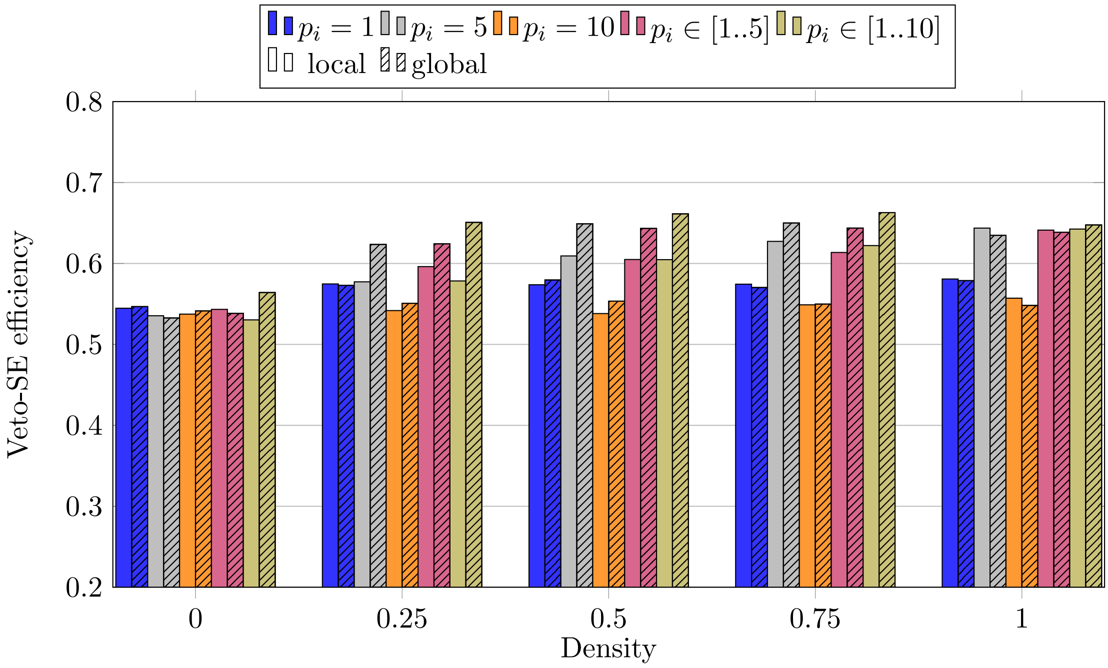

Figure 5.7: Veto-SE efficiency of local and global BBR_p-dynamics for different pivotal thresholds and densities of network

the network. This shows that, despite the fact that the three criteria we examine are conceptually different, the amount of information that the voters have always matters. Therefore, when the voters only have little information about the current profile, the bias in their deviations is important.

5.4.2.b Realistic networks

We now focus on more realistic networks by observing the quality of the reachable equilibria in Barabási-Albert graphs and in graphs with homophily. The results are presented in Figure 5.8 for the tree criteria used in the previous subsection, namely Condorcet efficiency, Borda closeness and Veto-SE efficiency.

Let us first observe the Condorcet efficiency of the dynamics. Generally, the frequency of electing a Condorcet winner is very low. Indeed, for Barabási-Albert graphs, this frequency does not exceed 40% and for graphs with homophily, the frequency is around 40% but with significantly better results for global dynamics, especially for homogeneous thresholds $p = (5, \dots, 5)$, and heterogeneous thresholds in [1..5] or in [1..10]. Like for convergence, the very low Condorcet efficiency in Barabási graphs can be explained by the low density of the graph. However, for graphs with homophily, the Condorcet efficiency for local dynamics is lower than in random graphs of density 0.5, which appears surprising. It seems that when agents are connected to agents with similar preferences, they fail to converge to an outcome acceptable for the whole society. It may be due to the fact that agents have a biased vision of the vote distribution and believe that the whole society have the same preferences as theirs. Nevertheless, the global dynamics in graphs with homophily perform clearly better than local dynamics

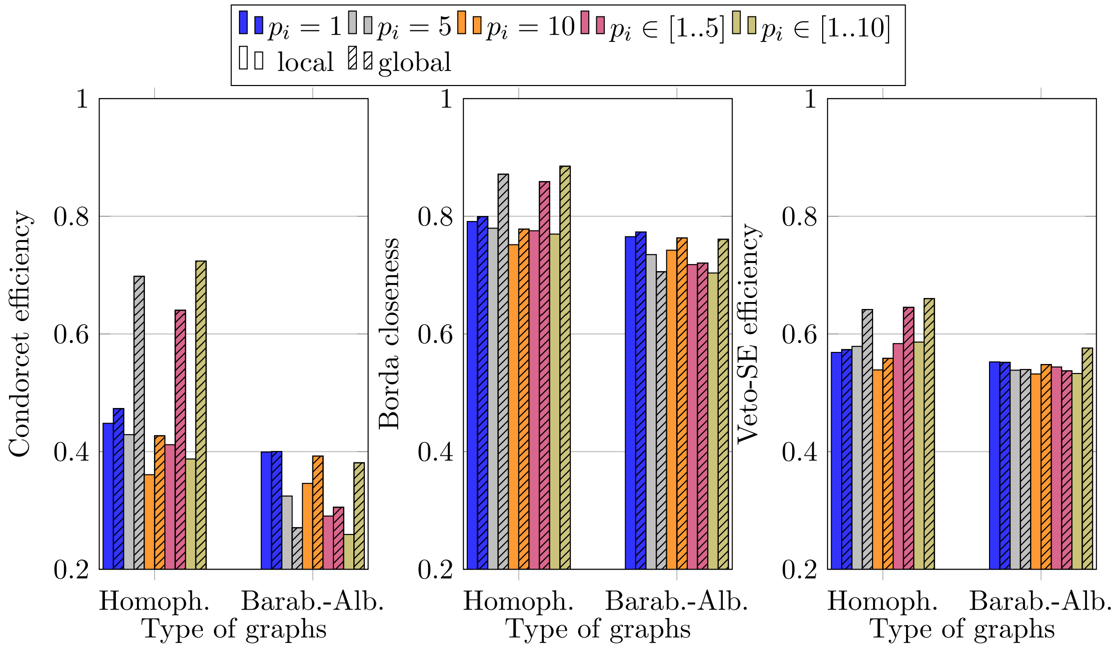

Figure 5.8: Quality of equilibria for local and global BBR$_p$-dynamics for different pivotal thresholds for graphs with homophily and Barabási-Albert graphs

and than global dynamics in random graphs of density 0.5. Therefore, if in addition of a local information from their connections in the graph, the agents regularly obtain a global information about agents with dissimilar preferences, then they are often able to reach an outcome acceptable for a large part of the society.

Regarding Borda closeness, Barabási-Albert graphs clearly behave like very sparse graphs and graphs with homophily like random graphs of density in between 0.25 and 0.75 (see Figure 5.6). The main observation for both graphs is the clear improvement in global dynamics compared to local dynamics.

Finally, the behavior of the dynamics for Veto-SE efficiency seems to be similar to the behavior of random graphs with corresponding density. Like the other criteria, global dynamics do better than local dynamics.

To summarize, graphs that are close to real social networks seem finally to mostly behave like random graphs of corresponding density, except graphs with homophily for some criteria. This highlights the bias in the deviations induced by the amount of information that agents have. Following the same idea, the quality of equilibria is clearly improved in global dynamics where agents regularly obtain a global information about the current state of the game. In order to eventually reach an acceptable outcome for the society, more than the identity or the preferences of the agents with who agents are connected, this is the number of connections which matters.

Actually, the quantity of information that the voters have from their successors in the social network, provided that these relatives are trusted agents, is essential in order to prevent election control from external agents, such as the polling institute.

5.5 Manipulation of the public opinion poll

In our model, the voters base their belief on the results of the poll. If we consider the polling institute as an agent who has her own preferences over the candidates, then the question of manipulation by the polling institute naturally arises. The polling institute as an agent is denoted by Π and her preferences are expressed via a linear order $\succ_\Pi$ over the candidates.

We focus on Plurality and assume in this section that all the voters have the same pivotal threshold p_u, i.e., $p = (p_u, \ldots, p_u)$, since the convergence results of the poll-confident BBR$_p$-dynamics under Plurality hold only under this condition. Furthermore, we make the assumption that the only possible manipulation moves for the agents are believed best responses.

Concretely, the manipulations of polling institute Π must be restricted in order to satisfy some likelihood conditions. For instance, Π could not announce that a candidate has no point if at least one voter has voted for it, otherwise this voter would know that the polling institute is lying. This credibility requirement is described more generally in the following definition.

Definition 5.5 (Likelihood condition) *The vector of scores Δ is a plausible communicated result for the poll on voting profile σ if $\Delta(x) \geq \max_{i \in N} Sc_i^\sigma(x)$, for every candidate $x \in M$.*

We assume that any manipulation performed by Π satisfies the likelihood condition. Let the *manipulation margin* be the number of points which are available, i.e., that polling institute Π is free to assign, after having fulfilled the likelihood condition.

We consider two different goals for manipulation from the polling institute. The first one is to make a precise candidate elected, whereas the latter is to obtain the best possible outcome at the next local equilibrium.

5.5.1 Enforcing the election of a candidate

We first ask whether it is possible for polling institute Π to enforce the election of a given candidate x.

Election Enforcing:

Instance:	Linked voting game instance $\langle N, M, \succ, \mathcal{F}_\triangleright, G\rangle$ with $\mathcal{F}$=Plurality, state $\sigma^t \in \mathcal{B_F}^n$, homogeneous pivotal thresholds $p = (p_u, \ldots, p_u)$, and candidate $x \in M$
Question:	Is there a poll score vector Δ such that the local poll-confident BBR$_p$-dynamics starting from state σ^t converges to a (Δ, BBR_p)-equilibrium electing x?

We will prove that Election Enforcing is computationally hard even when the social network is a DAG. We denote by $PW(\Delta)$ the set of potential winners announced by Δ, according to homogeneous pivotal threshold p_u. The set $PW(\Delta)$ includes all the candidates for which the addition of at most p_u votes is sufficient to win the election, according to Δ.

Although the poll-confident believed best response dynamics under Plurality appears very simple, it actually turns out that manipulating the opinion poll in order to enforce the election of a given candidate is computationally hard, even when the network is very simple and the voters deviate only if they are strictly pivotal, i.e., $p_u = 1$.

Theorem 5.12 ELECTION ENFORCING *is* NP*-hard even when the social network is a DAG and* $p = (1, \ldots, 1)$.

Proof: The reduction is from 3-SAT, a restriction of the satisfiability problem known to be NP-complete (Theorem 1.2). We consider an instance $\langle \mathcal{C}, X \rangle$ of 3-SAT composed of a set $\mathcal{C} = \{C_1, \ldots, C_s\}$ of s clauses over a set $X = \{x_1, \ldots, x_v\}$ of v variables, where each clause contains exactly three literals.

Let us construct a linked voting game instance $\mathcal{I}_\tau = \langle N, M, \succ, \mathcal{F}_\rhd, G, \tau \rangle$, where $\mathcal{F}$=Plurality, as follows.

We construct a set of agents N including: a set Y of s clause-agents Y_i for $i \in [s]$, sets Y' and L of $3s$ literal-agents Y_{ij} and L_{ij} for $i \in [s]$ and $j \in [3]$, a set L' of agents L'_{ijk} for $i \in [s]$, $j \in [3]$ and $k \in [3]$ referring to 3 copies of the literals, sets Z^1 and Z^2 (whose union is denoted by Z) of respective agents Z^1_i and Z^2_j for $i \in [4s-1]$ and $j \in [11]$, a set X of agents X_i for $i \in [11]$, and agents X'_1 and X'_2.

The set of candidates M contains the candidates z, x' and x, the set $C\ell$ of clause-candidates c_i for $i \in [s]$, and the set Lit of literal-candidates ℓ_{ij} for $i \in [s]$ and $j \in [3]$.

In the network G, there is an arc from Y_i to L_{ij} and from Y'_{ij} to L'_{ijk}, for all $i \in [s]$, $j, k \in [3]$. There is an arc from L_{ij} to $L_{i'j'}$ if $i > i'$ and the j^{th} literal of i^{th} clause is the opposite of the j'^{th} literal of i'^{th}. In the sets Z^2 and X, there is one agent, say respectively Z^2_1 and X_1, who is connected to all the agents in the set.

Let ρ be a linear order over $C\ell \cup Lit$ and ρ^{-1} its reverse order. For an agent L_{ij}, let $\{\overline{\ell_{ij}}\}$ be an arbitrary order over the set of literal-candidates $\ell_{i'j'}$ associated with literals opposite to ℓ_{ij} and such that $i' < i$. The preferences $\succ$ are as follows ($i+1 = 1$ if $i = s$):

$$
\begin{array}{rlrl}
Y_i: & z \succ c_i \succ x \succ \rho & Z^j_i: & z \succ \rho^{-1} \succ x \\
Y'_{ij}: & z \succ \ell_{ij} \succ x \succ \rho^{-1} & X_i: & x \succ z \succ \rho \\
L_{ij}: & c_i \succ \ell_{ij} \succ \{\overline{\ell_{ij}}\} \succ z \succ \rho^{-1} \succ x & X'_1: & x' \succ z \succ \rho \\
L'_{ijk}: & \ell_{ij} \succ \ell_{i+1\ k} \succ z \succ \rho \succ x & X'_2: & x' \succ z \succ \rho^{-1}
\end{array}
$$

where ρ and ρ^{-1} are assumed to be defined, in the preferences, without the candidates already mentioned in the order. Let A_ρ and $A_{\rho^{-1}}$ be the sets of agents having, respectively, order ρ and ρ^{-1} in their preferences.

The turn function τ is such that $\{L, L'\} \succ_\tau \{Y, Y'\}$, and $L_{ij} \succ_\tau L_{i'j'}$ for $i' > i$. The tie-breaking rule $\rhd$ is such that $z \rhd x \rhd \ell_{11} \rhd \cdots \rhd \ell_{13} \rhd \cdots \rhd \ell_{s1} \rhd \cdots \rhd \ell_{s3}$.

We claim that all the clauses in $\mathcal{C}$ are satisfiable if and only if polling institute Π can enforce the election of candidate x in the local poll-confident BBR$_p$-dynamics starting from the truthful profile σ^0 and where $p = (1, \ldots, 1)$.

At σ^0, each L_{ij} and L'_{ijk} vote respectively for c_i and ℓ_{ij}, and agents Y_i and Y'_{ij} respectively observe it. So, by the likelihood condition, each candidate $y \in C\ell \cup Lit$ must have at least 3 points in the poll, i.e., $\Delta(y) \geq 3$. Each agent in X votes for x and it is visible for X_1, thus $\Delta(x) \geq 11$. All the sets Y, Y' and Z vote for candidate z, but only agent Z_1^2 has a non-null out-degree for observing that, and thus $\Delta(z) \geq 11$, whereas the real score of z is $8s+10$. X'_1 and X'_2 vote for x' so $\Delta(x') \geq 1$. To summarize, the margin of manipulation for institute Π is $8s$. We can prove that the only strategy for Π to make x win is to assign 8 more points to one literal-candidate associated with each clause, and the new points cannot be given to literal-candidates corresponding to opposite literals. In this poll, z is announced winner.

Suppose there exists a truth assignment ϕ of the variables such that all the clauses are satisfiable. Let polling institute Π give in Δ, 8 more points to one candidate ℓ_{ij} associated to a literal true in ϕ chosen for each $i \in [s]$. Let us denote by Φ this set of chosen literal-candidates. Candidate z is done winner by Δ^0, thanks to the tie-breaking. The agents L_{ij} and $L'_{i-1\ k\ j}$ $(i-1 = s$ if $i = 1)$, for $k \in [3]$, have incentive to change their vote to candidate ℓ_{ij}, that they prefer to z, if ℓ_{ij} is able to win. This is the case if $\ell_{ij} \in \Phi$. Such deviations bring their predecessor agents Y_i and $Y_{i-1\ k}$, for $k \in [3]$, to deviate in consequence to candidate x because x has one point less than ℓ_{ij} but it is better in the tie-breaking. Since ϕ satisfies all the clauses, all the agents in $Y \cup Y'$ see one agent in $L \cup L'$ deviating and then they all deviate to x. Thus, candidate x obtains $4s+11$ points, whereas candidate z only gets $4s+10$ points, and so x becomes the new winner.

Suppose now that Π is able to enforce the election of x. Denote by ω the winner given by Δ. Observe first that every agent necessarily deviates to a candidate in $PW(\Delta)$. This is trivially the case for the first deviation by definition of a believed best response under Plurality. Moreover, since we start from σ^0 that is truthful, the voters actually supporting ω do not deviate, therefore the winner score does not decrease. Let us suppose that this property holds until the k^{th} step. If at the $(k+1)^{\text{th}}$ step, an agent deviates to $y \notin PW(\Delta)$, then the score of the winner has previously decreased since by induction assumption the score of y did not increase. But it is impossible since all the voters deviated to candidates in $PW(\Delta)$ by assumption and $p = (1, \ldots, 1)$, contradiction.

The previous observation implies that x should belong to $PW(\Delta)$, and that $PW(\Delta) \setminus \{\omega\} \neq \emptyset$, otherwise nobody deviates.

- Let us first suppose that $\omega \in C\ell \cup Lit$ and $|PW(\Delta) \setminus \{\omega\}| = 1$. Then, the only possibility is to give points to x (or ω if at a moment x wins over ω and an agent can observe it). Therefore, no agent from $L \cup L' \cup Z \cup X$ deviates because all these agents prefer ω to x, or they already voted for x (in the case of agents in X). All the agents of $Y \cup Y'$ deviate to x, except one, the agent Y_i or Y'_{ij} associated with $\omega = c_i$ or $\omega = \ell_{ij}$, and also exactly one agent between X'_1 and X'_2 deviates to x (they do not have both incentive to deviate since, except for x' and z, they have reversed preferences). Therefore, x can only obtain $4s+11$ points while z still has $4s+11$ points. Consequently, x cannot win because z has a better position in the tie-breaking rule.
- Suppose now that $\omega \in C\ell \cup Lit$ and $|PW(\Delta)| > 1$, or that $\omega \in \{x, x'\}$. The agents in A_ρ or $A_{\rho^{-1}}$ (except potentially $Y \cup Y'$) can agree on a candidate $y \in PW(\Delta)$ to which they deviate. They are too numerous to enable the agents $Y \cup Y'$ to elect x, therefore x cannot win.
- Hence, ω must be candidate z. No agent in Z, nor in $X \cup \{X'_1\} \cup \{X'_2\}$ deviates because z is the first of second choice of these agents and they cannot observe the deviations of agents in other sets. Moreover, the agents in L and L' will never

deviate to x since this candidate is their last choice. Thus, all the agents in $Y \cup Y'$ must deviate to x in order to make x obtain more votes than z and so win. This implies that at least one agent among $\{L_{ij} : j \in [3]\}$ for all i must deviate to candidate ℓ_{ij}, and at least one agent among $\{L'_{ijk} : j \in [3]\}$ for all $i \in [s], j \in [3]$, must deviate to candidate $\ell_{i+1\ k}$ (or ℓ_{1k} if $i = s$). Thus, one candidate related to a literal of a clause must belong to $PW(\Delta)$ for each clause. They need to obtain 8 points in order to belong to $PW(\Delta)$ (because $\Delta(z) = 11$) and thus all the $8s$ available points are used. There cannot be two candidates ℓ_{ij} and $\ell_{i'j'}$ in $PW(\Delta)$ with $i < i'$ corresponding to two opposite literals because otherwise $L_{i'j'}$ would have observed the deviation of L_{ij} to ℓ_{ij} since $L_{ij} >_\tau L_{i'j'}$, and thus $L_{i'j'}$ could not deviate then to $\ell_{i'j'}$ since she believes that ℓ_{ij} is the winner and $\ell_{ij} \rhd \ell_{i'j'}$. Note that agent $L_{i'j'}$ believes that candidate z is a potential winner but has no incentive to deviate to it because she prefers literal-candidate ℓ_{ij} to z. Hence, by setting to true all the literals associated to the candidates ℓ_{ij} favored by Δ, we obtain a truth assignment of the variables satisfying all the clauses.

□

Beyond the computational difficulty of manipulating the opinion poll in the worst-case, we explore a heuristic perspective. We design a very simple heuristic and test its efficiency by simulations. For all the experiments of this section, we run 10,000 instances of 100 voters and 10 candidates, where the preferences are generated from impartial culture, and the graphs are randomly generated via Erdös and Rényi [1959]'s model for different densities. We have also tested our heuristics on graphs supposed to be closer to real social networks, such as Barabási-Albert graphs or graphs with homophily. However, we do not present results on these types of graphs because the results are similar to those for random graphs with corresponding density.

Our heuristic approach for Election Enforcing is based on the combination of two algorithms, namely Algorithms 5.1 and 5.2. Algorithm 5.1 constructs a poll score vector where, as much as possible, there are only two best candidates in the poll. The intuition is that, if Π wants to make a target candidate x elected then, as much as the manipulation margin allows, she predicts as the winner a "specter" candidate y, that Π wants to show as a threat to the voters, and x as a potential winner. More precisely, after having fulfilled the likelihood condition (l. 2-4), points are added to x until it becomes a potential winner (l. 5-7). Then, we rise the score of y until one more point to y would remove x from the set of potential winners (l. 8-10), in order to quickly increase the gap with other candidates while keeping x in $PW(\Delta)$. At this point, y is the winner in Δ and x is "at the limit" of $PW(\Delta)$. Now we simultaneously increase the scores of x and y (l. 11-14) in order to further eliminate other candidates from $PW(\Delta)$. If one point remains in the end, we assign it to x if that does not make it the winner, or otherwise to the last ranked candidate in Δ by safety (l. 15-20). The process stops if at some step the margin is not sufficient.

Algorithm 5.2 is a heuristic based on Algorithm 5.1 where we choose, as a specter candidate y, the most disliked candidate in the preferences of the agents compared to the target candidate x. By this way, we aim at creating a situation where only two candidates, x and y, are favorites in the election. One of them, the "specter" y, is mostly disliked by the population but is announced the winner with a slight lead over

Algorithm 5.1: Margin rebalance between two candidates

Input: Instance $\mathcal{I} = \langle N, M, \succ, \mathcal{F}_{\triangleright}, G\rangle$ with $\mathcal{F}$=Plurality, state σ^t, pivotal threshold p_u, target $x \in M$, specter $y \in M \setminus \{x\}$

Output: Δ: communicated score results of the poll from σ

$margin \leftarrow n$;
foreach $z \in M$ **do**
  $\Delta(z) \leftarrow \max_{i\in N} \mathrm{Sc}_i^t(z)$; `# likelihood condition`
  $margin \leftarrow margin - \Delta(z)$;
while $margin > 0$ *and* $x \notin PW(\Delta)$ **do**
  $\Delta(x) \leftarrow \Delta(x) + 1$; `# make` x `a potential winner`
  $margin \leftarrow margin - 1$
while $margin > 0$ *and* x *remains in* $PW(\Delta)$ *if* $\Delta(y)$ *is increased by 1* **do**
  $\Delta(y) \leftarrow \Delta(y) + 1$; `# make` y `the winner while`
  $margin \leftarrow margin - 1$; x `remains a potential winner`
while $margin > 1$ **do**
  $\Delta(x) \leftarrow \Delta(x) + 1$; `# widen the gap between` x, y
  $\Delta(y) \leftarrow \Delta(y) + 1$; `and the other candidates`
  $margin \leftarrow margin - 2$;
if $margin = 1$ **then**
  if $p_u > 1$ **then**
    $\Delta(x) \leftarrow \Delta(x) + 1$; `# last point to` x `if it does not make` x `win`
  else
    $last \leftarrow \arg\min_{z\in M} \Delta(z)$;
    $\Delta(last) \leftarrow \Delta(last) + 1$; `# otherwise to the last candidate in` Δ
return Δ;

x. One could think that, in reaction, a large part of the electorate will report her ballot to the other candidate x.

Algorithm 5.2: Heuristic for ELECTION ENFORCING

Input: Instance $\mathcal{I} = \langle N, M, \succ, \mathcal{F}_{\triangleright}, G\rangle$ with $\mathcal{F}$=Plurality, state σ^t, pivotal threshold p_u, target $x \in M$

Output: Δ: communicated score results of the poll from σ

Sort the candidates of $M \setminus \{x\}$ in increasing order of the number of voters who prefer x to the candidate ;
foreach *specter* $y \in M \setminus \{x\}$ **do**
  $\Delta \leftarrow$Algorithm 5.1$(\mathcal{I}, \sigma^t, p_u, x, y)$;
  if $\mathcal{F}(\Delta) = y$ *and* $x \in PW(\Delta)$ **then return** Δ;
return Algorithm 5.1$(\mathcal{I}, \sigma^t, p_u, x, \arg\max_{z\in M\setminus\{x\}} \max_{i\in N} \mathrm{Sc}_i^{\sigma}(z))$

We present some experiments in Figures 5.9 and 5.10, where the target candidate is either the Condorcet winner (restriction to the Condorcet domain), the Borda winner, the truthful winner (the winner of the truthful profile) or the best candidate in $\rhd$. The frequency of election of the target candidate is given in a context of poll manipulation via Algorithm 5.2 or no manipulation, both with global and local poll-confident BBR$_p$-dynamics. For global dynamics, the polling institute manipulates at each global step in order to make target candidate x elected at the equilibrium of the next local poll-confident BBR$_p$-dynamics. The voters are assumed to all have a pivotal threshold equal to $p = (1, \dots, 1)$ (Figure 5.9) or to $p = (5, \dots, 5)$ (Figure 5.10).

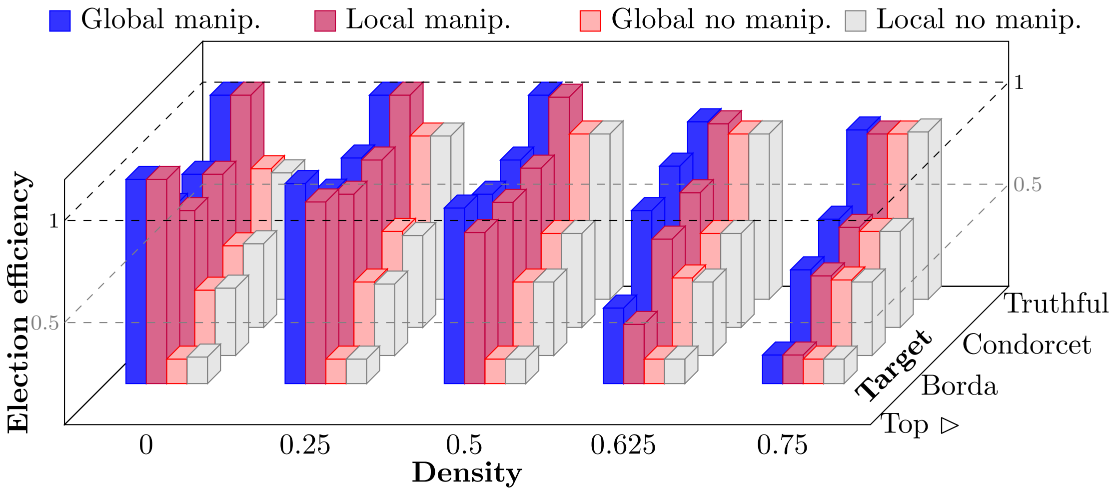

Figure 5.9: Algorithm 5.2 in local and global poll-confident BBR$_p$-dynamics where $p = (1, \dots, 1)$ for different densities of network

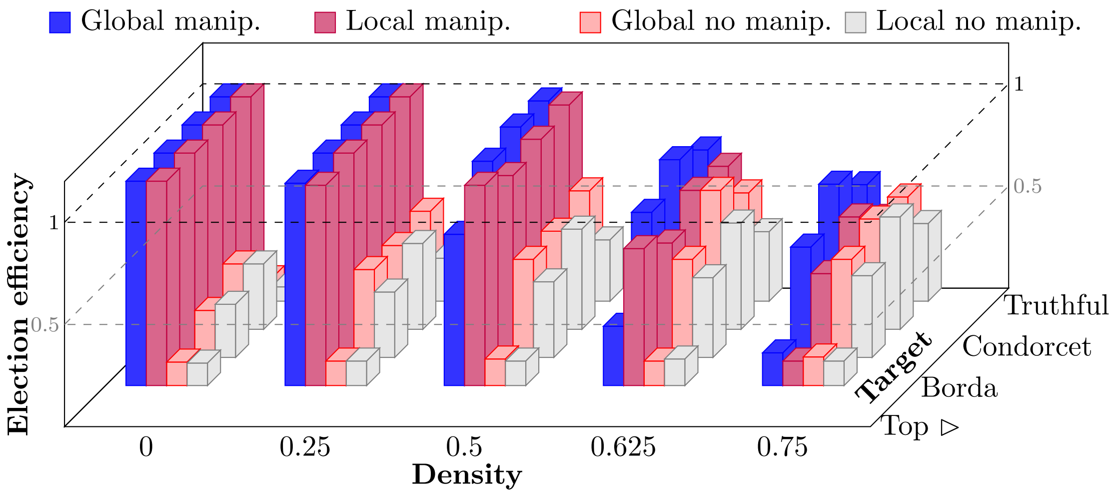

Figure 5.10: Algorithm 5.2 in local and global poll-confident BBR$_p$-dynamics where $p = (5, \dots, 5)$ for different densities of network

We can observe that this heuristic is very efficient, especially on sparse graphs. In fact, from a density equal to 0 to a density equal to 0.5, the frequency of the election of the target candidate is at least twice the frequency with no manipulation by the polling institute. The gap is even larger when $p = (5, \ldots, 5)$. This can be explained by the fact that, when $p = (1, \ldots, 1)$, the manipulation margin must be large enough to give to x and y, scores with a difference of one point, which can be difficult to achieve in Algorithm 5.1. However, for $p = (5, \ldots, 5)$, a gap of 5 points is allowed between the scores of the candidates x and y for x being a potential winner. So, this condition can be easier to fulfill in Algorithm 5.1, leading to better final results for electing the target candidate when $p = (5, \ldots, 5)$.

The gap between manipulation from the polling institute and no manipulation is particularly important when the target candidate is the top candidate in the tie-breaking rule, i.e., candidate a. This candidate is favored in case of ties in the outcome of the election but a priori, contrary to the Condorcet winner or the Borda winner or even the truthful candidate, a is not advantaged in the preferences of the voters. In fact, with no manipulation, the probability for a to be elected is around 0.12. Therefore, the profit of manipulation for this candidate is more visible than for the other target candidates since, in addition, the manipulation of the polling institute is very efficient in this case (here the advantage of a in $\rhd$ matters).

However, in general, when the density of the network increases, the benefit of manipulation for the polling institute is less significant. Indeed, the gap between the frequency of electing the target candidate, with manipulation from Π, compared to this frequency without manipulation, decreases in the same time as the density of the network increases. This is due to the decrease of the manipulation margin when the density of the network increases, because of the likelihood condition. Actually, when the network is a complete graph, the manipulation margin is equal to zero and thus there is no possible manipulation for the polling institute.

5.5.2 "Best response" dynamics of the polling institute

From a different perspective, instead of enforcing the election of a specific candidate, the polling institute Π could try to perform better responses at each global step, with immediate benefits in the associated local dynamics, in the same myopic spirit as a voter in iterative voting. In other words, Π manipulates at each global step in order to make elected her best possible candidate at the end of the next local dynamics. The optimization problem associated with the computation of such a myopic strategy for the polling institute is called Poll Best Response Computation.

Poll Best Response Computation:

Instance:	Linked voting game instance $\langle N, M, \succ, \mathcal{F}_{\rhd}, G\rangle$ with $\mathcal{F}$=Plurality, state σ^t, homogeneous pivotal threshold $p = (p_u, \ldots, p_u)$, preference ranking $\succ_\Pi$ over M
Problem:	Find a poll score vector Δ such that the local poll-confident BBR_p-dynamics starting from σ^t converges to a (Δ, BBR_p)-equilibrium electing the best possible candidate with respect to $\succ_\Pi$

As stated in Theorem 5.12, Poll Best Response Computation is computationally hard for the polling institute, even when the network is a directed acyclic graph, otherwise one could recognize positive instances of Election Enforcing.

Consequently, like for Election Enforcing, we adopt a heuristic approach by restricting the manipulation of Π to a simple move exposed in Algorithm 5.1: trying to favor only two candidates. We derive Algorithm 5.3 where a manipulated poll score vector Δ is built by Algorithm 5.1 for every ordered pair of candidates (x, y), and the associated sequence of local deviations is tested. We choose the poll score that leads to an equilibrium, in the tested local dynamics, electing the best candidate for Π.

Algorithm 5.3: Heuristic for Poll Best Response Computation

Input: Instance $\mathcal{I}_\tau = \langle N, M, \succ, \mathcal{F}_\triangleright, G, \tau\rangle$ with $\mathcal{F}$=Plurality, state σ^t, pivotal threshold $p = (p_u, \ldots, p_u)$, $\succ_\Pi$, acyclic or transitive spanning subgraph G' of G

Output: Δ: communicated score results of the poll from σ

```
foreach target x ∈ M do
    foreach specter y ∈ M \ {x} do
        Δ_y ← Algorithm 5.1(I, σ^t, p_u, x, y);
        R(y) ← F((Δ_y, BBR_p)-equilibrium) in instance I_{G'};
return Δ_y for which R(y) is the best in ≻_Π ;
```

However, in order to efficiently test the sequence of local deviations, we restrict the network to the cases where the dynamics is guaranteed to converge after a polynomial number of steps, namely the DAGs and the transitive graphs (Proposition 5.1 and cases of Corollaries 5.3 and 5.5). We thus consider a spanning subgraph G' of the network that is either empty, acyclic or transitive. Obviously, we would like to have an acyclic or transitive spanning subgraph that is as close as possible to the real network G. Unfortunately, the decision version of Maximum Acyclic Subgraph (i.e., the problem Feedback Arc Set of deciding whether there exists a subset of arcs of size k containing at least one arc from each dicycle in the graph) is an NP-complete problem [Karp, 1972], as well as the decision version of Maximum Transitive Subgraph [Yannakakis, 1978]. We consequently use simple classical approximate algorithms for computing these subgraphs. More precisely, for computing a spanning acyclic subgraph, we use a simple $\frac{1}{2}$-approximate algorithm [Korte, 1979] that works as follows: let $A_<$ be the set of arcs $(i, j) \in E$ such that $i < j$ and $A_>$ be the set of arcs $(i, j) \in E$ such that $i > j$, we return $G' := (N, \arg\max\{|A_<|, |A_>|\})$. For computing a spanning transitive subgraph, we actually use an algorithm computing a *maximal* transitive subgraph in $\mathcal{O}(n^3)$ [Chakraborty et al., 2015]: for every agent i and every agent j successor of i, if there exists some agent $\ell \neq j$ such that $(i, \ell) \notin E$, then remove arc (j, ℓ) if present in E, in addition if there exists some agent $\ell \neq i$ such that $(\ell, j) \notin E$, then remove arc (ℓ, i) if present in E.

Note that the local poll-confident BBR_p-dynamics that we test in the algorithm always converges, even in global dynamics where the initial state of the associated local dynamics may not be truthful and may be conditioned by the false informations of the poll. This is due to the fact that Proposition 5.1 as well as Corollaries 5.3 and 5.5 establish the convergence of the local dynamics from any initial state (for pivotal thresholds greater than one, we use a turn function satisfying the condition given in Corollary 5.5). Moreover, every agent believes in the same initial state coming from the manipulation of the polling institute, since the polling institute Π fulfills the likelihood condition in her manipulation.

We apply Algorithm 5.3 on local dynamics where Π can only manipulate the initial poll, as well as on global dynamics, where Π can manipulate the results of the poll at each global step. Observe that in this context, Observation 5.1 does not hold anymore, that is a global equilibrium may not be a Nash equilibrium even when $p = (1, \ldots, 1)$. Indeed, the information given by the poll may not correspond to the current state. However, there is no reason to have less convergent profiles than in the case where the polling institute is sincere because the manipulation of Π is driven by her preferences, represented as a linear order over the candidates.

The results given by Algorithm 5.3 are presented in Figure 5.11 when $p = (1, \ldots, 1)$ and in Figure 5.12 when $p = (5, \ldots, 5)$. The figures show the average of the rank in $\succ_\Pi$ of the final winner when the dynamics converges (lower is better). Note that the experimental convergence of the dynamics is not affected by the manipulation from the polling institute. Indeed, the frequency of convergence that we obtain is similar to the results given in Figure 5.3.

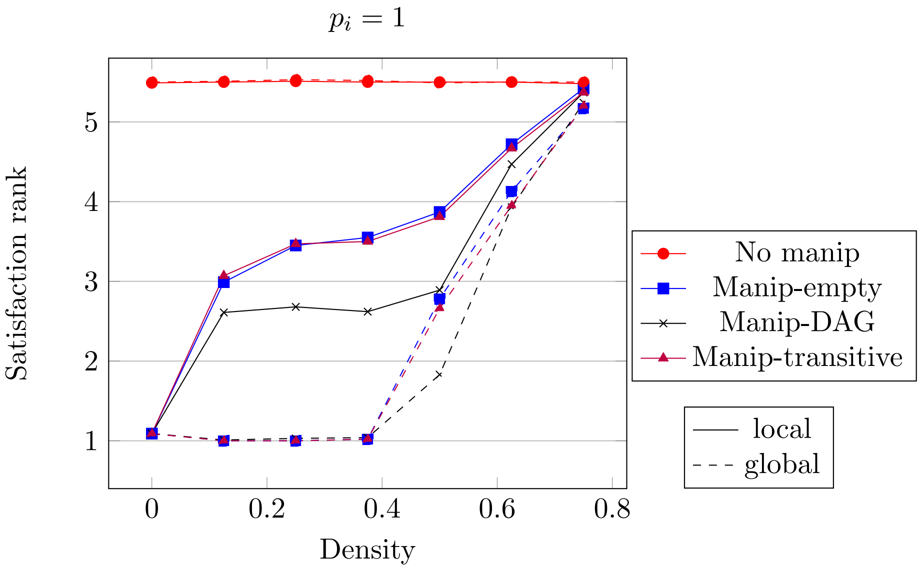

Figure 5.11: Algorithm 5.3 in local and global poll-confident BBR_p-dynamics where $p = (1, \ldots, 1)$ for different densities of network

The results are good for polling institute Π. For both types of dynamics, the rank with manipulation is clearly better than without manipulation, with a large gap for sparse graphs. The gap decreases with the increase of the density. The rank of the winner in the preferences of Π is clearly lower for global equilibria than for local equilibria, especially for sparse graphs until density 0.4. This implies that this is profitable for the polling institute to manipulate and communicate the results of several polls because she has more opportunities to manipulate, and this enables her to further influence the voters towards a specific direction.

The gap between the rank of the local equilibria and the global equilibria is less important when $p = (5, \ldots, 5)$ because, in this case, the quality of the local equilibria

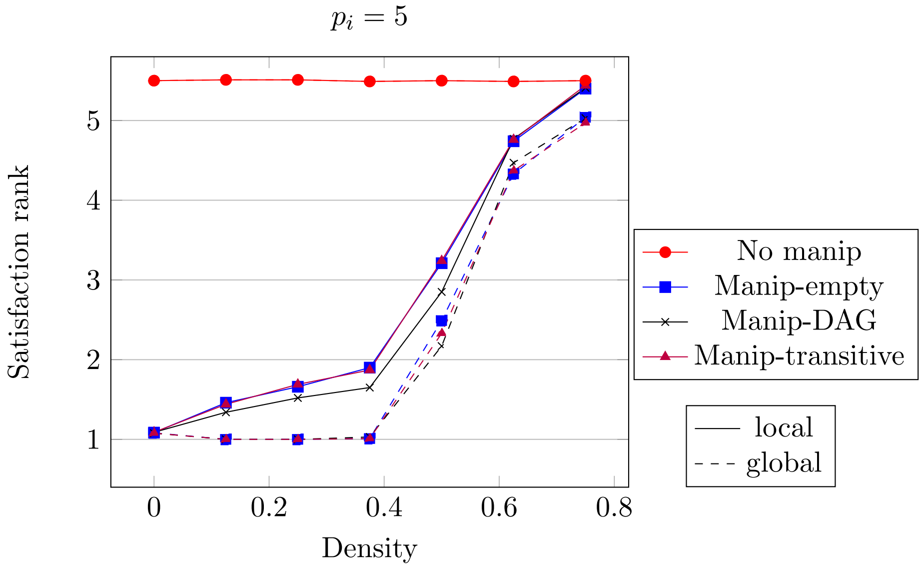

Figure 5.12: Algorithm 5.3 in local and global poll-confident BBR_p-dynamics where $p = (5, \dots, 5)$ for different densities of network

according to the preferences of Π is better. This is due to the fact, like in Algorithm 5.2, that the polling institute needs less manipulation margin to make the target candidate a potential winner when $p = (5, \dots, 5)$.

Like Algorithm 5.2, the results deteriorate with the increase of the density. This is related to the likelihood condition, making the manipulation margin decreasing with the increase of the density, since the agents have more information about the current profile. Indeed, because the agents are more informed in denser graphs, it becomes more difficult to mislead them. In fact, when the network is a complete graph, no manipulation is possible for Π since every agent has complete information about the current profile.

5.6 Concluding remarks

We have studied a best response dynamics in iterative voting where the voters aggregate the informations from opinion polls and social networks, and adopt a strategic behavior conditioned by pivotal thresholds. We have shown the convergence of the dynamics for some classes of graphs, notably the directed acyclic graphs and the transitive graphs, but in general it is difficult to recognize instances with cycles. Nevertheless, it turns out that the dynamics mostly converges in practice. The quality of the equilibria depends on the density of the network: better outcomes are found in dense graphs (there is more information). The equilibrium analysis allows to underline the bias produced by partial information and the dependency on the information sources, raising the question of election control.

Actually, a manipulation of the opinion poll can be hard to compute, even for simple sparse graphs. However, simple heuristics, based on the idea of announcing a “specter” candidate (which is particularly disliked) as the winner, are very efficient. Nevertheless,

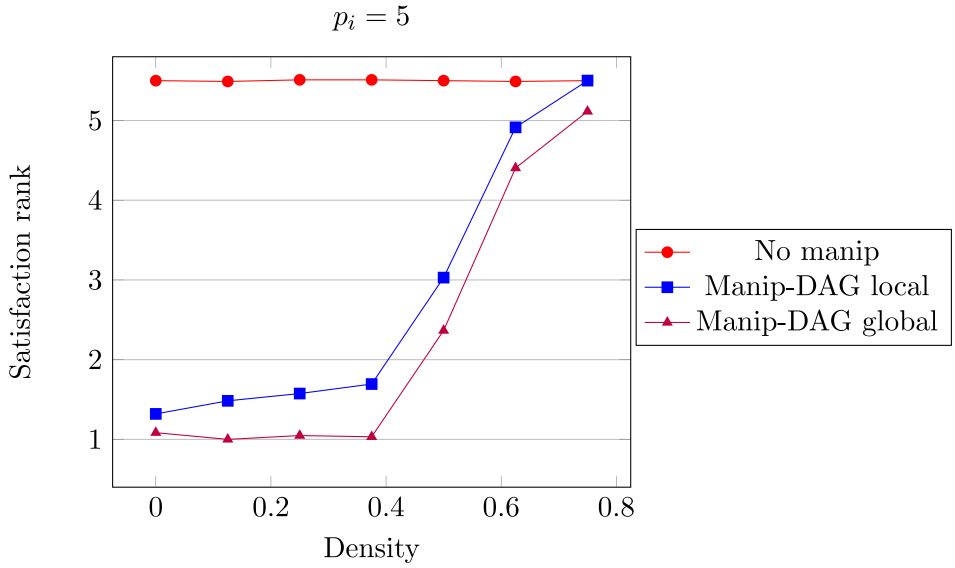

Figure 5.13: Algorithm 5.3 with an acyclic spanning subgraph G' and homogeneous thresholds $p = (5, \ldots, 5)$ on an instance where $p_i \in [1..10]$

it is more difficult to manipulate for the polling institute in dense graphs, where the knowledge of the voters is close to be complete. These heuristics are not too demanding regarding the knowledge of the network structure (we use subgraphs of the network), but they need to know the preferences and the pivotal thresholds of the agents. Nevertheless, assuming that the poll institute knows the preferences of the voters is not extravagant since the original task of the polling institute is to collect them. Moreover, we have conducted some experiments (see Figure 5.13) where the polling institute computes her strategy with thresholds that do not correspond to the real ones. Indeed, we suppose that the pivotal thresholds are uniformly distributed over the electorate with values in [1..10], but that the polling institute does not know these values. The polling institute then computes her strategy using Algorithm 5.3 by assuming that the pivotal thresholds are homogeneous and all equal to 5. Concretely, one could base on psychological studies in order to determine realistic pivotal thresholds to use in the algorithm. Even with a false assumption about the thresholds of the population, the results displayed by Figure 5.13 are very good for the polling institute and are similar to those obtained when the thresholds used in the algorithm coincide with the real thresholds of the agents.

Classical extensions for the model of poll-confident dynamics are the study of other voting rules, coalitional manipulation or more sophisticated heuristics. Actually, concerning the extension to other voting rules, we investigated the adaptation to the Veto rule, which is natural since the ballots are composed of only one candidate (the only disapproved candidate), like Plurality. A best response deviation based on vetoing the believed winner can be defined, which generalizes the direct best response (Definition 1.27) for Veto when the pivotal thresholds are all equal to 1 and the network is a complete graph. Based on this specific manipulation move, it is possible to adapt most of the proofs given in this chapter. More precisely, it is possible to prove the convergence of this dynamics under Veto from any initial profile, and for any pivotal thresholds even not homogeneous, when the network is a transitive graph, in the spirit of Obraztsova et al. [2016]'s Theorem 4. Furthermore, the Election Enforcing problem is also NP-hard for this dynamics, even when the social network is a DAG, in the same idea as Theorem 5.12. However, for the sake of clarity, we decided not to include the Veto rule in this chapter, because the interpretation of the deviations is less natural than for Plurality and the results are similar, except that the practical convergence rate of the dynamics is very low, especially for global dynamics.

Considering less classical extensions, one could think about voters who keep in memory the previous steps in the global dynamics, or relaxing the assumption of a poll on the entire electorate. The links with opinion diffusion in networks can also be examined. Indeed, in our model, the voters only get some information from their successors in the graph, but it can be natural to think that some voters try to influence their relatives to vote in a certain way.

Conclusion

We have presented several works dealing with classical social choice problems, in which interaction among agents is constrained by an accessibility relation, represented by a social network, which is a graph structure over the agents. We have especially studied how the possibility of collaboration or getting some information from the other agents can be limited according to the links of the social network. Our study is focused on two major problems in computational social choice: strategic voting and house allocation.

Summary of the contributions

Let us present in more detail the scope of our contributions.

In Chapter 2, we have explored coalitional manipulation in iterative voting by groups of voters that form fully connected components of the social network. Our main results show the existence of a considerate equilibrium for several well-known voting rules. A considerate equilibrium is a state immune to deviations from coalitions given by cliques of the social network, where in addition an altruistic condition is added regarding the neighbors of the clique in the network (consideration assumption). The goal of our usage of this solution concept was to find a trade-off between Nash equilibria and strong equilibria, based on the definition of realistic coalitions. Indeed, in a voting game, Nash equilibria always exist but can be numerous and thus irrelevant for capturing a plausible outcome of the game, whereas strong equilibria rarely exist (note that the Veto rule is an interesting special case where the existence of a strong equilibrium is guaranteed). However, the existence of a considerate equilibrium can be proved thanks to the consideration assumption or the fact that there is no coalition of size at least half of the voters, which implies that in fact many states are considerate equilibria, obtaining the same drawback as Nash equilibria. Nevertheless, the considerate equilibria that are reached by best response dynamics of the voting game are experimentally better, according to the quality of the outcome, than other weaker solution concepts like for instance the Nash equilibrium. Therefore, despite the large number of considerate equilibria, dynamics of the game can naturally filter acceptable outcomes for the whole society, by considering cliques of the network as possible coalitions and altruism among agents. Consequently, exploiting the social relations among the agents in coalitional manipulation enables to refine the game-theoretical analysis of a voting game. In fact, by defining realistic coalitions that are altruistic, the social network allows the iterative voting process to converge towards better plausible stable states.

We have continued to investigate some collaborative aspects modeled by a social network in Chapter 3 where, in a housing market, the only possible trades that can be performed from an initial allocation involve two connected agents in the graph. Surprisingly, this simple model of resource reallocation leads to computational hardness results, for predicting the allocations that can emerge from a dynamics of swaps between connected agents, even for social networks whose shape is relatively simple. The problems that we study are natural questions related to the reachability of certain objects for the agents and thus they should be useful both in the design of fair protocols and in the analysis of the dynamics of swaps. Nevertheless, for very simple structures of graphs, or some realistic social networks for modeling a possibility of collaboration, such as bounded degree graphs, and realistic bounds on the number of exchanges, we can provide tractable cases for these reachability questions. Hence, the social network, as a representation for the possibility of collaboration, has a strong impact on the computational properties of housing market, while it enables to model more realistic configurations.

In another perspective, we have investigated the impact of the social network on strategic voting and house allocation problems, where the network captures the information that is available to the agents. In Chapter 4, the social network represents the possibility of envy between two connected agents in a resource allocation problem where exactly one item must be assigned to each agent. We have proved that deciding the existence of an allocation that is envy-free with respect to this local notion of envy is computationally hard. However, we were able to provide tractable cases for very dense graphs or acyclic graphs, as well as approximation algorithms computing good allocations regarding the minimization of local envy. Moreover, the likelihood of existence of a locally envy-free allocation is high when the agents can only envy agents with preferences that are very different from theirs. This is a natural assumption in the sense that agents with similar preferences are more likely to have a close relationship that prevents envy between them. In contrast, the existence of an envy-free allocation in the standard meaning, without graph, is very rare in house allocation. Therefore, the social network enables to expand the possibility to design fair allocations regarding a realistic fairness criterion, while it strongly impacts the computational ability of constructing such allocations.

Finally, in Chapter 5, the social network is used as a local informative tool for strategic agents in iterative voting. In such a context, the voters manipulate according to a local information they have about the vote of their connections in the network and a global, but possibly out of date, information given by a public opinion poll. We have defined a specific dynamics of deviations, based on this strategic behavior, which is conditioned by the social network and the opinion polls. The dynamics of deviations is guaranteed to converge to a stable state for some restricted classes of graphs that can make sense in practice, and experiments even stress a large convergence rate of the dynamics for more general graphs. Furthermore, we have experimentally shown a basic intuition that the more the network is dense and the voters have information about the current state, the higher the quality of the outcome and the lower the possibility of manipulation for the polling institute. In such a context, the social network has a strong impact on the outcome of the procedure since the information that it provides influences the strategic behavior of the agents. Nevertheless, the social network enables to highlight the bias induced by partial information in strategic voting, which is actually present in real elections, because this is a configuration where voters, clearly, do not have complete information.

In a nutshell, the integration of a social network in social choice problems enables to consider more realistic scenarios. The outcome of the social choice procedures is deeply influenced by the structure of the social network and the computational properties of the models are also affected by the network topology. In fact, by assuming an accessibility relation over the agents that is not necessarily complete, many social choice procedures become intractable, as well as the satisfaction of desirable properties. However, this representation allows to focus on plausible interactions among the agents and thus to be more precise in the evaluation or the prediction of the outcomes of social choice processes.

Future work

All the questions we have explored have their own extensions. When talking about strategic voting, one could classically examine other voting rules, but the most interesting extension that we identify would be to combine our two approaches. Indeed, we have considered iterative voting by generalizing best response dynamics in two ways: by extending it to coalitions of voters given by the graph, or by relaxing the assumption of complete information to an information given by the network. Consequently, one could naturally consider a best response dynamics where connected agents in the social network could exchange information and join to collaborate in the elaboration of a strategy for manipulation. Regarding house allocation, it is possible to consider for instance initial partial allocations, or exchanges among more than two agents. Generalizing house allocation to a resource allocation problem with several resources per agent is also a direct extension, as well as other types of matching problems. Globally, for both settings, more attention should be paid on restricted domains for preferences such as single-peaked or single-crossing preferences. It is a natural extension for almost all the problems we have studied, and may lead to more positive results.

We do not claim that our work draws an exhaustive picture of the behavior of agents, in collective decision making procedures, whose interaction with other agents is conditioned by a social network. However, our goal was to point out some social choice problems for which modeling the possibility of interaction among agents by a social network makes sense. We especially focused on strategic voting and resource allocation but many other settings can be explored. By considering the swap dynamics model in house allocation where two connected agents can trade objects, one can think about models where the items to exchange are immaterial like information. In this case, the agents do not lose their previous “item” but progressively increase their knowledge along the exchanges. This is related to a famous problem in Artificial Intelligence, especially within the dynamic epistemic logic community, that is the *gossip problem* restricted to a communication network [Harary and Schwenk, 1974, Hedetniemi et al., 1988]. In such a problem, every agent has a piece of information, e.g., a secret, and can share this information with other agents by making phone calls. The goal is to find the minimum number of phone calls that are necessary to reach a state where every agent knows all the secrets of the other agents. Many variants of this problem exist. In the restriction involving a communication network, the phone calls are limited to connected agents in the graph. Concerning strategic voting, a natural model to study is *opinion diffusion in networks*. Indeed, while we explore strategic voting via coalitional manipulation and possibility of information that are given by a social network, a relevant extension would be to study strategic voters who try to convince their relatives in the network to vote in a certain way.

Other research directions regarding the representation of the possibility of interaction can be investigated. So far, we have only considered a social network as a graph over the agents for modeling an accessibility relation among agents. One can think about other ways to model this possibility of interaction among the agents, like for instance a hypergraph, where the edges are actually subsets of agents, allowing to directly take coalitions of agents into account. A computational difficulty arises from such a representation because it cannot be encoded in a compact way. A compromise should be found between the relevance and the compactness of the representation in order to get an appropriate representation. Furthermore, we have often examined dynamic processes in which the agents collaborate or collect information within several steps. It is pertinent to consider that the social network itself is dynamic according to these different steps. For instance, new links can be added among the relatives of two connected agents after their interaction.

Appendix

A Proof of equivalence in the reduction of Theorem 3.7

Theorem 3.7 RO-*max is* W[SAT]-*hard in a tree.*

Proof: Suppose first that there exists an assignment ϕ of weight k which satisfies formula φ. For each variable x_j in X set to true in ϕ, we decide to move the associated object z_j^{φ} by increasing order of the indices. Each such object needs to be "duplicated" at each agent R_i^0, representing a binary relation, into two objects $z_j^{\psi_{i,1}}$ and $z_j^{\psi_{i,1}}$ representing a copy of object z_j associated with the specific subformula.

The gadget of object duplication works as follows for agent R_i^0. Initially in the passage of z_j^{φ}, when object z_j^{φ} arrives to agent R_1^0, it first needs that agent R_1^0 gets object r_1^{j-1}. This is initially the case if $j = 1$. In that case, the swap can be performed. Otherwise, i.e., if $j > 1$, this can be done by exchanging with agent R_1^{j-1} who only accepts objects from $\{\psi_{i,2}^{j'} \cup r_i^{j'} : j' < j-1\}$. As further explained in the following, agent R_1^0 owns such an object if she has completely duplicated each previous object $z_{j'}^{\varphi}$ which passes by it where $j' < j$. Once agent R_1^0 gets object z_j^{φ}, object z_j^{φ} cannot go down in any member of the binary relation associated with R_1^0 because each of these agents does not prefer z_j^{φ} to their initial object. However, it can be exchanged with agent $Z_j^{\psi_{1,1}}$ against object $z_j^{\psi_{1,1}}$ allowing then the swap of agent R_1^0 with agent $\Psi_{1,1}$ in order to give to her, her own version of object z_j. It is possible then that agents $Z_j^{\psi_{1,1}}$ and $Z_j^{\psi_{1,2}}$ swap their objects and then agents $Z_j^{\psi_{1,1}}$ with agent R_1^0, and then agent R_1^0 with agent $\Psi_{1,2}$, in order to pass object $z_j^{\psi_{1,2}}$ to the corresponding target agent $\Psi_{1,2}$. Observe that if this second part is not accomplished, then no further swap can involve agent R_1^0 because no other agent than $Z_j^{\psi_{1,1}}$ accepts object $\psi_{1,1}^{j-1}$. By assuming that all the duplications of object z_j^{φ} have been made until an agent R_i^0, allowing the passage of object $z_j^{r_i}$ until agent R_i^0, the duplication of object $z_j^{r_i}$ at agent R_i^0 works following the same principle as for agent R_1^0, by replacing in the swaps and the preferences object z_j^{φ} by the specific object $z_j^{r_i}$. The whole duplication is necessary if agent R_i^0 needs to perform further swaps, i.e., if binary relation associated with R_i^0 is satisfied in formula φ when considering assignment ϕ.

At a moment, a version of object z_j cannot be duplicated any more because the neighbor $\Psi_{i,1}$ (or $\Psi_{i,2}$) of agent R_i^0 is a variable-agent $Y_{j',\ell}^0$. In this case, once agent $Y_{j',\ell}^0$ gets object $z_j^{y_{j',\ell}}$, either $j' \neq j$ and she will simply exchange it with an agent $Y_{j',\ell}^t$ for allowing the passage of the next object $z_t^{y_{j',\ell}}$ for $t > j$, or $j' = j$ and the passage of object $z_j^{y_{j,\ell}}$ enables to activate / validate this agent, like the truth assignment ϕ assigning true to variable x_j satisfies the corresponding occurrence of x_j. The validation is kept in memory with the swap of agents $Y_{j,\ell}^t$ and $T_{j,\ell}^1$, and the swap of agents $T_{j,\ell}^1$ and $T_{j,\ell}$ leading to the acquisition of validation-object $t_{j,\ell}$ by agent $T_{j,\ell}^1$.

In summary, until now, we have moved k objects within $\{z_j^{\varphi} : 1 \le j \le v\}$, by increasing order of indices, and duplicated them at each agent corresponding to a binary relation (at least for those which are satisfied by assignment ϕ), allowing to keep in memory the validation of all the variable-agents for which the corresponding occurrence of the variable is true in ϕ. Since the acquisition and the duplication of an object for an agent corresponding to a binary relation costs 6 swaps (or 5 for z_1^{φ}), the maximum number of swaps per agent until now is $6k$ (or $6k-1$) swaps.

After the move of these k objects, it suffices to activate the validation gadget for each variable-agent or relation-agent satisfied by truth assignment ϕ. This validation process leads for an agent B to the acquisition of object t_b. Every such agent performs now a swap with her v^{th} copy. This swap is possible for every agent R_i^0 who completely accomplished the duplication of each object $z_j^{r_i}$. Concerning the variable-agents, this swap is directly possible. Thereafter, all these satisfied variable-agents $Y_{j,\ell}$ can swap with agent $T_{j,\ell}^1$ for getting object $t_{j,\ell}$. By going up progressively along the syntax tree T_{φ} from the leaves and considering the associated relation-agent R_i^0, then at least one (respectively, both) of the agents $\Psi_{i,1}$ and $\Psi_{i,2}$ get respectively the objects $v_{\psi_{i,1}}$ and $v_{\psi_{i,2}}$ if $R_i^0 \in \mathscr{O}$ (respectively, $R_i^0 \in \mathscr{A}$), i.e., the associated binary relation is a disjunction (respectively, a conjunction). Therefore, if $R_i^0 \in \mathscr{O}$, then R_i^0 can swap with one of the agents $\Psi_{i,1}$ and $\Psi_{i,2}$ to obtain $t_{\psi_{i,1}}$ or $t_{\psi_{i,2}}$, which can be exchanged with agent T_i to obtain object t_i. Otherwise, that is if $R_i^0 \in \mathscr{A}$, then agent R_i^0 can swap with agent $\Psi_{i,1}$ to get object $t_{\psi_{i,1}}$ which then can be exchanged against t_i^1 with agent T_i^1. This object is the only one accepted by agent R_i^{v+1} to give object r_i^{v+1} to agent R_i^0, who needs it for agent $\Psi_{i,2}$ who can only accepts to give $t_{\psi_{i,2}}$ against r_i^v (already given to agent $\Psi_{i,1}$) or r_i^{v+1}. Thus, by performing the swaps between the agents T_i^1 and T_i, and then between T_i^1 and R_i^0, agent R_i^0 finally obtains object t_i. Observe that this validation process takes at most 6 swaps per agent (in case of a conjunction). One swap is added for the exchange with the predecessor agent.

Since truth assignment ϕ satisfies formula φ, after the validation phase, agent R_i^0 gets object t_1. After having previously exchanged her object with agent C', agent C can obtain object v_1 which is the only one that agent D accepts against object d. Finally, agent C obtains object d with at most $6k+7$ swaps per agent.

Suppose now that object d is reachable for agent C with at most $6k+7$ swaps per agent. Observe that agent D accepts to give object d only against object t_1. For that, it previously needs that agent R_1^0 herself obtains object t_1 (validation process). More generally, let us observe the conditions under which an agent R_i^0, modeling the i^{th} binary relation of formula φ, can obtain object t_i.

Firstly, let us consider the case where $R_i^0 \in \mathscr{O}$, i.e., the associated binary relation is a disjunction. Agent T_i accepts to give object t_i to R_i^0 only against object $t_{\psi_{i,1}}$ or object $t_{\psi_{i,2}}$, which are the validation-objects of the two agents $\Psi_{i,1}$ and $\Psi_{i,2}$ representing the two members of the disjunction. Hence, the only condition is that at least one of the agents $\Psi_{i,1}$ and $\Psi_{i,2}$ obtain respectively the objects $t_{\psi_{i,1}}$ and $t_{\psi_{i,2}}$, interpreting the fact that a disjunction is satisfied if at least one of its member is.

Now, suppose that $R_i^0 \in \mathscr{A}$, i.e., the associated binary relation is a conjunction. Object t_i must pass by agent T_i^1 to reach agent R_i^0 from its initial owner T_i. Agent T_i accepts to swap object t_i only against object $t_{\psi_{i,1}}$. Therefore, object $t_{\psi_{i,1}}$ must have previously passed by agent T_i^1 and by agent R_i^0 who must have obtained it via agent $\Psi_{i,1}$. Moreover, once T_i has accepted to swap object t_i with agent T_i^1, agent T_i^1 accepts to give object t_i only against object $t_{\psi_{i,2}}$. Therefore, object $t_{\psi_{i,2}}$ must have previously passed by agent R_i^0 who must have obtained it via agent $\Psi_{i,2}$. Hence, it is necessary that the agents $\Psi_{i,1}$ and $\Psi_{i,2}$ had previously obtained respectively the objects $t_{\psi_{i,1}}$ and $t_{\psi_{i,2}}$ to permit that agent R_i^0 obtains object t_i, reflecting the fact that a conjunction is satisfied if both of its members are satisfied.

Concerning a variable-agent $Y_{j,\ell}^0$, for obtaining object $t_{j,\ell}$, it needs that agent $T_{j,\ell}$ swap it with agent $T_{j,\ell}^1$, and then agent $T_{j,\ell}^1$ with agent $Y_{j,\ell}^0$. Agent $T_{j,\ell}$ accepts to swap object $t_{j,\ell}$ only against object $z_j^{y_{j,\ell}}$. Therefore, it firstly needs that object $z_j^{y_{j,\ell}}$ pass by agents $Y_{j,\ell}^0$ and $T_{j,\ell}^1$. This object corresponds to an occurrence of variable x_j. Hence, in summary, agent R_1^0 is validated, i.e., obtains object t_1, if there exist validated variable-agents such that the truth assignment of the variables setting to true the variables associated with the occurrences of the variables represented by these agents satisfies formula φ.

Now, the question is how each object $z_j^{y_{j,\ell}}$ is arrived until each validated variable-agent $Y_{j,\ell}^0$. Observe that initially object $z_j^{y_{j,\ell}}$ is owned by $Z_j^{y_{j,\ell}}$ who is only connected to agent $P(Y_{j,\ell}^0)$, the agent corresponding to the predecessor of this variable-agent in the syntax tree T_φ. Actually, object $z_j^{y_{j,\ell}}$ can go down to agent $Y_{j,\ell}^0$ only if $P(Y_{j,\ell}^0)$ has herself previously received object $z_j^{p(Y_{j,\ell}^0)}$ from agent $P(P(Y_{j,\ell}^0))$ who has herself previously received object $z_j^{p(p(Y_{j,\ell}^0))}$, and so on. By this way, it follows that the agent associated with the root of T_φ, R_1^0, must have received object z_j^φ, in order to permit a variable-agent $Y_{j,\ell}^0$ to receive object $z_j^{y_{j,\ell}}$. Agent R_1^0 can obtain object z_j^φ via agents C and Z_j^φ. Once she gets it, she is obliged to "duplicate" it in order to pass a version of z_j^φ to each of her neighbors $\Psi_{1,1}$ and $\Psi_{1,2}$ associated with the members of her related binary relation. If the complete duplication is not done by agent R_1^0 for every object z_j^φ which passes by her, then she cannot perform further swaps as explained more in details in the first part of the proof, and thus cannot be validated later by obtaining object t_1. Therefore, as detailed in the first part of the proof, for duplication of each object z_j^φ, agent R_1^0 uses 6 swaps (except for z_1^φ which costs 5 swaps). However, the maximum cost of a validation process is 7 swaps (as also mentioned in the first part of the proof) and occurs when the binary relation is a conjunction, as for agent R_1^0. Thus, since agent C obtains object d with at most $6k + 7$ swaps, agent R_1^0 can let pass only k objects within $\{z_j^\varphi : 1 \leq j \leq v\}$ (the gap of 1 swap given if z_1^φ is chosen is not sufficient to permit the passage of another object). By considering a truth assignment where we set to true only the variables associated with the k objects within $\{z_j^\varphi : 1 \leq j \leq v\}$ which pass by agent R_1^0, we get a truth assignment of weight k which satisfies formula φ. □

B Proof of equivalence in the reduction of Theorem 3.10

Theorem 3.10 RO-*sum and* RO-*makespan are* W[1]-*hard even for trees.*

Proof: Let us prove that there exists a clique of size k in graph $\mathcal{G}$ if and only if object x is reachable for agent Y within at most $k^3 + 4k^2 + k + 2$ swaps in total.

Suppose there exists a clique of size k in graph $\mathcal{G}$. Let V_C, with $|V_C| = k$, be the set of all vertices belonging to the clique, and C be the set of all edges of the clique. We consider the edges (v, w) in C with respect to the order over the edges previously assumed in the construction of the instance. Let us perform the rational swaps between the following pairs of agents for $(v, w) \in C$: $\{Y, Y^{[vw]}\}$, $\{Y, U_v^{vw}\}$, and $\{U_w^{vw}, U_v^{vw}\}$, that lead to give object u_v^{vw} to agent Y. Then, we decide to let object u_v^{vw} pass to branch A_v. At this moment, a further swap can be performed in the branch U^{vw} between Y and U_v^{vw}. Within the U^{vw}'s branches, we perform in total $\mathbf{2k(k-1)}$ **swaps**, considering all the $k(k-1)/2$ edges of the clique in C. Now let us focus on the passage of an object u_v^{vw} (respectively, u_w^{vw}) to branch A_v (respectively, A_w). To make the swaps rational within branch A_v, we need some auxiliary agents. We previously perform the swaps between agent A_v^ℓ and agent $A_v^{\ell[vw]}$, for each $1 \leq \ell \leq k - j$, if object u_v^{vw} in question is the j^{th} object to come into this branch. Then we perform the swaps along the path $[Y, A_v^1, \ldots, A_v^{k-j}]$ ($\mathbf{k^2(k-1)}$ **swaps** in total by considering all the objects u_v^{vw} and u_w^{vw} associated with an edge of the clique. After having performed all these swaps, all the agents A_v^ℓ, for $1 \leq \ell < k$ and $v \in V_C$, possess an object u_v^{vw} (or u_v^{wv}) associated with an edge (v, w) of the clique, and then can exchange with agent $A_v^{\ell *}$ in order to obtain object $a_v^{\ell *}$ and be "validated" (in total $\mathbf{k(k-1)}$ **swaps**). By rationality of the swaps, object a_v can now go to agent A_v^1 via path $[A_v, A_v^{k-1}, \ldots, A_v^1]$, for each $v \in V_C$ ($\mathbf{k(k-1)}$ swaps in total). Then, by increasing order of vertices v over V_C, let Y swap with $Y^{[v]}$, and then with A_v^1, in order to obtain object a_v (**2k swaps** in total by considering each $v \in V_C$). Let us focus now on the passage of object a_v to the branch of T^ℓ. Like in the A_v's branch, if a_v is the j^{th} object to come into this branch, then all the agents T^ℓ for $1 \leq \ell \leq k + 1 - j$ previously perform a swap with agent $T^{\ell[v]}$ in order to let object a_v pass to reach agent T^{k+1-j} ($\mathbf{k(k+1)}$ **swaps** in total). Then, agent T^{k+1-j} can exchange with agent T^{k+1-j*} to obtain object t^{k+1-j*}, and thus T^{k+1-j} is "validated" (**k swaps**). Since there are k validated A_v's branches, all the agents T^ℓ can be validated and thus, object t can go to agent Y via path $[T, T^k, \ldots, T^1, Y]$, for which the swaps are now rational ($\mathbf{k+1}$ **swaps**). Finally, agent Y can swap with agent X since object t is the only object that X prefers to object x (**one swap**), leading to the reachability of x by Y. Observe that we have exactly performed $k^3 + 4k^2 + k + 2$ swaps.

Suppose now that object x is reachable for agent Y within at most $k^3 + 4k^2 + k + 2$ swaps. The only way for agent X to give x in a rational swap is to obtain object t in return. Therefore, agent Y must previously get object t, initially owned by agent T, who only accept to give t against t^{k*}. Moreover, t must pass by all the agents T^ℓ and each of them accepts to give t to their neighbor only against object $t^{\ell-1*}$ or $y^{[t]}$. Since for agent T_1, this object is necessarily $y^{[t]}$, it must be $t^{\ell-1*}$ for all the others. Therefore, all the agents T^ℓ for $1 \leq \ell \leq k$ must obtain object $t^{\ell*}$ from their neighbor $T^{\ell*}$, who only accepts objects in $P := \{a_v : v \in V\}$. Thus, there must be k objects in total within P that move from their branch to the T branch in order to reach an agent T^ℓ. So far, the necessary swaps are those between X and Y (one swap), the swaps between each T^ℓ and

$T^{\ell*}$ (k swaps), and the swaps along the path $[T, T^k, \ldots, T_1, Y]$ ($k+1$ swaps), so in total $\mathbf{2k+2}$ **swaps**.

Consider an object $a_v \in P$ which must move to an agent T^ℓ. This object must follow the path $[A_v, A_v^{k-1}, \ldots, A_v^1, Y, T^1, \ldots, T^\ell]$. Consider first the subpath $[A_v^1, Y, T^1, \ldots, T^\ell]$, from the moment where object a_v reaches agent A_v^1. This agent only accepts to swap it against object $y^{[v]}$, therefore agent Y must previously perform a swap with agent $Y^{[v]}$ (**k swaps** in total by considering the k chosen objects in P). Observe that the j^{th} object in P entering in the T^ℓ's branch must reach agent T_{k+1-j}, otherwise an object should pass twice by the same agent, which contradicts the rationality assumption of the swaps. Therefore, by construction of the preferences, the objects in P must go into the T^ℓ's branch by increasing order of indices. By rationality of the swaps, agent Y accepts in exchange of object a_v in the swap with agent T^1 only objects coming from the T^ℓ's branch, and thus only objects in $\{t^{1[w]} : w \geq v\}$ that do not block the future swaps (typically $t^{1[v]}$ is appropriate). Therefore, T^1 must perform an exchange with one of the $T^{1[w]}$ before a_v. Observe that this also holds for the other agents in this branch and thus concerns all the agents T^ℓ for $1 \leq \ell \leq k+1-j$ if a_v is the j^{th} object of P to come into the branch. By counting the swaps between each such T^ℓ and one agent $T^{\ell[w]}$, and the swaps along the path $[Y, T_1, \ldots, T_{k+1-j}]$ for each object in P, we obtain $\mathbf{k(k+1)}$ **swaps**.

Now consider the first part where object a_v moves to agent Y from A_v along the path $[A_v, A_v^{k-1}, \ldots, A_v^1]$. The conditions are similar to those on the T^ℓ's branch. Each agent A_v^ℓ on the path only accepts object $y^{[v]}$ or $a^{\ell-1*}$ in return of giving object a_i, and agent A_v only prefers a^{k*} to a_v. Since for A_v^1 the preferred object is necessarily $y^{[v]}$, the other agents must obtain $a_v^{\ell-1*}$. However, all the agents $A_v^{\ell*}$, need an object within $D_v := \{u_v^{\delta^d(v)} : 1 \leq d \leq \delta(v)\}$ to give object $A_v^{\ell*}$ in a rational swap. It follows that each agent A_v^ℓ on the path must get object $a_v^{\ell-1*}$ and previously an object within D_v for letting pass object a_v to agent A_v^1. Thus, $k-1$ objects within D_v must be chosen to come into the A_v's branch. Once it is done, the remaining swaps are all the swaps between A_v^ℓ and $A_v^{\ell*}$ which lead to $\mathbf{k(k-1)}$ **swaps** in total, and the swaps for making object a_v reach agent A_v^1 along the path $[A_v, A_v^{k-1}, \ldots, A_v^1, Y]$, leading to $\mathbf{k^2}$ **more swaps**.

Now, consider an object u_v^{vw} (or u_v^{wv} depending on the order) which is chosen to come into the A_v's branch. Similarly as in the T^ℓ's branch, by construction of the preferences, the objects in D_v must arrive by increasing lexicographical order, and if u_v^{vw} is the j^{th} object in D_v which enters in the A_v's branch, then it must come to agent A_v^{k-j}. Moreover, each agent A_v^ℓ on the path $[A_v^1, \ldots, A_v^{k-j}]$ must previously make a swap with an auxiliary agent $A_v^{\ell[\delta^d(v)]}$ for $1 \leq d \leq \delta(v)$ that does not block the future swaps, typically with $A_v^{\ell[vw]}$ (or $A_v^{\ell[wv]}$ depending on the order). Therefore, by combining the swaps along $[Y, A_v^1, \ldots, A_v^{k-j}]$ and the swaps between each A_v^ℓ and one agent in $A_v^{\ell[\delta^d(v)]}$, we obtain in total $\mathbf{k^2(k-1)}$ **swaps**. To sum up, so far, we can count $k^3 + 2k^2 + 3k + 2$ necessary swaps. Therefore, it remains in the budget exactly $2k(k-1)$ swaps. By construction of the preferences, a previous swap between Y and the auxiliary agent $Y^{[vw]}$ is necessary to make a first swap between Y and U_v^{vw} occur. Observe that once an object u_v^{vw} is left from the U^{vw} branch, no other agent $U_{v'}^{v'w'}$ can swap with Y because the swap with the auxiliary agent is not possible. The only possibility is the swap between U_v^{vw} and U_w^{vw}, and then between U_v^{vw} and Y. This leads to the obligation of choosing both objects

u_v^{vw} and u_w^{vw} to make them pass to the branch of A_v, and we need four swaps to do it. Hence, with our remaining budget, we can only select $k(k-1)/2$ branches U^{vw} which correspond to an edge. Since the chosen objects allow validating $k-1$ incident edges of k vertices, the associated edges in $\mathcal{G}$ form a clique of size k.

The same reasoning holds for the makespan. Concerning the length of the sequence, it suffices to observe that the minimal sequence, by performing parallel swaps, is almost totally conditioned by the exchanges of agent Y. Obviously the number of swaps involving Y, precisely $5k(k-1)/2+3k+3$ swaps, is a lower bound for the makespan. But actually one can verify that all the other swaps can be performed in parallel of a swap involving Y. The only exception concerns the last swap between agent T^2 and agent T^1 for exchanging object t. Indeed, once Y gives to agent T^1 object a_v corresponding to the last vertex v of the clique (with respect to the order on vertices), agent T^1 can swap with agent T^{1*} to obtain object t^{1*}, while agent Y prepares in parallel the swap to obtain object t by swapping with agent $Y^{[t]}$. But agent T^1 still needs to swap with agent T^2 to obtain object t and Y has no swap to perform in parallel. Therefore, the makespan is $5k(k-1)/2+3k+4$.

The only possibility to answer true to a no-instance would be to choose more than $k(k-1)/2$ branches associated with an edge among the U^{vw}'s branches in order to validate more agents in the A_v's branches. However, it would imply at least three more swaps for agent Y and thus would increase the makespan, contradiction. □

Bibliography

Atila Abdulkadiroğlu and Tayfun Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.

Atila Abdulkadiroğlu and Tayfun Sönmez. House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233–260, 1999.

Rediet Abebe, Jon Kleinberg, and David C. Parkes. Fair division via social comparison. In *Proceedings of the Sixteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-17)*, pages 281–289, São Paulo, Brazil, May 2017. International Foundation for Autonomous Agents and Multiagent Systems.

David J. Abraham, Katarína Cechlárová, David F. Manlove, and Kurt Mehlhorn. Pareto optimality in house allocation problems. In *Proceedings of the Sixteenth International Symposium on Algorithms and Computation (ISAAC-05)*, pages 1163–1175, Sanya, Hainan, China, December 2005. Springer.

Karl A. Abrahamson, Rodney G. Downey, and Michael R. Fellows. Fixed-parameter tractability and completeness IV: On completeness for W[P] and PSPACE analogues. *Annals of Pure and Applied Logic*, 73(3):235 – 276, 1995.

Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows - Theory, Algorithms and Applications.* Prentice Hall, 1993.

R. Michael Alvarez, Frederick J. Boehmke, and Jonathan Nagler. Strategic voting in British elections. *Electoral Studies*, 25(1):1–19, 2006.

Elliot Anshelevich, Onkar Bhardwaj, and Martin Hoefer. Friendship and stable matching. In *Proceedings of the Twenty-First European Symposium on Algorithms (ESA-13)*, pages 49–60, Sophia Antipolis, France, September 2013a. Springer.

Elliot Anshelevich, Bugra Caskurlu, and Ameya Hate. Partition equilibrium always exists in resource selection games. *Theory of Computing Systems*, 53(1):73–85, 2013b.

Krzysztof R. Apt and Sunil Simon. A classification of weakly acyclic games. In *Proceedings of the Fifth International Symposium on Algorithmic Game Theory (SAGT-12)*, pages 1–12, Barcelona, Spain, October 2012. Springer.

Esteban Arcaute and Sergei Vassilvitskii. Social networks and stable matchings in the job market. In *Proceedings of the Fifth International Workshop on Internet and Network Economics (WINE-09)*, pages 220–231, Rome, Italy, December 2009. Springer.

Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach.* Cambridge University Press, 2009.

Kenneth J. Arrow. *Social Choice and Individual Values.* Wiley: New York, 1951.

Bengt Aspvall, Michael F. Plass, and Robert E. Tarjan. A linear-time algorithm for testing the truth of certain quantified boolean formulas. *Information Processing Letters*, 8(3):121–123, 1979.

Robert Aumann and Adam Brandenburger. Epistemic conditions for Nash equilibrium. In *The Language of Game Theory: Putting Epistemics into the Mathematics of Games*, pages 113–136. World Scientific, 2014.

Robert J. Aumann. Acceptable points in general cooperative n-person games. In Albert W. Tucker and Robert D. Luce, editors, *Contribution to the Theory of Games (AM-40)*, volume IV, chapter 16, pages 287–324. Princeton University Press, 1959.

Robert J. Aumann and Jacques H. Dreze. Cooperative games with coalition structures. *International Journal of Game Theory*, 3(4):217–237, 1974.

Haris Aziz and Bart De Keijzer. Housing markets with indifferences: A tale of two mechanisms. In *Proceedings of the Twenty-Sixth Conference on Artificial Intelligence (AAAI-12)*, pages 1249–1255, Toronto, Ontario, Canada, July 2012. AAAI Press.

Haris Aziz and Rahul Savani. Hedonic games. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 15, pages 356–376. Cambridge University Press, 2016.

Haris Aziz, Péter Biró, Jérôme Lang, Julien Lesca, and Jérôme Monnot. Optimal reallocation under additive and ordinal preferences. In *Proceedings of the Fifteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-16)*, pages 402–410, Singapore, May 2016a. International Foundation for Autonomous Agents and Multiagent Systems.

Haris Aziz, Ildikó Schlotter, and Toby Walsh. Control of fair division. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, pages 67–73, New York, NY, USA, July 2016b. AAAI Press / International Joint Conferences on Artificial Intelligence.

Haris Aziz, Jens L. Hougaard, Juan D. Moreno-Ternero, and Lars P. Østerdal. Computational aspects of assigning agents to a line. *Mathematical Social Sciences*, 90:93–99, 2017.

Haris Aziz, Sylvain Bouveret, Ioannis Caragiannis, Ira Giagkousi, and Jérôme Lang. Knowledge, fairness, and social constraints. In *Proceedings of the Thirty-Second Conference on Artificial Intelligence (AAAI-18)*, pages 4638–4645, New Orleans, Louisiana, USA, February 2018. AAAI Press.

Yoram Bachrach, Edith Elkind, and Piotr Faliszewski. Coalitional voting manipulation: A game-theoretic perspective. In *Proceedings of the Twenty-Second International Joint*

Conference on Artificial Intelligence (IJCAI-11), pages 49–54, Barcelona, Catalonia, Spain, July 2011. AAAI Press / International Joint Conferences on Artificial Intelligence.

Sophie Bade. Matching with single-peaked preferences. Technical report, Royal Holloway, University of London and Max Planck Institut for research on Collective Goods, Bonn, 2017.

Eyal Baharad and Shmuel Nitzan. The inverse plurality rule–an axiomatization. *Social Choice and Welfare*, 25(1):173, 2005.

Patrick Bahls, Sara Lake, and Andrew Wertheim. Gracefulness of families of spiders. *Involve, a Journal of Mathematics*, 3(3):241–247, 2010.

Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

John Bartholdi, Craig A. Tovey, and Michael A. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and welfare*, 6(2):157–165, 1989a.

John J. Bartholdi and James B. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

John J. Bartholdi, Craig A. Tovey, and Michael A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989b.

Cristina Bazgan, Bruno Escoffier, and Vangelis Th. Paschos. Poly-APX- and PTAS-completeness in standard and differential approximation. In *Proceedings of the Fifteenth International Symposium on Algorithms and Computation (ISAAC-04)*, pages 124–136, Hong Kong, China, December 2004. Springer.

Xiaohui Bei, Youming Qiao, and Shengyu Zhang. Networked fairness in cake cutting. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*, pages 3632–3638, Melbourne, Australia, August 2017. International Joint Conferences on Artificial Intelligence.

Sven Berg. Paradox of voting under an urn model: The effect of homogeneity. *Public Choice*, 47(2):377–387, 1985.

Piotr Berman, Marek Karpinski, and Alexander D. Scott. Approximation hardness of short symmetric instances of MAX-3SAT. Technical report, Institut des Hautes Études Scientifiques (IHES), 2004.

Aurélie Beynier, Yann Chevaleyre, Laurent Gourvès, Julien Lesca, Nicolas Maudet, and Anaëlle Wilczynski. Local envy-freeness in house allocation problems. In *Proceedings of the Seventeenth Conference on Autonomous Agents and MultiAgent Systems (AAMAS-18)*, pages 292–300, Stockholm, Sweden, July 2018. International Foundation for Autonomous Agents and Multiagent Systems.

Duncan Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.

André Blais. Strategic voting in the 2002 French presidential election. In Michael S. Lewis-Beck, editor, *The French Voter: Before and After the 2002 Elections*, chapter 5, pages 93–109. Springer, 2004.

Elizabeth Bodine-Baron, Christina Lee, Anthony Chong, Babak Hassibi, and Adam Wierman. Peer effects and stability in matching markets. In *Proceedings of the Fourth International Symposium on Algorithmic Game Theory (SAGT-11)*, pages 117–129, Amalfi, Italy, October 2011. Springer.

Anna Bogomolnaia and Matthew O Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.

Anna Bogomolnaia and Hervé Moulin. A new solution to the random assignment problem. *Journal of Economic theory*, 100(2):295–328, 2001.

Sylvain Bouveret and Jérôme Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *Journal of Artificial Intelligence Research*, 32:525–564, 2008.

Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible goods. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 12, pages 284–310. Cambridge University Press, 2016.

Sylvain Bouveret, Katarína Cechlárová, Edith Elkind, Ayumi Igarashi, and Dominik Peters. Fair division of a graph. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*, pages 135–141, Melbourne, Australia, August 2017. International Joint Conferences on Artificial Intelligence.

Steven J. Brams. Strategic information and voting behavior. *Society*, 19(6):4–11, 1982.

Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.

Steven J. Brams, Paul H. Edelman, and Peter C. Fishburn. Fair division of indivisible items. *Theory and Decision*, 55(2):147–180, 2003.

Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

Robert Bredereck, Andrzej Kaczmarczyk, and Rolf Niedermeier. Envy-free allocations respecting social networks. In *Proceedings of the Seventeenth Conference on Autonomous Agents and MultiAgent Systems (AAMAS-18)*, pages 283–291, Stockholm, Sweden, July 2018. International Foundation for Autonomous Agents and Multiagent Systems.

Markus Brill, Edith Elkind, Ulle Endriss, and Umberto Grandi. Pairwise diffusion of preference rankings in social networks. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, pages 130–136, New York, NY, USA, July 2016. AAAI Press / International Joint Conferences on Artificial Intelligence.

Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. On low-envy truthful allocations. In *Proceedings of the First International Conference on Algorithmic Decision Theory (ADT-2009)*, pages 111–119, Venice, Italy, October 2009. Springer.

Susumu Cato. Local strict envy-freeness in large economies. *Mathematical Social Sciences*, 59(3):319–322, 2010.

Katarína Cechlárová and Vladimir Lacko. The kidney exchange problem: How hard is it to find a donor? *Annals of Operations Research*, 193(1):255–271, 2012.

Katarína Cechlárová and Antonio Romero-Medina. Stability in coalition formation games. *International Journal of Game Theory*, 29(4):487–494, 2001.

Sourav Chakraborty, Shamik Ghosh, Nitesh Jha, and Sasanka Roy. Maximal and maximum transitive relation contained in a given binary relation. In *Proceedings of the Twenty-First Annual International Computing and Combinatorics Conference (COCOON-15)*, pages 587–600, Beijing, China, August 2015. Springer.

Yann Chevaleyre, Paul Dunne, Ulle Endriss, Jérôme Lang, Michel Lemaître, Nicolas Maudet, Julian Padget, Steve Phelps, Juan A. Rodrıguez-Aguilar, and Paulo Sousa. Issues in multiagent resource allocation. *Informatica*, 30(1):3–31, 2006.

Yann Chevaleyre, Ulle Endriss, Sylvia Estivie, and Nicolas Maudet. Reaching envy-free states in distributed negotiation settings. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, pages 1239–1244, Hyderabad, India, January 2007a. Morgan Kaufmann Publishers.

Yann Chevaleyre, Ulle Endriss, Jérôme Lang, and Nicolas Maudet. A short introduction to computational social choice. In *Thirty-Third International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM-07)*, pages 51–69, Harrachov, Czech Republic, January 2007b. Springer.

Yann Chevaleyre, Ulle Endriss, and Nicolas Maudet. Allocating goods on a graph to eliminate envy. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI-07)*, pages 700–705, Vancouver, British Columbia, Canada, July 2007c. AAAI Press.

Yann Chevaleyre, Ulle Endriss, and Nicolas Maudet. Distributed fair allocation of indivisible goods. *Artificial Intelligence*, 242:1–22, 2017.

Samir Chopra, Eric Pacuit, and Rohit Parikh. Knowledge-theoretic properties of strategic voting. In *Proceedings of the Ninth European Conference on Logics in Artificial Intelligence (JELIA-04)*, pages 18–30, Lisbon, Portugal, September 2004. Springer.

Silvano Colombo Tosatto and Marc van Zee. Social network analysis for judgment aggregation. In *Proceedings of the Thirteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-14)*, pages 1529–1530, Paris, France, May 2014. International Foundation for Autonomous Agents and Multiagent Systems.

Vincent Conitzer. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research*, 35:161–191, 2009.

Vincent Conitzer. Should social network structure be taken into account in elections? *Mathematical Social Sciences*, 64(1):100–102, 2012.

Vincent Conitzer. The maximum likelihood approach to voting on social networks. In *Proceedings of the Fifty-First Annual Allerton Conference on Communication, Control, and Computing (Allerton-13)*, pages 1482–1487, Monticello, IL, USA, October 2013. IEEE.

Vincent Conitzer and Tuomas Sandholm. Communication complexity of common voting rules. In *Proceedings of the Sixth ACM Conference on Electronic Commerce (EC-05)*, pages 78–87, Vancouver, Canada, June 2005. ACM.

Vincent Conitzer and Tuomas Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the Twenty-First Conference on Artificial Intelligence (AAAI-06)*, pages 627–634, Boston, Massachusetts, USA, July 2006. AAAI Press.

Vincent Conitzer and Toby Walsh. Barriers to manipulation in voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 6, pages 127–145. Cambridge University Press, 2016.

Vincent Conitzer, Tuomas Sandholm, and Jérôme Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM (JACM)*, 54(3):14, 2007.

Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating manipulations in voting with partial information. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence (AAAI-2011)*, pages 638–643, San Francisco, California, USA, August 2011. AAAI Press.

Stephen A. Cook. Characterizations of pushdown machines in terms of time-bounded computers. *Journal of the ACM (JACM)*, 18(1):4–18, 1971a.

Stephen A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing (STOC-71)*, pages 151–158, Shaker Heights, Ohio, USA, May 1971b. ACM.

Anastasia Damamme, Aurélie Beynier, Yann Chevaleyre, and Nicolas Maudet. The power of swap deals in distributed resource allocation. In *Proceedings of the Fourteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-15)*, pages 625–633, Istanbul, Turkey, May 2015. International Foundation for Autonomous Agents and Multiagent Systems.

Jean-François Daoust. Vote stratégique au Québec: Analyse de l'élection de 2012. *Politique et sociétés*, 34(2):3–15, 2015.

Marie Jean Antoine Nicolas de Caritat marquis de Condorcet. *Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*. L'imprimerie royale, 1785.

Bart de Keijzer, Sylvain Bouveret, Tomas Klos, and Yingqian Zhang. On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In *Proceedings of the First International Conference on Algorithmic Decision Theory (ADT-2009)*, pages 98–110, Venice, Italy, October 2009. Springer.

Morris H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the Eleventh ACM Conference on Electronic Commerce (EC-10)*, pages 347–356, Cambridge, Massachusetts, USA, June 2010. ACM.

Palash Dey, Neeldhara Misra, and Yadati Narahari. Complexity of manipulation with partial information in voting. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, pages 229–235, New York, NY, USA, July 2016. AAAI Press / International Joint Conferences on Artificial Intelligence.

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

John P. Dickerson, Jonathan R. Goldman, Jeremy Karp, Ariel D. Procaccia, and Tuomas Sandholm. The computational rise and fall of fairness. In *Proceedings of the Twenty-Eighth Conference on Artificial Intelligence (AAAI-14)*, pages 1405–1411, Québec City, Québec, Canada, July 2014. AAAI Press.

Franz Dietrich. A generalised model of judgment aggregation. *Social Choice and Welfare*, 28(4):529–565, 2007.

Rod G. Downey and Michael R. Fellows. Fixed-parameter tractability and completeness II: On completeness for W[1]. *Theoretical Computer Science*, 141(1-2):109–131, 1995a.

Rodney G. Downey and Michael R. Fellows. Parameterized computational feasibility. In Peter Clote and Jeffrey B. Remmel, editors, *Feasible Mathematics II*, pages 219–244. Springer, 1995b.

Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*, volume 4. Springer, 2013.

Anthony Downs. An economic theory of political action in a democracy. *Journal of Political Economy*, 65(2):135–150, 1957.

Jacques H. Dreze and Joseph Greenberg. Hedonic coalitions: Optimality and stability. *Econometrica*, 48(4):987–1003, 1980.

Paul E. Dunne and Yann Chevaleyre. The complexity of deciding reachability properties of distributed negotiation schemes. *Theoretical Computer Science*, 396(1-3):113–144, 2008.

François Durand. *Vers des Modes de Scrutin Moins Manipulables.* PhD thesis, Université Pierre et Marie Curie (Paris-VI), 2015.

David Easley and Jon Kleinberg. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World.* Cambridge University Press, 2010.

Florian Eggenberger and George Pólya. Über die statistik verketteter vorgänge. *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik*, 3(4):279–289, 1923.

Edith Elkind. Coalitional games on sparse social networks. In *Proceedings of the Tenth International Conference on Web and Internet Economics (WINE-14)*, pages 308–321, Beijing, China, December 2014. Springer.

Ulle Endriss. Judgment aggregation. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 17, pages 399–426. Cambridge University Press, 2016.

Ulle Endriss, Svetlana Obraztsova, Maria Polukarov, and Jeffrey S. Rosenschein. Strategic voting with incomplete information. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, pages 236–242, New York, NY, USA, July 2016. AAAI Press / International Joint Conferences on Artificial Intelligence.

Paul Erdös and Alfréd Rényi. On random graphs I. *Publicationes Mathematicae (Debrecen)*, 6:290–297, 1959.

Shimon Even, Alon Itai, and Adi Shamir. On the complexity of timetable and multicommodity flow problems. *SIAM Journal on Computing*, 5(4):691–703, 1976.

Patricia Everaere, Sébastien Konieczny, and Pierre Marquis. Belief merging versus judgment aggregation. In *Proceedings of the Fourteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-15)*, pages 999–1007, Istanbul, Turkey, May 2015. International Foundation for Autonomous Agents and Multiagent Systems.

Piotr Faliszewski and Ariel D. Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.

Piotr Faliszewski and Jörg Rothe. Control and bribery in voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 7, pages 146–168. Cambridge University Press, 2016.

Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner voting: A new challenge for social choice theory. In Ulle Endriss, editor, *Trends in Computational Social Choice*, chapter 2, pages 27–47. AI Access, 2017.

Robin Farquharson. *Theory of Voting*. Blackwell, 1969.

Timothy J. Feddersen, Itai Sened, and Stephen G. Wright. Rational voting and candidate entry under plurality rule. *American Journal of Political Science*, 34(4):1005–1016, 1990.

Michal Feldman and Moshe Tennenholtz. Partition equilibrium. In *Proceedings of the Second International Symposium on Algorithmic Game Theory (SAGT-09)*, pages 48–59, Paphos, Cyprus, October 2009. Springer.

Michael R. Fellows, Danny Hermelin, Frances Rosamond, and Stéphane Vialette. On the parameterized complexity of multiple-interval graph problems. *Theoretical Computer Science*, 410(1):53–61, 2009.

John A. Ferejohn. Decisive coalitions in the theory of social choice. *Journal of Economic Theory*, 15(2):301–306, 1977.

Michele Flammini, Mauro Manuel, and Tonelli Matteo. On social envy-freeness in multi-unit markets. In *Proceedings of the Thirty-Second Conference on Artificial Intelligence (AAAI-18)*, pages 1031–1038, New Orleans, Louisiana, USA, February 2018. AAAI Press.

Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Springer, 2006.

Duncan K. Foley. Resource allocation and the public sector. *Yale Economic Essays*, 7 (1):45–98, 1967.

Robert Forsythe, Roger B. Myerson, Thomas A. Rietz, and Robert J. Weber. An experiment on coordination in multi-candidate elections: The importance of polls and election histories. *Social Choice and Welfare*, 10(3):223–247, 1993.

Annika Fredén. Opinion polls, coalition signals and strategic voting: Evidence from a survey experiment. *Scandinavian Political Studies*, 40(3):247–264, 2017.

Linton C. Freeman. *The Development of Social Network Analysis: A Study in the Sociology of Science*. Empirical press, 2004.

Linton C. Freeman, John C. Mitchell, and Rolf Ziegler. Editorial. *Social Networks*, 1 (1):1–3, 1978.

Etsushi Fujita, Julien Lesca, Akihisa Sonoda, Taiki Todo, and Makoto Yokoo. A complexity approach for core-selecting exchange with multiple indivisible goods under lexicographic preferences. In *Proceedings of the Twenty-Ninth Conference on Artificial Intelligence (AAAI-15)*, pages 907–913, Austin, Texas, January 2015. AAAI Press.

David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

Peter Gardenfors. Assignment problem based on ordinal preferences. *Management Science*, 20(3):331–340, 1973.

Michael R. Garey and David S. Johnson. *Computers and Intractability: a Guide to the Theory of NP-Completeness*. Freeman, 1979.

William V. Gehrlein. Condorcet's paradox and the Condorcet efficienty of voting rules. *Mathematica Japonica*, 45:173–199, 1997.

Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41(4): 587–601, 1973.

Martin C. Golumbic. *Algorithmic Graph Theory and Perfect Graphs*. Academic Press, 1980.

Laurent Gourvès, Julien Lesca, and Anaëlle Wilczynski. Strategic voting in a social context: Considerate equilibria. In *Proceedings of the Twenty-Second European Conference on Artificial Intelligence (ECAI-16)*, pages 1423–1431, The Hague, The Netherlands, August-September 2016. IOS Press.

Laurent Gourvès, Julien Lesca, and Anaëlle Wilczynski. Object allocation via swaps along a social network. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*, pages 213–219, Melbourne, Australia, August 2017. International Joint Conferences on Artificial Intelligence.

Michel Grabisch, Antoine Mandel, Agnieszka Rusinowska, and Emily Tanimura. Strategic influence in social networks. *Mathematics of Operations Research*, 43(1):29–50, 2017.

Umberto Grandi. Social choice and social networks. In Ulle Endriss, editor, *Trends in Computational Social Choice*, chapter 9, pages 169–184. AI Access, 2017.

Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen B. Venable, and Toby Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and Borda score. In *Proceedings of the Third International Conference on Algorithmic Decision Theory (ADT-13)*, pages 181–192, Bruxelles, Belgium, November 2013. Springer.

Umberto Grandi, Emiliano Lorini, and Laurent Perrussel. Propositional opinion diffusion. In *Proceedings of the Fourteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-15)*, pages 989–997, Istanbul, Turkey, May 2015. International Foundation for Autonomous Agents and Multiagent Systems.

Umberto Grandi, Emiliano Lorini, Arianna Novaro, and Laurent Perrussel. Strategic disclosure of opinions on a social network. In *Proceedings of the Sixteenth Conference on Autonomous Agents and MultiAgent Systems (AAMAS-17)*, pages 1196–1204, São Paulo, Brazil, May 2017. International Foundation for Autonomous Agents and Multiagent Systems.

Davide Grossi and Gabriella Pigozzi. *Judgment Aggregation: A Primer*, volume 8 of *Synthesis Lectures on Artificial Intelligence and Machine Learning*. Morgan & Claypool Publishers, 2014.

Frank Harary and Allen J. Schwenk. The communication problem on graphs and digraphs. *Journal of the Franklin Institute*, 297(6):491–495, 1974.

Noam Hazon, Yonatan Aumann, Sarit Kraus, and Michael Wooldridge. Evaluation of election outcomes under uncertainty. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-08)*, pages 959–966, Estoril, Portugal, May 2008. International Foundation for Autonomous Agents and Multiagent Systems.

Sandra M. Hedetniemi, Stephen T. Hedetniemi, and Arthur L. Liestman. A survey of gossiping and broadcasting in communication networks. *Networks*, 18(4):319–349, 1988.

Lane A. Hemaspaandra. Computational social choice and computational complexity: BFFs? In *Proceedings of the Thirty-Second Conference on Artificial Intelligence (AAAI-18)*, pages 7971–7977, New Orleans, Louisiana, USA, February 2018. AAAI Press.

Martin Hoefer. Local matching dynamics in social networks. *Information and Computation*, 222(Supplement C):20–35, 2013.

Martin Hoefer, Michal Penn, Maria Polukarov, Alexander Skopalik, and Berthold Vöcking. Considerate equilibrium. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-11)*, pages 234–239, Barcelona, Catalonia, Spain, July 2011. AAAI Press / International Joint Conferences on Artificial Intelligence.

Harold Hotelling. Stability in competition. *The Economic Journal*, 39(153):41–57, 1929.

Aanund Hylland and Richard Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political Economy*, 87(2):293–314, 1979.

Ayumi Igarashi and Edith Elkind. Hedonic games with graph-restricted communication. In *Proceedings of the Fifteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-16)*, pages 242–250, Singapore, May 2016. International Foundation for Autonomous Agents and Multiagent Systems.

Ayumi Igarashi, Dominik Peters, and Edith Elkind. Group activity selection on social networks. In *Proceedings of the Thirty-First Conference on Artificial Intelligence (AAAI-17)*, pages 565–571, San Francisco, California, USA, February 2017. AAAI Press.

Matthew O. Jackson. *Social and Economic Networks*. Princeton University Press, 2008.

Samuel Karlin. *Total Positivity*, volume 1. Stanford University Press, 1968.

Richard M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Springer, 1972.

Bettina Klaus, David F. Manlove, and Francesca Rossi. Matching under preferences. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 14, pages 333–355. Cambridge University Press, 2016.

Jon M. Kleinberg and Éva Tardos. *Algorithm Design*. Addison-Wesley, 2006.

Jon M. Kleinberg and Éva Tardos. Balanced outcomes in social exchange networks. In *Proceedings of the Fortieth annual ACM Symposium on Theory of Computing (STOC-08)*, pages 295–304, Victoria, Canada, May 2008. ACM.

Hideo Konishi, Thomas Quint, and Jun Wako. On the Shapley–Scarf economy: the case of multiple types of indivisible goods. *Journal of Mathematical Economics*, 35 (1):1–15, 2001.

Aaron Koolyk, Tyrone Strangway, Omer Lev, and Jeffrey S. Rosenschein. Convergence and quality of iterative voting under non-scoring rules. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*, pages 273–279, Melbourne, Australia, August 2017. International Joint Conferences on Artificial Intelligence.

Bernhard Korte. Approximative algorithms for discrete optimization problems. In *Annals of Discrete Mathematics*, volume 4, pages 85–120. Elsevier, 1979.

Takashi Kurihara. A simple characterization of the anti-plurality rule. *Economics Letters*, 168:110–111, 2018.

Jérôme Lang and Lirong Xia. Voting in combinatorial domains. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 9, pages 197–222. Cambridge University Press, 2016.

Dominique Lepelley and Fabrice Valognes. Voting rules, manipulability and social homogeneity. *Public Choice*, 116(1-2):165–184, 2003.

Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. `http://snap.stanford.edu/data`, June 2014.

Omer Lev and Jeffrey S. Rosenschein. Convergence of iterative voting. In *Proceedings of the Eleventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS-12)*, pages 611–618, Valencia, Spain, June 2012. International Foundation for Autonomous Agents and Multiagent Systems.

Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the Fifth ACM Conference on Electronic Commerce (EC-04)*, pages 125–131, New York, NY, USA, May 2004. ACM.

Christian List and Philip Pettit. Aggregating sets of judgments: An impossibility result. *Economics & Philosophy*, 18(1):89–110, 2002.

Christian List and Clemens Puppe. Judgment aggregation: A survey. In Paul Anand, Prasanta Pattanaik, and Clemens Puppe, editors, *Handbook of Rational and Social Choice*, chapter 19, pages 457–482. Oxford University Press, 2009.

David F. Manlove. *Algorithmics of Matching under Preferences*, volume 2. World Scientific, 2013.

Nicholas Mattei and Toby Walsh. Preflib: A library for preferences `http://www.preflib.org/`. In *Proceedings of the Third International Conference on Algorithmic Decision Theory (ADT-13)*, pages 259–270, Bruxelles, Belgium, November 2013. Springer.

Reshef Meir. Plurality voting under uncertainty. In *Proceedings of the Twenty-Ninth Conference on Artificial Intelligence (AAAI-15)*, pages 2103–2109, Austin, Texas, January 2015. AAAI Press.

Reshef Meir. Iterative voting. In Ulle Endriss, editor, *Trends in Computational Social Choice*, chapter 4, pages 69–86. AI Access, 2017.

Reshef Meir. *Strategic Voting*, volume 12 of *Synthesis Lectures on Artificial Intelligence and Machine Learning*. Morgan & Claypool Publishers, 2018.

Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. Convergence to equilibria in plurality voting. In *Proceedings of the Twenty-Fourth Conference on Artificial Intelligence (AAAI-10)*, pages 823–828, Atlanta, Georgia, USA, July 2010. AAAI Press.

Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A local-dominance theory of voting equilibria. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation (EC-14)*, pages 313–330, Palo Alto, California, USA, June 2014. ACM.

Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. Iterative voting and acyclic games. *Artificial Intelligence*, 252:100–122, 2017.

Matthias Messner and Mattias K. Polborn. Robust political equilibria under plurality and runoff rule. *IGIER Working Paper*, 288, 2005.

Matthias Messner and Mattias K. Polborn. Strong and coalition-proof political equilibria under plurality and runoff rule. *International Journal of Game Theory*, 35(2):287–314, 2007.

James A. Mirrlees. An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, 38(2):175–208, 1971.

Hervé Moulin. *Fair Division and Collective Welfare.* MIT Press, 2004.

Roger B. Myerson. Graphs and cooperation in games. *Mathematics of Operations Research*, 2(3):225–229, 1977.

Roger B. Myerson and Robert J. Weber. A theory of voting equilibria. *The American Political Science Review*, 87(1):102–114, 1993.

John Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.

Mark Newman, Albert-Laszlo Barabasi, and Duncan J. Watts. *The Structure and Dynamics of Networks*, volume 19. Princeton University Press, 2011.

Trung Thanh Nguyen and Jörg Rothe. Minimizing envy and maximizing average Nash social welfare in the allocation of indivisible goods. *Discrete Applied Mathematics*, 179:54–68, 2014.

Svetlana Obraztsova, Evangelos Markakis, Maria Polukarov, Zinovi Rabinovich, and Nicholas R. Jennings. On the convergence of iterative voting: How restrictive should restricted dynamics be? In *Proceedings of the Twenty-Ninth Conference on Artificial Intelligence (AAAI-15)*, pages 993–999, Austin, Texas, January 2015. AAAI Press.

Svetlana Obraztsova, Omer Lev, Maria Polukarov, Zinovi Rabinovich, and Jeffrey S. Rosenschein. Nonmyopic voting dynamics: An optimistic approach. In *Proceedings of the Tenth Multidisciplinary Workshop on Advances in Preference Handling (MPREF-16)*, New York, NY, USA, July 2016.

Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

Vangelis Th. Paschos. A (delta/2)-approximation algorithm for the maximum independent set problem. *Information Processing Letters*, 44(1):11–13, 1992.

Prasanta K. Pattanaik and Manimay Sengupta. Outcomes of admissible Nash equilibria and sophisticated voting when decisions are based on pairwise comparisons. *Mathematical Social Sciences*, 2(1):39–54, 1982.

Bezalel Peleg. Consistent voting systems. *Econometrica*, 46(1):153–161, 1978.

Bezalel Peleg. *Game Theoretic Analysis of Voting in Committees*. Cambridge University Press, 1984.

Bezalel Peleg. Game-theoretic analysis of voting in committees. In Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, chapter 8, pages 395–423. Elsevier, 2002.

Bezalel Peleg and Hans Peters. *Strategic Social Choice: Stable Representations of Constitutions*. Springer, 2010.

Charles R. Plott. Axiomatic social choice theory: An overview and interpretation. *American Journal of Political Science*, 20(3):511–596, 1976.

Ariel D. Procaccia. Cake cutting algorithms. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 13, pages 311–330. Cambridge University Press, 2016.

Ariel D. Procaccia, Nisarg Shah, and Eric Sodomka. Ranked voting on social networks. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-15)*, pages 2040–2046, Buenos Aires, Argentina, July 2015. AAAI Press.

Michel Regenwetter, Bernard Grofman, Ilia Tsetlin, and Anthony A. J. Marley. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge University Press, 2006.

Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *Proceedings of the Eleventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS-12)*, pages 635–644, Valencia, Spain, June 2012. International Foundation for Autonomous Agents and Multiagent Systems.

Reyhaneh Reyhani and Mark C. Wilson. Best-reply dynamics for scoring rules. In *Proceedings of the Twentieth European Conference on Artificial Intelligence (ECAI-12)*, pages 672–677, Montpellier, France, August 2012. IOS Press.

Kevin W. S. Roberts. Voting over income tax schedules. *Journal of Public Economics*, 8(3):329–340, 1977.

Jack Robertson and William Webb. *Cake-Cutting Algorithms: Be Fair If You Can.* AK Peters/CRC Press, 1998.

Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. Kidney exchange. *The Quarterly Journal of Economics*, 119(2):457–488, 2004.

Agnieszka Rusinowska, Rudolf Berghammer, Harrie De Swart, and Michel Grabisch. Social networks: Prestige, centrality, and influence. In *International Conference on Relational and Algebraic Methods in Computer Science*, pages 22–39. Springer, 2011.

Abdallah Saffidine and Anaëlle Wilczynski. Constrained swap dynamics over a social network in distributed resource reallocation. In *Proceedings of the Eleventh International Symposium on Algorithmic Game Theory (SAGT-18)*, pages 213–225, Beijing, China, September 2018. Springer.

Tuomas W. Sandholm. Contract types for satisficing task allocation. In *Proceedings of the AAAI Spring Symposium: Satisficing Models*, pages 23–25, Palo Alto, California, USA, March 1998. AAAI Press.

Tuomas W. Sandholm. Distributed rational decision making. In Gerhard Weiss, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, chapter 5, pages 201–258. MIT Press, 1999.

Mark A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

Mark A. Satterthwaite and Hugo Sonnenschein. Strategy-proof allocation mechanisms at differentiable points. *The Review of Economic Studies*, 48(4):587–597, 1981.

John Scott. *Social Network Analysis.* Sage, 2017.

Murat Sertel and Remzi Sanver. Strong equilibrium outcomes of voting games are the generalized Condorcet winners. *Social Choice and Welfare*, 22(2):331–347, 2004.

Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23–37, 1974.

Lloyd S. Shapley. Simple games: an outline of the descriptive theory. *Behavioral Science*, 7(1):59–66, 1962.

Sigal Sina, Noam Hazon, Avinatan Hassidim, and Sarit Kraus. Adapting the social network to affect elections. In *Proceedings of the Fourteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-15)*, pages 705–713, Istanbul, Turkey, May 2015. International Foundation for Autonomous Agents and Multiagent Systems.

Arkadii Slinko and Shaun White. Nondictatorial social choice rules are safely manipulable. In *Proceedings of the Second International Workshop on Computational Social Choice (COMSOC-08)*, pages 403–414, Liverpool, England, September 2008.

Birgitte Sloth. The theory of voting and equilibria in noncooperative games. *Games and Economic Behavior*, 5(1):152–169, 1993.

Akihisa Sonoda, Etsushi Fujita, Taiki Todo, and Makoto Yokoo. Two case studies for trading multiple indivisible goods with indifferences. In *Proceedings of the Twenty-Eighth Conference on Artificial Intelligence (AAAI-14)*, pages 791–797, Québec City, Québec, Canada, July 2014. AAAI Press.

Hugo Steinhaus. The problem of fair division. *Econometrica*, 16(1):101–104, 1948.

Lars-Gunnar Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.

Alan D. Taylor. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press, 2005.

Alan D. Taylor and William S. Zwicker. *Simple Games: Desirability Relations, Trading, Pseudoweightings*. Princeton University Press, 1999.

Jan Tinbergen. *Redelijke Inkomensverdeling*. De Gulden Pers, Haarlem, 1946.

Taiki Todo, Haixin Sun, and Makoto Yokoo. Strategyproof exchange with multiple private endowments. In *Proceedings of the Twenty-Eighth Conference on Artificial Intelligence (AAAI-14)*, pages 805–811, Québec City, Québec, Canada, July 2014. AAAI Press.

Jeffrey Travers and Stanley Milgram. The small world problem. *Phychology Today*, 1 (1):61–67, 1967.

Jeffrey Travers and Stanley Milgram. An experimental study of the small world problem. In *Social Networks*, pages 179–197. Elsevier, 1977.

Alan Tsang and Kate Larson. The echo chamber: Strategic voting and homophily in social networks. In *Proceedings of the Fifteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-16)*, pages 368–375, Singapore, May 2016. International Foundation for Autonomous Agents and Multiagent Systems.

Alan Tsang, John A. Doucette, and Hadi Hosseini. Voting with social influence: Using arguments to uncover ground truth. In *Proceedings of the Fourteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-15)*, pages 1841–1842, Istanbul, Turkey, May 2015. International Foundation for Autonomous Agents and Multiagent Systems.

Alan Tsang, Amirali Salehi-Abari, and Kate Larson. Boundedly rational voters in large(r) networks. In *Proceedings of the Seventeenth Conference on Autonomous Agents and MultiAgent Systems (AAMAS-18)*, pages 301–308, Stockholm, Sweden, July 2018. International Foundation for Autonomous Agents and Multiagent Systems.

Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.

Hans Van Ditmarsch, Jérôme Lang, and Abdallah Saffidine. Strategic voting and the logic of knowledge. In *Proceedings of the Eleventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS-12)*, pages 1247–1248, Valencia, Spain, June 2012. International Foundation for Autonomous Agents and Multiagent Systems.

Hal R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974.

Vijay V. Vazirani. *Approximation Algorithms*. Springer, 2001.

Mark Voorneveld. *Potential Games and Interactive Decisions with Multiple Criteria*. Center for Economic Research, Tilburg University, 1999.

Toby Walsh. Generating single peaked votes. *CoRR*, abs/1503.02766, 2015.

Stanley Wasserman and Katherine Faust. *Social Network Analysis: Methods and Applications*, volume 8. Cambridge University Press, 1994.

Anaëlle Wilczynski. Poll-confident voters in iterative voting. In *Proceedings of the Thirty-Third Conference on Artificial Intelligence (AAAI-19)*, Honolulu, Hawaii, USA, January-February 2019. To appear.

Lirong Xia, Michael Zuckerman, Ariel D. Procaccia, Vincent Conitzer, and Jeffrey S. Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09)*, pages 348–352, Pasadena, California, USA, July 2009. IJCAI Organization / AAAI Press.

Mihalis Yannakakis. Node- and edge-deletion NP-complete problems. In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing (STOC-78)*, pages 253–264, San Diego, California, USA, May 1978. ACM.

Ryo Yoshinaka. Higher-order matching in the linear lambda calculus in the absence of constants is NP-complete. In *Proceedings of the Sixteenth International Conference on Rewriting Techniques and Applications (RTA-05)*, pages 235–249, Nara, Japan, April 2005. Springer.

H. Peyton Young. *Equity: In Theory and Practice*. Princeton University Press, 1995.

Reza Zafarani and Huan Liu. Social computing data repository at ASU. `http://socialcomputing.asu.edu`, 2009.

Lin Zhou. On a conjecture by Gale about one-sided matching problems. *Journal of Economic Theory*, 52(1):123 – 135, 1990.

Michael Zuckerman, Ariel D. Procaccia, and Jeffrey S. Rosenschein. Algorithms for the coalitional manipulation problem. *Artificial Intelligence*, 173(2):392–412, 2009.

William S. Zwicker. Introduction to the theory of voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 2, pages 23–56. Cambridge University Press, 2016.

Index of definitions

Résumé long en français

Introduction

La *théorie du choix social* repose sur l'étude de la prise de décision collective. Il existe de nombreuses situations dans la vie quotidienne où un groupe d'individus doit prendre une décision ensemble. On peut citer par exemple l'élection de représentants politiques, le choix d'une date pour un rendez-vous, le choix d'un menu commun dans un restaurant, la répartition de biens lors d'un héritage, ou encore la constitution de groupes de travail. La théorie du choix social a connu un véritable essor à partir de la moitié du XXème siècle, notamment à partir des travaux de Kenneth Arrow avec son célèbre théorème d'impossibilité [1951]. Cependant, des travaux fondateurs ont été réalisés dans le domaine dès le XVIIIème siècle avec en particulier les écrits de Nicolas (marquis) de Condorcet (1743-1794) et Jean-Charles (chevalier) de Borda (1733-1799).

Dans un problème de choix social, le but est d'agréger les préférences de différents agents portant sur un ensemble d'alternatives donné afin de sélectionner, comme décision finale, une ou plusieurs alternatives ou d'établir un ordre (pouvant être partiel ou un même un préordre) sur ces alternatives. Le but du choix social consiste alors à élaborer des procédures, les plus équitables et efficaces possibles, pour la prise de décision collective. Cependant, au-delà de la qualité des procédures, il est important de prendre en compte leur coût en termes de calcul et de communication. Cet axe de recherche, né à la fin du XXème siècle, initialement au sein de la communauté informatique, a émergé sous le nom de *choix social computationnel* [Brandt et al., 2016]. Ce domaine apparaît comme un champ de recherche avec des possibilités d'application concrètes, se trouvant à l'intersection entre l'économie, les mathématiques et l'informatique.

Plusieurs aspects peuvent être étudiés dans le cadre de problèmes de choix social. Premièrement, une étude axiomatique peut être faite en caractérisant les procédures d'agrégation de préférences en fonction des propriétés intéressantes qu'elles satisfont. L'aspect computationnel est également important, donnant lieu à une littérature vaste établissant la complexité algorithmique de mécanismes donnés. Par ailleurs, les préférences des agents, de par leur structure, peuvent influer de manière conséquente sur les procédures de choix social, que ce soit au niveau de leur calculabilité ou de la qualité de leur résultat. Il convient alors de déterminer des modèles de préférences réalistes. L'élicitation des préférences des agents peut ainsi devenir une question essentielle, notamment en ce qui concerne les coûts de communication des protocoles d'élicitation. D'un autre côté, il est possible que les agents ne soient pas sincères lorsqu'ils dévoilent leurs préférences, car ils peuvent avoir intérêt d'un point de vue stratégique à mentir, phénomène connu sous le nom de *manipulation*. Puisque la manipulation de la part des agents est en général un comportement que l'on cherche à prévenir, les questions relatives aux comportements stratégiques des agents constituent un enjeu fondamental en choix social.

Lorsque l'on parle de choix social, on pense souvent aux processus de vote. Néanmoins, le spectre des problèmes de choix social est bien plus large que ce seul domaine d'application. Il regroupe notamment la *théorie du vote*, l'*allocation de ressources* et le *partage équitable*, la *formation de coalitions* et l'*agrégation de jugements*.

Dans un problème de choix social, les agents sont conscients de l'existence et de l'implication d'autres agents. Généralement, dans la plupart des problèmes de décision collective, les agents peuvent interagir. Cette interaction peut prendre différentes formes : communication, collaboration, coopération, influence, collecte d'information, diffusion d'information et autres. Traditionnellement en choix social, les agents sont supposés pouvoir interagir avec tout autre agent. Cependant, dans la vie réelle, pour des raisons liées à des questions de communication, d'affinités entre agents ou à des problèmes liés à la distance, l'interaction avec certains agents peut ne pas être possible. La possibilité d'interaction entre agents peut être modélisée par une *relation d'accessibilité* qui n'est alors pas forcément complète par défaut. Dans cette perspective, la relation d'accessibilité reliant les agents pourrait également faire partie des données d'entrée d'un problème de choix social.

Il existe plusieurs manières de représenter la relation d'accessibilité entre agents. En raison de l'importance et de l'influence croissante des réseaux sociaux de nos jours, modéliser la relation d'accessibilité entre agents à l'aide d'un *réseau social* apparaît pertinent. Dans ce cadre, la relation d'accessibilité consiste en une relation binaire telle que deux agents peuvent interagir seulement s'ils sont connectés dans le réseau. Classiquement, on représente le réseau social à l'aide d'un graphe, c'est-à-dire par un ensemble de nœuds représentant les différents agents, et par des liens entre les nœuds symbolisant les connexions entre les agents dans le réseau. Le réseau social en lui-même peut modéliser une proximité géographique entre les agents, une relation entre pairs ou collègues, des relations amicales ou même un réseau social en ligne comme il en existe plusieurs de nos jours.

Les réseaux sociaux sont un domaine de recherche florissant et très en vogue ces derniers temps. Ils sont de plus en plus utilisés, notamment en économie, afin de comprendre les relations sociales qui contraignent ou guident les agents dans leurs choix et décisions [Jackson, 2008, Easley and Kleinberg, 2010]. L'introduction d'une structure de graphe afin de modéliser les relations entre agents a notamment été considérée en théorie des jeux coopératifs, avec l'utilisation d'un graphe de coopération [Aumann and Dreze, 1974, Myerson, 1977]. En choix social, l'utilisation d'un réseau social modélisant les interactions entre agents est plus récente et plus partielle. Cette approche a été principalement adoptée dans le contexte de problèmes de formation de coalitions, même si l'on peut aussi noter certains travaux analysant l'impact des réseaux sociaux dans des problématiques de vote ou d'allocation de ressources. Un état de l'art récent [Grandi, 2017] met en lumière certains travaux intégrant des réseaux sociaux dans des problématiques de choix social, et souligne des questions ou axes de recherche pertinents dans ce cadre.

Le but de cette thèse est de relâcher l'hypothèse classique, en choix social computationnel, qui consiste à supposer que tout agent peut interagir avec n'importe quel autre. Pour cela, on modélise grâce à un réseau social la possibilité d'interaction entre les agents. On se propose d'étudier l'impact de cette généralisation sur les interactions sociales qui peuvent avoir lieu dans des problèmes de choix social. On concentre notre étude sur deux types d'interaction particulièrement décisifs dans des processus de prise de décision collective : la *collaboration* entre agents et la *prise d'information*. De plus, notre travail concerne en particulier deux problèmes de choix social : le *vote stratégique*, modélisé sous la forme d'un jeu stratégique itératif appelé *vote itératif*, et l'*allocation de ressources*, avec l'étude d'un problème particulier reposant sur des biens indivisibles dans lequel chaque agent doit obtenir exactement une ressource ("house allocation"). Dans de tels problèmes de choix social, comment le réseau social influence-t-il les interactions entre les agents ? Plus précisément, comment la limitation sur l'information qui est accessible pour les agents ou sur les possibles agents collaborant entre eux affecte-t-elle le déroulement et le résultat des procédures de décision collective dans ces problèmes particuliers ?

Ces questions sont examinées selon deux axes : le type de problème étudié (vote stratégique ou allocation de ressources) et le type d'interaction sociale pris en considération (collaboration ou prise d'information). Le second axe détermine la structure de la thèse. Ainsi, après un chapitre préliminaire (Chapitre 1) introduisant les notations et concepts utilisés tout au long de la thèse, on étudie dans une première partie (Partie I), comment le réseau social peut modéliser la possible collaboration entre agents. Au sein de cette première partie, un premier chapitre (Chapitre 2) porte sur la manipulation par coalitions d'agents en vote itératif et un second chapitre (Chapitre 3) sur des échanges de biens indivisibles entre agents. Pour le chapitre 2, les possibles coalitions d'agents pouvant manipuler sont déterminées en fonction du réseau social et dans le chapitre 3, les seuls échanges possibles sont ceux impliquant des agents connectés dans le réseau social. Dans la seconde partie de la thèse (Partie II), on analyse comment le réseau social peut modéliser l'information qui est disponible pour les agents. En particulier, on étudie dans un premier chapitre (Chapitre 4) un problème d'allocation de ressources où la relation d'envie entre les agents est déterminée par un réseau social, puis dans un second chapitre (Chapitre 5) on se penche sur un problème de vote itératif sous incertitude dans lequel la connaissance des agents dépend du réseau social.

Nos contributions sont regroupées dans les parties I et II de la thèse, elles ont fait l'objet de publications dans des conférences internationales d'intelligence artificielle : ECAI-16 [Gourvès et al., 2016], IJCAI-17 [Gourvès et al., 2017], AAMAS-18 [Beynier et al., 2018], SAGT-18 [Saffidine and Wilczynski, 2018] et AAAI-19 [Wilczynski, 2019].

Chapitre 1 : Préliminaires et notations

Ce chapitre préliminaire présente les notions et notations qui sont utilisées tout au long de la thèse. Le cadre général des problèmes de choix social étudiés dans la thèse est introduit ainsi que deux problèmes particuliers sur lesquels on se concentre dans la thèse : le vote, plus précisément le vote stratégique, et l'allocation de ressources avec des biens indivisibles. Les concepts principaux sont rappelés ainsi que quelques éléments d'état de l'art.

Cadre général

Un problème de choix social est caractérisé par un ensemble d'*agents* et un ensemble d'*alternatives* pour lesquels les agents expriment des *préférences*. Les préférences des agents sur l'ensemble des alternatives peuvent prendre différentes formes.

Tout d'abord, les préférences des agents peuvent être *cardinales*. Dans ce contexte, les préférences d'un agent sont modélisées par l'intermédiaire d'une fonction d'utilité définie sur l'ensemble des alternatives, où chaque agent associe à chaque alternative une valeur numérique. De manière plus générale, les agents ne peuvent exprimer qu'une relation *ordinale* sur les alternatives, dont l'élicitation peut s'avérer plus aisée pour les agents. On suppose tout au long de la thèse que les agents expriment un ordre complet sur les alternatives où, pour certains modèles, l'indifférence entre deux alternatives peut être permise.

Afin de modéliser une possibilité d'interaction entre les agents, on introduit dans ce cadre classique de choix social un *réseau social* sur les agents. Ce réseau social est modélisé par une structure de graphe dont les nœuds symbolisent les agents et dont les liens entre les nœuds représentent une possibilité d'interaction entre les agents associés. L'accent est mis sur la relation binaire représentée par le graphe. Ainsi, un graphe non-orienté est vu comme un cas particulier de graphe orienté où la relation binaire associée est symétrique. Toutes les définitions que l'on utilise au long de la thèse suivent ce principe. Par exemple, une clique dans un graphe orienté est un sous-ensemble de sommets du graphe tel qu'il existe un arc entre chaque paire de sommets de ce sous-ensemble, c'est-à-dire la relation binaire associée est complète sur ce sous-ensemble d'éléments.

On s'attarde également sur la manière de générer des réseaux sociaux réalistes. Pour cela, certaines propriétés observables dans des réseaux réels et théorisées dans la littérature sont utiles. Dans ce cadre, les notions de *"petit monde"* [Travers and Milgram, 1967, 1977] et d' *"invariant d'échelle"* [Barabási and Albert, 1999] sont notamment rappelées, ainsi que la manière de générer ce type de réseaux. Par exemple, les graphes aléatoires de type Erdös-Rényi [1959] satisfont la propriété de "petit monde" en espérance, c'est-à-dire que la distance entre deux nœuds du réseau tend à être petite, tandis que les graphes aléatoires de type Barabási-Albert [1999] sont connus pour être invariants d'échelle, c'est-à-dire que les nœuds du réseau suivent une loi de puissance avec un mécanisme d'attachement préférentiel où les nœuds ont plus de chance d'être reliés dans le réseau à des nœuds possédant déjà de nombreuses connexions. Une autre notion importante dans un réseau social est l'*homophilie* : deux nœuds "similaires" ont tendance à être reliés dans le réseau. Dans un contexte où le réseau social concerne des agents tentant de prendre une décision collectivement, il apparaît pertinent de considérer que deux agents sont similaires s'ils ont des préférences proches. On établit donc un protocole de génération de graphe aléatoire avec homophilie où les agents sont reliés entre eux en fonction de la proximité de leurs préférences, sur la base d'une distance de Kendall-Tau entre les ordres de préférence des agents.

Enfin, on se concentre dans la thèse sur des problèmes de choix social où le but pour les agents est de sélectionner exactement une alternative parmi l'ensemble d'alternatives initial. La procédure pour y parvenir peut être *centralisée*, si une autorité externe détermine à partir de certaines règles quelle sera l'alternative choisie en fonction de préférences soumises par les agents ou, au contraire, elle peut être *distribuée*, si les agents construisent la décision finale par eux-mêmes en suivant un certain protocole. Certaines procédures peuvent également combiner ces deux aspects. Les procédures de choix social sont traditionnellement évaluées en fonction des axiomes qu'elles satisfont. Ces propriétés peuvent concerner le déroulement de la procédure en elle-même, comme par exemple le fait d'être *non-manipulable*, c'est-à-dire le fait qu'aucun agent n'ait intérêt à reporter des préférences qui ne sont pas les siennes. Mais également, elles peuvent s'appuyer sur la qualité de la solution obtenue. On peut s'intéresser par exemple à la sélection d'une solution qui soit *Pareto-efficace*, c'est-à-dire qu'il n'est pas possible à partir de cette solution d'améliorer la satisfaction d'un agent sans détériorer celle d'un autre.

D'autres critères sont propres au problème de choix social étudié. Dans la thèse, on se concentre sur le vote stratégique et l'allocation de biens indivisibles.

Théorie du vote

Le vote est incontournable dans la vie quotidienne. Outre les élections politiques, on peut citer des processus de vote afin de choisir une date ou un lieu pour une réunion, le choix d'un menu commun pour une sortie dans un restaurant, l'élection de représentants en entreprise ou dans des associations, et d'autres encore. Dans un modèle de vote, des agents, également appelés *électeurs* ou *votants*, ont des préférences par rapport à un ensemble d'alternatives, également appelées *candidats*. Le résultat de l'élection, c'est-à-dire le gagnant, est déterminé par une règle de vote. Chaque électeur est invité à soumettre son bulletin de vote à un système central qui calcule puis communique le gagnant de l'élection. La règle de vote prend en entrée un ensemble de bulletins de vote, un pour chaque électeur (en supposant que le modèle ne prenne pas en compte l'abstention), qui est appelé profil de vote. On se concentre dans la thèse sur des élections à vainqueur unique, impliquant l'utilisation d'une règle de départage des ex aequo. On choisit une règle de départage déterministe basée sur un ordre sur les candidats.

Plusieurs règles de vote ont été conçues dans la théorie du vote. Quasiment toutes les règles de vote que l'on étudie dans la thèse attribuent un score à chaque candidat et donnent comme gagnant le candidat qui maximise ce score. Toutes les règles de vote analysées requièrent la soumission d'un bulletin sous forme d'ordre sur les candidats (même si parfois moins d'information suffit). Plus précisément, on considère trois familles de règles de vote. La première, les *règles de notation positionnelles*, rassemble des règles pour lesquelles il existe un vecteur de points qui détermine le nombre de points remportés par chaque candidat en fonction de sa position dans l'ordre soumis dans chaque bulletin. Parmi ces règles, on compte notamment la règle de vote *Pluralité* où un candidat remporte un point pour chaque bulletin dans lequel il est placé en première position. Le profil de vote peut alors être simplifié en supposant que chaque bulletin porte la mention d'un seul candidat, celui qui est approuvé. La règle de vote *Véto* appartient également à cette famille mais cette fois-ci tous les candidats du bulletin remportent un point, sauf le candidat placé en dernière position du bulletin qui n'en remporte aucun. Ceci peut également se simplifier par un bulletin ne comportant que la mention du candidat contre lequel on veut exprimer un véto. La règle de vote k-*approbation* permet quant à elle d'attribuer, pour chaque bulletin, un point à exactement k candidats, qui sont censés être les k candidats préférés du votant. Cette règle généralise à la fois Pluralité et Véto, en choisissant des valeurs de k particulières. Enfin, la dernière règle sur laquelle on va se pencher au sein des règles de notation positionnelles est la règle de *Borda*, qui associe à chaque position de l'ordre soumis dans le bulletin un nombre de points différent, avec le même écart de points entre deux positions consécutives et le plus grand nombre de points pour la première position dans le bulletin.

La deuxième famille de règles de vote que l'on étudie consiste en des règles sur plusieurs tours ou dites par *éliminations successives*. Parmi elles, la règle de *Pluralité à deux tours*, qui est notamment utilisée en France, ne garde au deuxième tour que les deux candidats ayant le plus fort score de Pluralité dans le profil de vote. Puis, le candidat qui est préféré, à la majorité, parmi ces deux finalistes remporte l'élection. Suivant le même principe, la règle de *Vote simple transférable (STV)*, utilisée notamment en Australie et en Irlande, élimine à chaque tour le candidat qui a le plus faible score sous la règle de Pluralité. Le vainqueur est le candidat restant à la fin de ce processus d'éliminations successives ou, de manière équivalente, le candidat qui gagne à la majorité absolue à partir d'un certain tour.

Enfin, la dernière famille de règles de vote abordées repose sur la *comparaison par paires*. Parmi ces règles, on se concentre sur les règles de *Maximin* et de *Copeland* qui consistent à élire un candidat maximisant un certain score dont le calcul se base sur des comparaisons par paires avec les autres candidats. Le score de Maximin d'un candidat est le nombre minimum d'électeurs qui le soutiennent dans tout duel avec un autre candidat. Le score de Copeland d'un candidat est le nombre de candidats qu'il bat à la majorité absolue en comptabilisant tous les duels avec un autre candidat (il existe d'autres versions du score de Copeland).

Les règles de vote sont en général évaluées en fonction des axiomes qu'elles satisfont, reflétant des propriétés désirables pour le processus de choix en soi ou le candidat élu. On peut notamment citer le *critère de Condorcet* qui consiste à savoir si une règle permet d'élire le *vainqueur de Condorcet* lorsqu'il existe, le vainqueur de Condorcet étant un candidat battant tous les autres à la majorité absolue. Une règle élisant toujours le vainqueur de Condorcet est dite *Condorcet consistante*. Une manière empirique de se ramener au critère de Condorcet est d'observer l'*efficacité de Condorcet* d'un processus de vote, c'est-à-dire la fréquence d'élection du vainqueur de Condorcet lorsque celui-ci existe. De manière analogue, on peut étudier les scores de Borda du vainqueur d'un certain processus de vote pour voir à quel point ceux-ci se rapprochent du score de Borda du vainqueur sous la règle de vote Borda.

Le *vote stratégique* apparaît fréquemment lors d'élections réelles, et en particulier lors d'élections politiques. On dit que les votants manipulent lorsqu'ils soumettent un bulletin qui ne correspond pas à leurs préférences véritables. Les électeurs souhaitent en général, par ce biais, éviter l'élection d'un candidat qu'ils ne veulent pas voir élu. Malgré le fait que la manipulation puisse apparaître comme un comportement à prévenir, il n'existe pas de processus de vote où la manipulation serait impossible [Gibbard, 1973, Satterthwaite, 1975]. Une approche pour contourner le problème de la manipulation en vote consiste à élaborer des règles de vote difficiles à manipuler en termes de calcul. Une littérature riche a développé ce point, en analysant le coût de calcul d'une manipulation pour différentes règles de vote [Bartholdi et al., 1989b, Bartholdi and Orlin, 1991]. Cependant, cette approche n'est pas suffisante pour éviter les manipulations, comme le soulignent certains travaux récents [Conitzer and Sandholm, 2006, Faliszewski and

Procaccia, 2010]. Le principal inconvénient de cette approche repose sur l'analyse de la complexité dans le pire cas : même s'il est compliqué d'un point de vue computationnel de calculer une manipulation pour une règle de vote donnée dans le pire des cas, cela n'empêche pas la manipulation des électeurs dans la pratique. Une autre perspective consiste à explorer les aspects relevant de la théorie des jeux dans le comportement stratégique des votants en laissant les votants manipuler. La manipulation n'est alors pas nécessairement considérée comme un comportement à prévenir, mais plutôt comme un comportement à prendre en compte dans le processus de vote. Dans ce cadre, on peut alors étudier l'existence d'équilibres dans le jeu de vote. Différents concepts de solution peuvent être pris en compte en fonction des déviations stratégiques qui sont considérées. Par exemple, l'*équilibre de Nash* [Nash, 1951] est un état stable par rapport aux déviations unilatérales. Il a notamment été bien étudié dans le cadre de jeux de vote. De manière alternative, l'*équilibre fort* [Aumann, 1959] est un état stable par rapport aux déviations coalitionnelles où tout le monde doit gagner strictement dans la déviation. Les équilibres forts ont également été bien examinés dans des jeux de vote, par le biais de jeux admettant toujours un tel équilibre [Peleg, 2002, Peleg and Peters, 2010], ou de caractérisations faisant intervenir la notion de vainqueur de Condorcet [Sertel and Sanver, 2004].

Il est possible d'envisager une version dynamique d'un jeu de vote, ce que l'on appelle le *vote itératif*, où les électeurs peuvent changer de bulletin dans une perspective stratégique de façon itérative [Meir et al., 2010]. Il est communément supposé qu'un seul électeur peut dévier à chaque étape. Un votant dévie de son bulletin actuel vers un nouveau s'il estime que cette nouvelle stratégie lui permet de faire élire un candidat qu'il préfère. Le vote itératif peut être vu comme les réponses à une succession de sondages d'opinion où les électeurs peuvent observer les votes précédents et élaborer des stratégies en conséquence, ou il peut simplement décrire les changements stratégiques dans les intentions de vote des votants, qui peuvent évoluer au gré de ce qu'ils perçoivent des bulletins des autres votants. Les propriétés de convergence des dynamiques de jeu dans des processus de vote itératif ont été largement étudiées dans des travaux récents. Les résultats principaux dans le cadre classique du vote itératif sont la garantie de convergence vers un équilibre de Nash pour des déviations unilatérales implémentant des *meilleures réponses directes* sous les règles de vote Pluralité et Véto [Meir et al., 2010, Reyhani and Wilson, 2012]. Une meilleure réponse directe sous Pluralité consiste pour un votant à dévier vers un nouveau bulletin approuvant un candidat devenant le nouveau vainqueur et qui est le meilleur possible selon les préférences du votant en question. Pour Véto, la meilleure réponse directe possible consiste à faire élire un nouveau vainqueur que l'on préfère au précédent en déviant vers un bulletin contenant un véto contre l'ancien vainqueur. Il semble en particulier que seules les déviations unilatérales aient été étudiées jusqu'à présent en vote itératif.

Dans ce cadre classique du vote itératif, il est fait l'hypothèse que tous les agents connaissent le bulletin de tous les autres. Ceci apparaît comme une hypothèse assez irréaliste, en particulier pour des instances de grande taille. Par conséquent, de nombreux travaux élaborent des modèles intégrant de l'incertitude dans des jeux de vote. Certains modèles considèrent que les votants pensent en termes de probabilités afin de déterminer les profils de vote les plus vraisemblables [Myerson and Weber, 1993, Messner and Polborn, 2005, Hazon et al., 2008]. D'autres se servent des outils de la logique modale afin de formaliser le concept de manipulation en présence d'information incomplète [Chopra et al., 2004, Van Ditmarsch et al., 2012]. Dans l'approche de *dominance locale* [Meir et al., 2014, Meir, 2015], des seuils d'incertitude sont introduits afin de déterminer les profils de vote possibles, qui seraient tous ceux se trouvant à une distance au plus ce seuil du profil de vote réel en considérant une certaine métrique. Une autre approche consiste à prendre en compte une certaine fonction d'information représentant le type d'information communiquée aux votants, qui peut être le vainqueur, les scores des candidats ou autre [Reijngoud and Endriss, 2012, Endriss et al., 2016]. La susceptibilité de différentes règles de vote à être manipulables considérant un certain type de fonction d'information est étudiée en faisant l'hypothèse que les votants adoptent un comportement averse au risque sur l'ensemble des profils qui sont cohérents avec l'information qui leur est donnée. Une autre manière de gérer l'incertitude est de supposer que les votants considèrent tous les profils de vote cohérents avec des informations partielles sur le vote des autres agents [Conitzer et al., 2011, Dey et al., 2016].

Enfin, le vote stratégique combiné avec des réseaux sociaux commence à être étudié [Grandi et al., 2017]. Il est notamment fait l'hypothèse que les votants infèrent de leurs liens dans le réseau les bulletins courants de certains votants [Chopra et al., 2004, Sina et al., 2015, Tsang and Larson, 2016].

Allocation de ressources avec des biens indivisibles

L'allocation de ressources avec des biens indivisibles [Chevaleyre et al., 2006, Bouveret et al., 2016] est un champ de recherche primordial de l'intelligence artificielle, à l'intersection entre l'économie et l'informatique, avec de nombreuses applications dans la vie réelle. Dans un problème d'allocation de ressources, on dispose d'un ensemble d'agents et d'un ensemble d'objets (aussi appelés ressources ou biens) et les agents ont des préférences sur les objets. Le but est de répartir les objets entre les agents de la manière la plus efficace et la plus équitable possible. On suppose que les ressources ne sont ni partageables ni divisibles entre les agents, et que les agents ont des préférences ordinales sur les objets. Le cas de ressources divisibles, souvent connu sous le nom de problème de *partage de gâteau* ("cake cutting") a également fait l'objet d'intenses investigations [Steinhaus, 1948, Robertson and Webb, 1998, Procaccia, 2016].

Des mesures d'équité et d'efficacité peuvent être élaborées dans des problèmes d'allocation de ressources afin d'identifier les allocations souhaitables pour la société. Dans le cadre de préférences ordinales, une notion d'efficacité classique et pertinente est l'*efficacité de Pareto*. Une allocation est dite Pareto-efficace s'il n'est pas possible d'améliorer la satisfaction d'un agent sans détériorer celle d'un autre, c'est-à-dire s'il n'existe pas d'autre allocation telle que chaque agent est au moins autant satisfait que dans l'allocation de départ et au moins un d'entre eux est strictement plus satisfait. En ce qui concerne l'équité, plusieurs mesures ont été élaborées dans la littérature. On peut notamment citer la *proportionnalité* et l'*absence d'envie*. La proportionnalité [Steinhaus, 1948] est utilisée spécifiquement dans un contexte où les agents ont des préférences cardinales et requiert que chaque agent obtienne une part qu'il évalue comme au moins son utilité pour l'ensemble des objets divisée par le nombre d'agents. L'absence d'envie [Tinbergen, 1946, Foley, 1967, Varian, 1974] impose quant à elle qu'aucun agent n'envie ce qui est attribué à quelqu'un d'autre, c'est-à-dire que chaque agent préfère au moins sa part à celle des autres. Il est à noter que pour que la notion d'absence d'envie ait un sens, elle doit être combinée avec un critère minimal d'efficacité, comme par exemple la complétude de l'allocation considérée, qui implique que tous les objets soient distribués. Autrement, même l'allocation n'affectant aucun objet est sans envie. On se restreint d'ailleurs dans les problèmes d'allocation de ressources considérés à des allocations complètes.

On se penche dans la thèse sur un problème particulier d'allocation de biens indivisibles, appelé *allocation de maisons* ("house allocation") [Hylland and Zeckhauser, 1979, Abdulkadiroğlu and Sönmez, 1998, 1999], où chaque agent doit recevoir exactement un bien. On peut par exemple penser à des problèmes d'affectation d'étudiants à des appartements, ou l'allocation de créneaux horaires à des employés ou de rendez-vous dans un service administratif. Dans la littérature économique, ce problème est également connu sous le nom de *couplage unilatéral* ("one-sided matching") [Zhou, 1990], un type spécifique de couplage avec préférences [Manlove, 2013, Klaus et al., 2016] où seulement une des deux parties a des préférences sur l'autre, ici les agents sur les ressources. Les questions de Pareto-efficacité et d'équité ont notamment été examinées pour ce problème particulier [Bogomolnaia and Moulin, 2001, Abraham et al., 2005]. Le protocole simple de *dictature en série* est particulièrement pertinent dans ce contexte et est notamment utilisé dans des applications réelles. Étant donné un ordre sur les agents, chacun d'eux est appelé à son tour à choisir un objet parmi les objets disponibles restants. Si l'ordre est généré de manière aléatoire, on parle de *dictature en série aléatoire* ("random serial dictatorship") [Abdulkadiroğlu and Sönmez, 1998]. Ce protocole a le mérite d'être non-manipulable.

La définition de l'envie est très simple dans le cadre du problème d'allocation de maisons : un agent en envie un autre s'il préfère l'objet affecté à l'autre agent à son propre objet. Cependant, dans ce cadre, l'absence d'envie est une condition très exigeante et rare car il faudrait considérer une allocation où tous les agents obtiennent leur meilleur objet et donc les instances positives sont restreintes à celles pour lesquelles les préférences des agents sont telles que leur meilleur objet est différent.

Il existe un autre sous-problème spécifique dans le cadre de l'allocation de maisons, qui consiste à supposer que les agents ont initialement un objet, c'est-à-dire qu'il existe une allocation initiale. Le but est alors de réallouer les objets de la manière la plus efficace et la plus équitable possible au sein des agents. Ce problème est connu sous le nom de *marché du logement* ("housing market") [Shapley and Scarf, 1974]. Dans un tel contexte, la manière la plus naturelle pour réallouer les ressources repose sur des échanges entre les agents. Deux perspectives peuvent être adoptées pour l'échange d'objets. La première, distribuée, consiste à laisser les agents échanger par eux-mêmes et analyser quelles peuvent être les conditions d'échanges réalistes et prévoir vers quelle allocation la séquence d'échanges va converger. Dans la seconde, centralisée, un coordinateur externe peut guider les agents dans leurs échanges afin de les orienter vers une solution acceptable pour la communauté.

Des algorithmes centralisés très efficaces ont été conçus pour le problème de marché du logement, à savoir le célèbre algorithme "top trading cycle" (TTC) [Shapley and Scarf, 1974] et certaines de ses variantes [Abdulkadiroğlu and Sönmez, 1999, Aziz and De Keijzer, 2012]. L'algorithme TTC garantit l'obtention d'une allocation qui est, entre autres propriétés, Pareto-efficace. Le principal inconvénient de l'approche centralisée est l'obligation pour les agents de communiquer une partie de leurs préférences à une autorité externe à laquelle ils doivent faire confiance.

Alternativement, les agents peuvent effectuer la réallocation de manière distribuée en échangeant et négociant directement ensemble. Moins exigeante en termes de coûts de communication, cette approche présente l'avantage de l'indépendance des agents vis-à-vis d'une quelconque entité extérieure. Néanmoins, elle peut aboutir à des allocations moins intéressantes pour les agents à cause de leur myopie dans les échanges qui ne leur permet pas d'avoir une vision globale comme pourrait l'avoir une autorité externe. Certains travaux examinent ce processus d'échanges décentralisé, en déterminant des conditions réalistes pour les échanges et analysant la qualité des allocations qui peuvent être atteintes. Très étudié dans le cadre de plusieurs ressources par agent [Sandholm, 1998, Chevaleyre et al., 2017], cette approche n'a été que récemment introduite pour le problème plus précis de marché du logement [Damamme et al., 2015].

Théorie de la complexité

On étudie dans la thèse des problèmes de choix social sous plusieurs angles et principalement celui de la complexité computationnelle. Des concepts fondamentaux de la théorie de la complexité sont donc rappelés. On utilise des notions de complexité classiques, avec principalement des preuves de NP-difficulté, mais aussi des notions de complexité paramétrée. En ce qui concerne la complexité paramétrée, on se ramène notamment à des preuves d'appartenance à FPT, avec la présentation d'algorithmes où la difficulté est circonscrite à une fonction dépendant du paramètre, rendant alors l'algorithme efficace lorsque le paramètre en question est petit, des preuves d'appartenance à XP, avec des algorithmes dont la complexité devient polynomiale si le paramètre est une constante, et des preuves de difficulté dans les hiérarchies W et A, qui se situent entre les classes FPT et XP dans la hiérarchie des classes de complexité paramétrée.

Partie I : Le réseau social comme outil collaboratif

Tout d'abord, on se propose d'étudier comment le réseau social peut modéliser la possibilité de collaboration entre agents. La collaboration renvoie à l'idée d'agents œuvrant ensemble, sur un même pied d'égalité, en combinant leurs ressources, afin d'accomplir un même but ou des buts compatibles. Il s'ensuit que pour qu'il y ait collaboration entre deux agents, la relation de possibilité d'interaction entre ces deux agents doit nécessairement être symétrique. On fait donc l'hypothèse pour cette partie que le réseau social est représenté par un graphe non-orienté sur les agents.

Chapitre 2 : Manipulation par coalitions en vote itératif

Tout d'abord, on se concentre sur un problème de vote stratégique où la collaboration entre agents est matérialisée par des agents pouvant manipuler ensemble au sein de coalitions déterminées par le réseau social. Ce chapitre se réfère à notre article publié dans la conférence ECAI-16 [Gourvès et al., 2016].

Dans le contexte du vote stratégique modélisé en tant que jeu de vote, ne considérer que des déviations unilatérales dans les bulletins soumis, renvoyant ainsi au concept d'*équilibre de Nash*, comme il est communément admis dans le cadre classique du vote itératif, apparaît comme une hypothèse trop faible. Premièrement, le nombre d'états du jeu étant équilibres de Nash est trop important pour prédire convenablement l'issue possible du jeu. Dans un second temps, la restriction à des déviations unilatérales en elle-même ne permet pas de couvrir tous les cas possibles de manipulation. En effet, dans la vraie vie, il peut arriver que plusieurs électeurs s'entendent sur une stratégie commune à adopter. Les instructions de vote au sein de partis politiques en sont un exemple typique.

D'un autre côté, si l'on considère la manipulation par coalitions d'agents, supposer que tout sous-ensemble d'agents est une possible coalition en vue de voter stratégiquement, renvoyant ainsi au concept d'*équilibre fort*, semble une hypothèse trop forte. Premièrement, un tel équilibre existe rarement dans des jeux de vote. De plus, pour des raisons diverses, allant de la distance à des problèmes possibles de communication, en passant par la question des affinités entre individus, il n'est pas toujours possible qu'un groupe d'agents puisse collaborer. Par conséquent, en pratique, on peut raisonnablement exclure certaines coalitions, lors de la définition d'un état stable, dans l'ensemble des groupes d'agents pouvant dévier à prendre en compte. Par exemple, il est possible de ne considérer que des coalitions d'électeurs qui sont membres d'une partition donnée, comme dans le concept d'*équilibre de partition* [Feldman and Tennenholtz, 2009].

Une façon de prendre en compte des coalitions réalistes d'électeurs consiste à exploiter les réseaux sociaux. En effet, si les liens dans le réseau social caractérisent pleinement les relations sociales entre les agents, cela fait sens de considérer comme des coalitions possibles, des groupes d'agents déterminés à partir du graphe représentant le réseau social. En particulier, puisque l'établissement d'une stratégie commune implique la participation active de tous les membres de la coalition, supposer qu'une coalition possible est donnée par une clique du graphe apparaît pertinent. Pour aller plus loin, on peut également supposer qu'un agent est lié par des relations sociales qui le forcent à prendre en considération les autres participants. Par conséquent, on fait l'hypothèse qu'un agent est non seulement guidé par ses propres préférences sur l'ensemble des candidats, mais poursuit également l'objectif d'optimisation du bien-être des communautés auxquelles il appartient. Un concept de solution qui découle naturellement de cette configuration est l'*équilibre de considération* [Hoefer et al., 2011].

Un équilibre de considération est défini comme un état où aucune coalition d'agents correspondant à une clique du réseau social n'a intérêt à dévier au sein d'une stratégie commune qui ne nuit pas au bien-être de ses voisins dans le réseau social. Il est à noter qu'un équilibre de partition est un sous-cas de l'équilibre de considération pour lequel le graphe associé est un ensemble de cliques disjointes (graphe constitué de "clusters"). C'est également le cas pour l'équilibre de Nash, qui correspond à un équilibre de considération pour lequel le graphe associé est vide (pas d'arêtes dans le graphe). Il s'ensuit que si un équilibre de considération existe dans un certain jeu de vote quel que soit la structure du réseau social, alors un équilibre de partition ainsi qu'un équilibre de Nash sont garantis d'exister.

L'équilibre de considération, tout comme l'équilibre de partition, n'ayant été étudié que pour un cas particulier de jeu de congestion [Anshelevich et al., 2013b, Hoefer et al., 2011], on se propose d'étudier, pour différentes règles de vote classiques, l'existence d'un

		Règles de notation positionnelles				Eliminations successives		Comparaison par paires	
		Pluralité	Véto	k-approbation	Borda	STV	Pluralité à 2 tours	Copeland	Maximin
Eq. fort	Existence	×	✔	×	×	×	×	×	×
	Convergence	×	×	×	×	×	×	×	×
Eq. de considération	Existence	✔	✔	×	×	✔	✔	×	✔
	Convergence	×	×	×	×	×	×	×	×
Eq. de partition	Existence	✔	✔	×	×	✔	✔	×	✔
	Convergence	×	×	×	×	×	×	×	×
Eq. de Nash	Existence	✔	✔	×	×	✔	✔	×	✔
	Convergence	✔[1]	✔[2]	×[2]	×[2]	×	×	×	×

[1] [Meir et al., 2010]
[2] [Lev and Rosenschein, 2012, Reyhani and Wilson, 2012]

Tableau 3: Tableau récapitulatif des résultats d'existence d'équilibres (pour différents concepts de solution) et de convergence des dynamiques de jeu associées selon différentes règles de vote

équilibre de considération dans un jeu de vote stratégique et la capacité des dynamiques de jeu à converger vers un tel équilibre.

Différentes règles de vote classiques sont examinées : des règles de notation positionnelles telles que Pluralité, Véto, Borda et k-approbation, des règles à éliminations successives comme la Pluralité à deux tours et le Vote simple transférable (STV), et enfin des règles par comparaison de paires avec Maximin et Copeland. Bien que l'on prouve l'existence d'un équilibre de considération pour de nombreuses règles de vote, on montre que la dynamique de déviations n'est presque jamais garantie de converger vers un tel état stable. Ce résultat négatif de convergence est valable même lorsque l'on se restreint aux cas plus simples de déviations par des coalitions d'agents venant d'une partition ou de déviations unilatérales. Un récapitulatif des résultats d'existence et de convergence est présenté dans le Tableau 3.

Plus précisément, on prouve que dans les jeux de vote sous Pluralité, STV, Maximin ou Pluralité à deux tours, un équilibre de considération existe quelle que soit la structure du réseau social. Plus généralement, on prouve ce résultat d'existence pour des règles de vote dites "sensibles à la stricte majorité", qui englobent notamment les règles de votes spécifiques précédemment citées. Les règles sensibles à la stricte majorité sont telles que seules des coalitions de plus de la majorité des votants peuvent changer le résultat de tout état et seules des coalitions composées d'exactement la moitié des votants peuvent changer le résultat d'un profil unanime, où tous les votants donnent le même bulletin, sans pour autant pouvoir changer le résultat du profil unanime favorisant le candidat privilégié par la règle de départage des ex aequo. En revanche, on exhibe des contre-exemples pour les autres règles de vote étudiées, montrant que même un équilibre de Nash n'est pas garanti d'exister sous les règles de vote de Borda, de Copeland et k-approbation. Néanmoins, lorsque le nombre de votants est strictement supérieur à deux, un équilibre de Nash est garanti d'exister sous Borda et Copeland, ce qui n'est pas le cas pour un équilibre de partition. Mais, pour le cas d'un nombre impair de votants, un équilibre de considération existe toujours sous Copeland, résultat qui ne tient pas sous Borda.

Le cas de la règle de vote Véto est particulier. Bien que n'étant pas une règle sensible à la stricte majorité, un équilibre de considération est garanti d'exister pour toute structure de réseau social. Mais, on prouve également un résultat particulièrement intéressant : un équilibre fort existe toujours dans un jeu de vote sous Véto. La preuve est basée sur l'utilisation du mécanisme "feasible elimination procedure" introduit par Peleg [1978]. De plus, on montre que savoir si un candidat particulier peut être élu dans un état qui est un équilibre fort dans le jeu de vote sous Véto peut être déterminé en temps polynomial. On en déduit une mesure de qualité pour évaluer une règle de vote, qui est valable pour tout profil de préférence (contrairement au critère de Condorcet). Il s'avère empiriquement que cette mesure de qualité semble assez conflictuelle par rapport à l'efficacité de Condorcet.

Concernant la convergence des dynamiques de déviations, les résultats sont plutôt négatifs puisque l'on montre que la dynamique associée à un équilibre de partition peut comporter des cycles sous Pluralité et Véto. Ce résultat est valable même pour des déviations restreintes aux meilleures réponses directes, alors que sous cette restriction la dynamique des déviations unilatérales est connue pour converger [Meir et al., 2010, Reyhani and Wilson, 2012] sous ces règles de vote. Pour toutes les autres règles de vote étudiées, la convergence n'est pas garantie même pour des dynamiques associées à des déviations unilatérales.

Des résultats expérimentaux sont également présentés dans ce chapitre, évaluant la fréquence d'existence des différents concepts de solution analysés ainsi que la fréquence de convergence des dynamiques associées. Des simulations sont faites sur la base de préférences générées aléatoirement de manière uniforme ou aléatoirement "single-peaked", et de réseaux sociaux qui sont soit des graphes de type Erdös-Rényi générés pour des densités particulières, soit des graphes avec homophilie. La qualité des équilibres est évaluée en fonction de différents critères, dont l'efficacité de Condorcet, la proximité des scores de Borda et la possibilité pour le gagnant d'être élu dans un état étant équilibre fort sous Véto. Il s'avère que l'équilibre de considération a le même défaut que l'équilibre de Nash concernant le trop grand nombre d'états qui sont des équilibres de considération. En revanche, la qualité des équilibres atteints par la dynamique de déviations est clairement meilleure que celle des équilibres de partition ou des équilibres de Nash, et même clairement meilleure que ceux d'équilibres coalitionnels où les coalitions sont des cliques, comme dans l'équilibre de considération, mais où les agents n'ont pas de considération pour leurs voisins dans le réseau lorsqu'ils effectuent des déviations.

L'hypothèse de considération, qui rend les agents altruistes dans leurs déviations, est importante car c'est par elle que tiennent nos résultats d'existence des équilibres de considération, mais aussi par elle que les équilibres de considération sont si nombreux pour une instance donnée. L'équilibre de partition semble être un bon compromis au niveau du nombre d'états stables et vis-à-vis de la qualité des équilibres atteints, sans qu'une hypothèse de considération soit nécessaire.

Chapitre 3 : Dynamique d'échanges dans des problèmes d'allocation de maisons

Dans ce chapitre, on s'intéresse à une collaboration entre agents reliés dans le réseau social, au sein d'un processus d'allocation de ressources où les agents peuvent s'échanger des biens. Ce chapitre se réfère à nos articles publiés dans la conférence IJCAI-17 [Gourvès et al., 2017] et dans la conférence SAGT-18 [Saffidine and Wilczynski, 2018].

Ce chapitre étudie un problème spécifique d'allocation de ressources dans lequel chaque agent doit recevoir exactement un bien, problème connu sous le nom d'allocation de maisons ("house allocation"). Plus précisément, chaque agent est initialement doté d'un objet et donc le but est de réallouer les biens aux agents de manière efficace et équitable, en fonction de leurs préférences exprimées sous forme d'ordre sur les objets, ce qui renvoie plus particulièrement à un problème de marché du logement ("housing market"). Une manière naturelle de réallouer les objets dans ce cadre consiste à laisser les agents effectuer des échanges entre eux, de façon distribuée. Néanmoins, ce processus peut également être guidé de manière centralisée par un coordinateur externe. C'est le cas de l'algorithme "top trading cycle" qui permet d'obtenir, par le biais d'échanges successifs choisis entre agents, une allocation qui est entre autres Pareto-efficace. Dans ce processus, il est implicitement supposé que tout agent est à même d'échanger avec n'importe quel autre. Or, pour des raisons variées allant de problèmes de communication à des problèmes d'affinités en passant par des contraintes géographiques, tout échange n'est pas susceptible de survenir.

On se propose dans ce chapitre de représenter par un réseau social les échanges possibles entre agents. Ainsi, on considère qu'un échange n'est possible que s'il intervient entre deux voisins dans le réseau social, modélisé par un graphe non-orienté. On suppose également que tout échange est rationnel, dans le sens où les deux agents impliqués dans l'échange préfèrent l'objet actuellement détenu par l'autre agent par rapport à l'objet qu'ils possèdent. Dans cette perspective, les agents sont vus comme "myopes" car ils souhaitent obtenir directement un objet intéressant sans accepter de passer par un objet moins apprécié afin d'obtenir un meilleur objet dans une stratégie à long terme.

Afin d'illustrer un tel problème d'échanges, on peut penser à des plateformes d'échanges en ligne où chaque participant renseignerait l'objet qu'il possède actuellement et souhaiterait échanger, et exprimerait des préférences sur les objets qu'il désirerait obtenir. Dans ce cadre, le réseau social pourrait naturellement modéliser des contraintes géographiques car on peut raisonnablement croire que les agents ne seraient pas enclins à parcourir une distance importante pour l'échange. De plus, la rationalité des échanges fait également sens dans ce contexte puisque les agents ne sont pas garantis que d'autres ne quittent la plateforme entre-temps, pouvant ainsi mettre à mal certaines stratégies à long terme, et donc il peut sembler moins risqué de ne faire que des échanges directement améliorants.

Ce chapitre est dédié à l'étude de la *dynamique d'échanges*, dans une perspective distribuée, dans laquelle les agents, partant d'une allocation initiale, effectuent des échanges rationnels entre voisins dans le réseau social, jusqu'à l'obtention d'une allocation *stable* où aucun échange de ce type n'est plus possible. Dans ce cadre, les allocations *atteignables* par une séquence d'échanges sont particulièrement intéressantes puisque ce sont les seules allocations susceptibles de se produire.

Trois problèmes sont étudiés afin d'analyser la dynamique d'échanges : REACHABLE OBJECT (RO), sur la possibilité pour un certain agent d'obtenir un objet particulier par le biais d'une séquence d'échanges, REACHABLE ASSIGNMENT (RA), posant la question de l'atteignabilité d'une allocation donnée, et enfin GUARANTEED LEVEL OF SATISFACTION (GLS), posant la question de la garantie pour un certain agent d'obtenir toujours un objet parmi ses k premiers préférés dans toute allocation atteignable stable. La complexité de ces trois problèmes est examinée, soit en complexité classique avec des preuves de difficulté ou des algorithmes efficaces pour des cas particuliers des problèmes, soit en

		Pas de contraintes	Contraintes de budget (paramètre k)		
			max	sum	makespan
RO	Graphe général	NP-c	NP-c pour $k \geq 2$	W[1]-c	W[1]-c
	Arbre	NP-c	W[SAT]-difficile	W[1]-c	W[1]-c
	Degré borné	NP-c	NP-c pour $k \geq 2$	FPT	FPT
RA	Graphe général	NP-c	NP-c pour $k \geq 3$	FPT	W[1]-difficile
	Arbre	P	P	P	P
	Degré borné	NP-c	NP-c pour $k \geq 3$	FPT	?
GLS	Graphe général	co-NP-c	co-NP-c pour $k \geq 3$	co-W[1]-difficile / co-A[2]	co-W[1]-difficile / co-A[2]
	Arbre	co-NP-c	co-W[SAT]-difficile	co-W[1]-difficile / co-A[2]	co-W[1]-difficile / co-A[2]
	Degré borné	co-NP-c	co-NP-c pour $k \geq 3$	co-A[2]	co-A[2]

Tableau 4: Résultats globaux de complexité classique et paramétrée des problèmes RO, RA and GLS

complexité paramétrée en considérant des paramètres naturels vis-à-vis de la séquence d'échanges choisie. L'idée est qu'il paraît raisonnable de supposer que les agents ne sont pas prêts à effectuer un grand nombre d'échanges avant d'obtenir l'objet qu'ils désirent. Plus précisément, trois paramètres sont pris en compte : le paramètre *max* renvoyant au nombre d'échanges maximum fait par un agent dans la séquence, le paramètre *sum* concernant le nombre d'échanges total dans la séquence d'échanges, et enfin le paramètre *makespan* qui est la longueur de la séquence en considérant que tous les échanges qui peuvent être effectués en parallèle, car n'impliquant pas les mêmes agents, le sont effectivement.

On prouve que les problèmes RO, RA et GLS sont difficiles d'un point de vue computationnel dans le cas général et ce pour RO et GLS, même lorsque le réseau social est un arbre assez simple. En revanche, RA s'avère être résoluble en temps polynomial lorsque le réseau social est un arbre. Pour ce qui est de l'approche paramétrée, le paramètre max ne permet pas de contourner la difficulté computationnelle intrinsèque des problèmes RO et GLS puisqu'ils demeurent difficiles en fonction de ce paramètre, même dans un arbre. Néanmoins, les paramètres s'appuyant sur la taille de la séquence d'échanges, en l'occurrence sum et makespan, permettent de limiter cette difficulté computationnelle. Bien qu'ils ne permettent pas de rendre plus facile la résolution des problèmes dans le cas général, ils permettent de circonscrire leur difficulté à des classes de complexité paramétrée qui ne sont pas si élevées dans la hiérarchie. Ceci implique notamment la possibilité de résoudre ces problèmes efficacement lorsque le paramètre est une constante, hypothèse qui paraît assez naturelle si les agents ne sont pas prêts à attendre un long laps de temps avant d'obtenir l'objet qu'ils désirent. De plus, ces paramètres permettent de rendre RO résoluble efficacement du point de vue de la complexité paramétrée (FPT) quand le graphe est de degré borné, ce qui peut sembler une hypothèse naturelle si le réseau est supposé représenter les relations proches de chaque agent. En ce qui concerne GLS, ce problème s'avère être très proche du problème complémentaire de RO et donc les résultats de complexité de RO s'appliquent pour GLS vis-à-vis des classes de complexité complémentaires. Pour une vue globale des résultats de complexité sur RO, RA et GLS, se référer au Tableau 4.

Enfin, un dernier problème est étudié, évoquant un processus centralisé dans l'idée de l'algorithme TTC. Le but est de guider les échanges entre les agents afin d'atteindre une allocation Pareto-efficace au sein des allocations atteignables. Il s'agit d'un problème d'optimisation et non de décision car, autant une allocation Pareto-efficace n'est pas garantie d'être atteinte, autant une allocation atteignable qui est Pareto-efficace au sein de l'ensemble des allocations atteignables existe forcément. Il est clair qu'une telle allocation est forcément stable puisque l'on se restreint à des échanges rationnels. On prouve que le problème consistant à trouver une telle allocation atteignable est computationnellement difficile mais notamment, le cas où le graphe est un chemin ou une étoile sont résolubles en temps polynomial. Le cas un peu plus général de l'arbre reste ouvert.

En règle générale, les problèmes abordés s'avèrent difficiles à résoudre, même pour des structures de graphe assez simples. Des cas polynomiaux peuvent néanmoins être identifiés dans des réseaux exhibant des structures très particulières comme des chemins, des étoiles, des graphes à degré borné ou encore, pour certains problèmes (RA), des arbres. De plus, du point de vue de la complexité paramétrée, certains paramètres portant sur la taille de la séquence d'échanges permettent de limiter la difficulté computationnelle des problèmes.

Partie II : Le réseau social comme outil d'information

On étudie dans cette partie comment le réseau social peut déterminer l'information qui est disponible pour les agents. Ainsi, le réseau social traduit plutôt une relation de visibilité entre les agents. Cette relation n'a pas de raison d'être symétrique par défaut et donc, le graphe est supposé être simplement un graphe orienté.

Chapitre 4 : Absence d'envie locale en allocation de maisons

Ce chapitre s'intéresse à une mesure d'envie *locale* déterminée par un graphe, dans des problèmes d'allocation de ressources avec biens indivisibles où chaque agent doit recevoir exactement un objet. Une certaine relation de visibilité entre les agents donnée par le réseau social restreint les possibilités d'envie à prendre en compte entre les agents. Ce chapitre étend notre article publié dans la conférence AAMAS-18 [Beynier et al., 2018].

Dans le cadre d'un problème d'allocation de maisons sur lequel on se concentre, la notion d'envie qui consiste à imposer qu'aucun agent ne préfère l'objet affecté à un autre agent par rapport au sien, est très exigeante puisqu'elle consiste à donner à chaque agent son objet préféré. On se propose dans ce chapitre de relâcher cette mesure d'envie en considérant qu'un agent ne peut envier que ses successeurs dans le réseau social, représenté par un graphe orienté, et modélisant une relation de visibilité entre les agents. En effet, pour de nombreuses raisons, tous les agents ne sont pas nécessairement visibles pour tout agent. Mais la visibilité est cependant importante dans la définition de l'envie : peut-on vraiment dire que l'on envie le détenteur d'un objet donné, bien que l'on ne sache pas de qui il s'agit ? D'un autre côté, d'un point de vue centralisé, le graphe peut également modéliser les seuls liens dont le coordinateur extérieur se soucie, par exemple le fait de ne pas créer d'envie entre membres d'une même équipe ou d'employés aux compétences équivalentes.

Ainsi, on s'intéresse en particulier dans ce chapitre aux allocations *localement sans* *envie*, qui sont telles qu'aucun agent ne préfère l'objet qui est affecté à l'un de ses successeurs dans le graphe à l'objet qui lui est affecté. Plusieurs problèmes sont alors analysés.

Premièrement, la question de l'existence d'une allocation localement sans envie est traitée, par le biais du problème de décision DEC-LEF. La complexité de ce problème est examinée, au travers d'une étude de complexité classique et une approche paramétrée en fonction de paramètres du graphe, comme sa couverture de sommets minimale ou son degré. Les résultats de complexité pour ce problème sont présentés dans le Tableau 5 où n renvoie au nombre d'agents et k au paramètre considéré, dans une perspective de complexité paramétrée. Le problème est prouvé difficile même pour un graphe particulièrement simple reposant sur un couplage. De plus, le problème reste difficile même lorsque le graphe est seulement composé de deux cliques disjointes, ou dans des graphes relativement denses où chaque agent est relié à tous les autres agents sauf deux. En revanche, dans des graphes très denses où chaque agent est relié à tout autre ou à tout autre sauf un, le problème de l'existence d'une allocation localement sans envie peut être résolu en temps polynomial. Par ailleurs, il existe un protocole simple basé sur un processus de dictature en série ("serial dictatorship") qui permet de garantir l'existence et de construire une allocation localement sans envie dans des graphes orientés sans circuit. En ce qui concerne la taille de la couverture de sommets minimale du graphe, considérer ce paramètre permet d'obtenir un algorithme polynomial pour DEC-LEF si ce paramètre est une constante, c'est-à-dire que le problème est dans XP. Seulement, sous les hypothèses de complexité usuelles, il n'y a pas d'algorithme permettant de circonscrire la difficulté algorithmique du problème à une fonction dépendant uniquement de ce paramètre puisque l'on prouve que DEC-LEF est W[1]-difficile avec ce paramètre.

degré extérieur δ^+ du graphe	$\delta^+ \leq k$ ($k \geq 1$ fixé)	NP-c
	$\delta^+ \geq n - k$ ($k \geq 3$ fixé)	NP-c
	$\delta^+ \geq n - 2$	P
nombres de clusters c formant le graphe	$c = n/k$ ($k \geq 2$ fixé)	NP-c
	$c = k$ ($k \geq 2$ fixé)	NP-c
	$c = 1$ ou $c = n$	P
paramètre k sur la taille de la couverture de sommets dans un graphe supposé non-orienté		XP W[1]-difficile

Tableau 5: Résultats de complexité de DEC-LEF

Dans le cas où une allocation parfaitement sans envie n'existe pas, on peut vouloir trouver une allocation qui s'en approche le plus possible. La question de l'optimisation de l'absence d'envie locale est alors abordée, à travers deux mesures différentes pour l'envie locale : le nombre d'agents localement sans envie et une certaine moyenne d'envie locale de la société. On prouve que dans le sous-cas du graphe non-orienté, il existe des algorithmes d'approximation efficaces pour l'optimisation du nombre d'agents non-envieux, utilisant la notion d'ensemble stable dans un graphe, qui est un sous-ensemble de sommets n'ayant aucune arête reliant deux de leurs éléments. Pour l'optimisation d'un certain degré d'envie locale, on exhibe un algorithme d'approximation reposant sur la dérandomisation.

Puis, on s'intéresse à un problème où l'autorité extérieure a encore plus de pouvoir, dans le sens où elle n'affecte pas seulement les objets aux agents, mais affecte également les agents à des positions dans le graphe, de manière à trouver une allocation localement sans envie. L'idée est que la structure du graphe est donnée, mais sans affectation des agents à des places précises. Une interprétation directe pour ce problème serait que le graphe représenterait les plages horaires d'un emploi du temps conditionnant la visibilité des agents et sur laquelle un coordinateur externe pourrait avoir une certaine emprise. Ce problème s'avère difficile d'un point de vue computationnel, mais dans le cas d'un graphe non-orienté très dense, où tous les agents sont reliés à tous les autres ou tous les autres sauf un, le problème peut être résolu efficacement.

Enfin, la question de l'atteignabilité d'une allocation localement sans envie est également étudiée, où le graphe modélise à la fois une possibilité d'échange entre agents dans le sens du Chapitre 3 et une possibilité d'envie. On se demande alors si les agents, par le biais d'échanges rationnels successifs entre voisins du graphe, peuvent atteindre une allocation localement sans envie. Pour ce cas, on suppose également que le graphe est non-orienté car les échanges ne peuvent se produire que si la relation d'interaction entre agents est symétrique. Sur la base des preuves de difficulté pour l'atteignabilité du Chapitre 3, on prouve que ce problème est difficile à résoudre même lorsque le réseau est un arbre. En revanche, certains cas polynomiaux peuvent être exhibés si l'on restreint encore plus la structure du graphe. C'est le cas notamment de l'étoile.

Pour finir, des simulations sont effectuées afin d'avoir une idée plus claire du comportement de l'envie locale. En se basant sur des graphes générés aléatoirement, de type Erdös-Rényi avec différentes densités, réguliers avec différents degrés, ou de type Bárabasi-Albert, on confirme l'intuition selon laquelle la fréquence d'existence d'allocations localement sans envie est plus importante lorsque le graphe est peu dense. De plus, on remarque que lorsque le graphe est généré avec homophilie, il y a très peu et très peu souvent d'allocation localement sans envie car des voisins dans le réseau ont tendance à préférer les mêmes objets. En revanche, lorsque c'est le graphe complémentaire qui est généré avec homophilie, on trouve de nombreuses et fréquentes allocations localement sans envie.

Ainsi, même si les problèmes étudiés s'avèrent difficiles d'un point de vue computationnel, certains cas naturels peuvent être résolus efficacement comme par exemple le cas où le graphe complémentaire, que l'on peut qualifier de graphe de non-envie, est un couplage, qui peut représenter une situation où la seule personne qu'un agent ne peut pas envier est son partenaire. De plus, dans des configurations où les agents n'envient pas les personnes qui leur sont proches bien que celles-ci aient des préférences similaires (suivant la tendance à s'affilier à ses semblables), de nombreuses allocations sont localement sans-envie, confirmant alors une intuition assez naturelle.

Chapitre 5 : L'incertitude dans des problèmes de vote itératif

Ce chapitre se penche sur un problème de vote itératif dans lequel la visibilité des bulletins des autres agents est conditionnée par un réseau social, représenté par un graphe orienté dont les nœuds sont les agents. Ce chapitre est une extension de notre article publié dans la conférence AAAI-19 [Wilczynski, 2019].

On s'intéresse dans ce chapitre à un moyen simple d'agréger les différentes informations disponibles pour un votant dans un processus de vote itératif, afin de gérer l'incertitude. En s'inspirant des élections politiques, on se concentre sur le vote itératif comme un moyen de représenter les changements de stratégie dans les intentions de vote opérés par les votants, de manière parfois successives, dans une période pré-électorale. Dans ce cadre, pour faire face à l'incertitude concernant l'intention de vote des autres votants, les agents disposent notamment d'une information globale, donnée par des sondages d'opinion et diffusés à tout l'électorat, et d'une information locale, constituée des intentions de vote des membres de leur entourage. Pour représenter cette dernière information, on suppose que les agents connaissent l'intention de vote courante de leurs successeurs dans le réseau social, modélisé par un graphe orienté. Cette modélisation nous permet de généraliser le cadre classique du vote itératif [Meir et al., 2010] où l'on relâche la forte hypothèse d'information complète vis-à-vis de l'état courant. Intégrer l'incertitude au sein d'un processus de vote itératif apparaît pertinent afin de rendre le modèle plus réaliste, mais contrairement à la plupart des travaux dans cette ligne de recherche, on se concentre sur des sources d'information réalistes, à savoir les sondages d'opinion et les réseaux sociaux, et l'on suppose que les agents adoptent une manière simple pour agréger ces sources d'information. En effet, on suppose que les agents ne suivent pas forcément un comportement averse au risque car, étant donné la grande quantité d'incertitude en raison du nombre important de votants, si c'était le cas, il n'y aurait pas de manipulation dans les élections politiques, alors qu'il y a effectivement de la manipulation de la part des électeurs. De plus, il apparaît que les électeurs ne pensent pas en matière de probabilités, ce qui constituerait un effort cognitif trop important.

Dans ce modèle, on se concentre par souci de simplicité sur des déviations unilatérales et bien que certains résultats soient généralisés à plusieurs règles de vote, l'accent est mis sur l'étude de la règle de vote Pluralité.

Plus précisément, on considère un modèle où les votants possèdent des croyances vis-à-vis de l'intention de vote courante des autres votants. Initialement, un sondage d'opinion, qui est de notoriété publique, leur donne un aperçu de la distribution des votes, en donnant le score des candidats du profil de vote initial, qui est supposé véridique, c'est-à-dire reflétant les véritables préférences des électeurs. Les votants basent leur croyance initiale sur ce sondage car ils ont confiance dans les résultats donnés par le sondage. De plus, ils sont capables d'observer les changements stratégiques dans l'intention de vote de leurs successeurs dans le réseau social, c'est-à-dire que s'il existe un arc entre deux agents dans le graphe modélisant le réseau social, le premier agent (origine de l'arc) connaît l'intention de vote courante du second agent (destination de l'arc). Ainsi, les votants peuvent mettre à jour leur croyance de l'état actuel en supposant, de manière simple et myope, que pour tous les électeurs qui n'appartiennent pas à leurs connexions dans le réseau social, l'intention de vote n'a pas changé depuis le premier sondage d'opinion.

Dans ce cadre, on suppose que les agents effectuent des déviations suivant un principe de meilleure réponse assez simple : ils choisissent de voter, sous Pluralité, pour le candidat qu'ils préfèrent parmi ceux qu'ils considèrent susceptibles de gagner. La détermination des candidats susceptibles de gagner à un moment donné est modélisée grâce à des seuils de pivot propres à chaque votant, c'est-à-dire que si la différence entre le score d'un candidat et le score de l'actuel gagnant dans les croyances de l'agent est inférieure à ce seuil, alors ce candidat est considéré comme un gagnant potentiel et l'agent pourrait envisager de voter pour lui. On peut noter que ce modèle se ramène au cadre classique du vote itératif dans le cas où le réseau social est un graphe complet et tous les agents ont un seuil égal à un.

A partir de ce comportement stratégique de meilleure réponse de la part des votants et de leurs croyances par rapport à l'état actuel, déterminées par le réseau social et le sondage d'opinion, on se concentre sur deux types de dynamiques de déviations : une dynamique *locale* dans laquelle un seul sondage d'opinion initial donne lieu à des déviations stratégiques de la part des votants jusqu'à atteindre un état d'équilibre ou un cycle, et une dynamique *globale* caractérisée par plusieurs sondages d'opinion dans laquelle on considère plusieurs dynamiques locales dans le sens où si le processus de vote itératif atteint un équilibre local, les résultats de cet équilibre sont communiqués à travers un nouveau sondage public à l'ensemble des votants et une nouvelle dynamique locale recommence jusqu'à atteindre un cycle ou un équilibre dit global. On peut remarquer qu'un équilibre global est alors un équilibre de Nash puisqu'il y a information complète des scores des candidats à cet état.

Dans ce chapitre, ces deux dynamiques sont étudiées, vis-à-vis de leurs propriétés de convergence et de la qualité de leurs équilibres. On montre que si le réseau social est un graphe sans circuit, alors la dynamique locale est garantie de converger vers un équilibre, quels que soient la règle de vote choisie et les seuils de pivot considérés. D'un autre côté, si le réseau social est un graphe dont la relation binaire représentée est transitive, alors la dynamique locale converge lorsque les seuils de pivot sont homogènes au sein de l'électorat et si la stratégie de réponse considérée est garantie de converger lorsque le graphe est complet, c'est-à-dire si la convergence est assurée en cas d'information complète des scores des candidats. Pour un graphe général, on prouve que déterminer si une instance peut comporter un cycle dans sa dynamique locale est difficile computationnellement parlant. En ce qui concerne la convergence de la dynamique globale, on montre que même pour des graphes très simples, comme le graphe vide ne comportant aucun arc, la dynamique peut cycler. De plus, tout comme la dynamique locale, reconnaître les instances où la dynamique globale peut cycler dans un cas de graphe général est difficile.

Des simulations sont également effectuées afin d'avoir une idée de la convergence en pratique des dynamiques et de la qualité de leurs équilibres respectifs. On génère en particulier des graphes aléatoires de type Erdös-Rényi avec différentes densités, des graphes de type Bárabasi-Albert et des graphes avec homophilie. On remarque que malgré la possibilité d'obtenir des cycles pour des graphes généraux, la fréquence de convergence en pratique s'avère assez élevée, en particulier pour des graphes très denses ou très épars. Concernant la qualité des équilibres, plusieurs mesures de qualité ont été examinées, et notamment l'efficacité de Condorcet, les scores de Borda et la mesure de qualité élaborée au Chapitre 2 concernant la fréquence d'élection d'un candidat élu dans un équilibre fort sous la règle Véto. Pour toutes ces mesures de qualité, les deux observations principales sont que les équilibres globaux ont une nette tendance à être meilleurs que les équilibres locaux, et que les équilibres pour les deux dynamiques tendent à être légèrement meilleurs lorsque le graphe se densifie. Ces constats confirment donc l'intuition que la quantité d'information possédée par les votants joue un rôle important dans le processus de vote itératif, permettant ainsi d'aboutir à des équilibres satisfaisant mieux le bien-être social lorsque les votants sont mieux informés.

La question d'être bien informé et le fait que les votants ont confiance dans les résultats des sondages, puisqu'ils se basent dans notre modèle sur les résultats communiqués, amènent à la question de la manipulation par l'institut de sondage. En effet, que se passerait-il si l'institut de sondage avait lui-même des préférences sur les candidats et souhaiterait voir certains candidats élus plutôt que d'autres ? L'institut de sondage a-t-il intérêt à mentir en communiquant de faux résultats de sondage ?

On se penche donc sur l'étude de deux problèmes concernant la possibilité de contrôle de l'élection par l'institut de sondage. Le premier est un problème de décision consistant à déterminer s'il est possible de communiquer des scores de candidats de la part de l'institut de sondage de manière à faire élire un candidat donné. Bien entendu, afin de satisfaire des conditions minimales de vraisemblance, les résultats communiqués doivent être cohérents avec l'information détenue par les votants. On prouve que ce problème est difficile d'un point de vue computationnel même pour un graphe sans circuit et des seuils de pivot homogènes égaux à un. Cependant, au-delà de la difficulté computationnelle au pire cas, qui n'empêche pas forcément la manipulation de se produire, on adopte également une approche heuristique pour la résolution de ce problème. Pour cela, on établit une heuristique relativement simple basée sur l'idée d'annoncer, autant que possible, seulement deux candidats comme gagnants potentiels : l'un est le candidat cible que l'on souhaite voir élu et l'autre, annoncé vainqueur, est un candidat largement détesté au sein de l'électorat. On espère ainsi qu'une large part de l'électorat va voter en faveur du candidat cible car il est présenté comme la seule alternative pour éviter l'élection du candidat largement détesté. Cette heuristique, bien que très simple, s'avère efficace en pratique d'après nos simulations, en particulier sur des graphes qui ne sont pas très denses.

Au lieu de vouloir l'élection d'un candidat précis, un autre problème est regardé, dans une perspective d'optimisation. Celui-ci consiste à faire l'hypothèse que l'institut de sondage a, tout comme les votants, des préférences ordinales strictes sur les candidats et qu'il souhaite manipuler les résultats du sondage afin de faire élire dans l'équilibre correspondant le meilleur candidat possible en fonction de ses préférences. Il s'ensuit que ce problème est difficile, autrement le premier problème serait facile à résoudre. On propose également une heuristique reposant sur une idée proche de la première heuristique et on analyse ses performances. Tout comme la première heuristique, cette dernière est très efficace en pratique, en particulier sur les graphes épars. L'avantage de cette heuristique est qu'elle ne demande pas nécessairement de connaître l'intégralité du réseau social.

Ainsi, bien qu'il existe des cas où la manipulation par l'institut de sondage semble être difficile d'un point de vue computationnel, il apparaît que de simples heuristiques peuvent s'avérer efficaces pour calculer une telle manipulation. Néanmoins, lorsque les agents sont très bien informés, ces heuristiques deviennent incapables de guider les agents vers une issue désirée par l'institut de sondage.

Conclusion

Dans la thèse, on a considéré deux problèmes de choix social particuliers, le vote stratégique et l'allocation de biens indivisibles, et on a tenté par le biais d'un réseau social de relâcher l'hypothèse classique d'interaction possible au sein de tout groupe d'agents. On s'est penché spécifiquement sur deux types d'interaction : la collaboration et la prise d'information. Dans le premier cas, le réseau social est modélisé par un graphe non-orienté car la relation d'accessibilité entre agents n'a de sens que si elle est symétrique, alors que dans le second cas le réseau social est un graphe orienté traduisant une relation de visibilité entre agents. On montre à travers nos différents travaux que la structure du réseau social a un impact important sur les issues des procédures de choix social considérées mais également sur la complexité des problèmes abordés. De manière générale, l'introduction d'un réseau social permet une analyse plus fine des problèmes de choix social en ajoutant du réalisme à leur modélisation, tout en rendant parfois plus difficile la résolution de ces problèmes, en fonction de la structure du graphe.

De nombreuses extensions sont possibles pour ce travail. Premièrement, l'analyse de chaque problématique de choix social abordée peut être en elle-même étendue. Par exemple, concernant le vote itératif, on a généralisé la notion de déviation stratégique soit en considérant des déviations coalitionnelles plutôt que des déviations unilatérales ou des déviations avec incertitude plutôt qu'avec information complète. Une extension directe serait donc de combiner les deux généralisations. En ce qui concerne les problèmes d'allocation de ressources, on pourrait considérer des allocations non restreintes à un seul bien par agent mais aussi, par exemple, se pencher sur des biens qui soient plutôt du domaine de l'information ou de la connaissance. Le but serait alors d'enrichir sa connaissance au fur et à mesure des interactions, comme dans le problème de "gossip" bien connu en intelligence artificielle.

En prenant en compte le cadre général des problèmes de choix social abordés, on pourrait également penser à d'autres structures que le graphe afin de penser la représentation de la possibilité d'interaction. On pourrait notamment vouloir modéliser les groupes pouvant interagir ensemble de manière plus directe que via des cliques dans un graphe. La difficulté serait de trouver un compromis entre interprétabilité et compacité de la structure. Une autre voie serait de considérer que le réseau social puisse être dynamique. En effet, on a fait l'hypothèse tout au long de la thèse que la structure de réseau social était fixe mais il pourrait être pertinent de supposer que celle-ci évolue au gré des interactions entre les agents, surtout que la plupart des processus étudiés sont dynamiques. Par exemple, deux agents pourraient échanger leurs contacts lors de leur interaction et donc de nouveaux liens pourraient être créés.

Résumé

Le choix social repose sur l'étude de la prise de décision collective, où un ensemble d'individus doit convenir d'une solution commune en fonction des préférences de ses membres. Le problème revient à déterminer comment agréger les préférences de différents agents en une décision acceptable pour le groupe. Typiquement, les agents interagissent dans des processus de décision collective, notamment en collaborant ou en échangeant des informations. Il est communément supposé que tout agent est capable d'interagir avec n'importe quel autre. Or, cette hypothèse paraît irréaliste pour de nombreuses situations. On propose de relâcher cette hypothèse en considérant que la possibilité d'interaction est déterminée par un réseau social, représenté par un graphe sur les agents.
Dans un tel contexte, on étudie deux problèmes de choix social : le vote stratégique et l'allocation de ressources. L'analyse se concentre sur deux types d'interaction : la collaboration entre les agents, et la collecte d'information. On s'intéresse à l'impact du réseau social, modélisant une possibilité de collaboration entre les agents ou une relation de visibilité, sur la résolution et les solutions de problèmes de vote et d'allocation de ressources. Nos travaux s'inscrivent dans le cadre du choix social computationnel, en utilisant pour ces questions des outils provenant de la théorie des jeux algorithmique et de la théorie de la complexité.

Mots Clés

Choix social computationnel, réseau social, vote stratégique, allocation de ressources, théorie des jeux algorithmique, complexité

Abstract

Social choice is the study of collective decision making, where a set of agents must make a decision over a set of alternatives, according to their preferences. The question relies on how aggregating the preferences of the agents in order to end up with a decision that is commonly acceptable for the group. Typically, agents can interact by collaborating, or exchanging some information. It is usually assumed in computational social choice that every agent is able to interact with any other agent. However, this assumption looks unrealistic in many concrete situations. We propose to relax this assumption by considering that the possibility of interaction is given by a social network, represented by a graph over the agents.

In this context, we study two particular problems of computational social choice: strategic voting and resource allocation of indivisible goods. The focus is on two types of interaction: collaboration and information gathering. We explore how the social network, modeling a possibility of collaboration or a visibility relation among the agents, can impact the resolution and the solution of voting and resource allocation problems. These questions are addressed via computational social choice by using tools from algorithmic game theory and computational complexity.

Keywords

Computational social choice, social network, strategic voting, resource allocation, algorithmic game theory, computational complexity