



HABILITATION À DIRIGER LES RECHERCHES

Towards Fairer Collective Decisions

Sylvain Bouveret

Current version: January 27, 2025

COMMITTEE

Christian KLAMLER, Associate Professor, Institute of Public Economics, University of Graz, **reviewer**

Meltem ÖZTÜRK, Professeure, PSL, Université Paris Dauphine, **reviewer**

Bruno ZANUTTINI, Professeur, Université de Caen Normandie, **reviewer**

Nadia BRAUNER VETTIER, Professeure, Université Grenoble-Alpes, president

Evangelos MARKAKIS, Professor, Athens University of Economics and Business, examiner

Pierre MARQUIS, Professeur, Université d'Artois, examiner

Agnieszka RUSINOWSKA, Directrice de Recherche, CNRS, examiner

Ce manuscrit est dédié à Alice, qui a eu 18 ans le 26 janvier 2024.

Contents

Résumé	1
Abstract	3
I Introduction	5
II Fair enough: fairness beyond proportionality and envy-freeness	9
II.1 Technical preliminaries	10
II.1.1 Individual preferences	11
II.1.2 Fairness and efficiency	12
II.2 A scale of five fairness criteria	14
II.2.1 Five fairness criteria	14
II.2.2 A scale of criteria	18
II.2.3 Computational properties	20
II.2.4 Particular cases	22
II.3 Balancing envy between agents	22
II.3.1 Existing relaxations of envy-freeness	22
II.3.2 The minOWA criterion	24
II.3.3 Computing minOWA envy allocations	25
II.4 Epistemic envy	26
II.4.1 Epistemic envy-freeness	26
II.4.2 Envy-freeness with social constraints	27
II.4.3 A weaker social epistemic envy-freeness property	28
II.5 Envy approved by the society	29
II.5.1 K-approval envy-freeness	29
II.5.2 K-approval non-proportionality	31
II.5.3 Computation	32
II.6 Conclusion: the landscape of fairness properties	33

III The unreasonable fairness of picking sequences	35
III.1 An elicitation-free protocol for fair division	37
III.2 The fairest sequence problem	38
III.3 Computing optimal sequences	40
III.4 Manipulating picking sequences	42
III.5 Of strategyproof picking sequences	45
III.6 Picking sequences, swap deals and efficiency	49
III.6.1 Sequenceability as an efficiency criterion	49
III.6.2 Cycle-deals optimality	50
III.6.3 A scale of efficiency	51
III.6.4 Link with fairness notions	52
III.7 Conclusion	53
IV And the winner is... Alternative voting rules for fairer representation	57
IV.1 Technical preliminaries	58
IV.2 Voter Autrement	60
IV.2.1 Experimental setting	60
IV.2.2 Voting rules	62
IV.2.3 The ballots	64
IV.2.4 Results	66
IV.2.4.1 Of biases	66
IV.2.4.2 Global insights	67
IV.3 Simulations for the legislative election	72
IV.3.1 The voting rules simulated	72
IV.3.2 Predicting a second round	74
IV.3.3 Redistricting	76
IV.3.4 Simulations	78
IV.3.5 Conclusions	81
IV.4 Conclusion	82
V Conclusion: a fair and safe operating space for humanity	85
Bibliography	101

Résumé

Ce manuscrit d'habilitation à diriger les recherches s'intéresse à la notion d'équité dans les problèmes de décision collective. Par « problème de décision collective », nous entendons ici toute situation dans laquelle un ensemble d'agents (pouvant être des personnes ou des entités comme des entreprises ou des collectivités) doivent s'accorder pour choisir en commun une option parmi un ensemble d'options candidates. Dans ce manuscrit, nous nous focalisons sur deux problèmes de décision collective particuliers : les problèmes de partage de biens indivisibles d'une part, et les élections d'autre part.

Les travaux présentés dans ce document s'inscrivent dans la lignée des recherches effectuées dans le champ de la théorie du choix social, et plus précisément dans le domaine du choix social computationnel, discipline située à l'interface entre la microéconomie et l'informatique. Dans ce cadre, nous nous intéressons plus précisément à la notion d'équité, à ses déclinaisons dans les problèmes de partage et dans la théorie du vote, et à ses propriétés computationnelles.

Ce manuscrit est organisé en trois chapitres principaux dont les deux premiers sont dédiés aux problèmes de partage de biens indivisibles. Dans le premier chapitre en particulier, nous nous appuyons sur les formalisations standards de l'équité dans la théorie des problèmes de partage, et en particulier sur la notion de proportionnalité et d'absence d'envie, pour définir un ensemble de nouveaux critères s'appuyant chacun sur des principes différents, par exemple sur l'équilibre de l'envie entre les agents, sur la notion de connaissance limitée des agents entre eux, et sur l'approbation sociale de l'envie. Cela nous conduit également à définir une échelle de critères d'équité, permettant de trouver des solutions alternatives lorsque les critères standards ne peuvent être satisfaits, et permettant dans le même temps de caractériser la difficulté d'un problème de partage de ressources. Au final, cette exploration nous permet de dresser tout un paysage de critères d'équité, reliés entre eux, et permettant d'adopter une approche adaptée selon le problème à traiter.

Dans le deuxième chapitre, nous présentons une approche différente aux problèmes de partage équitable, adaptée aux contextes dans lesquels on ne dispose pas forcément d'une autorité centrale bienveillante en charge de recueillir les préférences des agents et de calculer une allocation. Cette approche se présente sous la forme d'un protocole dans lequel les agents choisissent chacun et chacune à leur tour l'objet qu'ils ou elles préfèrent parmi les objets restant. Dans ce chapitre, nous explorons plus particulièrement deux questions liées à ce protocole. La première question est celle de déterminer la séquence d'agents la plus équitable, ce qui, comme nous le montrons dans le chapitre, peut s'avérer une tâche délicate d'un point de vue computationnel. La deuxième question que nous explorons est celle de la manipulabilité de ce protocole : d'une part, nous cherchons à caractériser sous quelles conditions ce protocole peut être facilement manipulé par les agents, et nous explorons également les conditions particulières dans lesquelles ce protocole est insensible à la manipulation. Enfin, nous terminons le chapitre en mettant en évidence un certain nombre de liens entre cette famille de protocoles et les propriétés standards d'équité et d'efficacité des partages, ce qui nous amène finalement à conclure que ce protocole, en plus d'être simple à

comprendre et à mettre en œuvre en pratique, a de très bonnes propriétés théoriques.

Le troisième chapitre, quant à lui, est dédié entièrement à la problématique du vote, dans un contexte très particulier qui est celui des élections politiques. Les techniques que nous mettons en œuvre dans ce chapitre ne relèvent pas, comme aux deux chapitres précédents, de l'analyse théorique, mais s'appuient intégralement sur des approches expérimentales. Dans une première partie de ce chapitre, nous expliquons comment la mise en œuvre d'expérimentations *in situ* et en ligne permettant de faire tester à des participantes et participants volontaires des méthodes de vote alternatives peut permettre de comprendre l'influence de ces méthodes de vote sur le résultat d'une élection. Puis, dans la seconde partie, nous appliquons à la même question (mais sur un contexte électoral différent) un autre outil, celui des simulations informatiques. Notre objectif est que ces méthodes complémentaires permettent d'éclairer le débat public sur le choix des méthodes de vote pour les élections politiques en comprenant de manière plus fine comment l'utilisation de ces méthodes pourraient redessiner le paysage politique.

Le manuscrit s'achève en tâchant de remettre en perspective les travaux sur la théorie du choix social dans le contexte des nombreuses crises planétaires actuelles liées à l'effondrement des conditions environnementales. Plus précisément, nous montrons qu'un certain nombre de problèmes liés au respect des limites planétaires peuvent être naturellement vus sous le prisme du choix social. C'est le cas par exemple pour le partage de ressources disponibles en quantités de plus en plus limitées pour des raisons subies ou choisies, ou encore dans le cas d'infrastructures partagées gérées comme des communs. Cette exploration nous permet ainsi de replacer les techniques présentées dans le manuscrit dans le contexte d'une réflexion plus globale contribuant au développement d'une société désirable.

Abstract

This habilitation thesis is about fairness in collective decision making problems. Here, by "collective decision making", we refer to any kind of situation where a set of agents (that can be persons or entities like companies or local authorities) have to jointly reach an agreement in order to choose an option among several available ones. In this document, we particularly focus on two collective decision making situations: fair division of indivisible goods on the one hand, and elections on the other hand.

The works presented in this habilitation thesis belong to the field of social choice theory, and more precisely to the domain of computational social choice, which precisely lies at the interface between microeconomics and computer science. In this context, we especially investigate the notion of fairness, its variations in fair division problems and in voting theory, and its computational properties.

The document is organized in three main chapters, among which the first two ones are dedicated to fair division of indivisible goods. In the first chapter, we elaborate on the standard formal models of fairness in fair division theory, and especially on proportionality and envy-freeness. We develop a new set of fairness criteria, each of which conveying different principles. These principles might rely for instance on the idea of balancing envy among agents, on the idea of limiting the mutual knowledge of agents, or on the idea of introducing a social approval of envy. This exploration also leads us to define a whole scale of fairness criteria, allowing to find alternative solutions when standard criteria cannot be met, while also characterizing the difficulty of a fair division problem. In the end, we draw a whole landscape of interconnected fairness criteria, that can be used to adopt an adapted approach to each kind of problems we have to solve.

In the second chapter, we investigate a different approach to fair division problems. This approach is adapted to contexts in which there does not necessarily exist a benevolent central authority in charge of eliciting the agents preferences and computing an allocation. This approach is based on a simple protocol in which the agents are designated one by one. Every time an agent is designated, they pick one object among those that remain. In this chapter, we especially investigate two questions related to this protocol. The first one is to determine the fairest sequence of agents, which could be, as we show in this chapter, computationnally intricate. The second question concerns the manipulability of the protocol: on the one hand, we aim at characterizing under which conditions this protocol can be easily manipulated by the agents, and on the other hand, we also explore the particular conditions that make this protocol insensitive to manipulation (strategyproof). Finally, we end up the chapter by unveiling somewhat unexpected connections between this family of protocols and standard fairness and efficiency properties. That leads us to conclude that this protocol, on top of being simple in practice and easy to understand, has very good theoretical properties.

The third chapter is entirely dedicated to voting, in the very special context of political elections. Contrary to the first two chapters, the techniques proposed in this chapter are not based on theoretical analysis, but rather on experimental approaches. In a first part of the chapter, we explain how *in situ* and

online experiments, where voluntary participants can test alternative voting rules, help us understanding the influence of these voting methods on the election result. Then, in the second part, we apply to the same question (but in a different election context) another tool, namely, computer simulations. Our aim is to show how these complementary methods can enlighten the public debate on the choice of voting rules for political elections by showing how the use of these rules can redraw the political landscape.

The document ends up by trying to put those works on social choice theory in perspective of the current global crises related to the collapse of the environmental conditions. More precisely, we show that several problems related to keeping the society within the planetary boundaries can naturally be seen from the point of view of social choice. This is for instance the case for the division of resources available in smaller and smaller amounts, or for shared infrastructures managed as commons. This exploration hence leads us to show how the techniques presented in the document can fit in a more global reflection that hopefully aims at developing a more desirable society.

Chapter I

Introduction

Collective decision making deals with the situation where a set of agents (persons, companies, agencies, entities) have to jointly choose an option among several available ones. This might be for instance the case when a group of friends wants to choose a restaurant where to have dinner tonight, when brothers and sisters have to fairly divide up an inheritance, or when a nation of citizens elects a president. Beyond these examples, in a world where environmental problems are probably the most crucial challenges our society has to face, some studies (see *e.g.*, [Dugast and Soyeux, 2019](#)) suggest that most leverage points will not be individual, but will be based on decisions that involve a collective consensus. This raises several critical questions. Who should decide of the decisions to make in common? On which basis? How should these decisions be made? As we shall see, collective decision making is at the heart of society, and no transition to a sustainable world will be possible without a deep reflection on collective decision mechanisms.

The scientific field dedicated to the formal study of collective decision making methods is *social choice theory*. This field typically addresses problems that can be informally described as follows:

Given a set of *alternatives* and a finite set of *agents* expressing *opinions* over the alternatives, find a *collective opinion* on the alternatives, that could be for instance a ranking over these alternatives (stating which one is collectively preferred, which one comes second, and so on), or simply the collective choice of a unique alternative.

The underlying assumption behind this formulation of the problem is that basing the collective choice on the individual opinions is intrinsically desirable. As we shall see later, the axiomatic foundations of social choice theory also implicitly entail that collective decision making mechanisms that are based on a *fair* consideration of all the individual opinions are more desirable than the others. The core problem here is then to agree on what could be considered fair in a situation where the agents' individual opinions might be incompatible or even antagonistic. Before further discussion, it is important to notice here that fairness does not necessarily refer to a moral property. Following [Young, 1994b](#), a collective decision making is fair if it is accepted and considered appropriate by all the community members, with regards to the individual needs, statuses and contributions. Hence, in this sense, fairness is more a guarantee of stability and social peace than a moral principle.

This habilitation thesis is about fairness in collective decision making. We especially focus on two settings: *fair division* and *voting*. Collective decision making is not limited to these two problems, and includes other ones such as matching, coalition formation or judgment aggregation. Nevertheless, fair

division and voting are two prominent prototypical collective decision making problems having strong theoretical and philosophical roots dating back to antiquity.

Fair division In fair division, a common resource has to be shared among a set of agents that might have different preferences, needs, interests, or entitlements about the resource. As recalled by [Moulin, 2003](#), one of the roots of distributed fairness can be attributed to Aristotle in the *Nicomachean Ethics*: "*Equals should be treated equally, and unequals unequally, in proportion to relevant similarities and differences*". The fact that equals should be treated equally is clear to all appearances: two persons similar in all the dimensions relevant to the problem at stake should be treated interchangeably by the division procedure. The tricky question here is to determine which dimensions are relevant to the division problem: preferences, merits, statuses, entitlements, basic needs... Choosing among those dimensions goes well beyond the scope of this thesis, but it is important to bear in mind that any answer to this question embeds an implicit political vision of the world. Treating unequals unequally is even more complex. Once again, the question of determining the "relevant differences" pertains to the political vision we adopt.

The formal model on which our work is based is rooted in the theoretical foundations of microeconomic analysis. Here, the central assumption is that each agent embeds a preference relation, and that every individual choice consists in finding the outcome that maximizes it. The vision we adopt is also guided by the collective welfare approach, where a benevolent dictator proposes a fair division mechanism that is supposed to fairly take into account the similarities and differences, as well as the potentially conflicting interests of the stakeholders.

In this framework, we usually distinguish two kinds of fair division problems according to the essential nature of the resource to share.

1. *Cake-cutting* problems deal with the case where the resource is continuous. In this field, the continuous nature of the resource generally guarantees the existence of allocations with strong fairness properties like envy-freeness or proportionality,¹ which can be then calculated using standard optimisation approaches. This is why the major scientific questions are here more concerned with the search for distributed allocation procedures that guarantee these fairness properties.
2. If the resource is finite and discrete, then the problem comes down to allocating a set of indivisible (and usually rival) objects to the agents. In that case, the existence of envy-free or proportional allocations cannot be guaranteed anymore and finding some when they exist becomes extremely challenging. This is the setting we will focus on in the first part of this thesis.

Voting Voting refers to the situation where a group of agents have to choose among a set of predefined candidates one (in *single-winner voting*) or several ones (in *multi-winner voting aka committee voting*) that will represent them all. Even if voting also belongs to the set of collective decision making problems, the intrinsic nature of this problem is very different from fair division. In fair division, in some way antagonistic preferences can be reconciled to a certain extent, because the agents have a private use of the objects they receive, and different objects will be allocated to agents with different preferences. In voting, the result of the election is common to everyone: some agents will probably therefore be extremely disadvantaged by the decision. Fairness will therefore not be based on the same formal principles in voting

¹Although this may not be the case if we require additional properties like Pareto-efficiency ([Barbanel et al., 2009](#)) or if we add further constraints on the protocol used to cut the cake ([Stromquist, 2008](#)).

as in fair division: properties such as envy-freeness or proportionality, which are of major importance in fair division are essentially meaningless in the context of voting.

The mathematical study of voting systems dates back to the seminal book of [Condorcet, 1785](#), and the birth of the modern school of social choice and voting theory is in general attributed to the work of [Arrow, 1951](#). Until the late 1980's, the research agenda in this field was mainly axiomatic. [Bartholdi et al., 1989a](#); [Bartholdi et al., 1989b](#) were the first to import computational issues in the field of social choice, by analyzing questions like manipulation issues that turned out to be computationally difficult and required a careful analysis and tricky algorithms to be solved. Since these early works, the study of these questions have led to a fruitful collaboration between economics and computer science that have been at the origin of *Computational Social Choice* (Comsoc). In the last two decades, researchers in Comsoc have investigated questions that go well beyond those related to the computational complexity of strategic behaviour in voting ([Brandt et al., 2015](#)). The typical questions that arise in this field can relate to the computation of the voting rules themselves (as some of them turn out to be computationally complex), communication complexity of voting, or the development of compact preference representation languages for settings where the set of candidates is combinatorial ([Lang, 2004](#)).

In parallel to voting, other fields of social choice have benefited from the contributions of computer scientists. The domain of fair division has especially been the object of longstanding collaborations between economics and computer science. Even if the first formal works in fair division in the computer science community date back from the late 1990's, the early work of [Steinhaus, 1948](#) already has an algorithmic flavour since one of the core topics concerns the allocation procedures that guarantee some fairness properties. The pioneering work of [Rassenti et al., 1982](#) in the neighbouring domain of auctions was also one of the firsts to introduce the notion of compact representation in allocation problems of indivisible goods. This later led to a flourishing line of works in computational social choice.

Voting and fair division of indivisible goods, as structurally different as they might be, nevertheless raise similar questions. One of these questions, that will be the central topic of this thesis, can be formulated as follows:

How can fairness be formally defined, and how does the use of different fairness notions impact the collective decision and its computation in practice?

The form this question takes varies from one area to another. In fair division of indivisible goods, properly defining fairness is still a challenge in itself. As recalled earlier, the standard solution concepts like envy-free or proportional allocations are not guaranteed to exist. The question of how to relax them properly has no clear answer, and leads to a rich landscape of fairness properties that will be presented in details in [Chapter II](#). A central question behind the definition of these fairness properties is to compute allocations satisfying these properties. As we will see, this can be challenging. Even more, the simple question of determining whether such allocations exist or not for a given instance can be computationally complex as well.

A standard workaround for this potentially high complexity is to rely on a (semi-)distributed allocation procedure instead of using a centralized optimization algorithm. Here, fairness is *procedural*, that is, results from the application of a supposedly fair distributed procedure instead of being guaranteed by a central authority. Distributed fair division procedures also have other very appealing features: they can be used in contexts where the agents are reluctant to reveal their own preferences to a central authority, may it be for privacy reasons, or just because the preference elicitation process is tedious and imperfect.

As we will see, designing fair distributed procedures may also be computationally challenging. This will be the central topic of Chapter III.

In voting, the landscape of rules that can be reasonably considered fair for single-winner and multi-winner elections is now relatively clear and has benefited from decades of thorough axiomatic investigation.² Unfortunately, the set of impossibility results, starting with Arrow's seminal theorem (Arrow, 1951), shows that there is no such thing as a perfect voting rule,³ and that compromises have to be made. Therefore, the choice of a fair voting rule cannot only be the result of an axiomatic process, but must also be discussed on ethical bases and accepted by the members of the society. The part dedicated to voting in this thesis is a contribution of computer science to the society's debate on fairness of voting rules for political elections. Here, the central question is to know to which extent a change in the voting rule chosen changes the result of the election. The underlying idea is that people can only debate on the opportunity of changing a voting rule if they have an idea of the potential impacts it has on the result (e.g. candidates they promote). The approach we present to address this question is based on the design of voting experiments on the one hand, and on computer simulations on the other hand. This will be the topic of Chapter IV.

We can observe that the contributions of the two parts are of very different nature. The contributions are mostly theoretical in the first two chapters on fair division, and mostly experimental in the last chapter on voting. The objective was to give a broader point of view of the possible contributions of computer science to social choice, including experimental approaches, that clearly complement the theoretical contributions. Hopefully, these different contributions are still presented in a consistent way.

Note: The chapter titles are respective tributes to the papers "Fair Enough: Guaranteeing Approximate Maximin Shares" by Procaccia and Wang, 2014; Kurokawa et al., 2018, "The Unreasonable Fairness of Maximum Nash Welfare" by Caragiannis, Kurokawa, et al., 2016, "And the loser is... Plurality voting" by Laslier, 2011, and "A Safe Operating Space for Humanity" by Rockström et al., 2009.

²Although the discussion about the respective merits of voting rules is still subject to active debates in the community (see e.g. Laslier, 2011).

³At least if we stick to the standard model where the voters' preferences are expressed using linear orders and are not subject to interpersonal comparisons.

Chapter II

Fair enough: fairness beyond proportionality and envy-freeness

Chapter foreword

—

In this chapter, we investigate several notions of fairness in resource allocation problems.

The first part of the chapter, dedicated to the definition of a scale of fairness, is mostly based on a collaborative work with Michel Lemaître, rooted in the discussions we had with our former PhD student Charles Lumet, initially presented in 2014 at the international conference on Autonomous Agents and MultiAgent Systems (Bouveret and Lemaître, 2014), and further extended in 2015 to a version for the international journal on Autonomous Agents and MultiAgent Systems (Bouveret and Lemaître, 2015).

The second part of the chapter proposes three alternative views on fairness, each one driven by a different motivation. The first one, on balancing envy, corresponds to the first research question investigated by Parham Shams, a PhD student that I supervised with Aurélie Beynier and Nicolas Maudet, and presented in 2021 at the conference on *Algorithmic Decision Theory* (Shams et al., 2021). The second approach, on epistemic envy, is the result of a collaboration with Haris Aziz, Ioannis Caragiannis, Ira Giagkousi and Jérôme Lang that has been presented in 2018 at the AAAI conference on Artificial Intelligence (Aziz et al., 2018). Lastly, the third approach, on approval envy, corresponds to another research question investigated by Parham Shams, presented in 2020 at the *European Conference on Artificial Intelligence* (Shams et al., 2020) and later published in the *Journal of Artificial Intelligence Research* (Shams et al., 2022).

All the contributions of this chapter in terms of fairness concepts definitions are summarized in Figure II.1.

—

Aziz, Haris, Sylvain Bouveret, Ioannis Caragiannis, Ira Giagkousi, and Jérôme Lang (2018). “Knowledge, Fairness, and Social Constraints”. In: *Proceedings of the 32nd AAAI conference on Artificial Intelligence (AAAI’18)*. New Orleans, Louisiana, USA: AAAI Press. URL: <http://recherche.noiraudes.net/resources/papers/AAAI18.pdf>.

Bouveret, Sylvain and Michel Lemaître (May 2014). “Characterizing Conflicts in Fair Division of Indivisible Goods Using a Scale of Criteria”. In: *Proceedings of the 13th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS’14)*. Paris, France: ACM. URL: <http://recherche.noiraudes.net/resources/papers/AAMAS14.pdf>.

— (2015). “Characterizing conflicts in fair division of indivisible goods using a scale of criteria”. In: *Autonomous Agents and Multi-Agent Systems* 30.2, pp. 259–290. ISSN: 1573-7454. DOI: [10.1007/s10458-015-9287-3](https://doi.org/10.1007/s10458-015-9287-3). URL: <http://recherche.noiraudes.net/resources/papers/JAAMAS15.pdf>.

Shams, Parham, Aurélie Beynier, Sylvain Bouveret, and Nicolas Maudet (Aug. 2020). “Fair in the Eyes of Others”. In: *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI’20)*. Santiago de Compostela, Spain: IOS Press. URL: <https://arxiv.org/abs/1911.11053>.

— (Nov. 2021). “Minimizing and balancing envy among agents using Ordered Weighted Average”. In: *Proceedings of the 7th International Conference on Algorithmic Decision Theory (ADT’21)*. Toulouse, France: Springer Verlag. URL: <http://recherche.noiraudes.net/resources/papers/ADT21.pdf>.

— (2022). “Fair in the Eyes of Others”. In: *Journal of Artificial Intelligence Research (JAIR)* 75, pp. 913–951. URL: <https://jair.org/index.php/jair/article/view/13778/26864>.

Suppose you have to solve a fair division problem. After you have properly identified the objects and the agents at stake, and after you have convinced the agents to reveal their preferences, the next step will be to choose a fairness criterion and try to compute an allocation that best satisfies this criterion. In this chapter, we will focus on the question of how to properly define fairness concepts in the context of fair division of indivisible goods. More precisely, the work presented here is driven by the following question:

In fair division of indivisible goods, the standard fairness properties of envy-freeness and proportionality are not guaranteed to be satisfiable. What alternative fairness solution concepts could we possibly use in the case they cannot be satisfied?

In the first part of this chapter (Section II.2), we introduce three alternative fairness concepts and show (somewhat unexpected) connections between these concepts, proportionality, and envy-freeness. The main contribution is that, together, these criteria form a whole scale of fairness concepts of increasing strength that can help characterizing how much antagonistic the agents' preferences are.

In the second part of the chapter (Sections II.3, II.4 and II.5), we complete the landscape of fairness solution concepts by proposing three different approaches that all aim at relaxing envy-freeness.

The first one (Section II.3) is based on envy-minimization, a concept that was already introduced by Lipton et al., 2004, but that we extend with the idea of balancing envy among agents using aggregators from multiobjective decision making.

The second one (Section II.4) is based on a completely different approach. Here, the idea is that envy-freeness is inherently a *knowledge-sensitive* notion. More precisely, it strongly relies on the knowledge that the agents have on the share of the other agents. If we relax this implicit assumption that the agents' shares are common knowledge, we obtain a much weaker (and probably more realistic in some situations) notion of envy-freeness.

Finally, the third approach (Section II.5) questions the *subjective* flavour of envy-freeness. Namely, envy strongly relies on a subjective comparison of an agent's share with the shares of the other agents. We could imagine the case where a_1 envies a_2 but they objectively have no reason to do so (here, by "objectively" we mean that a majority of the agents other than a_1 judge that a_1 has no reason to envy a_2). This approach, inspired by the notion of *unanimous envy* introduced by Parijs, 1997, leads to yet other relaxations of envy-freeness.

All the contributions of this chapter in terms of fairness concepts definitions are summarized in Figure II.1 (page 34). This shows in a hopefully synthetic way the full landscape of fairness properties for additive multiagent resource allocation problems. What is striking is that all these properties are related to each other in some way.

Before diving into the main contributions, we will first give in Section II.1 a few technical preliminaries that will be useful for this chapter and the next one.

1 Technical preliminaries

This chapter and the next one will be about *fair division of indivisible goods*, that can be formally defined as follows.

Let $\mathcal{A} = \{a_1, \dots, a_n\}$ be a finite set of *agents* and $\mathcal{O} = \{o_1, \dots, o_m\}$ be a finite set of indivisible *objects* (that shall also be called *items*). An *allocation* of the objects to the agents is a vector $\vec{\pi} = \langle \pi_1, \dots, \pi_n \rangle$,

where $\pi_i \subseteq \mathcal{O}$ is the *bundle* of objects allocated to agent a_i , called agent a_i 's *share*. Even if in general any vector $\vec{\pi}$ could be considered to be a valid allocation, most work of fair division restrict the set of *feasible* allocations to those that satisfy the two following properties.

1. Non-shareable items: $i \neq j \Rightarrow \pi_i \cap \pi_j = \emptyset$
2. No free-disposal: $\bigcup_{a_i \in \mathcal{A}} \pi_i = \mathcal{O}$

We will denote by $\mathcal{F}_{\mathcal{A}, \mathcal{O}}$ (or simply \mathcal{F} when the context is clear) the set of feasible allocations. All the allocations considered in this document are implicitly feasible.

Fair division of indivisible goods is about finding a "good" allocation of the objects to the agents. As recalled in Chapter I, properly formalizing on what basis a good allocation could be defined is a long-standing philosophical question pertaining to the principles of distributive justice. The usual approach in social choice theory – and, by extension, computational social choice – is to consider that the agents are equipped with *preference orders* that can be used as proxies for their individual tastes (or needs) about the bundles they may receive. The objective will then be to find an allocation that best satisfies the agents with respect to these preference orders, whatever that means.

Note that an implicit assumption of our model is that the agents have *equal entitlements* about the resource, that is, that no agent can pretend to more resources than the others for reasons exogenous to the problem at stake. Some works have investigated the extension of the standard model of fair division to the case where agents have different entitlements represented by weights (see [Suksompong, 2025](#), for a recent survey on these works). This is however beyond the scope of this thesis, and we will stick to the standard model with equal entitlements, in which the final decision only relies on the expression of individual preferences.

1.1 Individual preferences

Ordinal and cardinal preferences There are two usual models that are used in the literature to formally define the agents' preferences. The first one is the *ordinal model*, in which we assume that the agents' preferences are expressed using a *binary relation* \leq on the bundles of items, that is at least reflexive and transitive.¹ Here, $\pi \leq \pi'$ should be understood as meaning "bundle π' is preferred to bundle π ". The ordinal model is appealing because it is conceptually simple and yet powerful enough to reason on individual preferences. However, it is intrinsically limited in the sense that even if we can perform pairwise comparison of bundles, there is no way to express intra- or inter-personal preference intensity. To overcome this limitation, a whole line of work in economics and computation consider *cardinal preferences*. In this model, the agents' preferences are expressed using a numerical *utility function* $u : 2^{\mathcal{O}} \rightarrow \mathbb{N}$ mapping each bundle to a utility.² Here, the implicit meaning is that the higher the utility of a bundle is, the more preferred this bundle is. Restricting utilities to the positive numbers also entails that we only deal with items that are beneficial to the agents; in other words, goods. The case where items have nega-

¹In some settings where it could be problematic to leave some pairs of bundles incomparable, an additional assumption could impose that the relation is total as well.

²Note that for formal reasons we assume here, as will be the case throughout the document, that the utilities are (positive) integer, despite it might not always be the case in the literature.

tive values – in other words, they are *chores* – is also the topic of many works in fair division,³ but clearly out of the scope of this document.

Additive preferences The careful reader might have noticed that the preference models we have defined above are based on abstract mathematical objects (a relation and a function), but do not say anything about how these objects are represented in practice. Let us consider for the sake of example a utility function u relative to a set of objects \mathcal{O} . An explicit representation of this utility function requires $2^{|\mathcal{O}|}$ values, which causes two main issues: (i) one needs to be able to store a number of values that is exponential in the number of objects, and (ii) assuming that this utility function reflects the preferences of an agent, a decision maker would have to ask a potentially exponential number of questions to this agent to fully elicit this utility function.

The classical way to overcome this difficulty is to use a *compact representation language*. Such a language serves two purposes: first, it provides a syntax that is supposed to help the users to express their preferences over bundles of objects in an intuitive (and concise) way, and second, the language semantics gives a way to compute the utilities of bundles (or to compare two bundles). The study of ordinal and cardinal compact preference representation languages for fair division of indivisible goods has led to a fruitful line of work (see *e.g.* Bouveret, Chevaleyre, and Maudet, 2016 for an overview), involving sometimes elaborate, expressive, and powerful languages – such as, *e.g.*, logic-based languages.

In the last decade, the trend in the community seems to give up intricate compact representation languages to focus on a much simpler preference setting (that can be also seen as a compact preference language, although it is rarely presented as such): *additive preferences*. In this setting, each agent only expresses utilities on *individual objects*. The utility of an agent for a given bundle is just the sum of the utilities of the items it contains. More formally, we assume that each agent a_i gives a weight $w_{i,j} \in \mathbb{N}$ that expresses the utility this agent attributes to object o_j . The *individual utility* of agent a_i is then:

$$u_i(\pi) \stackrel{\text{def}}{=} \sum_{o_j \in \pi} w_{i,j}.$$

Note that this preference setting is quite limited in the sense that it is impossible for an agent to express that two objects are *complements* ($u(\{o, o'\}) > u(\{o\}) + u(\{o'\})$) or *substitutes* ($u(\{o, o'\}) < u(\{o\}) + u(\{o'\})$). However, this setting is empirically often considered to offer a good compromise between expressive power and conciseness and is used in many studies, including seminal ones (see *e.g.* Demko and Hill, 1988; Beviá, 1998; Brams and Taylor, 2000; Herreiner and Puppe, 2002; Brams, Edelman, et al., 2003; Lipton et al., 2004; Bansal and Sviridenko, 2006; Asadpour and Saberi, 2007; de Keijzer et al., 2009; Markakis and Psomas, 2011, just to name a few).

In this document, unless stated otherwise, we consider that the agents' preferences are additive.

1.2 Fairness and efficiency

Let us wrap things up a little bit. In our formal model of the fair division problem, what we have so far is a finite set of agents $\mathcal{A} = \{a_1, \dots, a_n\}$ having preferences over a finite set of objects $\mathcal{O} = \{o_1, \dots, o_m\}$. These

³Although much less than fair division of indivisible *goods*. Since the introduction of the concept of chores in fair division by Gardner, 1978, not so many works truly deal with chores, in spite of the fact that chores are not the exact symmetrical counterpart of goods (see *e.g.* Bouveret, Cechlárková, et al., 2019, Related Work section)

preferences are modeled in the form of an additive utility function u_i for each agent a_i , conveying the expression of a_i 's evaluation for each bundle π . According to this model, and adapting the terminology from the survey by [Chevalleyre, Dunne, et al., 2006](#), we define an *additive MultiAgent Resource Allocation* instance (add-MARA instance for short) as a triple $\langle \mathcal{A}, \mathcal{O}, w \rangle$, where \mathcal{A} is a set of agents, \mathcal{O} is a set of objects, and $w \in \mathbb{N}^{n \times m}$ is a matrix specifying the weight $w_{i,j}$ given by each agent a_i to each object o_j .

In this document, we will sometimes investigate even more restricted settings than additive preferences. Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. This add-MARA will be said to have:

- *strict preferences on objects* if for all agent a_i , $w_{i,j} \neq w_{i,j'}$ for all pairs of objects $o_j \neq o_{j'}$;
- *strict preferences on bundles* if for all agent a_i , $u_i(\pi) \neq u_i(\pi')$ for all pairs of bundles $\pi \neq \pi'$;
- *same-order preferences* if for all pair of agents $a_i \neq a_{i'}$, $w_{i,j} \geq w_{i',j}$ if and only if $w_{i',j} \geq w_{i,j}$;
- *identical preferences* if $w_{i,j} = w_{i',j}$ for all pair of agents $(a_i, a_{i'})$;
- *0–1 preferences* if $w_{i,j} \in \{0, 1\}$ for all a_i and o_j .

Now that the problem input is formally defined, the objective is to find a suitable allocation of the objects to the agents that takes into account these revealed preferences and tries to satisfy the agents as most as possible accordingly. A minimal requirement is that the allocation is *efficient*, that is, that the resource is not wasted nor under-exploited. The very minimal transposition of this principle is no free-disposal, but most works on fair division also consider *Pareto-efficiency* as a necessary condition for an allocation to be efficient. Let us quickly recall its definition. Let $\vec{\pi}$ and $\vec{\pi}'$ be two allocations. $\vec{\pi}$ Pareto-dominates $\vec{\pi}'$ if and only if $u_i(\pi_i) \geq u_i(\pi'_i)$ for all $i \in \mathcal{A}$, and $u_j(\pi_j) > u_j(\pi'_j)$ for at least one $j \in \mathcal{A}$. An allocation is Pareto-efficient if and only if it is not Pareto-dominated by any other allocation.

Beyond the minimal requirements of efficiency, the crucial and intricate question is how to formalize *fairness*. Here, two standard options are available. The first one consists of defining a collective utility function (CUF) aggregating individual agents' utilities. If the CUF is well chosen, its outcome when applied to individual utilities reflects the fairness (and possibly other desirable properties) of a given allocation. The (benevolent) arbitrator just looks for an allocation maximizing this CUF. Three standard CUF are available:

- the classical utilitarian CUF $g^{ut} : (u_1, \dots, u_n) \mapsto \sum_{i=1}^n u_i$;
- the egalitarian CUF $g^{eq} : (u_1, \dots, u_n) \mapsto \min_{i=1}^n u_i$;
- the Nash CUF $g^{Na} : (u_1, \dots, u_n) \mapsto \prod_{i=1}^n u_i$.

The second option consists of defining, by means of a Boolean (logical) criterion, what is considered fair. This is the approach followed by most works in the cake-cutting literature and by several in the context of fair division of indivisible goods, starting from the seminal paper by [Lipton et al., 2004](#) that deals with envy-freeness. Note that there are connections between Boolean fairness criteria and maximizing a CUF. For instance, maximizing the egalitarian CUF is guaranteed to return an allocation that also satisfies proportionality (see Definition 1) if there exists one. For envy-freeness, things are less clear. [Brams and King, 2005](#), on the one hand, show that satisfying envy-freeness and maximizing the egalitarian CUF can be in some situations antagonistic. On the other hand, [Caragiannis, Kurokawa, et al., 2016](#)

also show unexpected connections between (a variant of) envy-freeness and the Nash CUF by showing that any allocation maximizing the Nash CUF is envy-free up to one good (EF1).

As mentioned earlier, in this chapter, we will focus on Boolean fairness criteria, although one of the envy-freeness relaxation approaches we will introduce (see Section II.3) somewhat reconciles both approaches by mixing the idea of envy-freeness with CUF maximization.

Computational concepts Even if we assume the reader to be familiar with the most standard complexity classes, we quickly recall a few definitions here. \mathbf{P} (resp. \mathbf{NP}) is the class of decision problems that can be decided in polynomial time by a deterministic (resp. non-deterministic) Turing machine. \mathbf{CONP} denotes the complementary (in the sense of complexity) of class \mathbf{NP} , that is, the set of problems for which we can verify in polynomial time that an instance is negative. $\Sigma_2^{\mathbf{P}}$ is the class of decision problems that can be decided in polynomial time by a non-deterministic Turing machine using an \mathbf{NP} oracle. Finally, \mathbf{PSPACE} denotes the class of decision problems that can be solved in polynomial space by a deterministic Turing machine.

Sometimes, the problems we investigate are optimization problems, of the form *e.g. given input I , find the highest k such that there is a solution f such that $f(I) = k$* . When we will deal with complexity classes of such problems, we will always refer to the decision version of them, that is, *e.g. given I and a constant K , is there a solution f such that $f(I) \geq K$* .

When we introduce combinatorial optimization models, like Mixed Integer Programs (MIP), we denote variables with bold letters (*e.g.* \mathbf{x}_i), to distinguish them from constants or other data.

Finally, we need a few notations from graph theory (especially in Section II.4). If $G = (V, E)$ is a graph (directed or not) and $u \in V$ is a vertex of this graph, we denote by $\text{Neigh}_G(u)$ the set of u 's (out-)neighbours, while $\text{deg}_G(u)$ denotes u 's (out-)degree. In both cases, we may omit G when the context is clear.

2 A scale of five fairness criteria

In the landscape of fairness criteria, two properties are unavoidable: envy-freeness and proportionality. It is a well known fact that envy-freeness is strongest than proportionality (in the sense that any envy-free allocation is also proportional), and that there can be add-MARA instances that do not even admit a proportional allocation. In this section, we complete these two properties by introducing three additional criteria. We also show the connections they have with envy-freeness and proportionality, and investigate their properties.

2.1 Five fairness criteria

Let us first formally introduce the five fairness criteria we deal with in this section, starting with proportionality. Among these criteria, some of them rely on classical notions in economics (CEEI), while some others are more original (max-min share and min-max share). Note that all these criteria have an appealing property: contrary to other approaches like the maximisation of egalitarian CUF, none of this criteria rely on any interpersonal comparison of utilities. In practice, it means that the decision maker does not have to bear the burden of directly comparing different individuals on the only basis of their utilities, which can be difficult and ethically sensitive.

Proportionality Let us start with one of the most prominent criteria used in the literature on fair division to convey fairness properties: *proportionality*. The idea is quite simple. From the point of view of proportionality, an allocation should be considered to be fair if and only if each agent values their own share at least as much as the equal share. Initially defined directly on the resources themselves (Steinhaus, 1948), many authors have given a natural utilitarian interpretation of this notion:

Definition 1 (Proportional share) Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. The proportional share of agent a_i for this instance is:

$$u_i^{PS} \stackrel{\text{def}}{=} \frac{1}{n} u_i(\mathcal{O}) = \frac{1}{n} \sum_{o_j \in \mathcal{O}} w_{i,j}.$$

An allocation $\vec{\pi}$ satisfies the criterion of proportional share, denoted by $\vec{\pi} \models PS$, if $u_i(\pi_i) \geq u_i^{PS}$ for each agent $a_i \in \mathcal{A}$.

This notion conveys the idea that any allocation procedure should guarantee to each agent a lower bound of welfare, which corresponds to the utility this agent would get in an equal share of the resources.

Note that for obvious reasons this notion is often called, e.g. by Thomson, 2015, *equal-division lower bound*.⁴ Here, we stick to the term proportionality that seems to be slightly more common in the computational social choice literature.

Max-min share The proportional share criterion, coined by Steinhaus, 1948, was initially introduced in the context of fair division of a continuous resource (*a.k.a* cake-cutting) to convey and generalize the idea of a divide-and-choose like procedure. Imagine a game with two players, Alice and Bob, that have to share a continuous but not necessarily homogeneous cake. A very simple procedure for doing so is the following: (i) Alice cuts the cake in two pieces; (ii) Bob chooses one of both pieces; (iii) Alice takes the remaining piece. This simple "I cut you choose" procedure has a very appealing property. Alice can guarantee herself at least half of the cake (from her point of view) by cutting it in exactly two equal parts (once again, from her own point of view). Bob can also guarantee himself at least half of the cake, by choosing the piece he prefers. In other words, if the agents play rationally, this simple procedure guarantees their proportional share. More generally, for more than 2 agents, it can be proved that it is always possible to find an allocation of the cake that satisfies proportionality, although this will require slightly more intricate procedures than "I cut you choose", e.g. the last-diminisher procedure (see e.g. Brams and Taylor, 1996, for a reference book on the topic).

Unfortunately, no procedure can guarantee proportional fair share in the context of fair division of indivisible goods, as we can obviously observe with an instance with two agents and one desirable object. This fairness notion has thus been adapted by Budish, 2011, that defines the *max-min share* as a relaxation of the proportional share:

Definition 2 (Max-min share) Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. The max-min share of agent a_i for this instance is:

$$u_i^{MmS} \stackrel{\text{def}}{=} \max_{\vec{\pi} \in \mathcal{F}} \min_{a_j \in \mathcal{A}} u_i(\pi_j)$$

An allocation $\vec{\pi}$ satisfies the criterion of max-min share, denoted by $\vec{\pi} \models MmS$, if $u_i(\pi_i) \geq u_i^{MmS}$ for each agent $a_i \in \mathcal{A}$.

⁴Moulin, 1990 also employs the term *equal split guarantee*.

Like proportional share, the max-min share criterion is also based on a lower bound of welfare, that the allocation procedure should guarantee. This time, this lower bound is computed for each agent independently, as the result of a virtual "I cut, I choose last" procedure. In other words, u_i^{MmS} is the highest utility an agent a_i can guarantee to themselves in a procedure where they propose a partition of the objects into n bundles (*i.e.* they cut), all the other agents choose a different bundle, and a_i takes the remaining one. Note that in the cake-cutting context, this lower bound is exactly equal to the proportional fair share. Also note that this criterion does not require any interpersonal comparison of utilities. Technically, it does also not even require the preferences to be cardinal.⁵

Interestingly, this criterion that was, to the best of our knowledge, completely unknown at the beginning of the 2010's has received since an ever growing attention from the community, especially dedicated to the problem of figuring out in which conditions an (approximate) MmS allocation exist. We will discuss this problem (and its associated literature) in Section II.2.2.

Min-max share The max-min share criterion is based on a strict application of the "I cut, I choose last" principle to the case of indivisible goods. In the paper at the basis of this chapter (Bouveret and Lemaître, 2015), we investigated the idea of reverting this principle to define another fairness criterion that would be instead based on the "someone cuts, I choose first" principle. This is the idea behind *min-max share*:

Definition 3 (Min-max share) Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. The min-max share of agent a_i for this instance is:

$$u_i^{mMS} \stackrel{\text{def}}{=} \min_{\vec{\pi} \in \mathcal{F}} \max_{a_j \in \mathcal{A}} u_i(\pi_j)$$

An allocation $\vec{\pi}$ satisfies the criterion of min-max share, denoted by $\vec{\pi} \models mMS$, if $u_i(\pi_i) \geq u_i^{mMS}$ for each agent $a_i \in \mathcal{A}$.

In other words, u_i^{mMS} is the highest utility an agent a_i can hope to get *for sure* in a procedure where some other agent proposes a partition of the objects into n bundles (*i.e.* this agent cuts) and a_i chooses among the proposed bundles the one that they prefer.

This criterion has not received as much attention from the community as max-min share, mostly because it probably does not have as many appealing properties. Nevertheless, as we shall see later, it still has its own interest as an intermediate level of the scale of fairness we will define.

Envy-freeness *Envy-freeness*⁶ is probably one of the most prominent fairness concepts introduced for fair division problems. Initially defined by Tinbergen, 1953, this criterion conveys the idea that a society can be considered fair if and only if nobody would like to swap their place for the place of anybody else. This strong vision of envy-freeness has then been adapted by Foley, 1967 to a weaker utilitarian notion that can be formally stated as follows:

Definition 4 Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. An allocation $\vec{\pi}$ satisfies the criterion of envy-freeness (or, is simply envy-free), denoted by $\vec{\pi} \models EF$, if $u_i(\pi_i) \geq u_i(\pi_j)$ for each pair of agents $(a_i, a_j) \in \mathcal{A}^2$.

⁵In the initial work by Budish, 2011, it was even defined on ordinal preferences.

⁶Also sometimes called, *e.g.* by Thomson, 2015, the *no-envy* criterion.

Even before revealing the formal connections between different criteria, we can already feel that this criterion is more demanding in terms of the information it relies on. Previous criteria only relied on the determination of individual thresholds that could even be computed by the agents themselves on the basis of their individual utilities and the number of agents at stake. Here, envy-freeness relies on a complete pairwise comparison of the agents shares, which can only be made by the agents if they exactly and exhaustively know what the others receive. This limitation has led some authors to investigate relaxed versions of this criterion, as we will discuss in Section II.3.

Competitive Equilibrium from Equal Incomes The notion of equilibrium is probably one of the most important concepts in economic theory. The Competitive Equilibrium from Equal Incomes we present here is related to the seminal *Fisher model* of equilibrium (Fisher, 1892). This model formalizes an ideal market involving a set of buyers (the agents) and a set of goods whose quantities are limited, but that can be bought in fractional quantities. Each buyer is endowed with a (not necessarily additive) utility function, and with some fixed amount of money. The problem is to find unit prices for the goods so that if all the agents buy a bundle with maximal utility among those they can afford, then all the goods are sold. In other words, the aim is to find a set of prices that "clear the market": no good is over- nor under-demanded.⁷

The CEEI criterion corresponds to a particular case of the Fisher model, and differs from the initial model in the following sense: (i) the goods are indivisible and only one unit of each object is available; (ii) the agents' utilities are additive; (iii) each agent is endowed with the exact same amount of money (that we can consider w.l.o.g to be 1 euro). More formally, the CEEI can be defined as follows:

Definition 5 Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. A competitive equilibrium from equal incomes (CEEI) for this instance is a pair $(\vec{\pi}, \vec{p})$, where $\vec{\pi}$ is an allocation, and $\vec{p} \in [0, 1]^m$ is a price vector, such that:

$$\pi_i \in \arg \max_{\pi \subseteq \mathcal{O}} \left\{ u_i(\pi) \mid \sum_{o_j \in \pi} p_j \leq 1 \right\}, \text{ for all agent } a_i \in \mathcal{A}.$$

An allocation $\vec{\pi}$ is a CEEI, denoted by $\vec{\pi} \models \text{CEEI}$, if $\exists \vec{p} \in [0, 1]^m$ such that $(\vec{\pi}, \vec{p})$ is a CEEI.

As seen earlier, CEEI is deeply rooted in economic theory (see, e.g. Moulin, 2003, page 177), but had not been much considered in computer science, before the works by Budish, 2011 and Othman et al., 2010 on course allocation. Here, fairness comes from the fact that everyone is exactly endowed with the same amount of money at the beginning of the procedure.⁸

Interestingly, this notion also conveys a certain form of efficiency, as the following proposition shows:

Proposition 1 When the agents' preferences are strict on shares, any CEEI is Pareto-efficient. However, this is not necessarily true if the preferences are not strict.

The fact that a competitive equilibria from equal incomes is Pareto-efficient when the preferences are strict seems to be like a folk theorem in economics.⁹ However, curiously, the fact that a CEEI can fail

⁷See, e.g. Brainard and Scarf, 2005 for a summary of the Fisher model's equations and a detailed discussion about I. Fisher's initial work.

⁸Note that this fairness concept can be straightforwardly adapted to the case where, for exogenous reasons, agents should have different entitlements, which could be ensured by giving unequal incomes to them.

⁹Although this result seems to be regularly rediscovered, in the form, for instance, of the First Welfare Theorem by Babaioff et al., 2021.

being Pareto-efficient when preferences are not strict does not seem to be as well-known. We will come back to the connection between CEEI and efficiency in the next chapter, Section III.6.4.

2.2 A scale of criteria

In Section II.2.1, we have introduced five properties that all seem to convey different but reasonable aspects of fairness. Now how do these criteria relate to each other? It was already known in the literature that any CEEI allocation is also envy-free, and that any envy-free allocation is also proportional. These folk-theorems follow from quite straightforward arguments:

1. In any CEEI allocation, each agent a_i gets a bundle of maximal utility among the affordable ones. Since all the agents have the exact same income, they can afford exactly the same bundles. Hence, a_i 's bundle is necessarily at least as good as the bundle of any other agent. Thus, this allocation is envy-free.
2. In any envy-free (EF) allocation, each agent a_i gets the bundle of maximal utility for them among all the bundles received by all the agents (otherwise they would envy someone else). This value is necessarily greater than or equal to the average utility of these bundles, which is exactly equal to $\frac{u_i(\mathcal{O})}{n}$ since the allocation is complete. Hence, this allocation is proportional.

Hence, we can see that CEEI is a stronger criterion than EF, which is in turn stronger than proportionality. What is more surprising is that for additive preferences these relations extend to the other properties, and that forms a complete scale of fairness criteria, from the less demanding to the most demanding one.

Proposition 2 *Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance, and let $\vec{\pi}$ be an allocation. Then:*

$$(\vec{\pi} \vDash \text{CEEI}) \Rightarrow (\vec{\pi} \vDash \text{EF}) \Rightarrow (\vec{\pi} \vDash \text{mMS}) \Rightarrow (\vec{\pi} \vDash \text{PS}) \Rightarrow (\vec{\pi} \vDash \text{MmS}).$$

In other words, if $\mathcal{F}_{|X}$ denotes the set of feasible allocations that satisfy property X , we have:

$$\mathcal{F}_{|\text{CEEI}} \subset \mathcal{F}_{|\text{EF}} \subset \mathcal{F}_{|\text{mMS}} \subset \mathcal{F}_{|\text{PS}} \subset \mathcal{F}_{|\text{MmS}} \subset \mathcal{F},$$

all these inclusions being strict.

Since the criteria form a linear scale, given an add-MARA instance, the decision-maker has probably an interest in trying to find an allocation that satisfies the most demanding one, namely CEEI (since this allocation will also automatically satisfy the other ones). However, this is not always possible. As we have seen earlier, there are add-MARA instances that do not admit any proportional allocation for instance (hence, they obviously do not admit allocations satisfying more demanding criteria either). Therefore, Proposition 2 also has an implication in terms of add-MARA instances, that we formalize below:

Proposition 3 *Let \mathcal{I} be the set of all add-MARA instances, and $\mathcal{I}_{|X}$ the set of all add-MARA instances that admit at least an allocation $\vec{\pi}$ satisfying property X . Then we have:*

$$\mathcal{I}_{|\text{CEEI}} \subset \mathcal{I}_{|\text{EF}} \subset \mathcal{I}_{|\text{mMS}} \subset \mathcal{I}_{|\text{PS}} \subset \mathcal{I}_{|\text{MmS}} \subset \mathcal{I},$$

all these inclusions being strict.

Let us spend some time discussing the implications of this proposition. In the context of fair division of indivisible goods, social choice theorists have defined a set of nice properties, among which envy-freeness and proportionality, that serve as baseline fairness criteria. However, in this context, there are many instances where these properties are simply unsatisfiable, because the agents preferences are too conflicting. Proposition 3 provides a practical way to solve this issue. Assuming that satisfying the more demanding criterion is always better in terms of fairness, the (central) decision maker can thus look for an allocation that satisfies the highest possible criterion in the scale, and explain to the potentially unsatisfied (e.g. envious) agents that it is not possible to do better.

The careful reader might however have noticed that the rightmost strict inclusion of Proposition 3 entails that some add-MARA do not even admit allocations that satisfy the minimal criterion of max-min share. This is a problem, since for these instances, the decision maker could be left with no fallback fairness criterion to satisfy. However, in practice, it turns out that such instances are very rare. Actually, after having failed to find such an instance in thousands of randomly generated instances, we have even conjectured that $\mathcal{I} = \mathcal{I}_{\text{MmS}}$. It turns out that this conjecture was proved incorrect by Procaccia and Wang, 2014; Kurokawa et al., 2018, that came out with a very intricate example that becomes invalid with the slightest perturbation.¹⁰

Since then, a number of follow-up works have further investigated the question of when it is possible to ensure that an MmS allocation exists, and which fraction of the MmS threshold it is possible to guarantee when such an allocation does not exist. Obviously, for $n = 2$, we can always find an MmS allocation, just by applying the "I cut, you choose" protocol starting from any of the two agents. For $n \geq 3$, Kurokawa et al., 2016 proved that an MmS allocation always exists if $m \leq n + 4$, but if $m \geq 3n + 4$ there can be instances without such an allocation. Concerning approximations of MmS, Procaccia and Wang, 2014 already proved (on top of their seminal aforementioned counter-example) that it is always possible to find an allocation guaranteeing $\frac{2}{3}$ of the MmS threshold to every agent. Their result was later improved by Amanatidis, Markakis, et al., 2017 that proposed a simple algorithm to compute such an allocation, and also proved that a higher approximation ratio of $\frac{7}{8}$ could be guaranteed for the 3-agent case. After that, many follow-up works further improved the approximation factor (Ghodsi et al., 2018; Garg, McGlaughlin, et al., 2019; Barman and S. K. Krishnamurthy, 2020; Garg and Taki, 2021; Akrami, Garg, et al., 2023). The current best known lower bound for this approximation factor is $\frac{3}{4} + \frac{3}{3836}$ (Akrami and Garg, 2024).

On the negative side, Feige et al., 2021 give an example with three agents and nine goods for which no allocation gives to any agent a_i a better utility than $\frac{39}{40} u_i^{\text{MmS}}$. For $n \geq 4$, their construction can also be adapted to provide an instance for which no allocation is better than $(1 - \frac{1}{n^4})$ -MmS. This is the tightest upper bound known so far for the existence of an approximate MmS allocation. We can observe that the gap with the lower bound is still large.

To conclude this discussion about the existence of (approximate)-MmS allocations, we can observe that the interest of these results is mostly theoretical. As mentioned earlier, our personal observation was that the instances for which no MmS allocation exists are very rare. Our impression was confirmed in theory by Kurokawa et al., 2016 and Amanatidis, Markakis, et al., 2017 that proved that an MmS allocation exists with a very high probability for randomly drawn instances.

¹⁰Incidentally, they also obtained the best EC'14 paper award for the work presenting this example and proving (in a few pages) that it is correct.

2.3 Computational properties

Let us now focus on the computational properties of these criteria, with in mind the idea of determining whether desirable allocations can be easily computed or not. Here, we will focus on two main problems:

1. **X-COMP**: given an add-MARA instance $\langle \mathcal{A}, \mathcal{O}, w \rangle$ an agent a_i , and an integer K , do we have $u_i^X \geq K$ (here, X is a threshold-based criterion among MmS, PS, and mMS)?
2. **X-EXISTS**: given an add-MARA instance $\langle \mathcal{A}, \mathcal{O}, w \rangle$ and an agent a_i , does there exist an allocation $\vec{\pi}$ satisfying X (here, X is a criterion among MmS, PS, mMS, EF and CEEI)?

Max-min share Let us start with the first criterion of the scale. Obviously, even MMS-COMP is **NP**-complete: for a given agent a_i , computing u_i^{MmS} essentially comes down to solving an instance of the **PARTITION** problem. Now let us consider the MMS-EXISTS problem. What we can say for sure is that the problem belongs to Σ_2^P , because it can be solved by the following non-deterministic algorithm:

1. guess an allocation $\vec{\pi}$;
2. compute $u_i(\pi_i)$ for each agent a_i ;
3. check with an **NP** oracle that $u_i(\pi_i) \geq u_i^{MmS}$ for each agent a_i .

We strongly suspect that it is at least as hard as MMS-COMP, hence **NP**-hard, but to the best of our knowledge, no one has ever succeeded in proving it. [Nguyen and Rothe, 2023](#) even consider this problem as "*probably one of the most important open questions in the area of fair allocation*".

Open question 1 *What is the complexity of the problem of determining, for a given add-MARA instance, whether there exists an allocation satisfying the max-min share criterion?*

Proportional share Contrary to max-min share, PS-COMP is obviously easy. However, PS-EXISTS is **NP**-complete for a similar reason as for MMS-COMP: there is a direct reduction from the **PARTITION** problem to MMS-COMP with agents having identical preferences.

Min-Max share The case of min-max share is essentially similar to the max-min share, with the notable difference that mMS-COMP is now **coNP**-complete. Like MMS-EXISTS, mMS-EXISTS belongs to Σ_2^P , but its precise complexity is also unknown (although this open question has received far less attention from the community):

Open question 2 *What is the complexity of the problem of determining, for a given add-MARA instance, whether there exists an allocation satisfying the min-max share criterion?*

Envy-freeness As we have seen earlier, there is a long record of work concerning envy-freeness, and its complexity is well characterized. Checking whether an allocation is EF is polynomial, but [Lipton et al., 2004](#) have proved in their seminal paper that EF-EXISTS is **NP**-complete. Note that we also know ([de Keijzer et al., 2009](#)) that the problem of deciding whether an envy-free and Pareto-efficient allocation exists is Σ_2^P -complete.

Competitive equilibrium from equal incomes While the computation of a CEEI appears to be tractable in the continuous setting¹¹, it is not the case anymore in the discrete setting. More precisely, CEEI-EXISTS was finally proved to be NP-hard by Brânzei, Hosseini, et al., 2015, who also proved that it was CONP-hard to determine whether a given allocation $\vec{\pi}$ is a CEEI.

To the best of our knowledge, no practical method exists to compute a CEEI for an add-MARA instance, if we except the work of Othman et al., 2010 that is devoted to the computation of *approximate* CEEI. In practice, we use an algorithm based on the enumeration of all admissible allocations. For each allocation $\vec{\pi}$, determining whether this allocation is a CEEI comes down to solving the following program:

$$\begin{aligned} \text{find } & (\mathbf{p}_1, \dots, \mathbf{p}_m) \\ \text{s.t. } & 0 \leq \mathbf{p}_j \leq 1 & \forall o_j \in \mathcal{O} & \quad (\text{II.1}) \end{aligned}$$

$$\sum_{o_j \in \pi_i} \mathbf{p}_j \leq 1 \quad \forall a_i \in \mathcal{A} \quad (\text{II.2})$$

$$\sum_{o_j \in \pi'} \mathbf{p}_j > 1 \quad \forall a_i \in \mathcal{A} \text{ and } \pi' \text{ s.t. } u_i(\pi') > u_i(\pi_i). \quad (\text{II.3})$$

Constraints are self-explaining: Constraints II.2 express the fact that each agent can afford their share, while Constraints II.3 model the fact that a better share would be unaffordable. The difficulty here is that the last set of constraints is not linear. If we want to delegate to an LP solver the question of determining whether this allocation is a CEEI, we should transform a little bit the model by allowing the prices to be unbounded:

$$\begin{aligned} \text{find } & (\mathbf{p}_1, \dots, \mathbf{p}_m, \mathbf{d}) \\ \text{s.t. } & 0 \leq \mathbf{p}_j & \forall o_j \in \mathcal{O} & \quad (\text{II.4}) \end{aligned}$$

$$\sum_{o_j \in \pi_i} \mathbf{p}_j \leq \mathbf{d} \quad \forall a_i \in \mathcal{A} \quad (\text{II.5})$$

$$\sum_{o_j \in \pi'} \mathbf{p}_j \geq \mathbf{d} + 1 \quad \forall a_i \in \mathcal{A} \text{ and } \pi' \text{ s.t. } u_i(\pi') > u_i(\pi_i). \quad (\text{II.6})$$

As we have shown (Bouveret and Lemaître, 2016, Appendix), both models are equivalent in terms of satisfiability. This one being linear, we can use a LP solver to determine whether an allocation is CEEI. Note that it does not contradict the fact that the latter problem is CONP-complete, as the number of constraints in this program is exponential.

We implemented this approach in practice in a Python library dedicated to the computation of fair allocations.¹² Our experiments show that it can solve instances up to 4 agents and 10 objects in reasonable time. For bigger problems, other approaches (like, *e.g.* approximation algorithms) will probably be necessary.

¹¹ Actually, such an equilibrium can be found in polynomial time using the Eisenberg-Gale convex program, which comes down to maximizing the Nash CUF under linear constraints (Vazirani, 2007).

¹² This library is available online (<https://gricad-gitlab.univ-grenoble-alpes.fr/bouveres/fairdiv>).

2.4 Particular cases

Considering restricted settings of add-MARA instance can radically change the landscape of fairness criteria presented above. The first setting we consider is the one where all the agents have identical preferences. In some sense, these instances could be considered to be the ones where a fair solution is the most unlikely to exist, because the agents' preferences are extremely conflicting (basically, all the agents want the same objects). Actually, this is not completely the case. First, in that setting, there always exists an allocation satisfying max-min share. Second, it is easy to observe that any proportional allocation, if such an allocation exists, should give exactly the same amount of utility ($u_i(\mathcal{O})/n$) to all agents. It implies that:

- no proportional allocation can exist if, on top of being identical, preferences are also strict on shares (the agents should receive different bundles giving them the same utility, which is impossible in that case);
- any proportional allocation also satisfies mMS, EF, and CEEI. So basically in that case, our scale collapses to a two level scale.

Interestingly, other restrictions can make the problem of finding an allocation satisfying max-min share easy. This is for instance the case for 0–1 preferences, where such an allocation always exist and can be found by a picking sequence (see Chapter III). This is also the case for more intricate reasons if m does not exceed $n + 3$. And of course, this is the case if the number of agents is 2. The interested reader can refer to [Bouveret and Lemaître, 2015](#) (Section 6) for more details.

3 Balancing envy between agents

In the latter section we have introduced five fairness criteria that form a complete scale of properties, from the most demanding one (CEEI) to the weakest one (max-min share). The starting point of our study was motivated by the fact that in the context of indivisible goods, the two prominent fairness criteria – proportionality and envy-freeness – may be unsatisfiable in theory (and often are indeed in practice). Introducing a complete scale of criteria allows us first to characterize the degree of difficulty of a given instance (how conflicting or compatible the agents' preferences are), and second to be still able for most conflicting instances to propose a fallback solution to the fair division problem, even if no proportional or envy-free allocation can be found.

This section and the next two ones are still about trying to define intuitive and usable models of fairness mostly as weaker alternatives to envy-freeness. As we will see, the fairness notions we will propose do not form a scale as clearly as the ones we have introduced in Section II.1, but they are nevertheless related to each other, although based on completely different rationales.

3.1 Existing relaxations of envy-freeness

Before introducing our own relaxations of envy-freeness, let us start by observing that this question of how to properly relax envy-freeness has led to several proposals in the fair division literature, that we quickly review below.

First and foremost, [Lipton et al., 2004](#) have proposed to relax the binary notion of envy-freeness by defining several ways to measure envy as a quantity (supposed to be zero for any envy-free allocation).

This quantity is based on a measure of the *pairwise envy* between agents:

$$pe(a_i, a_j, \vec{\pi}) \stackrel{\text{def}}{=} \max\{0, u_i(\pi_j) - u_i(\pi_i)\}. \quad (\text{II.7})$$

The *individual envy* of an agent a_i can then be defined as the maximal pairwise envy¹³ between a_i and any other agent a_j :

$$e(a_i, \vec{\pi}) \stackrel{\text{def}}{=} \max_{a_j \in \mathcal{A}} (pe(a_i, a_j, \vec{\pi})). \quad (\text{II.8})$$

The implicit idea behind defining a quantified version of envy-freeness is to turn the fair division problem into an optimization problem, whose aim would be to find an allocation that minimizes in some way these individual envies. While [Lipton et al., 2004](#) focus on the problem of minimizing the sum of individual envies, we have proposed ([Shams et al., 2021](#)) another approach that consists in balancing as much as possible the vector of individual envies. That will be the topic of Section [II.3](#) below.

Defining a quantified version of envy is motivated by the fact that if an envy-free allocation does not exist, the agents should be ready to accept a small amount of envy. This is also the motivation behind the notion of *envy-freeness up to one good* (EF1). This notion was formally introduced by [Budish, 2011](#), but was already implicitly present in the seminal paper of [Lipton et al., 2004](#) (see *e.g.* Theorem 2.1). EF1 was then further refined by [Caragiannis, Kurokawa, et al., 2016](#) in the form of *envy-freeness up to any good* (EFX). The idea behind EF1 and EFX is that there exist situations where no envy-free allocation exists, but where envy disappears if we "forget" some items when comparing the agents' shares. More formally, an agent a_i is *EF1-satisfied* by an allocation $\vec{\pi}$ if $\forall a_j \in \mathcal{A}$, there is an item $o \in \pi_j$ such that $u_i(\pi_i) \geq u_i(\pi_j \setminus \{o\})$. Likewise, an agent a_i is *EFX-satisfied* by an allocation $\vec{\pi}$ if $\forall a_j \in \mathcal{A}$, $u_i(\pi_i) \geq u_i(\pi_j \setminus \{o\})$ for all $o \in \pi_j$ such that $u_i(o) \geq 0$. An allocation $\vec{\pi}$ is said to be EF1 (resp. EFX) if every agent is EF1-satisfied (resp. EFX-satisfied). An even more demanding version of EFX, called EFX₀ ([Plaut and Roughgarden, 2018](#)) requires no agent a_i envies the share π_j of another agent after removing from it any item (may it be strictly positively valued, or of utility 0). It is not difficult to see that these criteria also form a scale, as clearly:

$$(\vec{\pi} \models \text{EF}) \Rightarrow (\vec{\pi} \models \text{EFX}_0) \Rightarrow (\vec{\pi} \models \text{EFX}) \Rightarrow (\vec{\pi} \models \text{EF1}).$$

An interesting result is that the existence of an EF1 allocation is guaranteed. Even more interesting: such an EF1 allocation can be easily generated by executing a particular version of the allocation protocol that we will introduce in Chapter [III](#): picking sequences ([Caragiannis, Kurokawa, et al., 2016](#)). Concerning EFX, things seem more complicated, and the question of whether an EFX allocation always exists is still open:

Open question 3 *Does there always exist, for any given add-MARA instance with $n \geq 4$ agents, an EFX allocation?*

This open question has led to an intense research activity in the domain. In particular, it has been proved that an EFX allocation always exists in the following cases: 2 agents ([Plaut and Roughgarden, 2018](#); [Plaut and Roughgarden, 2020](#)), 3 agents ([Chaudhury et al., 2024](#)), 3 agents with at least one additive ([Akrami, Alon, et al., 2022](#)), identical agents ([Plaut and Roughgarden, 2018](#); [Plaut and Roughgarden, 2020](#)), additive valuations with two types of items ([Gorantla et al., 2023](#)), additive valuations with

¹³As a variant, the sum of pairwise envies can also be used, although using the maximal pairwise seems more usual.

two distinct possible values (Amanatidis, Birmpas, Filos-Ratsikas, et al., 2021). However, in spite of this active line of work, the general case still resists the attempts of the research community.

3.2 The minOWA criterion

As we have seen earlier, one of the most natural ideas to relax envy-freeness is to turn this binary criterion into a quantified version, that we seek to minimize. This is the idea proposed by Lipton et al., 2004, who focus on a quantity of envy defined as the sum of individual envies:

$$e(\vec{\pi}) = \sum_{a_i \in \mathcal{A}} e(a_i, \vec{\pi}), \quad (\text{II.9})$$

where $e(a_i, \vec{\pi})$ is the individual envy defined by Equations (II.7) and (II.8).

This approach has nevertheless a major drawback: focusing on the total measure of envy does not prevent this envy to be quite unequally shared among agents, exactly like classical utilitarianism is insensitive to unequal distribution of welfare. An extreme case could be that a single agent bears the entire envy, while the others do not suffer from any.

To fix this issue, we propose to aggregate individual envies using an operator that somewhat promotes equality, that is, an operator which is compatible with the *Pigou-Dalton inequality reduction principle*, while still maintaining some kind of efficiency flavour. To that aim, we chose to focus on *Ordered Weighted Averages* (OWA), that have been introduced by Yager, 1988 in the context of multiobjective optimization:

Definition 6 Let $\vec{\alpha} = (\alpha_1, \dots, \alpha_n)$ be a vector of weights. The ordered weighted average parameterized by $\vec{\alpha}$ is the function $\text{OWA}^{\vec{\alpha}} : \vec{x} \mapsto \sum_{i=1}^n \alpha_i \times x_i^\downarrow$, where \vec{x}^\downarrow denotes a permutation of \vec{x} such that $x_1^\downarrow \geq x_2^\downarrow \geq \dots \geq x_n^\downarrow$.

Among the operators of this family, we will only focus on those indeed reducing inequalities, that is, those such that $\vec{\alpha}$ is non increasing (we will call these vectors *fair*), hence giving more importance to the agents experiencing a higher amount of envy. Our solution concept can thus be expressed as the minimization of the quantity of envy defined using an OWA aggregator:

Definition 7 Let $(\mathcal{A}, \mathcal{O}, w)$ be an add-MARA instance, and $\vec{\alpha}$ be a non-increasing vector. An allocation $\vec{\pi}$ is an $\vec{\alpha}$ -minOWA Envy allocation if:

$$\vec{\pi} \in \arg \min_{\vec{\pi} \in \mathcal{F}} (\text{OWA}^{\vec{\alpha}} (e(a_1, \vec{\pi}), \dots, e(a_n, \vec{\pi}))).$$

A major advantage of this solution is that it always exists. Also note that it is compatible with envy-freeness: if an envy-free allocation exists, then its vector of envies will be $(0, \dots, 0)$, which will be minimal in the sense of $\text{OWA}^{\vec{\alpha}}$, no matter what $\vec{\alpha}$ is.

We might wonder to which extent this relaxation of envy-freeness is compatible with other allocations like EF1. It turns out that, unfortunately, it might be incompatible in theory:

Proposition 4 There exist add-MARA instances for which no allocation is at the same time EF1 and $\vec{\alpha}$ -minOWA Envy, no matter what $\vec{\alpha}$ is.

In other words, this result implies that these two fairness solution concepts might be incompatible, and that we should choose between the two. However, this can be tempered by the fact that (i) it is not the case anymore for binary utilities (in this particular case, there always exists a vector $\vec{\alpha}$ such that there is an allocation that is at the same time $\vec{\alpha}$ -minOWA Envy and EFX₀), and (ii) in practice, our experiments show that for real-world instances (Spliddit) and randomly-generated ones, EF1 and $\vec{\alpha}$ -minOWA Envy are almost always compatible (see [Shams et al., 2021](#), for more details on the experiments).

3.3 Computing minOWA envy allocations

We will now give some hints on how to compute such allocations efficiently. First, it is not hard to see that this minimization problem is NP-complete. Since most allocation problems can be naturally modeled as a Mixed Integer Program, and then efficiently handled in practice using an MIP solver, it is natural to model this problem as an MIP as well:

$$\text{minimize } \text{OWA}^{\vec{\alpha}}(\mathbf{e}_1, \dots, \mathbf{e}_n) \quad (\text{II.10})$$

$$\text{s.t } \mathbf{e}_i \geq \sum_{j=1}^m w_{i,j} \times (\mathbf{z}_i^j - \mathbf{z}_j^j) \quad \forall a_i, a'_i \in \mathcal{A} \quad (\text{II.11})$$

$$\mathbf{e}_i \geq 0 \quad \forall a_i \in \mathcal{A} \quad (\text{II.12})$$

$$\sum_{i=1}^n \mathbf{z}_i^j = 1 \quad \forall o_j \in \mathcal{O} \quad (\text{II.13})$$

$$\mathbf{z}_i^j \in \{0, 1\} \quad \forall a_i \in \mathcal{A}, \forall o_j \in \mathcal{O} \quad (\text{II.14})$$

Here, variables \mathbf{z}_i^j model the allocation: $\mathbf{z}_i^j = 1$ if and only if object o_j is allocated to agent a_i . \mathbf{e}_i is agent a_i 's individual envy, computed by Equations (II.11) and (II.12), as the maximal pairwise envy between a_i and any other agent a_j .

A major problem of this program is that the minimization criterion, given by Equation (II.10) is not linear. Linearizing the OWA operator is not completely straightforward, but can be done using a trick proposed by [Ogryczak and Śliwiński, 2003](#), that uses an alternative definition of the OWA operator using the Lorenz vector.

Let \vec{e} be a vector. Its *Lorenz vector* $\vec{L}(\vec{e})$ is the vector such that $L_k(\vec{e}) = \sum_{i=1}^k e_i^\downarrow$. It is easy to see that the OWA aggregation of a vector can be alternatively defined as a weighted sum of its Lorenz components:

$$\text{OWA}^{\vec{\alpha}}(\vec{e}) = \sum_{i=1}^n \alpha_i \times e_i^\downarrow = \sum_{i=1}^{n-1} (\alpha_i - \alpha_{i+1}) \times L_i(\vec{e}) + \alpha_n \times L_n(\vec{e}). \quad (\text{II.15})$$

The interesting thing is that the Lorenz components can be characterized as a linear program. Namely, the k^{th} component $L_k(\vec{e})$ can be computed as the result of the following LP:

$$\text{maximize } \sum_{i=1}^n \mathbf{a}_i^k \times e_i \quad (\text{II.16})$$

$$\text{s.t } \sum_{i=1}^n \mathbf{a}_i^k = k \quad (\text{II.17})$$

$$0 \leq \mathbf{a}_i^k \leq 1 \quad \forall i \in \llbracket 1, n \rrbracket \quad (\text{II.18})$$

Unfortunately, we cannot inject this linear program directly into the previous one, since the optimization directions (maximize / minimize) do not match. However, we can use the dual formulation of the latter program to make it compatible. Putting things together, if we respectively denote by \mathbf{r}_k and \mathbf{b}_i^k the dual variables of Equation (II.17) (resp. Equation (II.18)), the whole program can be expressed as follows:

$$\begin{aligned} \text{minimize } & \sum_{k=1}^n n - 1(\alpha_k - \alpha_{k+1}) \times (k\mathbf{r}_k + \sum_{i=1}^n \mathbf{b}_i^k) + \alpha_n \times (n\mathbf{r}_n + \sum_{i=1}^n \mathbf{b}_i^n) \\ \text{s.t } & \mathbf{r}_k + \mathbf{b}_i^k \geq \mathbf{e}_i \quad \forall i, k \in \llbracket 1, n \rrbracket \\ & \mathbf{b}_i^k \geq 0 \quad \forall i, k \in \llbracket 1, n \rrbracket \\ & + \text{Equations (II.11), (II.12), (II.13) and (II.14)} \end{aligned}$$

Experiments show that this approach performs quite well for real-world instances using a commercial-off-the-shelf MIP solver (see again [Shams et al., 2021](#), for more details on the experiments).

4 Epistemic envy

The idea behind relaxing envy-freeness using a measure of envy as proposed above is a quite standard approach in computer science, that basically consists in turning a Boolean criterion into a numerical measure that is supposed to reveal how close we are to satisfying the criterion. The most original contribution in this respect is the combination of the OWA operator and the measure of envy.

The relaxation we will propose in this section is based on a completely different approach. Envy-freeness, as well as the measure of envy introduced above, are inherently *knowledge-sensitive* concepts: envy, as it is defined, only makes sense if each agent exactly knows how the items are allocated. In other words, it is based on the assumption that the knowledge of the allocation among the agents is *maximal*. In practice, this is often not the case, either because the agents' shares are private knowledge, or because the agents may not even directly know who the other agents are.

4.1 Epistemic envy-freeness

Here, we propose to relax the assumption of complete knowledge and introduce an envy-freeness concept based on the alternative assumption that the agents have *minimal* knowledge about the allocation: this is what we call *epistemic envy-freeness* (EEF).

Definition 8 Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA. An allocation $\vec{\pi}$ satisfies the criterion of epistemic envy-freeness (or, is simply epistemic envy-free), denoted by $\vec{\pi} \vDash \text{EEF}$, if for each agent $a_i \in \mathcal{A}$ there exists an allocation $\vec{\pi}^i$ such that $\pi_i^i = \pi_i$ and $u_i(\pi_i) \geq u_i(\pi_j^i)$ for each agent $a_j \in \mathcal{A}$.

In other words, an allocation is EEF if for every agent a_i , even if a_i does not know what the other agents receive, they can think of an allocation $\vec{\pi}^i$ of the rest of items to the rest of agents such that in this allocation, they do not envy anyone. Of course, the true allocation $\vec{\pi}$ could be different from $\vec{\pi}^i$, but it does not matter to a_i because they do not know the actual allocation.

Obviously, an envy-free allocation is also epistemic envy-free, but the converse is not true. Even more: there exist add-MARA instances admitting an EEF allocation but no EF allocation (Aziz, Bouveret, Caragiannis, et al., 2018, Proof of Theorem 5).

More interestingly, this new concept is linked to another property of our scale of fairness criteria introduced in Section II.2:

Proposition 5 Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance, and let $\vec{\pi}$ be an allocation. Then:

$$(\vec{\pi} \vDash \text{EF}) \Rightarrow (\vec{\pi} \vDash \text{EEF}) \Rightarrow (\vec{\pi} \vDash \text{mMS}).$$

Moreover, this entailment is strict: there indeed exist add-MARA instances admitting an mMS allocation but no EEF allocation (see Aziz, Bouveret, Caragiannis, et al., 2018, Proof of Theorem 6). This shows that, interestingly, epistemic envy-freeness constitutes an intermediate level of our now 6-level scale of fairness properties.

Epistemic envy-freeness up to any good This notion of epistemic envy-freeness has recently been extended by Caragiannis, Garg, et al., 2023 to integrate the concept of envy-freeness up to any good (mentioned in Section II.3.1). Namely, an allocation $\vec{\pi}$ is *epistemic envy-free up to any item* (EEFX) if for each agent $a_i \in \mathcal{A}$ there exists an allocation $\vec{\pi}^i$ such that $\pi_i^i = \pi_i$ and agent a_i is EFX-satisfied with $\vec{\pi}^i$.

In the same paper, Caragiannis, Garg, et al., 2023 also define the notion of *minimum EFX share* (mXS) of an agent as the minimum utility an agent a_i can get in an allocation where they are EFX-satisfied. In other words,

$$u_i^{\text{mXS}} = \min_{\vec{\pi} \in \mathcal{F} \text{ s.t. } a_i \text{ is EFX-satisfied by } \vec{\pi}} u_i(\pi_i).$$

Similarly to max-min share, an allocation $\vec{\pi}$ satisfies mXS if $u_i(\pi_i) \geq u_i^{\text{mXS}}$ for all a_i .

What is interesting is that these notions are related to each other, and also relate to other fairness criteria, as we have: $(\vec{\pi} \vDash \text{MmS}) \Rightarrow (\vec{\pi} \vDash \text{EEFX}) \Rightarrow (\vec{\pi} \vDash \text{mXS}) \Rightarrow (\vec{\pi} \vDash \text{PROP1})$ – here, PROP1 is the counterpart of EF1 for proportionality (Caragiannis, Garg, et al., 2023).

4.2 Envy-freeness with social constraints

The latter notion of epistemic envy-freeness clearly corresponds to a pessimistic view of the knowledge an agent has on the allocation (where traditional envy-freeness corresponds to an optimistic point of view). Intermediate points of view are possible. For instance, we may assume that an agent only knows a subset of the other agents enough to know what bundle they have received. More precisely, we can assume here that there is directed graph called the *social graph* modeling the acquaintances between agents. Each node

corresponds to an agent. If there is an arc between a and a' , then it means that a knows the items allocated to a' . This notion of social graph naturally leads to the definition of graph-epistemic envy-freeness that follows:

Definition 9 Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA and G be a social graph over \mathcal{A} . An allocation $\vec{\pi}$ is G -epistemic envy-free, denoted by $\vec{\pi} \vDash G\text{-EEF}$, if for each agent $a_i \in \mathcal{A}$ there exists an allocation $\vec{\pi}^i$ such that $\pi_j^i = \pi_j$ for all $a_j \in \{a_i\} \cup \text{Neigh}_G(a_i)$ and $u_i(\pi_i) \geq u_i(\pi_j^i)$ for each agent $a_j \in \mathcal{A}$.

Obviously, $G\text{-EEF}$ corresponds to EEF if G has no edge, and corresponds to EF if every node in G has out-degree at least $n - 2$.¹⁴ It is also not difficult to observe that for any subgraph H of a graph G and any allocation $\vec{\pi}$, $(\vec{\pi} \vDash G\text{-EEF}) \Rightarrow (\vec{\pi} \vDash H\text{-EEF})$. Intuitively, since the agents know more agents in G than in H , $G\text{-EEF}$ puts more constraints to the allocation than $H\text{-EEF}$.

Interestingly (and less obviously), this implication is strict, in the sense that if we consider two graphs G and H such that edge (u, v) is in G but not in H and $\deg_H(u) \leq n - 3$, there exist add-MARA instances that have $H\text{-EEF}$ allocations, but no $G\text{-EEF}$ one. Hence, the set of social graphs forms, w.r.t inclusion, a very rich hierarchy of fairness concepts between EF and EEF.

Note that our proposition of $G\text{-EEF}$ is not the first one trying to integrate social constraints into the definition of fairness properties. For instance, Abebe et al., 2017; Bei et al., 2017; Chevaleyre, Endriss, et al., 2017 also assume that the agents are linked by a social network, and define fairness accordingly. However, the main difference here is that the objects allocated to the unknown (non-neighbours) agents are completely ignored.

4.3 A weaker social epistemic envy-freeness property

In some sense, $G\text{-EEF}$ is still quite demanding because each agent a_i must be non envious of their neighbours, and have to imagine that an allocation of the remaining objects to the unknown agents is possible so that a_i is not envious of those agents either. Instead of asking a_i to imagine such a fictive allocation, we could just relax this assumption and just assume that a_i will be happy if they obtain their proportional fair share among the remaining agents. This is the idea behind $G\text{-PEF}$:

Definition 10 Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA and G be a social graph over \mathcal{A} . An allocation $\vec{\pi}$ is G -proportional envy-free, denoted by $\vec{\pi} \vDash G\text{-PEF}$, if for each agent $a_i \in \mathcal{A}$ $u_i(\pi_i) \geq u_i(\pi_j)$ for all $a_j \in \text{Neigh}_G(a_i)$, and:

$$u_i(\pi_i) \geq \frac{1}{n - 1 - \deg_G(i)} \times \sum_{a_j \notin \text{Neigh}_G(a_i)} u_i(\pi_j).$$

Obviously, here, $G\text{-PEF}$ corresponds to PROP if G has no edge, and corresponds to EF if every node in G has out-degree at least $n - 2$. We can also observe that this property is weaker than $G\text{-EEF}$:

Proposition 6 Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance, G be a social graph over \mathcal{A} , and let $\vec{\pi}$ be an allocation. Then:

$$(\vec{\pi} \vDash G\text{-EEF}) \Rightarrow (\vec{\pi} \vDash G\text{-PEF}).$$

¹⁴This condition is obviously satisfied for a complete graph.

Once again, this implication is strict as soon as G contains a node of out-degree at most $n - 3$. And once again, if H is a subgraph of G , then for any allocation $\vec{\pi}$, $(\vec{\pi} \vDash G\text{-PEF}) \Rightarrow (\vec{\pi} \vDash H\text{-PEF})$. This implication is again strict, as soon as G contains an edge (u, v) with $\deg_H(u) \leq n - 3$ that does not belong to H .

Finally, we have mentioned earlier in Section II.4.2 the weaker notion of graph-envy-freeness, introduced by Chevalleyre, Endriss, et al., 2017. We can easily observe that this notion is even weaker than G -PEF, in the sense that any G -PEF allocation is also G -EF. Interestingly, we can prove that G -EF is actually strictly weaker than G -PEF, as long as G is not complete of course.

5 Envy approved by the society

Epistemic envy was based on the fact that envy is a *knowledge-sensitive* notion. The relaxation we will introduce in this section is rooted in the fact that envy is also an inherently *subjective* notion. To introduce this notion, consider an instance where no envy-free allocation exists. Suppose that there nevertheless exist two allocations $\vec{\pi}$ and $\vec{\pi}'$ such that in both allocation only one agent a_i is envious (say of a_j in both cases). Now suppose that in $\vec{\pi}$ all agents concur with the fact that a_j 's bundle is indeed better than a_i 's, while in $\vec{\pi}'$ everyone disagrees with a_i . According to Parijs, 1997, it is hard to justify that $\vec{\pi}$ is a better allocation to $\vec{\pi}'$: in the latter allocation, the envy of a_i towards a_j is socially unsupported. Taking social consensus as a proxy for objectivity, we could say that in the case of $\vec{\pi}$, a_i 's envy is justified whereas it is not in the case of $\vec{\pi}'$. In other words, $\vec{\pi}$ exhibits *unanimous envy*.

5.1 K-approval envy-freeness

Unanimous envy requires that the envy of an agent towards another one is supported by all agents in the society to consider that this envy is justified. We could think of a less extreme version of the same idea, where only a given threshold K of the agents is required to support envy.

Definition 11 Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance, let $\vec{\pi}$ be an allocation, a_i and a_j be two different agents and $K \in \llbracket 1, n \rrbracket$ be an integer. We say that a_i K -approval envies (K -app envies for short) a_j if there is a subset \mathcal{A}_K of K agents including a_i such that:

$$\forall a_k \in \mathcal{A}_K, u_k(\pi_i) < u_k(\pi_j).$$

In other words, a_i envies a_j and $K - 1$ other agents think that this envy is justified. If a_i n -app envies a_j , we say that a_i unanimously envies a_j . Obviously, a_i 1-app envies a_j if and only if a_i envies a_j .

We can define the counterpart of envy-freeness for this notion of approval envy:

Definition 12 Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. An allocation $\vec{\pi}$ is (K -app envy)-free, denoted by $\vec{\pi} \vDash (K\text{-app envy})\text{-free}$, if a_i does not K -app envies a_j for all pairs of agents (a_i, a_j) .

Observe that for an allocation to be (K -app envy)-free, for all pairs of agents (a_i, a_j) , either a_i or at least $n - K + 1$ agents have to think that a_i does not envy a_j . Notice that this is different from requiring

that at least K agents think that this allocation is envy-free. This explains the parenthesis around (K -app envy): this notion means "free of K -app envy", which is different from " K -app-(envy-free)".

Also observe that this notion of K -approval envy can be interpreted as a vote: when an agent a_i declares that they envy a_j , then a vote is organized to determine whether this envy is justified or not. If the collective consensus decides that this envy is not justified, then it is discarded and we do not take it into account when evaluating the current allocation. Here, several voting procedures could be used to determine whether this envy is justified. However, since the only two options are yes and no, the most reasonable voting rules are quota rules (Perry and Powers, 2010): a_i envies a_j if and only if at least a threshold of agents think so.

Properties of K -approval envy First, we can observe that if a_i K -app envies a_j , then a_i $(K - 1)$ -app envies a_j as well. Hence, if an allocation \vec{w} is (K -app envy)-free, then it is $((K + 1)$ -app envy)-free as well. However, the converse is not true: for $K \geq 3$, if an allocation \vec{w} is (K -app envy)-free, then it is not necessarily $((K - 1)$ -app envy)-free as well. Furthermore, there exist add-MARA instances that admit a (K -app envy)-free allocation, but no $((K - 1)$ -app envy)-free one (see Shams et al., 2022, Propositions 1 and 2), which proves that K -approval envy defines a strict and non-trivial hierarchy of fairness concepts for $K \geq 3$.

The careful reader might have noticed that the hierarchy described above is only valid for $K \geq 3$. What happens for $K = 2$? Interestingly, the following proposition shows that (2-app envy)-freeness is essentially a vacuous notion:

Proposition 7 *Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. If this instance admits a (2-app envy)-free allocation, then it also admits an envy-free allocation.*

The proof of this result is based on a technique similar to the *bundle reallocation technique* introduced by Lipton et al., 2004 and used later in many contexts involving envy. The idea is to perform a reallocation of the bundles by following the reverse direction of a cycle in the envy graph (where each agent is a node, and there is an edge between a_i and a_j if a_i envies a_j). Unfortunately, performing a cyclic reallocation of bundles does not guarantee that the K -app envy will decrease (the resulting allocation can even be worse in terms of approval envy). However, we can prove that for $K = 2$, it is always possible to find a weakly improving cycling reallocation that still guarantees (K -app envy)-freeness for $K \leq 2$. Proposition 7 can then be deduced from this observation.

Let us end this section about (K -app envy)-freeness by giving some hints about the complexity of decision problems related to this notion. First of all, since envy-freeness is equivalent to (1-app envy)-freeness, the problem of finding the minimum K for which there exists a (K -app envy)-free allocation is at least as hard as determining whether an envy-free allocation exists (that is, NP-complete). Moreover, for agents having identical preferences, observing that the entire hierarchy of (K -app envy)-freeness collapses (and that (1-app envy)-freeness and (unanimous envy)-freeness coincide), we can deduce that deciding whether an instance admits no (unanimous envy)-free allocation is CONP-complete.¹⁵

¹⁵See Shams et al., 2022, Corollary 2. Unfortunately, the way this corollary is stated in the paper is incorrect, because (i) it deals with allocations instead of instances, and (ii) it claims that it is NP-complete instead of CONP-complete. The correct version of the aforementioned corollary is "Deciding whether an instance exhibits unanimous envy is CONP-complete."

5.2 K-approval non-proportionality

As we have seen in Section II.2, envy is not the only fairness notion essentially based on a subjective evaluation. This is also the case for instance for proportionality. Several relaxations of proportionality have been proposed in the literature, like PROP1, the counterpart of EF1 for proportionality (Aziz, Moulin, et al., 2020; Barman and S. Krishnamurthy, 2019; Brânzei and Sandomirskiy, 2019; Conitzer et al., 2017). It is thus natural to propose a relaxation of proportionality based on the same idea as approval envy-freeness:

Definition 13 Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance, let $\vec{\pi}$ be an allocation, a_i be an agent, and $K \in \llbracket 1, n \rrbracket$ be an integer. We say that π_i is K -approval non proportional (K -app non prop for short) if there is a subset \mathcal{A}_K of K agents including a_i such that:

$$\forall a_k \in \mathcal{A}_K, u_k(\pi_i) < \frac{u_k(\mathcal{O})}{n}.$$

$\vec{\pi}$ will be said to be (K -approval non proportional)-free if and only if no π_i is K -app non prop.

Once again, the interpretation of this property is that an allocation is free of K -app non-prop: each agent a_i either thinks they receive a proportional share, or, if it is not the case, no more than $K - 2$ other agents agree with them. Here, we focus on non-proportionality rather than proportionality to be consistent with our definition of K -app envy.

Properties of K -approval non-proportionality K -approval non-proportionality exhibits similar properties than K -approval envy. First, if π_i is K -app non-prop, then it is $(K - 1)$ -app non-prop as well. Hence, if an allocation $\vec{\pi}$ is (K -app non-prop)-free, then it is $((K + 1)$ -app non-prop)-free as well, the converse not being true (like before): for $K \geq 3$, if an allocation $\vec{\pi}$ is (K -app non-prop)-free, then it is not necessarily $((K - 1)$ -app non-prop)-free as well. Furthermore, there exist add-MARA instances that admit a (K -app non-prop)-free allocation, but no $((K - 1)$ -app non-prop)-free one (see Shams et al., 2022, Propositions 6 and 7), which proves that, just like K -approval envy, K -approval non-proportionality defines a strict and non-trivial hierarchy of fairness concepts for $K \geq 3$.

Very much like K -approval envy, (2 -app non-prop)-freeness is also essentially a vacuous notion:

Proposition 8 Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. If this instance admits a (2 -app non-prop)-free allocation, then it also admits a proportional allocation.

This proposition can be again proved by using weakly improving cycles, with the word of caution that only some of such cycles can be improving in terms of (K -app non-prop)-freeness.

And finally, we have exactly the same complexity results for (K -app non-prop)-freeness than those we had for (K -app envy)-freeness: the problem of finding the minimum K for which there exists a (K -app non-prop)-free allocation is NP-complete. And again, the problem of deciding whether an instance admits no (unanimous non-prop)-free allocation is CONP-complete,¹⁶ which can be again seen by observing that the hierarchy of (K -app non-prop)-freeness properties collapses if all the agents have the same preferences.

¹⁶See Shams et al., 2022, Corollary 4, which contains the same mistakes as Corollary 2 (see Footnote 15).

Links with (K -app envy)-freeness As we have seen earlier, envy-freeness implies proportionality. The same holds for their respective relaxations, namely, EF1 and PROP1. It is thus natural to ask whether it is also the case for approval envy and approval non-proportionality. Unfortunately, the same relations do not seem to hold for these relaxations:

Proposition 9 *There exist add-MARA instances that have (3-app envy)-free allocations, but no (unanimous non-prop)-free ones.*

As a corollary of this proposition, for $K \geq 3$, (K -app envy)-freeness does not imply (K -app non-prop)-freeness. Moreover:

Proposition 10 *There exist add-MARA instances that admit proportional allocations, but no (unanimous envy)-free ones.*

These two propositions together imply that K -app envy and K -app non-proportionality seem to form two somewhat unrelated hierarchies of fairness properties (if we except the link between (1-app envy)-freeness and (1-app non-prop)-freeness).

5.3 Computation

As we have seen earlier, the problem of finding, for a given add-MARA instance, the minimum K such that the instance admits a (K -app envy)-free (resp. (K -app non-prop)-free) allocation is NP-complete. It does not mean that we cannot solve this problem efficiently in practice, *e.g.* by modeling this problem as a Mixed Integer Program, and by using MIP solvers.

Our MIP formulation uses the same variables \mathbf{z}_i^j as Section II.3 to model the allocation: $\mathbf{z}_i^j = 1$ if and only if object o_j is allocated to agent a_i . We also need the same constraint defined by Equation (II.12) that each item is allocated to exactly one agent.

On top of that variables and constraints, we need to introduce 0–1 variables and constraints that model the notion of envy (resp. proportionality) of a_i as perceived by another agent a_k :

- $\mathbf{e}_{kih} = 1$ if and only if a_i envies a_b according to a_k 's preferences;
- $\mathbf{p}_{ki} = 1$ if and only if a_i 's bundle is non-proportional according to a_k 's preferences.

The constraints that link variables \mathbf{e}_{kih} and \mathbf{p}_{ki} to allocation variables \mathbf{z}_i^j can simply be expressed as follows:

$$\sum_{j=1}^m w_{k,j} (\mathbf{z}_h^j - \mathbf{z}_i^j) > 0 \Leftrightarrow \mathbf{e}_{kih} = 1 \quad \forall k, i, h \in \llbracket 1, n \rrbracket \quad (\text{II.19})$$

$$\sum_{j=1}^m u(k, j) \cdot \mathbf{z}_i^j < \frac{\sum_{j=1}^m w_{k,j}}{n} \Leftrightarrow \mathbf{p}_{ki} = 1 \quad \forall k, i \in \llbracket 1, n \rrbracket \quad (\text{II.20})$$

Finally, we have to write the constraints that convey the fact that the allocation we want is (K -app envy)-free (resp. (K -app non-prop)-free):

$$\mathbf{e}_{iib} = 0 \vee \sum_{k=1}^n \mathbf{e}_{kih} \leq \mathbf{K} - 1 \quad \forall i, b \in \llbracket 1, n \rrbracket \quad (\text{II.21})$$

$$\mathbf{p}_{ii} = 0 \vee \sum_{k=1}^n \mathbf{p}_{ki} \leq \mathbf{K} - 1 \quad \forall i \in \llbracket 1, n \rrbracket \quad (\text{II.22})$$

All these constraints can be easily linearized using standard techniques. The global MIP is then obtained by minimizing \mathbf{K} under constraints defined by Equation (II.12) and linearizations of Equations (II.19) and (II.21) for approval envy (resp. (II.20) and (II.22) for approval non proportionality). See [Shams et al., 2022](#) for the complete linear model.

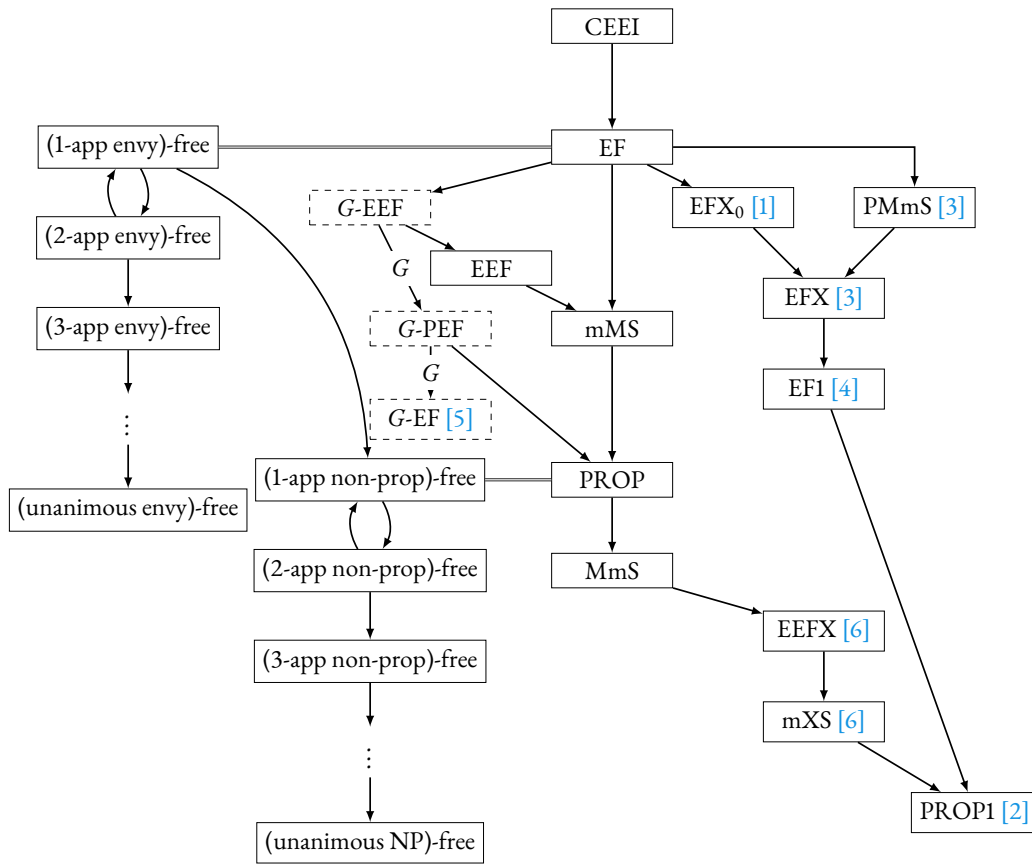
6 Conclusion: the landscape of fairness properties

In this chapter, we have introduced several original fairness criteria. These criteria complete existing fairness properties (on top of which proportionality and envy-freeness) to form a complex landscape of connected criteria. This landscape is summarized in Figure II.1.

How can we use such a landscape as a (central) decision maker? First of all, as proposed in Section II.2.2 we can look for an allocation that satisfies the highest possible criterion in the hierarchy, with the idea in mind that the strongest criterion is satisfied, the fairest the allocation is. This is the idea behind the scale of fairness introduced in Section II.2, but also the idea between the minOWA-envy approach.

Now suppose that the highest satisfiable property is low in the scale (or that minimizing OWA envy does not lead to very close bounds). Then, depending on the form of the problem at stake, we argue that opting for other branches of the landscape may be reasonable. In contexts where the problem involves a lot more items than agents, looking for allocations that satisfy EFX, EF1 or some derivatives (right branch of Figure II.1) is relevant, all the more that EF1 is always satisfiable and can be found either by maximizing the Nash CUF, or by applying a simple round-robin algorithm ([Caragiannis, Kurokawa, et al., 2016](#)). In contexts where the agents only partially know each other, then an approach based on epistemic envy-freeness definitely makes sense. Conversely, in contexts where all the agents know each other and have the possibility to discuss in order to determine whether the subjective perception of unfairness is justified or not, then it seems very appropriate to choose an approach based on approval envy-freeness and non-proportionality.

Let us go back to our initial question, presented in the introduction of this document, of defining what is exactly fairness in the contexts of fair division and voting. We can see here that there is no unique answer, and that the appropriate rule to apply may depend on the context. That will be exactly the same in the context of voting (see Chapter IV), where different rules convey different notions of "fairness", and choosing the right one depend on the situation at stake, and can probably only be decided by a collaborative approach between educated stakeholders.



- [1] [Plaut and Roughgarden, 2018](#)
- [2] [Aziz, Moulin, et al., 2020](#); [Barman and S. Krishnamurthy, 2019](#); [Brânzei and Sandomirskiy, 2019](#); [Conitzer et al., 2017](#)
- [3] [Caragiannis, Kurokawa, et al., 2016](#)
- [4] [Budish, 2011](#)
- [5] [Chevaleyre, Endriss, et al., 2017](#)
- [6] [Caragiannis, Garg, et al., 2023](#)

Figure II.1 — Landscape of (binary) fairness properties. An arrow indicates an implication link. A double edge indicates an equivalence. Dashed properties indicate hierarchies induced by a graph G . A label G on an arrow indicates that the implication is valid only for the same graph G .
 Note that the minOWA-envy approach does not appear in this figure since it is not strictly speaking a binary criterion.

Chapter III

The unreasonable fairness of picking sequences

Chapter foreword

This chapter is dedicated to the study of a particular fair division protocol, namely picking sequences.

The starting point of this chapter, addressing the question of finding the fairest sequence, is essentially based on a work with Jérôme Lang (Bouveret and Lang, 2011). The natural question that arises afterwards is the question of strategyproofness. The study of this question has led to some follow-up works on manipulation that resulted from a collaboration with Haris Aziz (Bouveret and Lang, 2014; Aziz et al., 2017). On the other hand, the study of strategyproof picking sequences led to a collaboration with Jérôme Lang, Hugo Gilbert and the master student Guillaume Méroüé (Bouveret, Gilbert, et al., 2023), that will soon be submitted to a journal.

In the final part of the chapter we exhibit unexpected connections between Pareto-efficiency, distributed fair division, and picking sequences. This contribution is the result of a close cooperation with Aurélie Beynier, Nicolas Maudet, and our former PhD student Parham Shams, and with Michel Lemaître and Simon Rey, that has been presented in 2019 at the international conference on Autonomous Agents and MultiAgent Systems (Beynier et al., 2019).

We also mention in the conclusion the notion of price of elicitation freeness, formally defined in a paper resulting from a collaboration with Dorothea Baumeister, Jérôme Lang, Trung Thanh Nguyen, Jörg Rothe and Abdallah Saffidine (Baumeister et al., 2017), and the extension of picking sequences to voting, which is the topic of a paper resulting from a collaboration with Yann Chevaleyre, François Durand and Jérôme Lang (Bouveret, Chevaleyre, et al., 2017).

Aziz, Haris, Sylvain Bouveret, Jérôme Lang, and Simon Mackenzie (2017). “Complexity of Manipulating Sequential Allocation”. In: *Proceedings of the 31st AAAI conference on Artificial Intelligence (AAAI’17)*. San Francisco, California, USA: AAAI Press. URL: <http://recherche.noiraudes.net/resources/papers/AAAI17.pdf>.

Baumeister, Dorothea, Sylvain Bouveret, Jérôme Lang, Nhan-Tam Nguyen, Trung Thanh Nguyen, Jörg Rothe, and Abdallah Saffidine (May 2017). “Positional scoring-based allocation of indivisible goods”. In: *Autonomous Agents and Multi-Agent Systems* 31.3, pp. 628–655. ISSN: 1573-7454. DOI: [10.1007/s10458-016-9340-x](https://doi.org/10.1007/s10458-016-9340-x). URL: <http://recherche.noiraudes.net/resources/papers/JAAMAS16.pdf>.

Beynier, Aurélie, Sylvain Bouveret, Michel Lemaître, Nicolas Maudet, Simon Rey, and Parham Shams (May 2019). “Efficiency, Sequenceability and Deal-Optimality in Fair Division of Indivisible Goods”. In: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS’19)*. Montreal, Canada: IFAAMAS, 9 pages. URL: <http://recherche.noiraudes.net/en/cycle-deals.php>.

Bouveret, Sylvain, Yann Chevaleyre, François Durand, and Jérôme Lang (Aug. 2017). “Voting by Sequential Elimination with few Voters”. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-17)*. Melbourne, Australia. URL: <http://recherche.noiraudes.net/resources/papers/IJCAI17-sequences.pdf>.

Bouveret, Sylvain, Hugo Gilbert, Jérôme Lang, and Guillaume Méroüé (2023). *Thou Shalt not Pick all Items if Thou are First: of Strategyproof and Fair Picking Sequences*. arXiv: [2301.06086](https://arxiv.org/abs/2301.06086) [cs.GT]. URL: <https://arxiv.org/abs/2301.06086>.

Bouveret, Sylvain and Jérôme Lang (July 2011). “A general elicitation-free protocol for allocating indivisible goods”. In: *Proceedings of the 22st International Joint Conference on Artificial Intelligence (IJCAI’11)*. Barcelona, Spain. URL: <http://recherche.noiraudes.net/resources/papers/IJCAI11.pdf>.

— (Aug. 2014). “Manipulating picking sequences”. In: *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI’14)*. Prague, Czech Republic: IOS Press. URL: <http://recherche.noiraudes.net/resources/papers/ECAI14.pdf>.

In the previous chapter, we mainly focused on providing actionable definitions of fairness, that could be used as a basis for practical fair division protocols. The underlying idea was to help a central (and assumed to be benevolent) decision-maker to find desirable allocations by first eliciting the agents' preferences and then solving a combinatorial problem to automatically compute an allocation that satisfies these fairness criteria.

Such an approach obviously has several drawbacks. First, it relies on a central authority which the agents should trust, both to protect their privacy, and to strictly apply ethical principles of distributive fairness without favouring any of the agents. This central authority should also possess a high computation power, because as we have seen earlier (*e.g.* in Section II.2), most problems that have to be solved are NP-hard. Finally, even eliciting preferences can be a complex task, as some agents may only have a rough idea of what their opinions on the items are, and that can only hardly translate to a precise (and additive) utility function.

In situations where relying on a central authority is not desirable or not even feasible, other approaches are possible. The first one is *distributed fair division*. Here, the idea is to start from an initial allocation (*e.g.* a random one), and let the agents negotiate to swap their objects. This approach has led to a fruitful line of work (see, *e.g.* Sandholm, 1998; Endriss et al., 2006) that has in particular investigated the convergence properties of different restrictions of bundle exchanges (for instance swaps restricted to two agents, exchanges along a graph,...). We mention this approach in Section III.6.

Another possibility is to use an *interactive protocol*, whose objective is to build an allocation by asking the agents a sequence of questions. This kind of approach is common in fair division of divisible goods (cake-cutting), but has also been studied in the context of indivisible goods. In particular, when there are two agents, the *Adjusted Winner* procedure (Brams and Taylor, 2000) is able to efficiently compute an equitable (both agents enjoy the same utility), envy-free and Pareto-efficient allocation, provided that one good can be split between both agents (that is, a fraction of this good can be allocated to each agent), or, equivalently, that some money is available as a compensation for an unbalanced utility. Later, Brams, Kilgour, et al., 2012 have proposed another procedure that guarantees, still for two agents and under mild conditions, an envy-free allocation, without the presence of fractional goods. This procedure, called the *Undercut Procedure*, takes as input only ordinal information from the agents, and informally works as follows. During a first phase, the agents are asked to name their preferred item. If both items differ, they are allocated to the agents asking them. Otherwise, the item is put aside on a *contested pile*. The second phase starts when all items have been either allocated to one agent, or put in the contested pile. During this second step, the items of the contested pile are allocated by asking each agent to name their *minimal bundles*. Here, a "minimal bundle" for agent i is any bundle π such that (i) $\pi \succeq_i \mathcal{O} \setminus \pi$, and (ii) any π' satisfying condition (i) is such that $\pi' \succ \pi$. When the minimal bundles have been named, the protocol randomly chooses one of them as a proposal (for instance π_1 for agent 1). Then, the other agent (2) has the opportunity either to accept $\mathcal{O} \setminus \pi_1$ or to undercut π_1 by taking for herself a bundle strictly less preferred than π_1 . Interestingly, it can be shown that if the minimal bundles differ, then the protocol guarantees to yield an envy-free allocation.

What can we do for more than two agents? Herreiner and Puppe, 2002 have proposed a very simple protocol that returns a Pareto-efficient and rank-maxmin-optimal allocation: the *Descending Demand Procedure*. This protocol works by asking the agents to name one by one their preferred bundles, until a feasible allocation can be found. The underlying assumption is that the agents have in mind a linear order over all the bundles of items, which can soon become unrealistic if the number of items grow. Another possibility (when ordering all the bundles becomes unfeasible) is to take inspiration from the generation phase of the undercut procedure and allocate goods incrementally. Of course, if the number of

agents grows and they all name their preferred item simultaneously, all goods will be likely to be contested, making the protocol essentially pointless. An alternative solution is to fix beforehand a sequence among agents and ask them to name their preferred item in turns. This is the idea behind *picking sequences*, that are the central topic of this chapter.

Informally, *picking sequences* work as follows. Suppose we have m objects to allocate to n agents. The central authority defines a sequence of agents of length m (with agents appearing several times in the sequence if $m > n$). Every time an agent is designated, they pick one object out of those that remain. For instance, if $n = 3$ and $m = 5$, the sequence $a_1a_2a_3a_3a_2$ means that agent a_1 picks an object first; then a_2 picks an object; then a_3 picks two objects in a row; and a_2 takes the last object. Here, fairness can be ensured by the choice of an appropriate sequence of agents by the central authority. For instance, for $n = 3$ and $m = 5$, we feel that the aforementioned sequence ($a_1a_2a_3a_3a_2$) is fairer than would be the sequence $a_1a_1a_2a_2a_3$, simply because in the first case, the fact that agent a_1 can pick their most preferred object is compensated by the fact that they will pick only once (instead of twice for the other agents). The central question of this chapter is then:

| Given a number of agents n and a number of objects m , what is the fairest sequence?

The starting point and the first part of this chapter (Sections III.1, III.2 and III.3) are entirely dedicated to the study of this question. The second part of the chapter (Sections III.4 and III.5) deals with a natural follow-up problem, namely, the question of knowing whether such a protocol can be easily manipulated by the agents and whether it is possible to find a version of the protocol that is at the same time fair and strategyproof.¹ The chapter then ends up with a contribution that shows nice and unexpected connections between Pareto-efficiency, distributed fair division, and picking sequences, in the form of a *scale of efficiency* that takes a similar form to the scale of fairness introduced in Chapter III. That completes the landscape of fairness criteria discussed in Figure II.1 in a nice way by adding some complementary efficiency properties, presented in Figure III.3.

1 An elicitation-free protocol for fair division

Let us start by formally defining the picking sequences protocol.

Definition 14 Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance. A picking sequence (or simply sequence when the context is clear) is a vector of \mathcal{N}^m . We will denote by $S(I)$ the set of possible sequences for instance I .

Let $\vec{\sigma} \in S(I)$ be a sequence of agents. We will denote by σ_t the t^{th} agent of the sequence. We will also say that $\vec{\sigma}$ generates an allocation $\vec{\pi}$ if $\vec{\pi}$ can be obtained as one of the possible results of the protocol where we go through the sequence $\vec{\sigma}$ of agents and each agent chooses *sincerely* at their turn one object among those of maximal utility for them. Note that several (potentially an exponential number of) allocations can be generated from the same sequence $\vec{\sigma}$, since several objects can have the same (maximal) utility for a given agent.

¹This question is also relevant in the context of centralized fair division, in spite of the fact that we have completely eluded it in Chapter II.

This protocol has arguably appealing properties. It is simple to understand, can be explained in a few words, and is preference elicitation-free. Technically, it does not even require any particular (*e.g.* additive) assumption on the agents' preferences, since the agents choose themselves which object they pick at their turns without explicitly expressing their preferences in the form of a preorder or utility function. However, we might notice that this protocol is not very well adapted to situations where the preferences over objects are complementary. In such situations where, *e.g.*, two objects are complements, it might be difficult and risky for an agent to pick one of the two without being sure that the second one will still be available the next time they pick. In spite of this limitation, this simple protocol is used in practice in a lot of everyday life situations, ranging from board games (see *e.g.* the *Settlers of Catan* for the attribution of initial positions) to draft mechanisms in sport and course allocation, *e.g.* at the Harvard Business School (Budish and Cantillon, 2012).

Historically, the formal study of this protocol seems to date back to Kohler and Chandrasekaran, 1971, in its simplest version, namely, the strict alternation for two agents ($a_1a_2a_1a_2\dots$). Later, Brams and Taylor, 2000 extended this study by also paying attention to balanced alternation of agents ($a_1a_2a_2a_1\dots$). Picking sequences have also been considered by Brams and King, 2005 in a situation where the agents have ordinal preferences. They make some interesting parallel between picking sequences (what they call *products of sincere choices*) and Pareto-efficiency, on which we will come back in Section III.6. Budish and Cantillon, 2012 also studied a randomized variant of this model, that they applied to the problem of allocating courses to students. They also showed that not only this protocol is manipulable in theory, but that the students also manipulate in practice. To the best of our knowledge, our paper (Bouveret and Lang, 2011) was the first to propose a complete formalization of picking sequences focusing on the problem of finding the fairest sequence of agents.

2 The fairest sequence problem

Imagine you are in the position of the central authority, whose role in this chapter, is just to choose and propose a sequence to the agents. Which sequence would you choose? Obviously, you cannot ask the agents to explicitly reveal their preferences and base your choice on them – otherwise, using a picking sequence would be essentially useless. The approach we propose in this chapter is the following:

1. we assume that the actual agents' preferences are unknown, but are assumed to be drawn according to some fixed probabilistic distribution;
2. for each sequence, we can compute an expected utility based on this probabilistic distribution;
3. the best sequence is one that maximizes this expected utility.

Utilities as scoring functions The crucial parameter here is the probabilistic distribution that the agents' preferences are supposed to follow, and that will depend on the problem at stake. To limit the range of possible weight functions w , we restrict them to those that are generated by a *scoring function*:

Definition 15 A scoring function over $\llbracket 1, m \rrbracket$ is a decreasing function² g from $\llbracket 1, m \rrbracket$ to \mathbb{N} : if $i > j$ then $g(i) < g(j)$.

²Here, contrary to the classical definition of scoring functions, we assume that the function is strictly decreasing to avoid having to deal with non-deterministic executions of picking sequences.

Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. w is said to have been generated by a scoring function g if there exists a profile $(\succ_1, \dots, \succ_n)$ of linear orders over \mathcal{O} such that for all a_i and o_j , $w_{i,j} = g(\text{rank}_{\succ_i}(o_j))$, where $\text{rank}_{\succ}(o)$ denotes the rank of object o in relation \succ .

In other words, the decision maker assumes that the agents have ordinal preferences in mind (which is enough for them to know which items to pick), and that these ordinal preferences are converted into utilities using the same scoring function for all of them. We focus on three prototypical functions:

- *Borda*: for any k , $g_B(k) = m - k + 1$;
- *lexicographic*: for any k , $g_L(k) = 2^{m-k}$;
- *quasi-indifferent (QI)*: for any k , $g_{QI}(k) = 1 + \varepsilon \times (m - k)$.

While the choice of Borda is quite standard in collective decision making (although initially tailored for voting situations), the lexicographic scoring function is adapted to contexts where there is a huge difference between items. On the opposite, the QI scoring function suits the case where all items almost have the same value, and where what matters most is the number of objects received.

Uncertainty over profiles Since we assume that the agents' preferences are generated by a common scoring function, we only have to define a probability distribution Ψ over the possible profiles $R = (\succ_1, \dots, \succ_n)$ to come up with a probability distribution over the utilities. In what follows, we will mostly work with two probabilistic models.

- The *full independence (FI)* model³, that assumes that the rankings of the agents are sampled uniformly at random from the set of all possible rankings independently from one another.
- The *full correlation (FC)* model, that assumes that all the agents have exactly the same ranking R (that is itself drawn uniformly at random from the set of all possible permutations).

In Section III.5, we will also work with a third model that we introduce here for simplicity: the model FI-FC $_{\lambda}$ assumes that the rankings are sampled according to FI with probability λ and FC with probability $1 - \lambda$, with $\lambda \in [0, 1]$ being a parameter that allows to adjust the level of correlation between the preferences.

Expected social welfare Now that we have defined the set of probabilistic distributions we use as the basis of our model, we can evaluate the quality of a given sequence in terms of expected fairness. For that, we will use an approach based on collective utility functions, by focusing on the two classical ones: the egalitarian and the utilitarian social welfare functions.

More precisely, let \mathcal{A} is a set of agents and \mathcal{O} a set of objects. Let g be a scoring function and Ψ be a probabilistic distribution over profiles (FI, FC, or FI-FC $_{\lambda}$). Finally, let $\vec{\sigma}$ be a picking sequence. We will denote by $EU_{g, \Psi}^{\vec{\sigma}}(a)$ the expected utility obtained by agent a , given scoring function g , in the allocation generated by $\vec{\sigma}$ if the utility profile is drawn according to Ψ . Finally, the expected social welfare will be defined by aggregating the expected individual utilities according to the classical utility functions:

³Also known as the *impartial culture assumption* in the social choice literature.

- Expected utilitarian SW: $SW_{g,\Psi}^{ut}(\vec{\sigma}) = \sum_{a_i \in \mathcal{A}} EU_{g,\Psi}^{\vec{\sigma}}(a_i)$;
- Expected egalitarian SW: $SW_{g,\Psi}^{eg}(\vec{\sigma}) = \min_{a_i \in \mathcal{A}} EU_{g,\Psi}^{\vec{\sigma}}(a_i)$.

In Section III.5, we will also investigate the Expected Nash SW $SW_{g,\Psi}^{Na}(\vec{\sigma})$, that will simply be defined as the product of expected individual utilities.

The fairest sequence problem We are now in position of putting things together to give a formal definition of the fairest sequence problem.

OPTSEQ- g - Ψ - sw

Given: A set of agents \mathcal{A} , a set of objects \mathcal{O} , a scoring function g , a probability distribution over profiles Ψ , a collective utility function sw

Find: A sequence $\vec{\sigma}$ maximizing $SW_{g,\Psi}^{sw}(\vec{\sigma})$.

3 Computing optimal sequences

Now that we have properly formalized our initial problem of finding the fairest sequence, we will give some insight on how optimal sequences can be computed, and whether there is hope that such optimal sequences can actually be computed in reasonable time. Since we have 2 possibilities for Ψ , 3 possibilities for g and 2 for sw , that makes 12 combinations to study. We will start with the easiest case of full correlation.

Full correlation Under full correlation, all the agents have the exact same utility function. Hence, a first easy observation is that no matter what agent will pick at a given turn, the object taken will always be the same. That simplifies a lot the problem of finding an optimal sequence, because at turn k , the object picked will always be the k^{th} best object. This has a direct implication for utilitarian social welfare: since each increase in individual utility provides exactly the same contribution to the collective utility no matter what the agent is, we have the following proposition:

Proposition 11 *Under full correlation, every sequence has the same utilitarian social welfare, which is $\sum_{k \in \llbracket 1, m \rrbracket} g(k)$. Therefore, every sequence is optimal.*

Things are different for egalitarianism, because here, we have to partition the picking rounds so that the agents utilities are almost balanced. This comes down to partition the set of values $g(\llbracket 1, m \rrbracket)$ in n balanced subsets. In other words, this is a generalization of the PARTITION problem.

Proposition 12 *The problem OPTSEQ- g -FC-eg is NP-complete. However, it can be solved in polynomial time if $g(1)$ is polynomially bounded.*

The NP-hardness part of the proposition comes from a simple reduction from the PARTITION problem. The second part of the proposition can be deduced by observing that the problem can be solved by a dynamic programming algorithm where we have to fill a table of $m \times K$, where K is the number of

possible utility values for an agent. If $g(1)$ is polynomially bounded, then so is K . As a corollary, we can observe that the problem can be solved in polynomial time for the Borda scoring function.

Now let us consider the QI case. Here, since $g(1)$ is polynomially bounded, by Proposition 12, this case can be solved in polynomial time by dynamic programming. However, we can also simplify the problem by observing that the optimal sequences have a very special form. Since what matters first is the number of items each agent receives, then each agent should receive either $\lfloor \frac{m}{n} \rfloor$ or $\lfloor \frac{m}{n} \rfloor + 1$ objects. Those who receive $\lfloor \frac{m}{n} \rfloor + 1$ objects will not constrain the egalitarian social welfare, so they can be put at the end of the sequence. Hence, if $q = m - n \times \lfloor \frac{m}{n} \rfloor$, w.l.o.g the optimal sequences will look like this:

$$\pi = \sigma(\underbrace{\{a_1, \dots, a_1\}}_{\lfloor m/n \rfloor}, \underbrace{\{a_2, \dots, a_2\}}_{\lfloor m/n \rfloor}, \dots, \underbrace{\{a_{n-q}, \dots, a_{n-q}\}}_{\lfloor m/n \rfloor}), \tau(\underbrace{\{a_{n-q+1}, \dots, a_{n-q+1}\}}_{\lfloor m/n \rfloor + 1}, \dots, \underbrace{\{a_n, \dots, a_n\}}_{\lfloor m/n \rfloor + 1}),$$

where σ and τ are two permutations.

The last q agents all receive strictly more items than any of the first $n - q$ agents. So once again, they will not constrain the egalitarian social welfare, no matter what τ is. Now let us consider the first $n - q$ agents. Any of this agent a_i will have utility $\lfloor \frac{m}{n} \rfloor + x_i \times \varepsilon$, x_i depending on a_i 's picking turns (that is, depending on permutation σ). Hence, to find the best sequence we have to find the one that balances the x_i s as most as possible. This exactly comes down to solving the Borda case under the constraint that the sequence is balanced, that can be, again, tackled by dynamic programming on q agents only (compared to the n initial agents).

Now let us finish this review of the full correlation and egalitarian setting by evoking the lexicographic scoring function. Here, the problem is easy to solve. Since what matters first for an agent is to have the earliest pick in the sequence, the agent whose first pick comes last will be the one that determines the egalitarian social welfare. Hence, this agent should receive the highest possible number of items. In other words, the optimal sequences in this case are exactly of this form:

$$\tau(a_1)\tau(a_2) \dots \tau(a_{n-1})\tau(a_n)\tau(a_n) \dots \tau(a_n), \text{ where } \tau \text{ is a permutation of } \mathcal{A}.$$

Full independence The FI case seems to be much more intricate than the FC case. When we published our first paper on the topic (Bouveret and Lang, 2011), we even suspected that simply computing the expected utility for a given sequence was NP-complete. It turned out that this was not the case: it was later proved by Kalinowski, Narodytska, and Walsh, 2013 that this problem is actually polynomial. They also characterized the optimal sequence for Borda and utilitarian social welfare:

Proposition 13 (Kalinowski, Narodytska, and Walsh, 2013) *Under full independence, the alternating sequence $(a_1 a_2 a_1 a_2 a_1 a_2 a_1 a_2 \dots)$ is optimal for the Borda scoring function and utilitarian social welfare.*

However, to the best of our knowledge, beyond this particular case, the precise complexity of computing the best sequence for the full independence case is still an open problem:

Open question 4 *What is the complexity of OPTSEQ- g -FI-sw, when $(g, sw) \neq (g_B, ut)$?*

Now let us end this full review of the problem by giving some examples of optimal sequences for the probably most interesting setting in practice, namely: FI, Borda scoring function, egalitarian CUF. These sequences have been computed by a brute-force algorithm, and are presented in Table III.1.

m	$n = 2$	$n = 3$
4	$a_1 a_2 a_2 a_1$	$a_1 a_2 a_3 a_3$
5	$a_1 a_1 a_2 a_2 a_2$	$a_1 a_2 a_3 a_3 a_2$
6	$a_1 a_2 a_1 a_2 a_2 a_1$	$a_1 a_2 a_3 a_3 a_2 a_1$
8	$a_1 a_2 a_2 a_1 a_2 a_1 a_1 a_2$	$a_1 a_1 a_3 a_3 a_2 a_2 a_3 a_2$
10	$a_1 a_2 a_2 a_1 a_1 a_2 a_1 a_2 a_2 a_1$	$a_1 a_2 a_3 a_1 a_2 a_2 a_3 a_1 a_3 a_3$

Table III.1 — Optimal sequences for small numbers of agents and objects under full independence, Borda scoring function, egalitarian CUF.

4 Manipulating picking sequences

One desirable property of allocation protocols that we did not discuss yet is *strategyproofness*. Informally, a protocol is said to be strategyproof (or incentive-compatible) if the dominant strategy for an agent is to act consistently with their true preferences. In a non-strategyproof allocation protocol, some agents could have an incentive to lie to obtain a better bundle.

Is the picking sequence protocol strategyproof? We can answer negatively by exhibiting a very simple example. Suppose that we have 4 objects and 2 agents having the following preferences:

- $a_1: o_1 > o_2 > o_3 > o_4$;
- $a_2: o_2 > o_3 > o_4 > o_1$.

Now assume that the sequence chosen by the central authority is $\vec{\sigma} = a_1 a_2 a_2 a_1$. Supposing the agents tell the truth, this sequence generates allocation $(\{o_1 o_4\}, \{o_2 o_3\})$.

But what if a_1 knows a_2 's preferences and acts maliciously? They can manipulate by picking o_2 instead of o_1 at first step. They will then obtain $\{o_1 o_2\}$ instead of $\{o_1 o_4\}$, which is obviously much better (no matter what the scoring function is).

The bad news is that picking sequences are always manipulable⁴, *unless they are non-interleaving*. Here, by non-interleaving, we mean sequences that are made of contiguous blocks of agents (for instance $a_1 \dots a_1 a_2 \dots a_2 a_3 \dots a_3$). The study of such sequences will be the topic of Section III.5.

Since almost all picking sequences are manipulable, the question we will ask here is to which extent it is easy for an agent to compute such a manipulation. More formally, given a set of objects, a set of agents and a policy, we want to know whether it is easy for a cheating agent to build an optimal picking strategy (that is, to obtain the best bundle they can obtain). Here, we suppose that only one agent manipulates, and that this agent perfectly knows (i) the sequence, (ii) the other agents' picking strategies (that is, what each agent will pick in any situation). Obviously, the only possible cheating action for the manipulative agent is to choose at a given step *not* to pick their preferred object.

Complexity of the manipulation problem The first observation we can make is that when there are only two agents – a manipulating agent, say a_1 and a sincere one, say a_2 –, a_1 can easily figure out

⁴And some practical observations (Budish and Cantillon, 2012) tend to show that manipulation is not only theoretical, and that real users exhibit strategic behaviours.

whether a given bundle π is accessible to them or not. Moreover, if this bundle is accessible, a_1 can build in polynomial time a strategy to get it. The idea is simple: since a_1 exactly knows in which order a_2 will pick the objects, then a_1 should pick π 's objects exactly in that order. If at some point an object is not available anymore, then it means that no strategy could possibly guarantee π to a_1 .

The interesting thing is that this observation generalizes to any number of agents. Assuming a_1 is still the manipulating agent, all they have to do is to merge the other agents $a_2 \dots a_n$ into a single (fake) agent (by also merging their picking strategies) and act like in the two agents case.

For instance, let us consider the following 3 agents, 6 items setting:

- $a_1: o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5 \succ o_6$;
- $a_2: o_2 \succ o_3 \succ o_4 \succ o_5 \succ o_6 \succ o_1$;
- $a_3: o_3 \succ o_4 \succ o_2 \succ o_5 \succ o_1 \succ o_6$.

Suppose that the sequence chosen by the central authority is $\vec{\sigma} = a_1 a_2 a_3 a_3 a_2 a_1$. Is there a way for agent a_1 (the cheating agent) to get for sure bundle $\{o_1, o_2\}$? To figure out, we will transform agents a_2 and a_3 into a single (fake) agent a_f by applying the following procedure (see [Bouveret and Lang, 2011](#), Algorithm 2):

- The first adversarial picker is a_2 (second round), and they will pick o_2 first, if it is still available.
- Then, a_3 comes twice in a row, and they will pick o_3 and o_4 .
- Then comes a_2 again, that will pick o_5 (the first item still available).
- After that round, no adversarial agent picks. Hence, a_f has the following picking strategy: $o_2 \succ o_3 \succ o_4 \succ o_5$, that can be completed arbitrarily with o_1 and o_6 .

Now what will be the strategy for a_1 to get their target bundle $\{o_1, o_2\}$? Their strategy is simply to pick those two objects in the exact same order of a_f 's preference order, namely: pick o_2 first, and o_1 second. Here, as we can see, a_1 will pick o_2 in the first round, then a_2 will pick o_3 , a_3 will pick o_4 and o_5 , a_2 will pick o_6 and o_1 will still be available for a_1 , which means that their target bundle was indeed achievable.

Now let us turn to our initial manipulation problem, where the cheating agent tries to find an optimal manipulation strategy:

Proposition 14 *The optimal manipulation problem is NP-complete if $n \geq 3$. However, if $n = 2$, the optimal manipulation strategy can be computed in polynomial time.*

For two agents, the proof relies on a greedy construction of the best achievable bundle: first a_1 (the cheating agent) tries to determine what is the best object o_i such that $\{o_i\}$ is achievable. Then they try to find the best object o_j such that $\{o_i, o_j\}$ is achievable, and so on. Since determining whether a bundle is achievable can be done in polynomial time (see above), the best bundle can also be constructed in polynomial time.

However, this construction does not work for more than 2 agents (and we cannot use the conversion trick used above to merge the sincere agents into a single fake one), and a reduction from 3-SAT shows

that the problem is NP-complete. Note that this is probably good news because it shows that picking sequences are somewhat resistant to manipulation.⁵

More than one manipulator So far, we have considered that only agent a_1 manipulated (and that they knew all the other agents' preferences, which is a standard assumption in the study of manipulation in social choice – see *e.g.* Zwicker, 2015). However, there can be situations where a single agent cannot obtain any additional utility by manipulating alone, but could obtain better by cooperating with another agent. Let us consider for instance the following setting with 6 objects and 3 agents having the following preferences:

- $a_1: o_1 > o_2 > o_5 > o_4 > o_3 > o_6$
- $a_2: o_1 > o_3 > o_5 > o_2 > o_4 > o_6$
- $a_3: o_2 > o_3 > o_4 > o_1 > o_5 > o_6$

Now assume that the sequence chosen by the central authority is $\vec{\sigma} = a_1 a_2 a_3 a_1 a_2 a_3$. Supposing the agents tell the truth, this sequence generates allocation $(\{o_1 o_5\}, \{o_3 o_4\}, \{o_2 o_6\})$. It can be easily observed that if A and B manipulate alone (independently of each other), they cannot obtain better. However, if they cooperate and agree to leave o_1 aside during their first picking stage, they can obtain $(\{o_1 o_2\}, \{o_3 o_5\}, \{o_4 o_6\})$, which is strictly better for both of them.

This setting is called *coalitional manipulation*, and is pretty standard in voting theory. However, a notable difference between voting and fair division is that in the first case, the outcome is the same for everyone, whereas in fair division, the agents receive different shares and can trade their bundles after the allocation. In our work (Bouveret and Lang, 2014), we have considered three situations.

1. No post-allocation trade allowed: after the execution of the sequence, each agent leaves the game with the bundle they obtained, and no one can swap any good or bundle with each other.
2. Post-allocation exchange of goods allowed between the manipulators: after the execution of the sequence, the cheating agents can exchange goods with each other, provided that these exchanges are mutually beneficial.
3. Post-allocation exchange of goods and side-payments allowed between the manipulators: basically the same as above, but here, a manipulating agent can accept a non beneficial trade, provided that a side-payment is given to them by the other agent.

It turns out that the problem is NP-complete for the first two cases. In the last case, since every kind of post-allocation is allowed, it means that (i) in the final allocation (after the exchanges), each object will be allocated to the agent of the coalition that values it the most, and (ii) the optimal joint picking strategy is the one that maximizes the total utilitarian social welfare of the coalition. As a consequence, solving this case comes down to solving the case where there is a single manipulator, which has been proved to be NP-complete in the general case, but in P for the case where only one agent is sincere.⁶

⁵This is however not the case anymore for restricted cases like 0–1 preferences, where the optimal manipulation problem can be solved in polynomial time (see Aziz, Bouveret, Lang, et al., 2017, for details).

⁶Note that in our initial paper (Bouveret and Lang, 2014), we claim that this problem is in P in the general case. This claim was later proved incorrect.

Everyone manipulates Let us finish this review of manipulation issues in picking sequences by giving some insights on the case where all agents are selfish manipulators and try to maximize their own utilities independently of each other. This is typically a game-theoretic setting. Kohler and Chandrasekaran, 1971 were the first to study questions related to strategic behaviour in picking sequences by giving a precise characterization of the *subgame perfect Nash equilibrium*⁷ (SPNE) for two agents and strict alternations. Namely, they proved that in this restricted case, the result of every SPNE is exactly the same as the allocation generated by a setting where the picking sequence is reversed, the agents' preferences as well, and the agents pick sincerely. Later, Kalinowski, Narodytska, Walsh, and Xia, 2013 further extended this result to any two-agent sequence, and also proved that for an unbounded number of agents, the computation of an SPNE is PSPACE-hard.

5 Of strategyproof picking sequences

In Section III.4, we investigated the complexity of manipulating picking sequences, with in mind the idea (common in voting theory) that a high complexity could be a computational barrier to manipulation. In this section, we adopt another point of view by proposing to restrict picking sequences to those that are strategyproof. As we have seen earlier, the only such sequences are the *non-interleaving ones*, namely, those that are made of contiguous blocks of agents (for instance $a_1 \dots a_1 a_2 \dots a_2 a_3 \dots a_3$).

The fairest non-interleaving picking sequence problem Since these sequences are strategyproof, there is no point studying manipulation issues in this context. However, the question studied in Sections III.2 and III.3 of finding the *fairest* such sequence is still relevant here. Can we still ensure fairness in such a highly-constrained setting where agents can only appear as contiguous blocks in the sequence?

It turns out that we can indeed ensure a kind of fairness, just by playing on the number of items that each agent receives (which is the only possibility we have in this context). For instance, in a setting with 3 agents and 10 objects, we feel that the sequence $a_1 a_1 | a_2 a_2 a_2 | a_3 a_3 a_3 a_3 a_3$ is fairer than $a_1 a_1 a_1 a_1 | a_2 a_2 a_2 | a_3 a_3$ (we add vertical bars for the sake of readability), because in the first case, the fact that a_3 gets the last turn is compensated by the fact that they receive more objects than a_1 and a_2 .

Here, the problem is much simpler to formalize than the general case. We can assume w.l.o.g that non-interleaving sequences always have the following form:⁸

$$\underbrace{a_1 \dots a_1}_{k_1} \underbrace{a_2 \dots a_2}_{k_2} \dots \underbrace{a_n \dots a_n}_{k_n}$$

Hence, a sequence is simply defined by a vector (k_1, \dots, k_n) such that $\sum_{i=1}^n k_i = m$, where k_i is the number of picking turns in a row allocated to agent a_i . The question is then to find the fairest such vector.

A dynamic programming algorithm The first observation we can make is that while there are dramatically fewer non-interleaving sequences than the n^m unconstrained ones, there are still $\binom{n-1}{n+m-1}$ such

⁷Since picking sequences look more like a sequential game than a one-shot game, this equilibrium concept is more suited to this context than the regular Nash equilibrium concept.

⁸Also note that since we do not know the actual agents' preferences, and that we only consider expected utilities, we can assume w.l.o.g that agent a_1 comes first in the sequence, followed by agent a_2 , and so on...

sequences, which is still much too many to find the fairest one by a brute-force approach. However, rather unexpectedly, it turns out that the OPTSEQ problem can be solved in polynomial time.

Proposition 15 *Restricted to non-interleaving sequences, the OPTSEQ-g-FI-FC $_{\lambda}$ -sw problem can be solved in time $O(m^2 \max(n, m))$ for any $\lambda \in [0, 1]$, any $sw \in \{ut, eg, Na\}$, and any g .*

In the specific case where $\lambda = 0$ (full correlation), the problem can be solved in time $O(m^2 n)$ for $sw \in \{eg, Na\}$ and is trivial for $sw = ut$ since in this case all vectors lead to the same social welfare.

We will now give some intuition of how the dynamic programming algorithm that solves these problems works (see [Bouveret, Gilbert, et al., 2023](#), for a formal presentation of the result). The idea is the following.

Let us consider a_1 , who is the first picker. We must choose how many items they pick. If they pick k_1 items, they will receive utility v_1 and there will be $m - k_1$ remaining items. We must choose the value of k_1 that maximizes the aggregation (+, min or \times , depending on the CUF chosen) of v_1 and the best expected social welfare that we can obtain from the remaining $(n - 1)$ agents and $(m - k_1)$ items. Here, v_1 only depends on k_1 and is exactly equal to $\sum_{i=1}^{k_1} g(i)$ (the utility of the best k_1 objects).

Now consider any a_i , the i^{th} picker. Here, we must choose how many items they pick among the $m - \sum_{j=1}^{i-1} k_j$ remaining items. If they pick k_i items, they will receive utility v_i and there will be $m - \sum_{j=1}^i k_j$ remaining items. We must choose the value of k_i that maximizes the aggregation of v_i and the best expected social welfare that we can obtain from the remaining $(n - i)$ agents and $(m - \sum_{j=1}^i k_j)$ items. Here, obviously, v_i depends on k_i but also on the items received before. The success of our dynamic programming approach depends on whether we are able to compute v_i efficiently. It turns out that this is the case:

Proposition 16 *For any $\Psi = FI - FC_{\lambda}$ and for any agent a_i , v_i only depends on:*

- *the number k_i of items received by a_i ;*
- *the number t of items picked before;*

and can be computed in time $O(m^3)$.

Note that in the FC case, the computation of v_i is even simpler (since a_i simply picks the items between ranks $\sum_{j=1}^{i-1} k_j$ and $\sum_{j=1}^i k_j$). In the FI case, we can once again use dynamic programming to compute v_i , as was already noticed by [Kalinowski, Narodytska, and Walsh, 2013](#). This algorithm can be generalized to the case of $FI - FC_{\lambda}$ for any $\lambda \in [0, 1]$ by using linearity of expectation.

Putting things together, we just give below the recursive formula that can be used to compute the optimal expected social welfare. Let $SW_{g, \Psi}^{sw}(i, \ell)$ be the best expected utility we can obtain for agents a_i, \dots, a_n if ℓ items have already been picked. This expected utility can be computed as follows:

$$\begin{aligned} SW_{g, \Psi}^{sw}(i, \ell) &= \max_{k \in \llbracket 0, m - \ell \rrbracket} F(v(k, \ell), SW_{g, \Psi}^{sw}(i + 1, \ell + k)), \\ SW_{g, \Psi}^{sw}(n, \ell) &= v(m - \ell, \ell), \end{aligned}$$

where $F = +$ (resp. \min, \times) if $sw = ut$ (resp. eg, Na).

Table III.2 shows some examples of optimal sequences for various values of n and m , and for the three social welfare functions considered, under the Borda scoring vector. In the same vein, Figure III.1 plots the optimal sequences for $n = 5$ and m varying from 5 to 300, also for the three social welfare functions and under the Borda scoring vector. A first observation is that for the egalitarian social welfare, the expected utilities of the agents tend to equalize when the number of items increases, which confirms that our sequences can indeed ensure a form of ex-ante fairness. This is ensured by the fact that, for egalitarian social welfare (ESW), agents coming later in the sequence pick more items than agents coming earlier, which can also be observed on the graphs and on the examples of Table III.2.⁹ Interestingly, these examples also show that it is not the case anymore for utilitarian social welfare (USW), where the vector (k_1, \dots, k_n) is not necessarily increasing, simply because here, the compensation for later picking does not make much sense anymore.

n	m	$sw = eg$	$sw = ut$	$sw = Na$
2	10	(4, 6)	(5, 5)	(4, 6)
2	20	(8, 12)	(10, 10)	(8, 12)
3	10	(3, 3, 4)	(4, 3, 3)	(3, 3, 4)
3	35	(9, 10, 16)	(13, 11, 11)	(10, 11, 14)
5	10	(2, 2, 2, 2, 2)	(2, 2, 2, 2, 2)	(2, 2, 2, 2, 2)
5	70	(12, 12, 12, 13, 21)	(18, 16, 14, 11, 11)	(13, 13, 13, 13, 18)
8	20	(2, 2, 2, 2, 2, 3, 3, 4)	(3, 3, 3, 3, 2, 2, 2, 2)	(2, 2, 2, 2, 3, 3, 3, 3)
8	100	(11, 11, 11, 11, 11, 12, 13, 20)	(18, 16, 15, 13, 12, 10, 8, 8)	(12, 12, 12, 12, 12, 12, 12, 16)

Table III.2 — Examples of optimal non-interleaving sequences for various n , m and sw and Borda scoring rule.

Discussion: of strategyproof allocation mechanisms We have investigated in this section non-interleaving picking sequences. The reason was that these sequences are the only ones to be strategyproof. However, it turns out that this nice property goes far beyond the simple case of picking sequences. In fact, various characterization theorems state that, under mild conditions (see *e.g.* Pápai, 2001), all strategyproof allocation mechanisms *in general* (*i.e.* not restricted to picking sequences) are within the family of *serial dictatorships*. A standard serial dictatorship (SSD) is a protocol that works as follows: a central arbitrator chooses a permutation of agents. At their turn in the permutation, the designated agent picks *all the items they like* among the remaining ones.¹⁰ A *constrained serial dictatorship (CSD)*, also called quota serial dictatorship, is a variant, where at each step, the designated agent chooses a predefined number of items (Pápai, 2000).

Clearly, SSD is unacceptable in terms of fairness. CSD does better, but this is at the price of Pareto-efficiency. The theoretical results seem to show that we cannot have all properties satisfied at the same time. Either we should sacrifice Pareto-efficiency (and a bit of fairness), or strategyproofness. From this point of view, Constrained Serial Dictatorships seem to be a reasonable compromise between those properties.

⁹We can even prove this observation formally (Bouveret, Gilbert, et al., 2023).

¹⁰Clearly, if a given agent has a strictly positive utility for every item, then they will pick them all and leave nothing for the other agents.

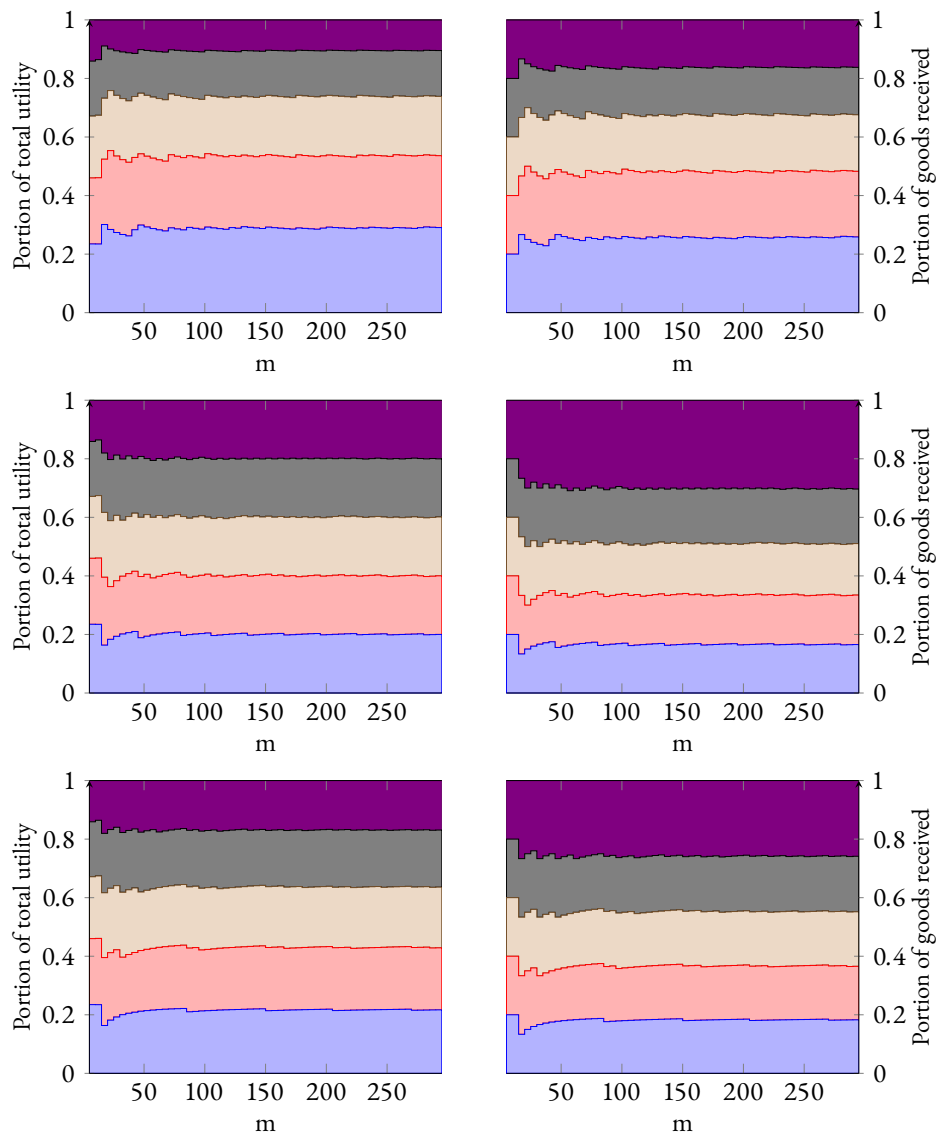


Figure III.1 — Portion of total utility (plots on the left) and of goods (right) received by each of 5 agents with m increasing from 5 to 300 in steps of 5. Maximizing USW (plots at the top), NSW (bottom), or ESW (middle), using Borda scoring vector and FI. – Adapted from a work by Guillaume M erou e (Bouveret, Gilbert, et al., 2023)

Now let us go back to picking sequences. The careful reader might have observed that CSDs exactly correspond to what we call non-interleaving picking sequences. Hence, if strategyproofness is a property that we want to guarantee, that leaves us no choice but to use non-interleaving sequences. Unfortunately, [Amanatidis, Birmpas, Christodoulou, et al., 2017](#) have shown that this family of protocols is incompatible with fairness (for several notions of fairness). In other words, it is essentially impossible to find an allocation protocol that both guarantees strategyproofness and fairness, which seems contradictory with our claim in this section. So how do we reconcile fairness and strategyproofness? The idea here is that we simply consider *ex-ante* fairness, while impossibility results concern *ex-post* fairness. The allocations resulting from the application of our picking sequences might actually be very unfair in the end. What we only guarantee is that the protocol is fair *in expectation*. This is the way we escape the impossibility results and somewhat reconcile fairness and strategyproofness.

6 Picking sequences, swap deals and efficiency

In this chapter, we have tried to show (hopefully in a convincing way) that picking sequences is a natural and appealing protocol to allocate indivisible objects to agents. We will complete this praise of picking sequences by showing unexpected connections with negotiation-based distributed fair division protocols on the one hand, and efficiency properties on the other hand.

6.1 Sequenceability as an efficiency criterion

Characterization of sequenceable allocations In Section III.1 we have formally defined the picking sequence protocol by giving a definition of the execution of a sequence. As we have seen, each sequence $\vec{\sigma}$ generates at least one allocation, and several ones (potentially even an exponential number) if the preferences are not strict on objects. Conversely, obviously, not every allocation is sequenceable, as the following simple example shows:

Example 1 Consider the add-MARA instance defined by the following weight matrix:

$$w = \begin{bmatrix} 2 & \textcircled{1} \\ \textcircled{1} & 2 \end{bmatrix}.$$

The (circled) allocation $\vec{\pi} = \langle \{o_2\}, \{o_1\} \rangle$ is not sequenceable.

In Example 1, $\vec{\pi}$ is not sequenceable, because neither a_1 nor a_2 receives their top object, whereas in any sequence, the agent that chooses first for sure picks their most preferred item. Such an allocation, where no agent receives their top object, will be called *frustrating*. Does this notion of frustrating allocation precisely characterize the set of sequenceable allocations? Not exactly, as the following example shows:

Example 2 Consider the add-MARA instance defined by the following weight matrix:

$$w = \begin{bmatrix} \textcircled{9} & 8 & 2 & \textcircled{1} \\ 2 & \textcircled{5} & \textcircled{1} & 4 \end{bmatrix}.$$

In allocation $\vec{\pi} = \langle \{o_1, o_4\}, \{o_2, o_3\} \rangle$, each agent receives their top object. However, after o_1 and o_2 have been allocated (they must be allocated first by all sequences generating $\vec{\pi}$), the dotted sub-allocation

remains. This sub-allocation is obviously non-sequenceable because it is frustrating. Hence $\vec{\pi}$ is not sequenceable either.

It turns out that this concept of frustrating sub-allocation provides a nice characterization of sequenceable allocations:

Proposition 17 *Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance and let $\vec{\pi}$ be an allocation. $\vec{\pi}$ is sequenceable if and only if no sub-allocation of $\vec{\pi}$ is frustrating, that is, in for any nonempty subsets of agents $\mathcal{A}' \subset \mathcal{A}$ and objects $\mathcal{O}' \subset \mathcal{O}$, at least one agent receives their top object in the restriction of the instance to \mathcal{A}' and \mathcal{O}' .*

The nice thing is that as a corollary of this result, we can obtain an algorithm that decides in time $O(n \times m^2)$ if an allocation $\vec{\pi}$ is sequenceable or not.

Also note that if the agents have same-order preferences, then obviously no frustrating sub-allocation can exist in any allocation $\vec{\pi}$. Hence in that case, all the allocations are sequenceable.¹¹

Of sequenceability and Pareto-efficiency This characterization result is nice, but we might wonder how useful it is to know whether an allocation is sequenceable or not. Our argument here is that sequenceability conveys a form of "local" efficiency: indeed, at each turn of the execution of a sequence the choice of an object is "locally" optimal, since it corresponds to the allocation of the best object to the agent at stake. To which extent is this form of efficiency compatible with the classical notion of Pareto-efficiency? It turns out that sequenceability can be seen as a weak form of Pareto-efficiency:

Proposition 18 *Every Pareto-optimal allocation is sequenceable (but the converse is not true).*

Actually, this result was already partially known in a weaker form. In a seminal work, [Brams and King, 2005](#) prove the equivalence between sequenceability (there referred to as product of sincere choices) and Pareto-efficiency. The main difference with our setting is that their notion of Pareto-optimality is much weaker than ours. In their setting, the agents' preferences are given as linear orders over objects, which is then lifted to preferences over bundles using the *responsive set extension* (see [Barberà et al., 2004](#), for an extensive discussion about preference lifting). This extension leaves many bundles incomparable. The notion of Pareto-efficiency used by [Brams and King, 2005](#) corresponds to what could be called *possible Pareto-efficiency*, which is indeed, in this setting, equivalent to sequenceability.

These various forms of Pareto-efficiency and their link with sequenceability have been extensively discussed in parallel and independently in our paper ([Bouveret and Lemaître, 2016](#)) and by [Aziz, Kalinowski, et al., 2016](#).

6.2 Cycle-deals optimality

We have seen earlier that picking sequences, which is an essentially distributed allocation protocol, can to some extent convey some flavour of efficiency. Now let us consider another kind of distributed protocol: improving deals.

As reminded in the introduction of this chapter, the idea is the following. We start from an initial allocation, that can be essentially random. Then, we let the agents negotiate with each other to exchange

¹¹ Actually, we have proved (see [Beynier et al., 2019](#), Proposition 6) that the opposite is also true: if an add-MARA is such that all the allocations are sequenceable, then it means that the agents have same-order preferences.

bundles of objects. The interesting thing is that if all the exchanges are mutually improving (that is, no agent loses utility in the trade), and if we do not impose any restriction on the kind of deals allowed (a deal can involve any number of agents and goods at the same time), then the protocol will eventually reach a Pareto-optimal allocation (Sandholm, 1998).

This result shows that we can still guarantee a strong form of efficiency even without central intervention of any kind. However, this holds only if we allow all forms of exchanges. What if we restrict ourselves to specific kinds of deals? Can we still guarantee some weaker form of efficiency?

It turns out that it is indeed the case for a special and very natural kind of deals: cycle-deals. Here, if $\langle \mathcal{A}, \mathcal{O}, w \rangle$ is an add-MARA instance and $\vec{\pi}$ an allocation, a N -cycle-deal applied to $\vec{\pi}$ is a sequence

$$\mu_1 \xrightarrow{\omega_1} \mu_2 \dots \mu_{N-1} \xrightarrow{\omega_{N-1}} \mu_N \xrightarrow{\omega_N} \mu_1,$$

such that for all i , μ_i is an agent and ω_i is an object of π_i . This sequence means that μ_1 will give object ω_1 to μ_2 , who will herself give object ω_2 to μ_3 , and so on until μ_N , who will receive object ω_{N-1} from μ_{N-1} and give back ω_N to μ_1 (so as to complete the cycle).¹²

Among the set of possible cycle-deals, some of them are obviously more interesting, namely, those that are (weakly) *improving*, that is, that are such that each agent involved is better-off afterwards (with one of them at least being strictly better-off). Since we have an underlying notion of improving allocations, we can naturally define an associated notion of optimality: here, an allocation will be said *N -cycle optimal* if it cannot be improved by any N -cycle deal.

Note that this notion of improving cycle deals is very close to similar notions like *trading cycles* that have been introduced and used outside the context of distributed fair division, for instance by Varian, 1974 and Lipton et al., 2004 in the context of envy-freeness.

6.3 A scale of efficiency

An easy observation is that N -cycle-optimality implies $(N - 1)$ -cycle optimality (but the converse is not true). Cycle optimality hence forms a chain of properties of decreasing strength, from n -cycle optimality to 2-cycle optimality (also called *swap-optimality*).

What is especially interesting and somewhat unexpected is that cycle-deals actually have a strong link with sequenceability:

Proposition 19 *Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance and $\vec{\pi}$ an allocation. $\vec{\pi}$ is sequenceable if and only if it is n -cycle-optimal.*

Proposition 19 shows that restricting the set of admissible exchanges in negotiation protocols indeed leads to weaker notions of efficiency. More interesting, this notion of efficiency perfectly fits in the scale Pareto-efficiency / sequenceability / non-efficiency that was implicitly implied by Proposition 18, and completes the landscape of efficiency notions in a similar way as fairness notions introduced in Section II.2.2. Adopting the same terminology as the scale of fairness introduced in Proposition 2:

¹²Actually, our paper (Beynier et al., 2019) is slightly more general than this and also considers the case where bundles of objects can be involved in the cycle.

Proposition 20 *Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance, and let $\vec{\pi}$ be an allocation. Then:*

$$\begin{aligned} (\vec{\pi} \models \text{Pareto-optimality}) &\Rightarrow (\vec{\pi} \models \text{Sequenceability}) \Leftrightarrow \\ &(\vec{\pi} \models n\text{-cycle optimality}) \Rightarrow (\vec{\pi} \models n-1\text{-cycle optimality}) \Rightarrow \\ &\dots (\vec{\pi} \models \text{swap optimality}). \end{aligned}$$

In other words, if \mathcal{F}_X denotes the set of feasible allocations that satisfy property X , we have:

$$\mathcal{F}_{\text{Pareto-optimality}} \subset \mathcal{F}_{\text{Sequenceability}} = \mathcal{F}_{n\text{-cycle optimality}} \subset \mathcal{F}_{n-1\text{-cycle optimality}} \subset \dots \subset \mathcal{F}_{\text{swap optimality}} \subset \mathcal{F}.$$

Moreover, there exist instances for which all these inclusions being strict.¹³

That completes the landscape of additive multiagent resource allocation problems, now made of two orthogonal scales, of fairness and efficiency. Note however that a significant difference with the scale of fairness is that finding an allocation that satisfies the highest level of efficiency is trivial here: giving all the objects to the same agent already yields a Pareto-efficient allocation. The problem will thus be to find an allocation that satisfies the highest level of fairness while also satisfying the highest level of efficiency, with the idea that trade-offs between both scales could be necessary.

6.4 Link with fairness notions

As presented above, finding a trade-off between fairness and efficiency might be necessary. Let us therefore end this landscape by giving some hints on the known links between efficiency and fairness concepts. The objective here will be to determine to which extent these concepts are compatible with each other.

Threshold-based fairness properties Let us start with an easy observation. For any fairness property X based on the individual utility being higher than a threshold (that is: max-min share, min-max share, proportionality), then: (i) $\vec{\pi} \models X$ does not entail that $\vec{\pi}$ is Pareto-efficient, and (ii) if $\exists \vec{\pi}$ such that $\vec{\pi} \models X$, then there also exists an allocation $\vec{\pi}'$ (possibly equal to $\vec{\pi}$) such that $\vec{\pi}' \models X$ and $\vec{\pi}'$ is Pareto-efficient. Hence for these three properties, as soon as we have found a suitable allocation, there is a guarantee that either it is Pareto-efficient, or we can find another one that is.¹⁴

Envy-freeness Now let us deal with envy-freeness. A well-known fact is that envy-freeness does not imply Pareto-efficiency, and Pareto-efficiency does not more imply envy-freeness. It is also known that the problem of deciding whether there exists a Pareto-efficient and envy-free allocation is Σ_2^P -complete (de Keijzer et al., 2009).

Unfortunately, envy-freeness does not even imply sequenceability (see Beynier et al., 2019, Example 4.2).

However, as we quickly mentioned in the introduction of Section II.3, there is an interesting link between envy-freeness and picking sequences, as an EF1 allocation can be obtained as the result of a particular picking sequence, the *round-robin* one, like $a_1 a_2 \dots a_n a_n a_{n-1} \dots a_1 a_1 a_2 \dots a_n \dots$ (Caragiannis, Kurokawa, et al., 2016).

¹³Except if the agents have single-peaked preferences, in which case the N -cycle optimality hierarchy collapses at the second level (see Beynier et al., 2019, Proposition 8).

¹⁴Although it might be CONP-complete to find this new allocation from the first one (Aziz, Biro, et al., 2016).

CEEI The case of Competitive Equilibrium from Equal Incomes is also interesting. We have seen in Proposition 1 of Section II.2.1 that if the preferences are strict on shares, any CEEI allocation is Pareto-efficient. We have also seen that a CEEI allocation can fail being Pareto-efficient if the preferences are not strict. It turns out that in this latter case, we can still guarantee sequenceability:

Proposition 21 *Every CEEI allocation is sequenceable.*

7 Conclusion

In this chapter, we introduced an interactive fair division protocol for indivisible goods. This protocol has arguably many appealing properties. It is mostly distributed: after having chosen a sequence (that can be made offline, even before knowing the agents and objects at stake), the role of the central authority is very light. It only relies on a minimal interaction (the communication complexity is low), as the agents do not have to communicate their entire preferences. It is quite intuitive, in the sense that it can probably be easily understood and trusted by any person even without deep mathematical knowledge. Moreover, while not being strategyproof in its general form, our results also show that (i) manipulating this protocol turns out to be computationally complex, and (ii) if strategyproofness is not negotiable, non-interleaving sequences provide this guarantee. While being inherently less fair, fairness can to some extent still be ensured by a proper distribution of the number of items allocated to the agents.

Price of elicitation-freeness Note that, as already mentioned earlier, picking sequences convey a very particular form of fairness, namely, *ex-ante* fairness (that is, fairness in expectation). A natural question is whether the allocation resulting from the allocation of a picking sequence can be arbitrary bad in terms of utility, compared to the allocation maximizing the social welfare with full preference elicitation. That led us to define the concept of Multiplicative Price of Elicitation Freeness ($MPEF_{sw}$, with $sw \in \{ut, eg, Na\}$) as the worst-case ratio in social welfare between an optimal allocation and a sequential allocation (Baumeister et al., 2017). Interestingly, for a balanced sequence like $(a_1 a_2 \dots a_n)^*$, this ratio is upper bounded both for the utilitarian SW and the egalitarian SW. Namely, both $MPEF_{ut}$ and $MPEF_{eg}$ are upper-bounded by $2 - \frac{1}{n} + \Theta(\frac{1}{m})$ when $m \rightarrow +\infty$. In practice, we also observe that this price of elicitation-freeness stays bounded, and tends to the neighbourhood of 1 when m grows (see Figure III.2).

Completing the landscape of criteria Finally, it turns out that, somewhat unexpectedly, picking sequences also convey some form of local efficiency related to what Pareto-efficiency is at the global level. That completes the landscape of desirable fairness and efficiency criteria that we have started in Chapter II. This completed landscape is shown in Figure III.3.

To finish this chapter and make a connection with the next one, we ask the question of whether such an appealing fair division protocol can be adapted to the context of voting, where centralized collective decision procedures are usually the standard solution. It turns out that it can: such a protocol can take the form of a procedure where voters *eliminate* candidates instead of picking items. Of course, this only works in situations where there are more candidates than voters, which is very limiting in the context of voting, and in particular totally rules out this kind of protocols in political elections. The study of such an elimination voting protocol (both for interleaving and non-interleaving strategyproof sequences) is the central topic of a work that resulted from a collaboration with Yann Chevaleyre, François Durand

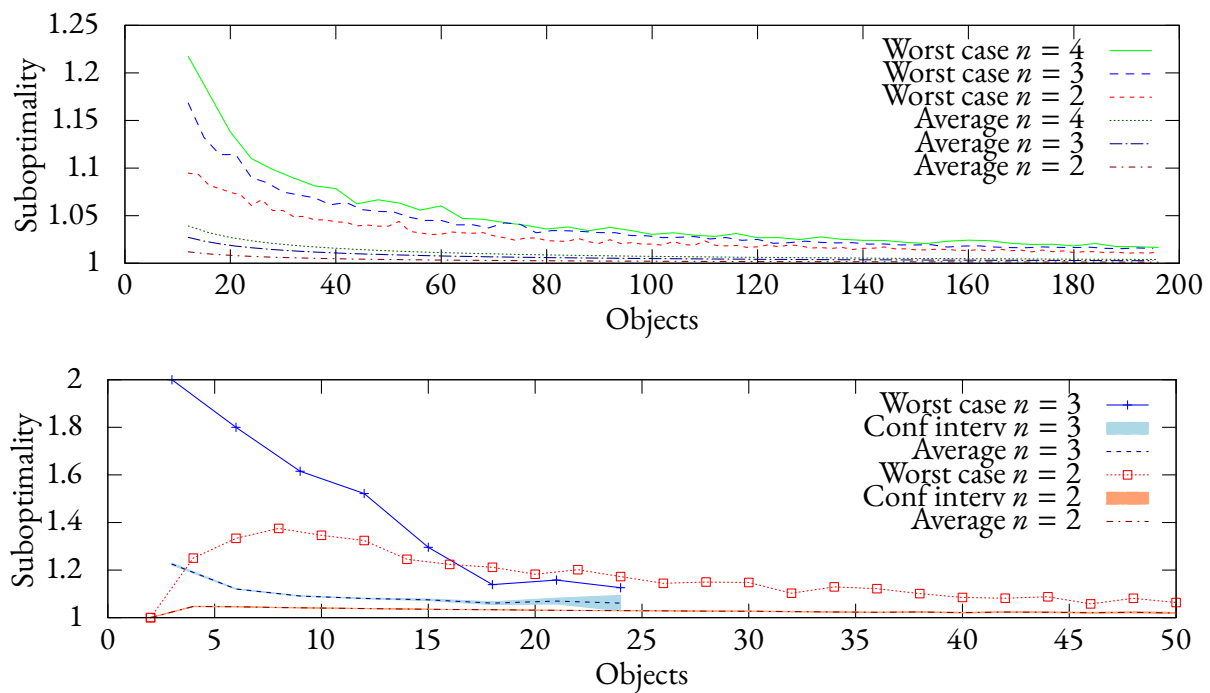
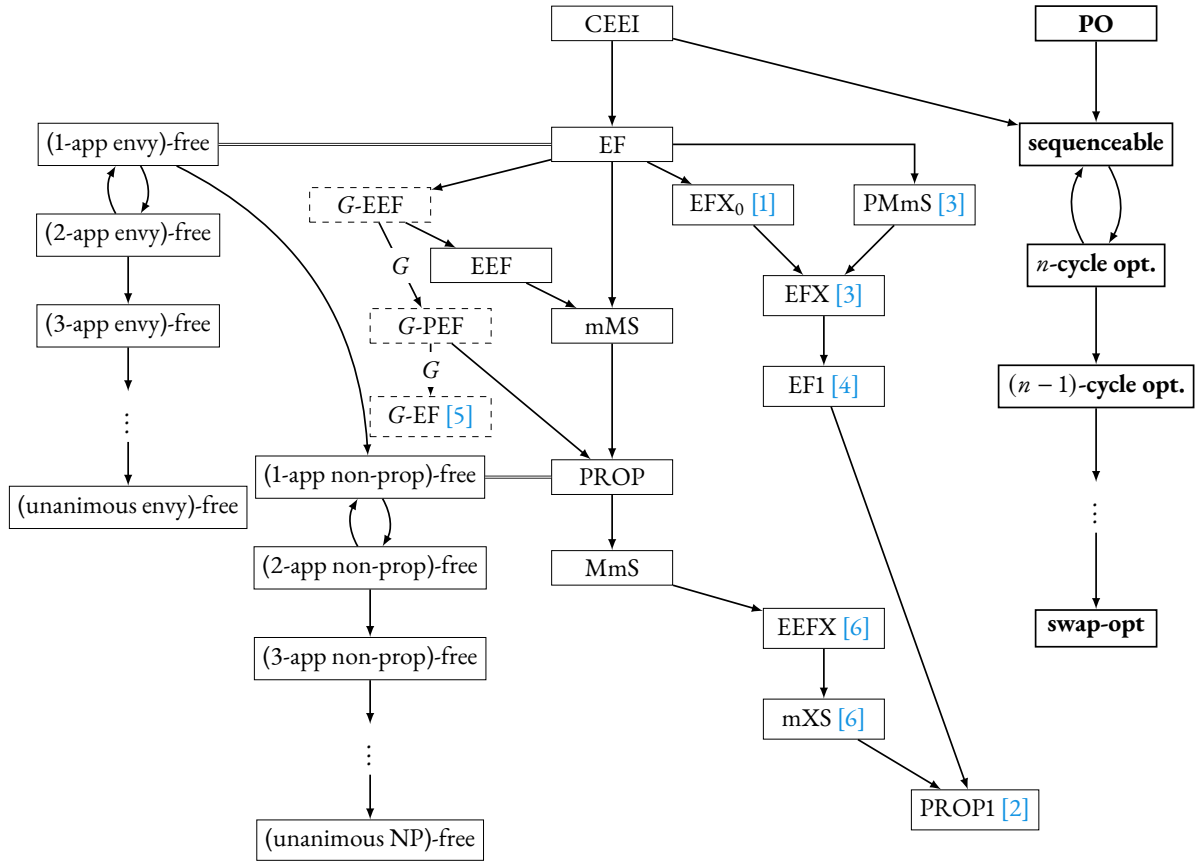


Figure III.2 — MPEF for utilitarian (top) and egalitarian (bottom) SW, under Borda scoring, and balanced picking sequences for various number of agents. Average and worst-case MPEF are given, for a set of profiles generated uniformly at random. Extracted from our paper ([Baumeister et al., 2017](#)), to which we refer for details on the experimental setting.

and Jérôme Lang that led to a publication at the International Joint Conference on Artificial Intelligence ([Bouweret, Chevaleyre, Durand, et al., 2017](#)).



- [1] [Plaut and Roughgarden, 2018](#)
- [2] [Aziz, Moulin, et al., 2020](#); [Barman and S. Krishnamurthy, 2019](#); [Brânzei and Sandomirskiy, 2019](#); [Conitzer et al., 2017](#)
- [3] [Caragiannis, Kurokawa, et al., 2016](#)
- [4] [Budish, 2011](#)
- [5] [Chevalyre, Endriss, et al., 2017](#)
- [6] [Caragiannis, Garg, et al., 2023](#)

Figure III.3 — Landscape of (binary) fairness (and efficiency) properties, completed after Figure II.1. An arrow indicates an implication link. A double edge indicates an equivalence. Dashed properties indicate hierarchies induced by a graph G . A label G on an arrow indicates that the implication is valid only for the same graph G . Here, the bold part indicates the new efficiency concepts introduced in this chapter.

Chapter IV

And the winner is... Alternative voting rules for fairer representation

Chapter foreword

This chapter explores alternative voting rules for political elections. It is composed of two different parts, each highlighting a different investigation technique.

The work presented in the first part of this chapter, dedicated to *in situ* and online experiments, is the result of a collaboration with Antoinette Baujard, Renaud Blanch, François Durand, Herrade Igersheim, Jérôme Lang, Annick Laruelle, Jean-François Laslier, Isabelle Lebon, Vincent Merlin and Théo Delemazure. In this collaboration, I was especially in charge of coordinating with Renaud Blanch the experiment in Grenoble, and in charge of the entire online experiment. The main output of this collaboration was the publication of two datasets (Bouveret et al., 2018; Bouveret et al., 2019), that respectively gather the data on 37 739 and 6358 participants.

The second part of this chapter is the result of a collaboration with Renaud Blanch that led to the publication of a technical report presented at the law commission of the national assembly (Blanch and Bouveret, 2018).

Even if none of these two topics resulted in a formal publication in a peer-reviewed journal or conference, they led to alternative significant scientific contributions and visibility: datasets that are at the basis of several publications in economics (see <https://www.gate.cnrs.fr/vote/> for a non-exhaustive list), substantial media coverage and dissemination (see <https://vote.imag.fr/about#media>), and a political visibility for the part concerning the legislative election. Moreover, they nicely complement the last two chapters on fair division by exhibiting experimental techniques that can be used on top of a thorough theoretical analysis to investigate the properties of collective decision making procedures.

Blanch, Renaud and Sylvain Bouveret (May 2018). *Simulations de systèmes de vote avec proportionnelle pour les législatives*. Tech. rep. Laboratoire d'Informatique de Grenoble. URL: <http://recherche.noiraudes.net/resources/2018-05-28-rapport.pdf>.

Bouveret, Sylvain et al. (July 2018). *Voter Autrement 2017 - Online Experiment*. Dataset and companion article on Zenodo. DOI: [10.5281/zenodo.1199545](https://doi.org/10.5281/zenodo.1199545). URL: <https://doi.org/10.5281/zenodo.1199545>.

— (Nov. 2019). *Voter Autrement 2017 for the French Presidential Election - The data of the In Situ Experiment*. Dataset and companion article on Zenodo. DOI: [10.5281/zenodo.3548574](https://doi.org/10.5281/zenodo.3548574). URL: <https://doi.org/10.5281/zenodo.3548574>.

We continue our journey through the different expressions of fairness by exploring a completely different collective decision making setting: *voting*. Voting refers to the situation where a group of agents, called the *voters*, express their opinions about a set of *candidates* (options, issues, alternatives...) using *ballots* that may be of various forms. These ballots are used as the input of a *voting rule*, that is used to reach a collective decision, that is, to choose a single candidate (the *winner*), or a committee of candidates (the *winning committee*), depending on the context.

In practice, such a situation occurs in quite different contexts. Choosing a restaurant where to have dinner with a group of colleagues, the social activity at the next lab days, or the name of a team are examples of what could be considered to be quite low stake situations. Ranking a set of candidates for a permanent position at the university, electing the representatives of a lab, or deciding which measures to apply to reduce the carbon emissions of a university can be reasonably considered to be middle stake contexts.¹ But the prototypical voting settings that involve the highest and most sensitive stakes are probably political elections. This will be the topic of this chapter.

Elections by universal suffrage are a crucial moment in democratic life. In France most political elections, on top of which the most two important ones, namely legislative and presidential elections, are based on two round majority. However, there are many other ways to conduct an election, each of which conveying different symbolic principles and having different political consequences. It starts from the way voters express their political opinions on the ballots. We can ask them, like in the French context, to which single candidate they choose to join up. We can ask them, like *e.g.* in the Irish case, to rank the candidates, so that a voter does not have to choose a unique candidate. Or we can ask them to be in the shoes of an evaluator that gives a score to each candidate, independently of the others. And even for a given ballot type, various voting rules can be used to determine the winner or the winning committee.

Changing the voting rule of a political election can have a lot of implications. The political landscape will be reconfigured. The voters will behave differently. The result will probably differ. The questions of how the political landscape reconfigures and how the voters adapt to this landscape are out of reach of our analyses. Hence, in this chapter, we will only focus on the following question:

How does the use of an alternative voting rule change the result of an election and to which extent could it be considered to lead to a fairer representation?

To investigate this question, we focus on two particular election contexts: the French presidential election of 2017 (with some insights on the 2022 election) and the French legislative election of 2017 (with some insights on the 2024 election). These two contexts illustrate two different techniques that can be used to answer to the initial question: *in situ* and online experiments on the one hand, and computer simulations on the other hand. Not that even if this work is partially rooted in the current dissatisfaction of a part of the French electorate regarding the current voting rules², our study is not prescriptive.

1 Technical preliminaries

Most work on voting theory historically focuses on *Social Choice Functions*, namely, functions that take as input ranked ballots, and investigate axiomatic properties of these functions. Voting rules based on

¹All these examples are on purpose related to the academic context, but of course they can be transposed anywhere else.

²See <https://vote.imag.fr/results/online-2022> and https://www.gate.cnrs.fr/wp-content/uploads/2024/07/VA22_CR_Enquete.pdf.

ranked ballots have a major advantage: just like the criteria discussed in Chapters II and III for fair division, they do not rely on any interpersonal comparisons of utilities. In other words, to operate, these voting rules do not require to compare the benefits and losses of the voters with each other. Unfortunately, as recalled in the introduction of this document, Arrow’s seminal theorem (Arrow, 1951) states that there does not exist any reasonable voting procedure only based on rankings. In spite of this impossibility theorem (and a few other ones in social choice theory), a substantially large number of works in social choice theory is still dedicated to axiomatic properties of social choice functions.

In this chapter, our focus will be somewhat different to this standard literature. First, we will consider different kinds of ballots (not only rankings), and secondly, we will not investigate axiomatic properties at all, but rather study how these voting rules behave in practical situations. The reader interested by theoretical properties of voting rules can for instance read the first part of the survey book on Computational Social Choice (Brandt et al., 2015), especially the introductory chapter on voting (Zwicker, 2015).

The only technical preliminary we need for this chapter concerns the notion of *voting rule*. Our definition is not completely standard, in the sense that it makes the ballot expression language (here, what we will call the *opinion language*) explicit, where usual definitions typically mainly focus on the function computing the winner.

Definition 16 (Opinion Language) *Let C be the set of candidates. An opinion language is a set \mathcal{L}_C of well-formed formulas depending on C .*

In our terminology, a voting rule is simply an opinion language with an aggregation function that specifies how to compute the election winner.

Definition 17 (Single-winner voting rule) *Let C be a set of candidates. A single-winner voting rule is a pair (\mathcal{L}_C, f_C) , where \mathcal{L}_C is an opinion language and f_C is a function mapping any vector \vec{v} of well-formed formulas from \mathcal{L}_C to a unique winner $f_C(\vec{v})$.*

In other words, a voting method is made of a preference representation language that depends on the set of candidates, and an aggregation function whose role is to compute a single winner from any vector of well-formed ballots (such a vector will be called voting profile).

The voting rules we use in this experiment are based on aggregation functions that are irresolute, in the sense that they can map a given profile to several winning candidates. To be well defined, these aggregation functions must be associated to a *tie-breaking* rule t_C that maps each set of co-winning candidates \mathcal{W} to a unique candidate $w \in \mathcal{W}$.

The difference between an opinion language and a voting rule is thus the formal specification of the method used to compute the election winner. In the experiments we describe in the first part of this chapter, we expose participants to voting rules and opinion languages. In both cases, the participants can give their opinions about candidates using a particular language (ranking, scores, etc.). However, in the case of a voting rule, they exactly know how their opinions will be aggregated to compute the winner, and they can take that information into account when they vote. This is why we distinguish the two situations in the underlying formalism of our experiments.

In the second part of this chapter, presented in Section IV.3, we will also consider multi-winner voting rules (also known as committee voting rules) that will be defined as follows:

Definition 18 (Multi-winner voting rule) *Let C be a set of candidates. A multi-winner voting rule is a pair (\mathcal{L}_C, f_C) , where \mathcal{L}_C is an opinion language and f_C is a function mapping any pair made of a vector \vec{v} of well-formed formulas from \mathcal{L}_C and an integer $K \in \mathbb{N}^*$ to a winning committee $f_C(\vec{v}, K)$ of size K .*

Like before, multi-winner voting rules we use can be based on aggregation functions that are irrelative, and that must be used in conjunction with a tie-breaking rule to match our definition of a voting rule.

2 Voter Autrement

The first part of this chapter will be devoted to *in situ* and online experiments that were carried out in April 2017 during the French presidential election. More precisely, these experiments were the third edition of the series "Voter Autrement" (the fourth edition, that we slightly evoke at the end of the section, was conducted in April 2022), whose objective is to test alternative voting rules for the French presidential election. This series of experiments have been conducted by researchers in economics and computer science (with a special interest in social choice), mainly from Caen, Grenoble, Paris, Saint-Étienne and Strasbourg.

Conducting these experiments required a tedious organization work, that involved a close collaboration not only with the set of colleagues aforementioned, but also with the local authorities, and the involvement of dozens of volunteers. In this section, we carefully describe the experimental protocol, and draw some conclusions that can be made from these experiments with respect to the initial question of this chapter.

2.1 Experimental setting

Context In April 2017 a presidential election took place in France. 11 candidates competed for this election: Nicolas Dupont-Aignan (NDA), Marine Le Pen (MLP), Emmanuel Macron (EM), Benoît Hamon (BH), Nathalie Arthaud (NA), Philippe Poutou (PP), Jacques Cheminade (JC), Jean Lassalle (JL), Jean-Luc Mélenchon (JLM), François Asselineau (FA) and François Fillon (FF).³ Even if our objective is not to draw any political conclusion, we recall that the political context was somewhat unusual for a French presidential election. First, the two major parties (PS and LR) organized primaries, respectively for the first and the second time in the party's history. Second, whereas the candidate of the right-wing party (LR), François Fillon, was predicted by opinion polls to be the winner of the election with a comfortable margin, a scandal that happened a few weeks before the election ruined all his aspirations and in the end he did not even manage to qualify for the second round. Third, this election was characterized by a substantial increase of the extremist votes, both at the left and the right sides. Finally, a new man (Emmanuel Macron) that was not from the historical parties emerged almost from nowhere and in the end won the elections.

The official results of the first round of this election are recalled in Table IV.1 (page 66). The experiments we ran were carried out in parallel with the official election.

Objectives Our experiment has several objectives. First, the aim is to increase the scientific knowledge about electoral behaviour, and investigate the influence of the voting rule on the result of the election. Another output is to determine whether the voting rules are usable in practice: whether they are well

³The candidates are presented using the official order randomly chosen by the French Constitutional Council. The associated colours are those chosen by the French newspaper *Le Monde* to reflect their political parties.

understood by voters, whether they are easy to tally, to verify, and so on. Another objective is to build a whole practical voting dataset that could be used for further research on experimental social choice. And finally, an indirect objective is pedagogical: without promoting a single voting rule (like other experiments do), our aim was to advertise voting rules as alternatives to the plurality with runoff rule used for the presidential election. Hence, our aim was to participate, as scientists, to the French public debate on electoral democracy, by providing insights on the potential use of alternative voting rules for high stake elections.

The online and in situ experiments In 2017, two parts of the *Voter Autrement* experiment were carried out in parallel: an online part and an *in situ* part.

First, an electronic experiment was organised on the web, to test alternative voting methods for the presidential election using a web application. The online experiment started on March, the 29th, and lasted until the second round of the official election, namely May, the 7th. A total of 37739 voters took part in the online part of the experiment. The participation to this experiment was voluntary and free. Anyone could take part just by going to the dedicated website.⁴ The website was first advertised using mostly academic mailing lists and social networks. In the course of the experiment, the information was then relayed by several press media⁵ and via several groups of interests, including activists from several political parties. During the whole experiment, we also received several hundreds of e-mails from participants to the experiment, including comments, suggestions and questions, mostly benevolent.

In parallel with this online experiment, we organized an *in situ* experiment in five cities in France: Hérouville Saint-Clair (Calvados), Strasbourg (Bas-Rhin), and 3 cities in Isère (Allevard les Bains, Crolles and Grenoble). This *in situ* experiment took place on April the 23rd, namely, the day of the first round of the official election. For each city, we had chosen several weeks before an appropriate polling station to run our experiment (these stations were chosen mainly for practical reasons). The day of the experiment, the voters had to first vote in the official polling station. Then, we asked them whether they wanted to participate in our experiment. If they agreed, they were then guided to another part of the room, where we had replicated the polling station, with polling cabins, urns, sign-off sheets, and of course our own paper ballots. By replicating the official voting protocol (including sign-off sheets that are useless to us) we aimed at inducing the smallest possible perturbation that could change the way the voters express their opinions, compared to the official election.

Note that the organization of such an *in situ* experiment is complex. First, it needs a close cooperation with city officials for the logistic organization. Secondly, it requires a formal authorization from the prefect.⁶ Thirdly, each registered voter of the polling stations concerned must be contacted individually by mail a few weeks before the election to inform them that an experimentation will take place. Fourthly, an information meeting should be organized a few weeks before the experiment. Fifthly, each registered voter should receive an individual mail after the experiment presenting the results. Finally, of course, dozens of volunteers are needed the day of the experiment in the polling stations.

In total, 6358 voters participated in this *in situ* experiment in the five polling stations.

⁴<https://vote.imag.fr/>

⁵See the press release on the website of the experiment, <https://vote.imag.fr/about#media>.

⁶In France, the prefect is the chief administrator in a department, and a representative of the central government.

2.2 Voting rules

Before speaking of the voting rules we tested, let us first recall how the president is elected in France. The official voting rule used is *plurality with runoff*. This two-round voting rule informally works as follows. For the first round, each voter casts a ballot that must contain a single candidate name to be valid and non-blank. If no candidate receives an absolute majority of votes, a second round is organized, where only the two candidates having received the highest numbers of votes compete. Like for the first round, each voter is asked to cast a single-name ballot. The candidate receiving the highest number of votes wins the election.

In our experiment, we consider two alternative families of voting rules respectively based on candidate evaluation and ranking.

Evaluative voting The first voting rules we use are variants of *Evaluative voting* (also called in English Range Voting, Utilitarian Voting, or Score Voting), a single round voting rule where each voter gives a grade from a predefined numerical scale s to each candidate. Given a voting profile, the score of a candidate is simply the sum of all the grades given by the voters to this candidate. The candidates with the highest score win.

In our version of range voting, for simplicity of preference elicitation, abstention is allowed, in the sense that a voter is not forced to give an evaluation to all the candidates. The interpretation we give to abstention is that if a voter does not evaluate a given candidate c , then c receives the lowest score from this voter.

In our experiment, we use the following scales:

- Approval voting (0; 1) (in all polling stations, and online);
- Approval / disapproval voting with scales (-0.5; 0; 1) (Allevard les Bains), (-1; 0; 1) (Allevard les Bains) and (-2; 0; 1) (Allevard les Bains);
- Evaluation voting with scales (0; 1; 2) (Strasbourg and online), (-1; 0; 1; 2) (online), (-1; 0; 1) (Strasbourg and online), (0; 1; 2; 3) (Strasbourg and online), (-1; 0; 1; 2) (Strasbourg and Hérouville Saint-Clair), (0; 1; 2; 3; 4; 5) (Hérouville Saint-Clair);
- Continuous evaluative voting (Grenoble).

For continuous evaluative voting, we use a virtually continuous scale as a proxy to the interval $[-1, 1]$. For the paper ballots used in Grenoble, it takes the form of a line on which a voter is asked to put a mark (see Figure IV.1 in Section IV.2.3 below).

In Crolles, the voters were also asked to give their opinions on the form of an evaluation on the scale $\llbracket 0, 20 \rrbracket$ (which is a standard grading scale in France). Finally, we also tested (also in Crolles) a mix of plurality with runoff and evaluative voting, namely, 2 round approval voting. We only mention it here for completeness, but we will not give much details on this rule in the chapter.

Ranking-based rules The second kind of rules we use are closer to the traditional setting of Arrowian social choice theory, where it is assumed that the ballots come in the form of linear orders (rankings) over the candidates. In practice, asking for complete rankings would not be realistic: there are eleven

candidates for the French election, and ranking all of them would be tedious and irrelevant. This is why we just ask for partial orders, or more exactly for possibly truncated rankings. More formally, in what follows, we call a k -ranking (resp. $(\geq k)$ -ranking) a ballot where the voters are asked to rank exactly (resp. at least) k candidates, that is (i) designate a subset of exactly (resp. at least) k candidates they approve of, and (ii) give a linear order on these candidates.

For practical and consistency reasons, the voting rules based on rankings have only been tested online and not *in situ*. In polling stations, the ballots used were mostly paper ballots, and asking for rankings is much harder on paper than on a digital interface (that can be designed to prevent misuses of the ballot and to reduce the risk of invalid votes). Note that using ranking-based paper ballots is not impossible though, and some countries use this modality for high stake political elections. For instance, Ireland uses this kind of ballots for the presidential election, that is based on Instant Runoff Voting. Needless to say that the tallying process can last much longer than in France.

The first rule of this kind we test in our experiment is Borda- k , which is a variant of the seminal Borda rule – one of the most well-known rules from the family of *scoring rules* – where voters only rank k candidates. This rule takes k -rankings as inputs. A candidate ranked i^{th} in a ballot wins $k+1-i$ points. In other words, in the Borda- k method, a candidate wins k points each time they are ranked first by a voter, $k-1$ each time they are ranked second, and so on. They do not win any additional point when they do not appear in a voter's list. Like with evaluative voting, the candidate with the highest number of points wins the election.

The second kind of voting method we use is Instant-Runoff Voting (IRV), which can be seen to be a variant of plurality with runoff with as many rounds as candidates. More formally, the $(\geq k)$ -IRV voting rule takes $(\geq k)$ -ranking ballots as input. Each round of the procedure works as follows.

1. The plurality loser is computed, that is, the candidate being ranked first by the lowest number of voters.
2. The plurality loser is eliminated from the election, that is, from the set of candidates, and from the voters' rankings.
3. Another round starts with the new set of voters and rankings.

The candidate reaching the last round is elected. Note that if at some round a candidate is ranked first by more than half voters, the process can stop as this candidate is guaranteed to be the winner.

Opinion on candidates On top of the aforementioned voting rules, we also tested in the online experiment two different opinion languages: continuous opinion and pairwise comparisons.

Like for continuous evaluative voting, in the continuous opinion language, each voter is asked, for each candidate, either to give a grade or to abstain. The main difference with evaluative voting (which explains that continuous opinion is not a voting rule) is that we do not explicitly state how a potential winner is computed; in particular this opens the door to several interpretations of abstention about a candidate.

To simulate a continuous scale, we used a slider with 100 positions (which we judged empirically enough for the scale to look continuous to the users).

The second opinion language we tested is the k -pairwise comparison language. Here, a set of k successive pairs of candidates (c, c') are presented to a voter. For each pair, the voter is asked which of the two

candidates they prefer. A voter might also abstain for a given pair presented. Once again, here, we do not explain how a potential winner could be computed, which explains why this could not be considered to be a voting rule.

2.3 The ballots

In the previous section, we have described the voting rules we tested in our experiment. We will now give a few more insights on the experimental protocol by describing the ballots used in the experiment and the web interface used for the online version.

The online experiment For the online experiment, each participant was faced with several voting rules and opinion languages, and a final questionnaire. This questionnaire contained various items concerning the participant's personal information (*e.g.* age, gender, socio-professional category...), and their perception of the voting rules encountered. We also ask them which candidate they voted (or intend to vote) for at the official election.

More precisely, each participant had to answer to:

1. 2 types of evaluative voting;
2. 1 type of ranking based voting rule among 4-Borda, ≥ 1 -IRV and ≥ 4 -IRV;
3. 1 continuous opinion;
4. 8 random duels;
5. 1 questionnaire.

Each step could be skipped. We do not further describe the interface here, as it can still be tested online⁷ for both 2017 and 2022 presidential elections.

Note that we do not prevent multiple participation, since we do not have any way of formally checking the identity of the participants.⁸

The *in situ* experiment in Grenoble For the *in situ* experiment, in most polling stations, paper ballots were used. The main difficulty is that the ballot should be self-contained: it should be easy to understand, and should precisely describe the voting rule used.⁹

We give in Figure IV.1 a copy of the paper ballot used for Grenoble. This ballot is original compared to the other *in situ* settings in the sense that a continuous scale has been used (see the right part of the front side of the ballot). This required to develop an improved tallying process, because tallying these ballots requires to be able to measure the distance between the left side of the continuous scale and the mark put by the voter. In other words, this requires making $11 \times m$ measures (where m is the number of ballots, 1080 for Grenoble) which is tedious in practice.

This is why we developed a computer-assisted tallying process based on three steps:

⁷<https://vote.imag.fr/>

⁸Besides, checking and storing the identity would not be desirable and legally difficult to implement for privacy reasons.

⁹Even if some volunteers could give some explanations the day of the experiment, we also wanted to test whether such a ballot was usable in a real context were the voters should understand it by themselves.



Figure IV.1 — The ballot used for the election in Grenoble, front (top) and back (bottom) side.

1. The ballots are scanned.
2. The ballots are skimmed through by a software tool we developed, that uses basic image processing approaches to detect marks.¹⁰
3. We manually check the accuracy of the detection by superimposing ballots and the marks found by the image processing tool, which can be quickly done.

2.4 Results

IV.2.4.1 Of biases

Before giving a few insights on the results of the experiments, let us evoke the potential biases of our sample of participants. The fact that we picked particular polling stations for the *in situ* experiment, and that the participation is on a voluntary basis anyway for both the *in situ* and the online experiments entails that our population is highly non representative of the French electorate. This can be formally verified by analyzing the answers to the final questionnaire. For instance, for the 2017 online experiment, about 60% of participants declared to be men, and 42% declared to be between 18 and 29 years old; 87% declared to have a higher education degree.¹¹ Things are not better concerning diversity, for the 2022 online experiment, where 66% of the participants declared to be men, and 91% declared to have a higher education degree.¹²

For the *in situ* experiment, the samples are more reasonably mixed in terms of gender and age, but they are biased in favor of highly educated people and (like for the online experiment) against conservative voters.

This lack of representativeness is a problem if we want to be able to compare the results for various tested voting rules to the official election. Table IV.1 for instance shows the difference between the result at the first round of the official election, and what the voters that participated in our online experiment declared as their vote for the official election. Let us consider Marine Le Pen for instance: while 21.3 % voted for her for the first round of the official election, only 1.57 % of our participants declared to do so.

	NDA	MLP	EM	BH	NA	PP	JC	JL	JLM	FA	FF	Blk	Abs
Off. (%)	4.7	21.3	24.01	6.36	0.64	1.09	0.18	1.21	19.58	0.92	20.01	2.55	22.23
Decl. (%)	1.08	1.57	22.11	18.54	0.34	2.12	0.16	0.58	42.37	1.36	4.38	2.38	3.02
Corr.	4.35	13.57	1.09	0.34	1.88	0.51	1.12	2.09	0.46	0.68	4.57	1.07	7.36

Table IV.1 — Official results of the 2017 experiment (first line) compared to the vote intentions declared by the participants to the online experiment (second line) – limited to those who have tested approval voting. The correction coefficient (last line) is simply given by the ratio between both scores.

Aligning our population of participants to the population of French voters is complex. The best we can do so far is to use the ratio between the official result and the declared vote to give different weights to the participants. Hence, using such a method, according to Table IV.1, last line, a participant who has

¹⁰The careful reader might have noticed the black discs on the ballots, that are used to calibrate the image processing tool.

¹¹See <https://vote.imag.fr/stats>.

¹²See <https://vote.imag.fr/results/online-2022>.

declared to vote for Marine Le Pen will count for 13.57 voters, while a voter for Benoît Hamon will only count for 0.34 voter. Note that for the *in situ* experiments, several alignments are possible. We could align our population to the national population (national correction), or align it using the local results at the polling station (local correction). For consistency reasons, in this document, we only use the national correction, but the results can vary a lot between the raw scores, the nationally corrected scores and the locally corrected scores.

IV.2.4.2 Global insights

We will now give a few insights on the results of the experiment. The complete results and interpretations can be found on the dedicated webpage.¹³

Different voting rules yield different rankings First and foremost, the election result closely depends on the voting rule used. As we can see from Figure IV.2 and Table IV.2, that respectively show the results of the online experiment and the results of the *in situ* experiment, the collective ranking varies with the voting rule. We might also observe that the result also significantly differs from the result of the official election, but this observation should not be overinterpreted. As reminded above, our population is highly non representative, and the method we use to correct the biases is far from perfect. Hence, the only relevant comparisons we can make are those between results given by applying identical correction techniques.

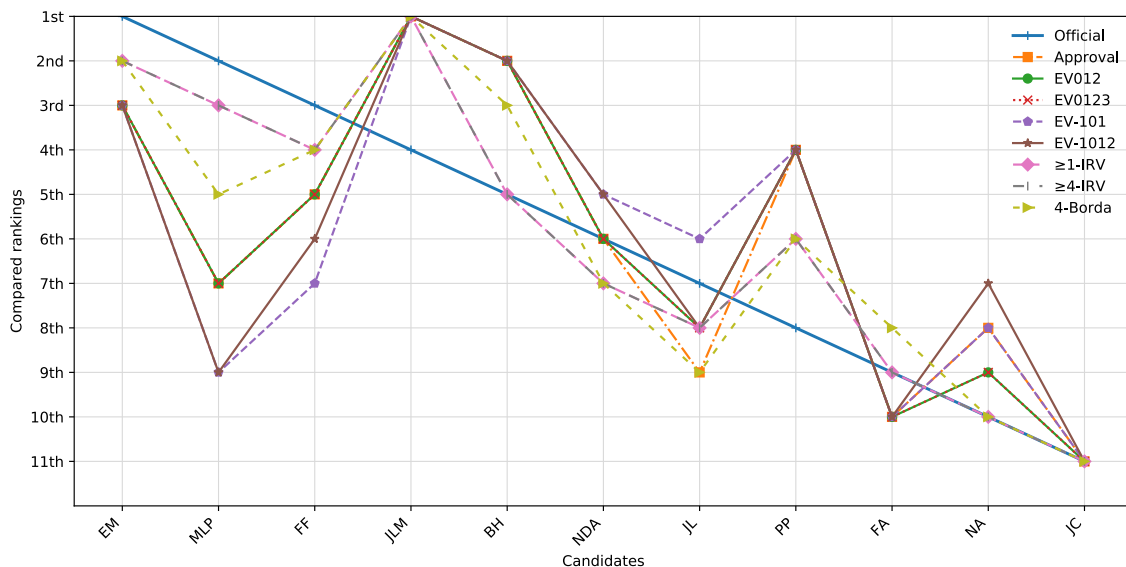


Figure IV.2 — Comparison of the ranks obtained by the candidates for different voting rules tested in the online experiment (nationally corrected data). A dynamic version of this graph, and many other results can be found here: <https://vote.imag.fr/results/online-2017>.

We can also observe in Figure IV.2 that several families of voting rules seem to emerge. On the one hand, the evaluative voting rules (including Approval Voting) yield similar rankings, while IRV voting rules (and to some extent Borda) form another group. We can also observe that the rankings obtained

¹³<https://www.gate.cnrs.fr/vote/>

		Official	Allevard les Bains			Crolles		Grenoble	
			EV(-0.5,0,1)	EV(-1,0,1)	EV(-2,0,1)	AV	OPI	AV	EV
raw scores	1st	EM	JLM	JLM	JLM	EM	EM	BH	BH
	2nd	MLP	BH	BH	BH	BH	BH	JLM	JLM
	3rd	JLM	EM	EM	EM	JLM	JLM	EM	EM
corr. scores	1st	EM	EM	JLM	JLM	EM	EM	BH	EM
	2nd	MLP	JLM	EM	EM	JLM	BH	EM	JLM
	3rd	JLM	BH	BH	BH	BH	JLM	JLM	BH

		Hérouville St Clair			Strasbourg				
		AV	EV(0-3)	EV(0-5)	AV	EV(0-2)	EV(-1,0,1)	EV(0-3)	EV(-1,0,1,2)
raw scores	1st	JLM	JLM	JLM	JLM	JLM	BH	JLM	JLM
	2nd	BH	EM	BH	BH	BH	JLM	BH	BH
	3rd	EM	BH	EM	EM	EM	EM	EM	EM
corr. scores	1st	EM	EM	EM	EM	JLM	JLM	EM	JLM
	2nd	JLM	JLM	JLM	JLM	EM	EM	JLM	EM
	3rd	BH	BH	BH	BH	BH	BH	BH	BH

Table IV.2 — Top three candidates for different voting rules tested in the 2017 *in situ* experiment, raw scores (top), and nationally corrected scores (bottom).

with IRV-like voting rules are closer to the official first round ranking, even though, as noted earlier, any comparison between our voting rule and the official one should be interpreted cautiously.

This visual observation is confirmed by a numerical analysis of the rankings. If we compute the Kendall-Tau distance matrix between the rankings yielded by different voting rules, and use it as the input of a clustering method (with 3 clusters), we obtain the following clusters for the online voting method:

- Cluster #1: AV / EV012 / EV0123 / EV-101 / EV-1012
- Cluster #2: ≥ 1 -IRV / ≥ 4 -IRV / 4-Borda
- Cluster #3: Official

Of course we cannot observe the same groups within the rankings of the *in situ* experiment, because all the rules tested are evaluative voting rules. However, among these evaluative voting rules, two groups also seem to emerge: on the one hand the rules including only positive scores, and on the other hand the rules including negative scores (with the exception of EV(-0.5,0,1)).

Applying the same clustering technique as before for the rankings of the *in situ* experiment gives:

- Cluster #1: All-EV(-0,5;0;1) / Cro-AV / Gre-EV / HSC-AV / HSC-EV(0-3) / HSC-EV(0-5) / Str-AV / Str-EV(0-2) / Str-EV(0-3)
- Cluster #2: Cro-OPI / Gre-AV

- Cluster #3: All-EV(-1;0;1) / All-EV(-2;0;1) / Str-EV(1,0,1) / Str-EV(-1,0,1,2)

If we have a look at the results of the 2022 online experiment (see Table IV.3), we can make similar observations. Clustering the different voting rules according to the rankings they yield also clearly parts them into similar subgroups (but the distinction between positive/negative EV and positive only EV is less clear here):

- Cluster #1: AV / EV(-1, 0, 1) / EV(0, 1, 2) / EV(-1, 0, 1, 2) / EV(0, 1, 2, 3)
- Cluster #2: IRV 4+ / IRV
- Cluster #3: Borda / MJ5 / MJ7
- Cluster #4: Official

	Official	AV	≥ 4-IRV	≥ 1-IRV	EV(-1, 0, 1)	EV(0, 1, 2)
1st	EM	YJ	EM	EM	YJ	YJ
2nd	MLP	JLM	JLM	MLP	EM	JLM
3rd	JLM	EM	MLP	JLM	JLM	EM

	EV(-1, 0, 1, 2)	EV(0, 1, 2, 3)	Borda	MJ5	MJ7
1st	YJ	YJ	JLM	YJ	YJ
2nd	JLM	EM	YJ	JLM	JLM
3rd	EM	JLM	EM	EM	EM

Table IV.3 — Top three candidates for different voting rules tested in the 2022 online experiment, nationally corrected scores. Here, "MJ5" (resp. "MJ7") stands for Majority Judgment with a 5-level scale (resp. 7-level scale).

Evaluative vs single-name voting rules Another observation we can make is that different voting rules tend to favour different types of candidates. While the official (single-name) voting rule (and, to some extent, the voting rules based on partial rankings) tend to favour more polarizing candidates, evaluation-based voting rules tend to favour more consensual candidates. Even if it is hard to evaluate to which extent a candidate is polarizing or consensual, we can evaluate it for instance by analyzing the distribution of the scores given by the participants to this candidate on the continuous evaluation scale. This distribution is shown in Figure IV.3. On this graph, we can see that a candidate like Marine Le Pen is extremely polarizing: a very high rejection score, a quite high score on the right side of the scale compared to the other candidates, and a quite low number of intermediate scores (as can be judged by the quite horizontal progression of her curve). On the opposite spectrum, candidates like Emmanuel Macron or Benoît Hamon appear to be much more consensual, as can be judged from their respective curves.

Looking at evaluative voting rules also gives a completely different interpretation of the political landscape, as several candidates that have a very low score in the official election can benefit from the fact that

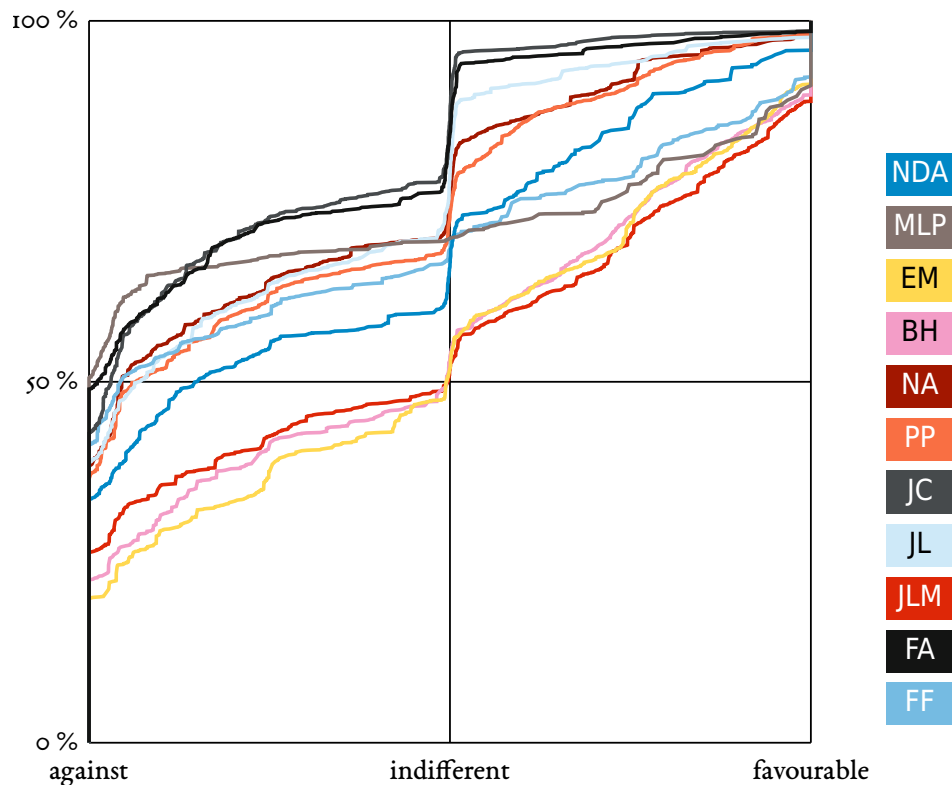


Figure IV.3 — Distribution of the continuous scores of each candidate (nationally corrected scores). The X -axis represents the continuous scale. The curves show, for each candidate c and each x , the percentage of participants that gave an evaluation $\leq x$ to this candidate c . (Adapted from a work by Renaud Blanch)

the voters expression is not limited to a single name. This can be seen *e.g.* on the results of the 2022 on-line experiment,¹⁴ where, for instance, the approval scores of candidates like Yannick Jadot or Philippe Poutou are dramatically higher than their scores at the official election. These "smaller" candidates were potentially victims of the *vote utile* phenomenon, where voters cast their ballots in favour of the candidate that has the highest chance to get to the second round among the ones they might approve.

This dichotomy between "big" and "small" candidates can also be revealed by the observation of the number of approvals given by the supporters of each candidate.¹⁵ What we observe is that the average number of approvals given by the participants is around 3 candidates / 11. However, this number tends to be lower in the population of participants that approved bigger candidates (like Emmanuel Macron, Marine Le Pen or François Fillon), and a non-negligible number of them give only one approval. Conversely, smaller candidates (typically Nathalie Arthaud, Jean Lassalle, François Asselineau and Jacques Cheminade) are very often co-approved with more than 3 other candidates (even 5 or 6).

Note that this classification of candidates and the observation that some families of rules tend to favour some types of candidate is not original and had already been observed, for instance by [Baujard, Gavrel, et al., 2014](#) (on the 2012 French presidential election). [Darmann et al., 2017](#) also provided similar observations on the 2015 parliament election in the Austrian federal state of Styria. They even propose

¹⁴See <https://vote.imag.fr/results/online-2022>.

¹⁵See *e.g.* <https://vote.imag.fr/results/online-2017#detail-approval> and <https://vote.imag.fr/results/grenoble#general>.

a classification of parties into four different types: (i) popular (strong support from a specific segment of the society, and seen positively by a large proportion of the voters), (ii) unpopular (strong support from a small segment of the society, and seen negatively by a large proportion of the voters), (iii) medium (acceptable to a large proportion of the society), (iv) polarizing (induces strong views, either positive or negative).

Perception of the voting rules The last conclusions we might draw from these experiments concern the perception of the voting rules by the participants. First, when they have the possibility to approve (or evaluate) several candidates (instead of only one for the official election), the experiment shows that the participants take this opportunity. This can be seen from the number of approved candidates (around 3 on average, with a minority of single-approval votes).

For evaluative voting, the participants seem to use the whole scale of available grades, but the use of the continuous scale shows that a large majority of evaluations concentrates around -1, 0 and 1. In the continuous scale tested in Grenoble, several participants even give grades that exceed the scale boundaries (especially negatively). This is interesting: it shows that the voters probably tend to exaggerate their preferences when the scale is longer, as was already observed by [Baujard, Gavrel, et al., 2018](#); [Baujard, Igersheim, et al., 2021](#).

Finally, our analysis of the questionnaires show that most people seem to appreciate alternative voting rules compared to the official voting rule. For instance, for the *in situ* experiment in Grenoble, on a scale of four values {0, 1, 2, 3, 4} (0 corresponding to "very bad" and 4 to "very good"), the average appreciations for approval voting and continuous evaluative voting are respectively 2.73 and 2.62, while it is only 2.06 for the official rule. There also seems to be an effect of the age of the participant on their appreciation of the official rule: on the same data, the average appreciation of the official rule is 2.32 for participants older than 70, and only 1.97 for 18-29 years old participants. This seems to be confirmed by the results of the 2022 online election,¹⁶ where Approval Voting seems to be the most appreciated voting rule, and we see a clear effect of the age on the appreciation of the official rule.

Nevertheless, these results on the perception of the voting rules by the participants must be considered cautiously, because once again our population is not representative of the actual French population. A recent survey carried out by a polling institute with a representative sample of the population shows much less enthusiasm for the alternative voting rules.¹⁷ Besides, it is worth noting that in this survey, the results seem to be much more consistent between the official election and evaluation methods: for instance, the top-three candidates remain the same (in spite of the fact that they include the polarizing candidate Marine Le Pen, which was highly disadvantaged by evaluation voting in our experiments).

Once again, these contrasted results prompt us to be cautious in what we can conclude from them. It is hard to compare the results given by the alternative voting rules to the ones given by the official rule, and in any case, it seems hard to tell anything about what would be the political landscape if the official rule were to be changed. Political parties, candidates, and the voters themselves would probably adapt to this new voting rule, just like the current political landscape (including the party primaries) is adapted to plurality with runoff.

¹⁶ See again <https://vote.imag.fr/results/online-2022>.

¹⁷ See https://www.gate.cnrs.fr/wp-content/uploads/2024/07/VA22_CR_Enquete.pdf.

3 Simulations for the legislative election

In this chapter, we investigate the question of determining what a change in the voting rules of political elections would induce in terms of results. In Section IV.2, we gave some insights by analyzing results of real-world experiments on the French presidential election. In this section, we will use another approach: computer simulations.

The starting point of this work is a request that was made to me and my colleague Renaud Blanch by the French office in charge of evaluating scientific and technological choices.¹⁸ In France, the national assembly (one of the two houses of the French parliament) is composed of 577 deputies, one per district (*circonscription*, see Figure IV.5, left), elected by direct universal suffrage using a majority-based two-round system detailed below in Section IV.3.1. Changing the voting rule for a rule that involves more proportionality is a long-standing question in France. The first question we had to answer is to which extent such a change in the voting rule would have changed the composition of the national assembly after the 2017 legislative elections. The second question was to study the effect of a reduction of the number of deputies,¹⁹ possibly in combination with a change in the voting rule, on the composition of the assembly in 2017.

To sum up, the main research question was to test the effect of several changes in the electoral system on the results of the 2017 legislative election. Several datasets were used to make our simulations, starting from the official results of the election (Ministère de l'Intérieur, 2017a; Ministère de l'Intérieur, 2017b). We also used for the part involving to deal with geographic boundaries (redistricting) the datasets proposed by Atelier de Cartographie de Sciences Po, 2017 (electoral districts), by Institut National de l'Information Géographique et Forestière, 2016 (*cantons* – administrative districts) and by Institut National de l'Information Géographique et Forestière, 2015 (departments).

Our methodology and the results of our simulations have been described in a complete report available online (Blanch and Bouveret, 2018). All the code used for this simulation is available on a git repository.²⁰

Main challenges Replaying the 2017 legislative election with an alternative voting rule requires to overcome two major difficulties. First, reducing the number of deputies, or attributing some part of the seats using a proportional rule implies to reduce the number of districts as well, and hence to redefine the electoral boundaries. Second, once this redistricting is made, we need to simulate the results of a vote in each (virtual) new district. While this can be easily done for the first round by combining the results of the official election, the main difficulty is to simulate the results of second rounds that may have never happened. We have proposed several methods to deal with these issues, that we present in the rest of this chapter.

3.1 The voting rules simulated

Before describing the voting rules we simulated, we should clarify that numerous candidates are engaged in this election, affiliated to a large number of political parties. Hence, some of the officially declared

¹⁸Office parlementaire d'évaluation des choix scientifiques et technologiques (OPECST), whose president was the mathematician and deputy Cédric Villani at the time of this study.

¹⁹Which was a proposition made by Emmanuel Macron during his campaign for the 2017 presidential election.

²⁰<https://gricad-gitlab.univ-grenoble-alpes.fr/vote22/simulation>

parties are only represented in a small number of electoral districts (even a single district for some parties). For that reason, we cannot work with the officially declared parties, because we need a consistent political landscape across the electoral districts. For our simulations, we deal with the nomenclature proposed by the Ministry of Interior, that classifies the political parties into 17 ordered *nuances*, represented (with a colour) in Figure IV.4. In what follows, we will use the term "party" to actually denote the *nuance* according to this classification, and we will denote by \mathcal{P} this set of *nuances*.



Figure IV.4 — Political parties (*nuances*) engaged in the 2017 legislative election.

We also limit our simulations to single-name voting rules: mixes of the official election and proportional voting rules but where each person only vote for a single party (at each round).

Official majoritarian rule The official rule used for the legislative election is based on a single-name two-round majority system that is applied district by district (one deputy is elected in each district separately). During the first round:

- a candidate obtaining a number a votes that is both greater than 50 % of the valid votes and greater than 25 % of the registered voters is directly elected (no second round);
- otherwise, the two candidates with highest plurality scores automatically qualify for the second round. Any other candidate obtaining a greater number a votes than 12.5 % of the registered voters also qualifies for the second round.

The second round (traditionally organized one week after the first round) is simply based on plurality voting.

The official rule is thus purely majoritarian. The rule only guarantees geographic representation (one deputy per electoral district). The assembly can hence be very far from representing the proportions of different political parties as expressed by the votes cast during the first round of the election (see Section IV.3.4).

Mixed rules To introduce a touch of proportionality in the voting rule without turning the whole election into a proportional one, one can mix the official rule with a proportional one. To do so, there exist a lot of possibilities, including methods that separate the election into a party-list national election and a district-based direct local election. This is for instance the case in Germany, where the voters cast two ballots for the *Bundestagswahl*: one for the candidates in the electoral district (*Direktkandidaten im Wahlkreis*), and the other for a list at the *Land* scale (*Landesliste*). Here, we rule out this kind of separate voting rules and focus on mixed rules that first elect a subset of the deputies using the official majoritarian rule, and then complete this subset by giving additional seats to some parties, so as to guarantee some form of proportionality. Let us formalize things a little.

Let N be the size of the assembly (hence, the final number of deputies to elect), and $K \leq N$ be the number of deputies elected by the majoritarian rule. Let also suppose that n parties compete in the election, and that party i gets k_i deputies elected by the majoritarian rule (hence $K = \sum_{i=1}^n k_i$). We also

suppose that party i gets q_i votes at the first round and denote by Q the total number of votes, that is, $Q = \sum_{i=1}^n q_i$ (we rule out blank or invalid votes). The question is thus to attribute the remaining $N - K$ seats to the different parties so as to guarantee some form of proportionality, that is, to compute a vector $\vec{l} = (l_1, \dots, l_n)$ so that $\sum_{i=1}^n l_i = N - K$. Three major attribution methods exist:

- **Additive rule:** here, \vec{l} is computed completely independently of the number of deputies already elected for each party (vector \vec{k}) by ensuring that \vec{l} is proportional to \vec{q} .
- **Compensatory rule:** here, \vec{l} is computed so as to make it proportional to the vote deficit incurred by each party. This deficit d_i for party i is the difference between the number of votes that this party should have obtained in a purely proportional election, and the number of deputies k_i already elected, namely: $d_i = \max(0, \frac{q_i \times K}{Q} - k_i)$.²¹
- **Corrective rule:** the corrective rule can be seen as an intermediate between the additive rule and the compensatory rule. Here, \vec{l} is computed using the number of voters that are not represented by any majoritarian deputy. Formally, if \overline{D}_i is the set of districts in which candidate of party i has *not* been elected using the majoritarian rule, and if $q_i(d)$ is the number of votes for party i in district d , then, l_i should be proportional to $\sum_{d \in \overline{D}_i} q_i(d)$.

All these attribution methods can be associated to a *participation threshold*, that limits the participation to the proportional part to the parties that have reached this threshold.

Rounding methods The careful reader might have noticed that using the proportional attribution rules might lead to fractional numbers of seats. The usual way to overcome this difficulty is to round down these fractions and attribute the remaining seats. There are once again two methods to attribute these remaining seats:

- **Highest average method (D'Hondt method):** for each party, the ratio between the number of votes obtained by a party and its number of seats obtained so far is computed. The party having the highest such ratio obtains an additional seat. This procedure is iterated until the number of seats is exhausted.
- **Highest remainder method (Hare method):** here, for each party i we compute the difference between the proportion that i should obtain and the lower rounding of this proportion. Then, the remaining seats are attributed to the parties in decreasing order of this difference.

3.2 Predicting a second round

As explained earlier, one of the major challenges of our simulations was to predict the result of virtual second rounds that never happened, because changing the number of deputies elected with the majoritarian rule requires to change the number of electoral districts and hence to change the electoral boundaries as well. Suppose that in our simulation, a new (virtual) district d is composed of some voters from official

²¹As can be seen from the definition, parties that are already over-represented are excluded from this phase. For high rates of proportionality, this could even result in an inversion, where initially under-represented parties become over-represented.

districts d_1, \dots, d_k . Simulating the results of the first round for d is easy, because we know which proportion of voters have voted for the different parties in the official election. However, the candidates that qualify for the second round in d might be very different from the candidates that qualify in d_1, \dots, d_k . Hence, we have no information on the behaviour of these voters when they are faced with the second round in d .

To deal with this issue, we have proposed and compared four approaches.

Naive predictor The first method we have proposed is a naive predictor, that consists in declaring winner the candidate with the highest first round score, as if there all the persons that have voted in the first round for a candidate which is not present in the second round abstained. Of course, this is highly unrealistic. However, we keep this method as a baseline for comparison.

Single-peaked predictor The second method refines the naive predictor by considering that the voters preferences are single-peaked. Namely, we consider here that the different parties are distributed uniformly along a 17 level left-right axis with EXG (*EXtrême Gauche*) at the leftest level and EXD (*EXtrême Droite*) at the rightest (see Figure IV.4, where the parties are represented from left to right). We suppose that a voter whose first round candidate is present in the second round votes for this candidate (like for the naive predictor). However, we now suppose that a voter whose first round candidate is not present in the second round votes for one of the two closest parties in the 17 level scale, with a probability that is inversely proportional to the distance between the voter and these two parties. For instance, a voter for ECO that has to face a second round opposing SOC to LR will have a probability 1/3 to vote for LR in the second round and 2/3 to vote for SOC (because the distance between ECO and LR is twice as much as the distance between ECO and SOC).

Empirical predictor Our third prediction method uses a voice transfer matrix that has been built by Bruno Cautrès, CNRS researcher in political science at CEVIPOF. This matrix, based on a poll that has been made from 2015 to 2017, and has concerned 25000 voters in France, specifies, for each possible second round, how the persons that have voted for a party in the first round transfer their vote to another party for the second round. This predictor is probably the most promising one, but we could only develop a basic version of it: a careful tuning would have required a closer collaboration with the CEVIPOF (that was not possible, due to the very limited time we had).

Statistical predictor Our last predictor tries to learn the vote transfers from the real data we have. More precisely, this model assumes that, for a given second round configuration, the score of each party during the second round is a linear function of the score of all parties during the first round. Our method just tries to learn the coefficients of these linear functions (one linear function per second round configuration). Given that there can be plenty of different second round configurations (and some of them have too few examples, even no example at all), we actually group some of them into clusters of similar configurations (distinguishing "major" parties from the others). In the end, we only distinguished 8 configuration clusters, for which we learned the associated linear functions.

Comparison of the predictors Table IV.4 gives the prediction error rate on the 2017 real data, expressed as the number of districts inaccurately predicted, divided by the total number of districts. The

statistical linear predictor clearly gives the best results.²² This is why we will use only this predictor in our simulations. Note that the joint use of the linear and the empirical predictor could be a promising improving step.

Prediction model	Error rate
Naive predictor	23.22 %
Single-peaked predictor	26.17 %
Empirical predictor	21.32 %
Linear (statistical) predictor	11.09 %

Table IV.4 — Error rate of each vote transfer prediction method, in percentage of the total number of districts.

3.3 Redistricting

As we have mentioned earlier, reducing the number of deputies elected with the official majoritarian rule implies to change the number of electoral districts, and hence to change the electoral boundaries.²³ In our simulations, we have used three different approaches.

Geographic merging approach The first redistricting method we propose uses the geographic boundaries of the official electoral districts. The method operates by merging districts within the same department. More precisely:

1. for each department, we compute the target number by applying the reduction rate uniformly;²⁴
2. inside each department, we merge the old districts uniformly so as to obtain the target number of districts.

To obtain a realistic electoral map, during the merging step (2), we use the district adjacency graph to ensure that the merged districts are connected. In other words, we compute the adjacency graph of a department (each node being a district, and edges denoting adjacent districts, see Figure IV.5), and we partition this graph into t classes (t being the target number of districts). Since this is an instance of the NP-complete balanced k -connected graph partitioning problem, we use a greedy approximation algorithm to compute these partitions. In practice, for a 30 % reduction scenario, our algorithm creates in average 1.5 disconnected districts out of 404.

As a variant of this approach, we have also applied the exact same merging algorithm using a finer-grained administrative division, *cantons*, as the starting point, instead of electoral districts. This division is composed of 2090 territorial units (instead of 577 for the electoral districts). We can use them because we also have in the official datasets the results at the cantons scale.

²² Although a form of cross-validation should be used to confirm the performance of this predictor.

²³ Note that we have only simulated the cases where the number of electoral districts is *reduced*. These methods do not apply to the case where this number increases.

²⁴ We round down these target numbers and attribute the reminding districts (if any) uniformly at random among the departments).

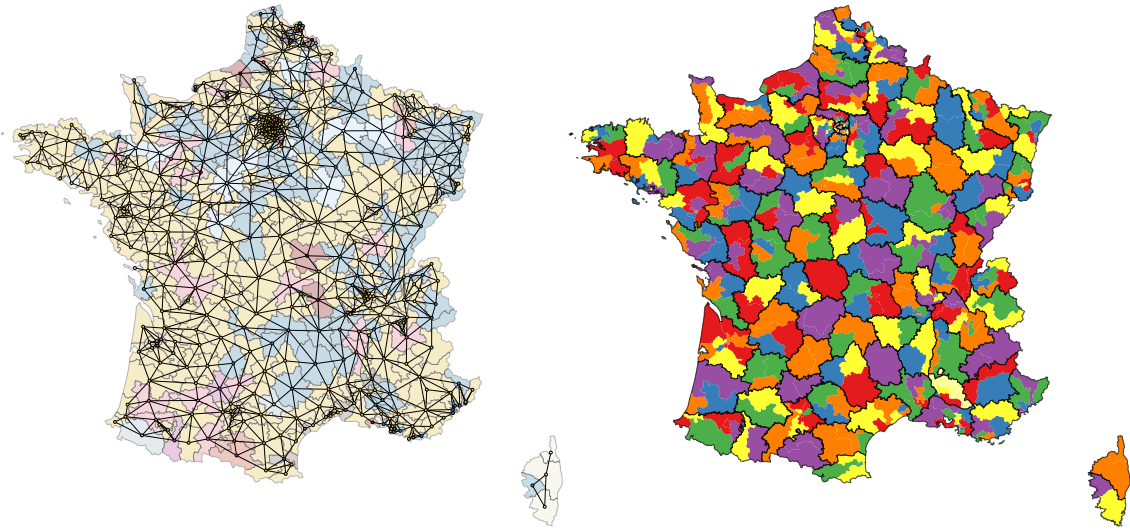


Figure IV.5 — Left: the French official electoral districts with underlying adjacency graph. Right: an example of electoral redistricting with 404 districts instead of 577, obtained by graph-based merging inside departments.

Statistical approach In the previous approach, the idea was to start from the actual districts (or cantons) and to merge them to obtain the desired number of districts. Here, we use a completely different approach. The idea is to create a realistic statistical model of the French electoral districts, and then to generate a completely artificial set of districts that fits our needs, according to this statistical model.

More precisely, we first establish a model of the typical electoral districts with 3 parameters per party p : (i) the probability z_p that no candidate of this party is present in this district, (ii) the average μ_p and (iii) the standard deviation σ_p of the scores (in percent) obtained by this party at the official election. Then, for each party p and each district to generate, we randomly draw a score according to a normal distribution with the same parameters μ_p and σ_p . We then withdraw this party from the district with probability z_p , and transfer its score to its allied parties.²⁵ Finally, we adjust the population size of each district by drawing it randomly according to a normal distribution, consistent with the distribution of the population in the actual districts.

Empirical approach The third redistricting method we use is based on a careful manual analysis of the possible district merging in several scenarios. A baseline scenario with 103 electoral districts has been proposed in a technical report by [Cohendet et al., 2018](#). Jérôme Lang, one of the authors of this report, has also proposed another version of this manual redistricting for a scenario with 344 districts. The main drawback of this approach is that it is valid only for a fixed target number of districts (contrary to the two previous approaches, that can be use to simulate any scenario). However, since it has been created using empirical political knowledge, this approach (for these two scenarios) has more chances to be realistic.

Comparison of redistricting methods Our experiments show that the different redistricting methods can make the results slightly vary. For instance, in a purely majoritarian scenario with 344 districts, the number of deputies obtained by REM varies between 190 and 210 depending on the redistricting

²⁵These potential alliances are determined using the conditional probability that a party is not represented in a district, given that another is represented.

method used. Without prescribing any of the method, we observe that the method based on geographical merging of *cantons* with connection constraint yields more realistic and homogeneous districts. This is why we chose to focus our simulations on this redistricting method only.

3.4 Simulations

We now give a few insights on the results of the scenarios we simulated for different voting rules. Each simulation depends on three parameters: (i) the electoral boundaries; (ii) a voting rule; (iii) a second round predictor. The number of possible scenarios being combinatorial, we will only give results for a few of them. More detailed results including other scenarios can be found in our technical report (Blanch and Bouveret, 2018).

For each scenario, we do not give exhaustively the number of votes for each party – the detailed results can be found in our report (Blanch and Bouveret, 2018) – but rather use two different representations. The first one is a graphics visualization proposed by Renaud Blanch, that represents the different parties as circular portions of different areas (see Figure IV.6). In this representation, the length of the circular arc of a party p is proportional to the number of seats $n_s(p)$ it obtains. As for the radius, it is proportional to the ratio $n_v(p)/n_s(p)$ between the number of voters for this party ($n_v(p)$) and the number of seats it obtains (in other words, this ratio is the number of voters per elected deputy for p). Hence, the total area of party p 's circular portion is proportional to the number of voters for p . The bold half-circle (the third one starting from the center) represents the purely proportional election: parties under this limit are *over-represented* (they have more seats than their proportion), and parties over this limit are *under-represented*. We see in particular in Figure IV.6 that in the 2017 official composition of the assembly, parties like ECO or FN are highly under-represented, while the presidential party REM is over-represented.

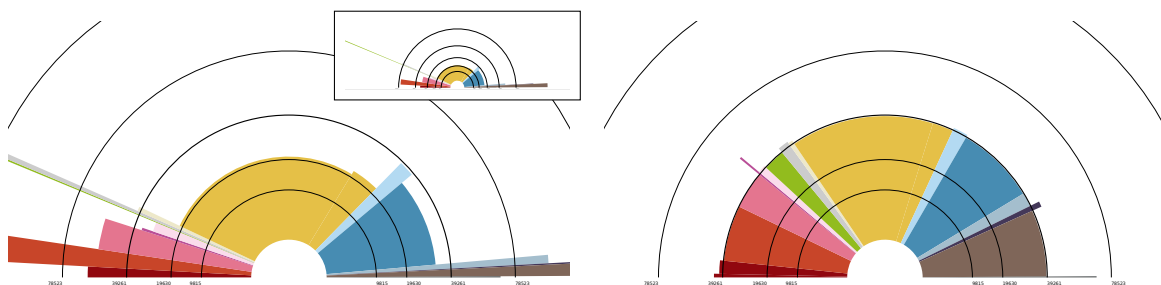


Figure IV.6 — Composition of the national assembly for the official results (left), and for a fully proportional voting rule on 577 electoral districts (right).

The second indicator we use to represent the results is an index introduced by Loosemore and Hanby, 1971 to quantify the distance to purely proportional representation. This index is formally defined as follows:

$$q \stackrel{\text{def}}{=} 1 - \frac{1}{2} \sum_{p \in \mathcal{P}} \left| \frac{n_v(p)}{\sum_{p' \in \mathcal{P}} n_v(p')} - \frac{n_s(p)}{\sum_{p' \in \mathcal{P}} n_s(p')} \right|$$

In this equation, the number inside the sum represents the distance between the fraction of seats that a party p should have received in a purely proportional representation, and the fraction of seats actually obtained by p . In a purely proportional representation, this index is equal to 1, whereas in an election

where n parties receive exactly the same number of votes, but only one party receives all the seats, q is close to 0 (namely, in this case, $q = 1/n$).

Equivalently, this index q represents the fraction of deputies that would keep their seats in we had to turn the assembly into a purely proportional one. Reciprocally, $1 - q$ represents the proportion of deputies that would lose their seats in such a situation.

Impact of the rounding method For each simulation, we have tested the difference between the two rounding methods introduced in Section IV.3.1. It turns out that the average difference between both methods is around 0.2 and 0.6 seats for each party (and the maximum difference observed for a given party does not exceed 2 seats). It thus seems that the rounding method only has a limited impact on the number of seats. This is why we chose to focus only on one of the two rounding methods: the d'Hondt method.

Mixed voting without redistricting The first scenario we explored is the case where we keep the number of electoral districts (577), but we simply increase the number of seats in the assembly by adding some deputies elected using the proportional method. The advantage of this scenario is that the simulation does neither depend on our redistricting methods nor on a second round predictor. Figure IV.7 shows the results obtained for 15 % of proportional seats, for the three mixed rules explained in Section IV.3.1 (additive, compensatory, corrective). The figure clearly shows that there is a significant difference between these voting rules, with a much higher effect of the proportional seats on the representation for the compensatory rule, that benefits to highly under-represented parties. This effect is observed for each amount of proportionality tested (10 %, 15 %, 20 % and 30 %).

We have also tested the influence of the participation threshold on the proportionality of the assembly. Table IV.5 gives the proportionality index for each voting rule and threshold configuration for a scenario with 15 % proportionality. We can observe that the introduction of a participation threshold has also a significant impact on the proportionality.

	0 % threshold			5 % threshold			10 % threshold		
	add.	corr.	comp.	add.	corr.	comp.	add.	corr.	comp.
q	0.72	0.75	0.80	0.71	0.74	0.79	0.70	0.73	0.78

Table IV.5 — Influence of the participation threshold for a scenario with 678 seats (15 % of proportional seats)

Reduction of the number of districts The second scenario we tested is exactly the opposite of the previous one. Here, we do not change the voting rule (no proportionality), but we only study the influence of a reduction in the number of districts. What we observe on the results shown in Table IV.6 is that the influence of reducing the number of seats on the proportionality is moderate. There is a slight reduction of proportionality with the reduction of the number of seats for artificially generated districts.

The 404 scenario The last scenario that we tested is a scenario with 404 deputies, including 344 that are elected using the majoritarian rule, and 60 that are elected using the proportional rule. It corresponds to a scenario with 30 % reduction of the number of deputies and 15 % proportionality that was judged

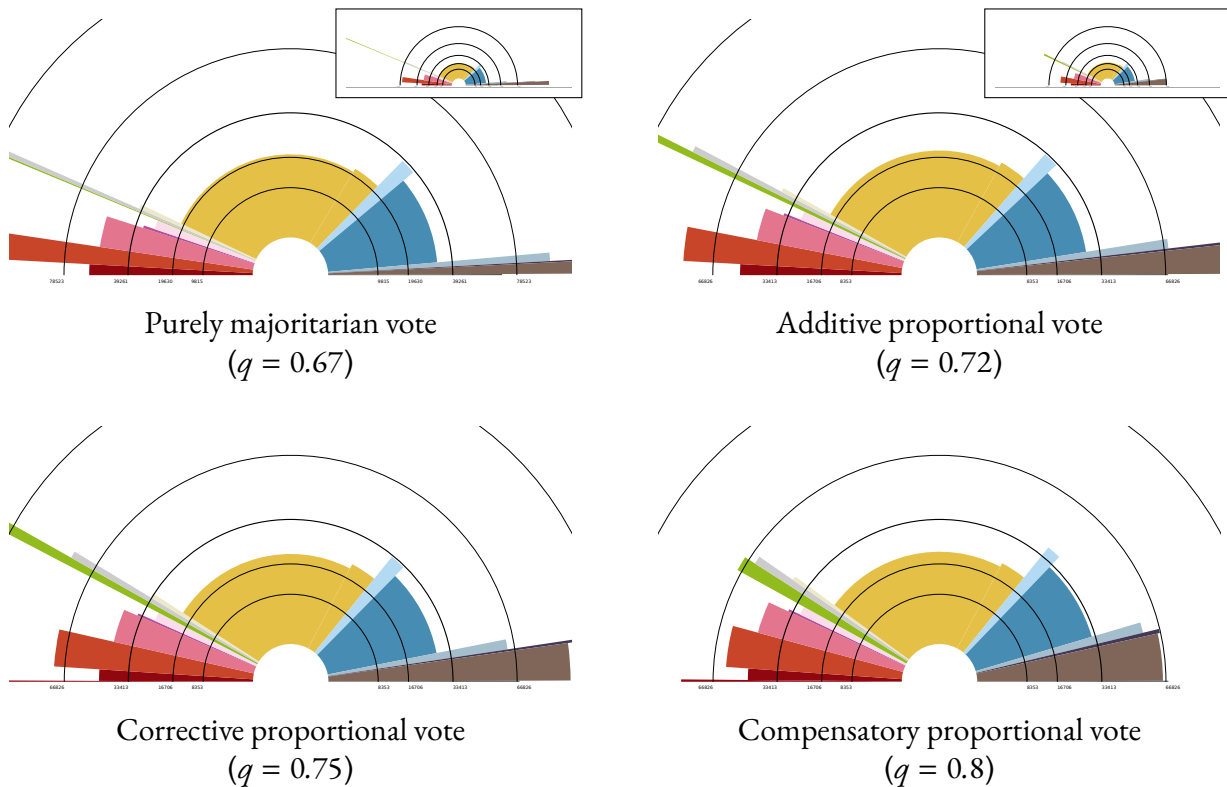


Figure IV.7 — Composition of the assembly for a scenario with 678 seats (15 % of proportional seats) for different mixed voting rules, and comparison with the actual assembly. No participation threshold used.

	577 dist.	404 dist.	364 dist.	344 dist.	323 dist.	303 dist.
q (geog. merging)	0.67	0.69	0.68	0.68	0.67	0.68
q (stat. generation)	0.67	0.66	0.65	0.65	0.65	0.65

Table IV.6 — Influence of the reduction of the number of districts, with two redistricting methods: geographical merging of *cantons* (with connection constraint), and statistical generation of artificial districts. Average on 1000 random draws.

credible after the governmental announces concerning this reform of the legislative election. Without developing much this scenario, our simulations show that the proportionality index q varies between 0.64 (manual redistricting) and 0.73 (statistical generation) for the additive rule, and between 0.71 (manual redistricting) and 0.77 (statistical generation) for the compensatory rule. In other words, we see two antagonistic effects of using a mixed rule that tends to increase the proportionality of the assembly, while reducing the number of electoral districts tends to reduce it. The manual redistricting elaborated by Jérôme Lang yields a much less proportional assembly, probably because the political analysis of plausible redistricting behaviours favours the big parties at the detriment of smaller (already under-represented) ones.

3.5 Conclusions

We can make several observations from our simulations. First, the rounding method has only a marginal effect that could be neglected in first approximation. Second, even with 25 % proportionality, the effect on the composition of the assembly is visible but quite moderate, especially with the additive rule. Third, the reduction of the number of seats entails a slight reduction of the proportionality of the resulting assembly. For the 404 scenario, these two effects combine with each other: increasing the proportionality of the rule mechanically increases the proportionality of the assembly, but reducing the number of seats decreases it.

What about 2024? The political landscape of the June 2024 anticipated legislative election is very different from the one in 2017. First and foremost, 23 *nuances politiques* were engaged instead of 17 (see Figure IV.8), according to the classification defined by the Ministry of Interior.²⁶ Second, contrary to the 2017 situation, the election was not dominated by a highly majoritarian party but the electorate was rather broken into several main blocks, among which the *Nouveau Front Populaire* (designated as UG) on the left wing, *Ensemble* (designated as ENS) – the presidential party –, and *Rassemblement National* (designated as RN) for the extreme right. The parties had a huge incentive to form coalitions like UG, and a lot of *nuances* appearing in Figure IV.8 just represent lonely candidates that refused to appear as a member of a coalition (for instance candidates from the greens – VEC – or *La France Insoumise* – FI – that refused the coalition UG).



Figure IV.8 — Political parties (*nuances*) engaged in the 2024 legislative election.

Figure IV.9 shows the official composition of the assembly and compares it with the computed composition for a mixed compensatory rule with 15 % proportionality and a mixed rule with 30 % proportionality (no redistricting). What is striking is that the official assembly is much more proportional ($q = 0.88$) than it used to be in 2017 ($q = 0.67$). And it can reach even higher rates of proportionality by adding a few additional proportional seats. We can also observe that the compensatory rule here produces an inversion for some parties: parties that were under-represented (resp. over-represented) with the official rule can become over-represented (resp. under-represented) by adding a slight touch of proportionality.

We still did not have time to simulate redistricting for the 2024 legislative election. Such a simulation will probably be much more delicate than for the 2017 election, due to the particular political situation where (i) a higher number of second rounds involved 3 parties or more, and (ii) in those situations, a lot of candidates (but not all) from UG, ENS or LR chose to withdraw to prevent RN to win the electoral district. Simulating such second rounds would probably require more than a simple linear predictor, and here, a close collaboration with political scientists would be very useful.

²⁶<https://www.resultats-elections.interieur.gouv.fr/legislatives2024/referentiel.html>

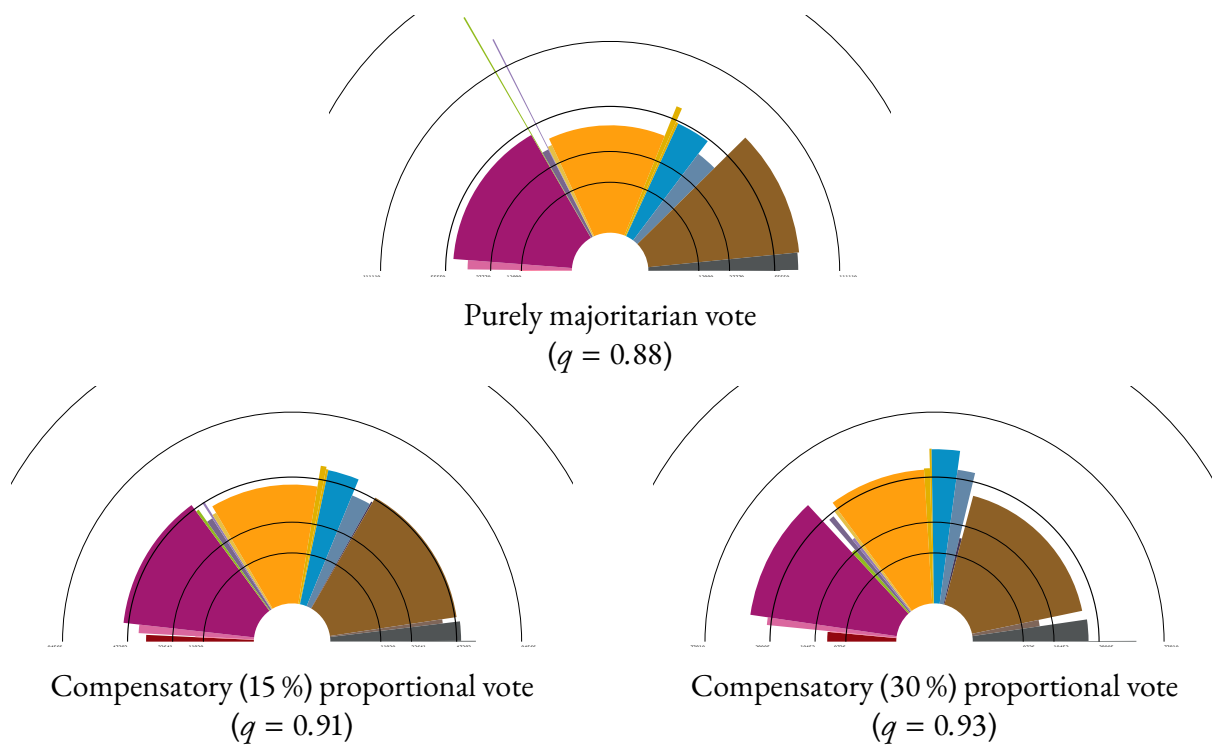


Figure IV.9 — Composition of the 2024 national assembly for the official results (top), for a compensatory proportional vote with 15 % proportionality (bottom left) and 30 % proportionality (bottom right). No redistricting.

4 Conclusion

In this chapter, we tried to make a contribution to the societal debate of knowing what kind of voting rules we want for the major political elections. Of course, our aim is not to prescribe anything. For that matter, it would be probably difficult to find any consensus on the best voting rule even in the scientific community.²⁷ We rather aim to show (i) the impacts of using alternative voting rules on the results of the election, and (ii) that alternative voting rules that, for some of them, have never been used before in political contexts, can actually be used in practice. Regarding the second point, our experiments conducted during the French presidential election, involving actual citizens, strikingly show that alternative voting rules can indeed be used in practice, they are well understood by voters, and might even be very much appreciated by them. Concerning the first point, both the simulations we have conducted on the French legislative election data and the experiments on the presidential election have shown that different voting rules indeed convey different forms of fairness principles, as they can result in quite different political landscapes. We hope that these studies will contribute to the political and societal debate, and will help designing a new form of democracy with which citizens feel much more in tune.

Even if we focused on high stake political contexts in this chapter, computational social choice can also be applied at low or intermediate stakes. For instance, designing new forms of platforms that help decision

²⁷ Actually, a vote was organized in the conference *Voting Power and Democracy* to determine which was the preferred voting rule by the community. The winner was Approval Voting – exactly the voting rule that was used to run the election (Laslier, 2011).

making for everyday life situations by proposing a range of voting rules adapted to those situations is a rich and still promising topic. This is the objective of the voting platform Whale,²⁸ that I have been developing for several years with the collaboration of Marie-Jeanne Natete, François Durand, Yann Chevalere and Alexis Hummel. The scientific questions raised by the development of such voting platforms on the web have been the central topic of a book chapter published in 2017 (Bouveret, 2017).

²⁸<https://whale.imag.fr/> and <https://whale5.noiraudes.net/>.

Chapter V

Conclusion: a fair and safe operating space for humanity

In this document, we investigated the concept of fairness in several different contexts of collective decision making. Here, as reminded in the introduction of this document, fairness was interpreted in a quite weak sense, namely, as the property of a collective decision (or a procedure) which is accepted and considered appropriate by all the community members, with regards to their individual needs, statuses and contributions (Young, 1994b). Hence, discussing the ethical and philosophical principles conveyed by our different notions of fairness is clearly relevant, but out of the scope of this Habilitation thesis. The question tackled in this thesis was rather to find collective decision making procedures that could be considered appropriate in different contexts and study their computational properties.

First, we have seen in Chapter II how this concept could be formally defined and declined to ensure endstate justice in the context of fair division of indivisible goods. That led us to sketch a somewhat comprehensive landscape of fairness properties, each property conveying a different interpretation of what a fair division should be, based *e.g.* on the satisfaction of a basic level of utility (Sections II.2 and II.3, based on Bouveret and Lemaître, 2014; Bouveret and Lemaître, 2015; Shams et al., 2021), on the limitation of the agents' knowledge (Section II.4, based on Aziz, Bouveret, Caragiannis, et al., 2018), or on a social consensus about individual envy (Section II.5, based on Shams et al., 2020; Shams et al., 2022). We have also shown that these properties are not independent, but related to each other by strong conceptual links.

Still in the same context of fair division of indivisible goods, we investigated in Chapter III a completely different approach, where fairness is not guaranteed *ex post*, but is ensured *ex ante* by the allocation procedure itself. The general procedure studied in this context, namely, picking sequences, is remarkably simple and intuitive, and yet ensures some form of fairness and efficiency while being partly distributed, just like a fully distributed allocated protocol like swap deals does. In this context, we investigated the problem of finding the fairest sequence, both for general ones (Section III.3, based on Bouveret and Lang, 2011), and non-interleaving strategyproof ones (Section III.5, based on Bouveret, Gilbert, et al., 2023), as well as the complexity of the manipulation problem (Section III.4, based on Bouveret and Lang, 2014; Aziz, Bouveret, Lang, et al., 2017). This finally led us to show that picking sequences are actually related to Pareto efficiency and cycle deals, and complete the landscape of fairness and efficiency properties for fair division (Section III.6, based on Beynier et al., 2019).

Lastly, we focused in Chapter IV on a completely different collective decision making context, namely

voting. Here, our approach was not to design new properties or procedures, as the landscape of the voting rules is already quite rich and extensive. We rather focused on the question of analyzing the effects of using alternative voting rules for past political elections. Our point of view is that understanding these effects is necessary in a society willing to consciously and collectively tackle the question of finding the right voting rule. This is a prerequisite for the proper functioning of democracy. To address this question, we used two different tools: first, online¹ and *in situ* experiments, for which we respectively gathered 37 739 and 6358 votes (Section IV.2, and Bouveret, Blanch, et al., 2018; Bouveret, Blanch, et al., 2019), and second, computer simulations (Section IV.3). These tools were respectively applied to the 2017 French presidential election and to the 2017 French legislative election. Our results show that changing the voting rules indeed has an effect on the political landscape, by favouring different kinds of candidates or parties.

Perspectives

Now, what could be the opportunities for future research on the topic of fairness in collective decision making? Before addressing this question, let us recall a few elements of context about the state of the world.

Humanity is currently facing a situation that can be considered unique in its history. Namely, the environmental impacts of human activities have grown so much since the beginning of the industrial revolution that they now represent a major driver of the Earth system, prompting the term *Anthropocene* to describe this new epoch (Steffen, Broadgate, et al., 2015).² Six of the nine *planetary boundaries* identified by the work of Stockholm Resilience Center’s researchers (Rockström et al., 2009; Steffen, Richardson, et al., 2015) and defined as *processes that are critical for maintaining the stability and resilience of the Earth system as a whole* are transgressed (Richardson et al., 2023), meaning that human-driven change not only influences the Earth system, but it also threatens to destabilize it. Whether we consider the on-going biodiversity collapse (Díaz et al., 2019) or global warming (Pörtner et al., 2022) — or in fact almost any planetary boundary — the situation is bleak.³

What should be the role of science in this complicated context? Here, several options are possible. The first one is to consider that science should continue as is independently of the world situation. The second one is to consider that we as scientists should question the way we practice it (for instance by reducing our professional carbon footprint) without changing the topics of our research. The third path followed by some researchers is based on the idea that not only the research practices should be questioned, but also the topics themselves.⁴ My personal opinion is that there are indeed strong arguments in favour of questioning the research practices as they are, but also that some research topics should be redirected, both because there is an urgent need to tackle environmental problems seriously (for that the contribution of everyone would be welcome) and because some research areas promote a model of society

¹<https://vote.imag.fr/>

²Even though the International Union of Geological Sciences (IUGS) has rejected the proposal to use the term *Anthropocene* as a formal unit of the Geologic Time Scale, it is still considered as relevant by Earth and environmental scientists, economists, and social scientists as a clear marker of human impact on the Earth system. It is also a well established term in the public discourse.

³This paragraph is extracted as is from the introduction of the formal project proposal *Anthropocène, Décroissance, Numérique* to which I contributed as a member of the future research team.

⁴Another option (not exclusive to the other ones), that I do not discuss here, is to redirect research to concrete actions that have actual impacts in the civil society, for instance in the form of non-violent civil resistance. This is *e.g.* the path followed by Scientist Rebellion (<https://scientistrebillion.org/>).

that clearly seems incompatible with planetary boundaries.

What about fairness issues in social choice in this context? Several problems related to developing a model of society that stays within the environmental limits can actually be naturally expressed as formal social choice problems. I will now give some instances of the kinds of collective decision making problems that may arise naturally when dealing with environmental issues. In this sense, those problems are natural opportunities for computational social choice to contribute to solving these issues.

Of fair division of scarce(r) resources

The first kind of problems I want to evoke is the problem of redistributing resources that become scarcer. Actually, a major part of climate change related problems (to only name one of the planetary boundaries) are due to the massive use of fossil energies (IPCC, 2021), namely, coal, oil and gas, that has never stopped increasing since the beginning of the industrial era (Frescoz, 2024). As we can see, the problem here is that these resources are much too abundant. More and more scientists, starting with those involved in the IPCC, state that the use of these resources should dramatically decrease if we want to have a chance to maintain the global temperature increase below 2 °C (IPCC, 2021). In the end, whether we consider carbon emissions, use of fossil energies, or some other material resource, the problem remains the same: humanity currently consumes too much of this resource; this consumption should decrease; hence the resource should become scarcer and we should decide how to reallocate it to the members of the society.

This resource reallocation problem is intricate, and fairness notions are not even clearly defined. Should the resource cut-offs be allocated uniformly in the society? Should richer agents (or countries) bear a larger share of the reduction? Should we take into account the current development needs? The past use of the resource? Deciding this at the global scale is a (geo)political problem that goes far beyond the scope of computational social choice. However, my claim is that this sort of problems will also happen more frequently at the local scale, where maybe social choice theory could help. Let us give an example inspired by a real application.

Fair division of carbon emissions The application concerns a research lab in computer science.⁵ It involves a finite set of agents $\mathcal{A} = (a_1, \dots, a_n)$ that represent the lab members. Each agent is characterized by a set of attributes that are supposedly either relevant to the problem, or sensitive. Here, the four attributes that we may consider are: gender, team, status (researcher, associate professor, PhD student...), permanent position (yes / no). The attributes may be dependent on each other. It could be for logical reasons related to hard constraints of the domain (e.g the permanent position attribute functionally depends on the status attribute), in which case we could even think of having a complex taxonomy characterizing the agents. But it may also be due to unwanted (and often harmful) correlations, like the fact that the observed proportion of women is negatively correlated to the status rise in academia.⁶

The lab members collectively decide to cut off the lab's carbon emissions by a certain amount. For a computer science lab, where a major part of carbon emissions are due to research trips, this comes down to drastically limiting the most impacting of those travels by imposing a quota. More precisely, to each scheduled trip μ_i is associated a total travel distance $d(\mu_i) \in \mathbb{N}$ and a means of transportation $t(\mu_i) \in Tr$, where Tr is the set of available means of transportation (bike, car, train, plane...). Moreover, to every

⁵The specific context of such a research lab is of course inspired by my own experience. However, note that this application can well be generalized to any (professional or not) micro-society having to face the same kind of issue.

⁶See e.g. [Ministère de l'Enseignement Supérieur, de la Recherche et de l'Innovation, 2021](#).

$t \in Tr$ is associated an impact factor $\omega(t)$ that gives the amount of carbon equivalent emitted in average by a person using transport t for one kilometer. For simplicity, we assume that the total amount of carbon emitted for trip μ_i only depends on the distance and the means of transportation used, and is given by the following expression: $d(\mu_i) \times \omega(t(\mu_i))$. Of course, each research trip is associated to one of the agents from \mathcal{A} . Finally, we may add to the research trips all the relevant characteristics for the problem at stake: the length of stay may be one of them, but there could also be some notion of priority, either for the lab, or for the concerned agent himself.

We assume that the lab has a total carbon emission quota Q defined for the entire year.⁷ The problem we have to solve can thus simply be defined as follows.

$$\left| \begin{array}{l} \text{Given a set of trips } \mathcal{M} = (\mu_1, \dots, \mu_m) \text{ and a global quota } Q, \text{ choose a subset of trips } \mathcal{S} \subset \mathcal{M} \\ \text{such that } \sum_{\mu_i \in \mathcal{S}} d(\mu_i) \times \omega(t(\mu_i)) \leq Q. \end{array} \right.$$

Note that since each trip is associated to one agent, we can either solve this problem by directly selecting the trips, or by splitting the quota and allocating a fraction of it to each agent (so that a trip will be selected only if the corresponding agent still has enough quota to support it). Hence, depending on the way we see it, this problem can be considered either to be a fair division problem (splitting the quota), or a selection problem (selecting the trips).

Of course, here, the difficulty is not just to find a feasible subset of trips. The difficulty is first and foremost to find a *fair* solution, which requires to be able to define what fairness is in this context. Unlike the notion of fairness investigated in this manuscript, here, it can obviously not be based on agents' opinions or preferences. Instead, it should be based on relevant features that all the lab members agree upon. For instance, it may be based on:

- agents' features, like seniority, with the idea that for instance junior researchers and PhD students should have a higher quota because travelling is more essential to them;
- research trips' features, like priority for the lab, with the idea that a higher priority trip should have more chances to be selected (for instance because the expected reward of this trip for the lab is higher), or length of stay, with the idea that the longer the stay is, the higher the chances are to yield fruitful collaborations.
- ...

We can also observe that, here, fairness should probably also be understood in the sense of *algorithmic fairness* (Kleinberg et al., 2018). Since the problem can be seen from the point of view of trip selection, ensuring that the result is not biased with respect to sensitive agent dimensions (e.g gender) is probably of primary importance. Note that depending on what lab members consider to be fair, some agent attributes (like, e.g status or seniority) may be considered either as relevant differences that should be taken into account in the allocation process, or as sensitive attributes with respect to which the selection process should be as little biased as possible.

A last aspect that should also be considered in the formalization of this problem is timing aspects. First, the problem is obviously repeated: the quota has to be allocated every year (and potentially decreases

⁷For simplicity, we may assume that all the research trips are associated to a given year (even if the stay overlaps the year change date).

every year). It probably makes the problem even more complicated, but is also an additional opportunity for fairness, since the agents' satisfaction can be averaged over different years (see *e.g.* [Lemaître et al., 1999](#), for an application of repeated fairness to the domain of space). We could even think of taking into account past emissions (before the cut-off) as an input of the problem, so as, *e.g.*, to give less entitlements to persons having emitted a lot in the past. Secondly, in the real application, the research trips are of course not known at the beginning of the year, but may randomly arrive during the year. The problem is thus not to select a subset of trips that are known beforehand, but rather to decide, for each new trip arriving, whether it is feasible given the quota constraint and will be accepted or not. This problem hence looks more like an *online* fair division (or selection) problem ([Aleksandrov and Walsh, 2020](#)).

As we can feel from this simple setting description, there is plenty of room for complex modeling and computational problems with a real application. However, being able to apply complex computational techniques to a formal model first necessitates to solve a quite more intricate problem: make people collaborate and agree upon the carbon emission cut-off implementation. In the context of a lab, a prerequisite is that people agree on a few facts that are stated below.

1. We have a serious environmental problem, and mankind is accountable for it.
2. Speaking only of climate change, reducing anthropogenic emissions is the major part of the solution to this problem.
3. Research labs must contribute and reduce their carbon footprint.
4. This necessary reduction will have a direct impact on everyone's research practices.
5. The way this reduction is borne by different members of the lab must be decided collectively.

Our experience shows that these assumptions are not self-evident. This is confirmed by a recent survey ([Blanchard et al., 2022](#)) that shows that even among a concerned research community, a lot of people are really not ready to accept the drastic changes that would nevertheless be necessary to globally reduce the sector's carbon footprint. However, collectively agreeing on them is a prerequisite for any supposedly fair allocation mechanism of carbon credits to be accepted by the stakeholders.

Managing the commons

As explained above, one of the big challenges I see for the anthropocene is the one of sharing resources whose available quantity is diminishing every day. One way to consider this issue is to adopt – like we did above – a private ownership perspective: resources (like carbon emissions) are considered to have private use, even though this use has (negative and positive) externalities. Another point of view would be that consuming fewer resources implies switching from private use to *common use*. Switching to a model of society which is neither directed by the market laws nor by an authoritarian state but where small communities of agents collectively and locally pool resources and decide how to use them in common ([Ostrom, 1990](#)) paves the way for a rich line of works in social choice.

The first question that arises when switching to the model of commons is the question of governance. Here, if we follow Elinor and Vincent Ostrom's principles for managing the commons, the model of governance that seems intrinsically adapted to this situation seems to be polycentric governance, that is, a form of governance that involves multiple semi-autonomous, interacting decision centers ([Carlisle](#)

and Gruby, 2019). This is particularly relevant for large Information and Communication Technology commons like Wikipedia (Forte et al., 2009) or other open source projects. The question of how to make decisions in these contexts arises naturally and is even a key success factor of these initiatives.⁸

Beyond the question of governance and collective decision making methods related to the management of the commons, another question that arises naturally when some resources are pooled (or some infrastructure is shared) within a group to be used in common is the question of dividing the costs of building, operating or maintaining these resources. The concrete example we may give is inspired by the work of our PhD student Nicolas Besson, supervised in common with Nadia Brauner and Nicolas Brulard. This work is rooted in a practical application to a decision support system for collaborative transportation in an association of farmers in a Short Food Supply Chain (see Besson et al., 2024, for a more complete description of the application). The association is composed of a group of farmers (denoted by A), sharing a vehicle (that we assume to have unlimited capacity for the sake of simplicity), and a client (denoted by a_0), where their production must be brought whenever they receive an order. The tour to pick up the production at the farms and to deliver it to the client is a Traveling Salesman Problem where $A \cup \{a_0\}$ is the set of locations to be covered. The transportation cost must be paid by the farmers only, on a fair basis.

We are clearly in a context where the resource (the vehicle) has not exactly a private use, but is pooled to be used in common. Simply, this resource has an operating cost, and this cost must be fairly split among the agents. Several approaches are possible. We can for instance suppose that the operating cost is a negative resource (a divisible chore) that must be shared among the agents and resort to usual solution concepts of fair division theory, like proportionality or envy-freeness. The difficulty here is that the agents have different utilities for using the resource (typically, this corresponds to the money they hope to earn by selling their products to the client), but also unequally contribute to the costs, in the sense that a farmer located further away from the client will induce higher costs to the association.⁹ Another possibility is to view the problem from the perspective of *cooperative game theory*, where it can be modeled as a *Transferable Utility game* (Shapley, 1953). In this setting, fairness is conveyed by stability notions: we will typically look for a payment scheme such that (i) no player has any incentive to leave the association and buy its own means of delivery (individual rationality), but also (ii) no sub-group of agents has an incentive to do so as well. This concept of stability corresponds to the seminal concept of the *core* of the game.

Modeling this problem gives rise to several difficult questions. The first one is computational: even if the concept of the core seems to be appropriate in a context where farmers should be incentivized to participate to the association, it is very complex to compute: this comes back to solving a linear program of n variables and $2^n + n + 1$ constraints (Rothe, 2015). Moreover, the core of a game can be empty, even in the special case of Traveling Salesman Games¹⁰ (Potters et al., 1992). This concretely means that this fairness notion is probably too strong in this context (just like envy-freeness for multiagent resource allocation problems), and may be inappropriate in practice. Several relaxations exist, like the *semicore* (Young, 1994a) and its relaxations. Although being less computationally complex, the question of computing it might also be an issue in practice, especially since the computation of the cost function itself

⁸This can be demonstrated by the active discussions on the voting methods to use in these communities that are often far ahead of the methods used for political elections at the state level. See for instance the seminal Schulze voting rule used in several open source communities, including the Debian project.

⁹The same problem arises in the context of building a power supply infrastructure for instance, where isolated users induce higher costs to the society – usually without being charged more than the others.

¹⁰A Traveling Salesman Game is a Transferrable Utility game where players are nodes in a weighted graph and the cost value associated to each coalition S of players is equal to the shortest Hamiltonian cycle between these nodes.

requires to solve several instances of the Traveling Salesman Problem.

But beyond these computational questions, the most fundamental ones are probably also to know whether the fairness (or stability) concepts introduced really make sense from an ethical point of view, and more basically whether they are considered appropriate by the stakeholders (the farmers) themselves. The underlying central question is the place we want to give to a digital decision making application in the collective decision process of the association of farmers. Would they accept the solutions the application computes and proposes? Should we explain those proposed solutions, or leave some information hidden to some of the farmers for the sake of privacy and acceptability? Should the tool rather be used as a mean for the farmers to understand the system in which they evolve and as a medium to help them discussing the possible solutions while delegating the computational burden to the machine? In the end, is it even desirable to make the farmers dependent on a digital tool or should we design the tool to be *self-obviating*?¹¹ All this questions (on top of computational questions) are also on the agenda of Nicolas Besson's PhD thesis.

What can we conclude from the discussion above? Obviously, social choice theory can have a role in the construction of a desirable society, that fits within the planetary boundaries but also ensures minimal social guarantees to everyone (Raworth, 2017). A lot of problems arising in this context actually seem to be coming down to resource sharing problems and more generally collective decision making problems.¹² This is good news for social choice theory, which, I believe, can enlighten the debate and actively contribute to the construction of such a desirable society.

However, our experience also shows that there can exist a big gap between theory and practice. A theoretical model imposed without a proper discussion with the stakeholders is essentially useless. One of the biggest challenges in this context will probably be to bridge this gap and find methodologies that both involve the researchers and the stakeholders at the design phase. In other words, methodologies that seek "*transformative change through the simultaneous process of taking action and doing research, which are linked together by critical reflection.*"¹³ Hopefully, this is a line of research that I will be able to develop in the future research team *Anthropocène, Décroissance, Numérique*, whose project is currently under active construction.

¹¹Self-obviating systems (Tomlinson et al., 2015) are systems where ICT technologies are used initially – typically for learning purposes – and then progressively discarded in everyday use by humans, which can still rely on the persistent results of previous computation.

¹²The implicit assumption is that an authoritarian society does not match the definition of a *desirable* society.

¹³Wikipedia entry on *Action Research*.

Bibliography

- Abebe, Rediet, Jon M Kleinberg, and David C Parkes (2017). “Fair Division via Social Comparison”. In: *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 281–289.
- Akrami, Hannaneh, Noga Alon, Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, and Ruta Mehta (2022). *EFX Allocations: Simplifications and Improvements*. arXiv: 2205.07638 [cs.GT]. URL: <https://arxiv.org/abs/2205.07638>.
- Akrami, Hannaneh and Jugal Garg (2024). “Breaking the 3/4 barrier for approximate maximin share”. In: *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, pp. 74–91.
- Akrami, Hannaneh, Jugal Garg, Eklavya Sharma, and Setareh Taki (2023). “Simplification and Improvement of MMS Approximation”. In: *32nd International Joint Conference on Artificial Intelligence, IJCAI 2023*, pp. 2485–2493.
- Aleksandrov, Martin and Toby Walsh (2020). “Online fair division: A survey”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 34. 09, pp. 13557–13562.
- Amanatidis, Georgios, Georgios Birmipas, George Christodoulou, and Evangelos Markakis (2017). “Truthful allocation mechanisms without payments: Characterization and implications on fairness”. In: *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 545–562.
- Amanatidis, Georgios, Georgios Birmipas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A Voudouris (2021). “Maximum Nash welfare and other stories about EFX”. In: *Theoretical Computer Science* 863, pp. 69–85.
- Amanatidis, Georgios, Evangelos Markakis, Afshin Nikzad, and Amin Saberi (Dec. 2017). “Approximation Algorithms for Computing Maximin Share Allocations”. In: *ACM Trans. Algorithms* 13.4. ISSN: 1549-6325. DOI: 10.1145/3147173. URL: <https://doi.org/10.1145/3147173>.
- Arrow, Kenneth (1951). *Social choice and individual values*. revised edition 1963. John Wiley and Sons.
- Asadpour, Arash and Amin Saberi (2007). “An approximation algorithm for max-min fair allocation of indivisible goods”. In: *Proc. of STOC-2007*.
- Atelier de Cartographie de Sciences Po (2017). *Carte des circonscriptions législatives 2012 et 2017*. <https://www.data.gouv.fr/fr/datasets/carte-des-circonscriptions-legislatives-2012-et-2017/>. Accessed on April 4th, 2018.
- Aziz, Haris, Peter Biro, Jérôme Lang, Julien Lesca, and Jérôme Monnot (2016). “Optimal Reallocation under Additive and Ordinal Preferences”. In: *International Conference on Autonomous Agents And Multiagent Systems (AAMAS’16)*, pp. 402–410.
- Aziz, Haris, Sylvain Bouveret, Ioannis Caragiannis, Ira Giagkousi, and Jérôme Lang (2018). “Knowledge, Fairness, and Social Constraints”. In: *Proceedings of the 32nd AAAI conference on Artificial Intelligence (AAAI’18)*. New Orleans, Louisiana, USA: AAAI Press. URL: <http://recherche.noiraudes.net/resources/papers/AAAI18.pdf>.

- Aziz, Haris, Sylvain Bouveret, Jérôme Lang, and Simon Mackenzie (2017). “Complexity of Manipulating Sequential Allocation”. In: *Proceedings of the 31st AAAI conference on Artificial Intelligence (AAAI’17)*. San Francisco, California, USA: AAAI Press. URL: <http://recherche.noiraudes.net/resources/papers/AAAI17.pdf>.
- Aziz, Haris, Thomas Kalinowski, Toby Walsh, and Lirong Xia (2016). “Welfare of Sequential Allocation Mechanisms for Indivisible Goods”. In: *22nd European Conference on Artificial Intelligence (ECAI 2016)*. Ed. by Gal A Kaminka, Maria Fox, Paolo Bouquet, Eyke Hüllermeier, Virginia Dignum, Frank Dignum, and Frank van Harmelen. IOS Press, pp. 787–794.
- Aziz, Haris, Hervé Moulin, and Fedor Sandomirskiy (2020). “A polynomial-time algorithm for computing a Pareto optimal and almost proportional allocation”. In: *Operations Research Letters* 48.5, pp. 573–578. ISSN: 0167-6377.
- Babaioff, Moshe, Noam Nisan, and Inbal Talgam-Cohen (2021). “Competitive Equilibrium with Indivisible Goods and Generic Budgets”. In: *Mathematics of Operations Research* 46.1, pp. 382–403.
- Bansal, Nikhil and Maxim Sviridenko (2006). “The santa claus problem”. In: *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pp. 31–40.
- Barbanel, Julius B, Steven J Brams, and Walter Stromquist (2009). “Cutting a pie is not a piece of cake”. In: *The American Mathematical Monthly* 116.6, pp. 496–514.
- Barberà, Salvador, Walter Bossert, and Prasanta K Pattanaik (2004). “Ranking Sets of Objects”. In: *Handbook of Utility Theory*. Vol. 2. Kluwer, pp. 893–977.
- Barman, Siddharth and Sanath Krishnamurthy (July 2019). “On the Proximity of Markets with Integral Equilibria”. In: *Proceedings of the AAAI Conference on Artificial Intelligence* 33, pp. 1748–1755. DOI: [10.1609/aaai.v33i01.33011748](https://doi.org/10.1609/aaai.v33i01.33011748).
- Barman, Siddharth and Sanath Kumar Krishnamurthy (2020). “Approximation algorithms for maximin fair division”. In: *ACM Transactions on Economics and Computation (TEAC)*. Vol. 8. 1, pp. 1–28.
- Bartholdi, John J, Craig A Tovey, and Michael A Trick (1989a). “The computational difficulty of manipulating an election”. In: *Social choice and welfare* 6, pp. 227–241.
- (1989b). “Voting schemes for which it can be difficult to tell who won the election”. In: *Social Choice and welfare* 6, pp. 157–165.
- Baujard, Antoinette, Frédéric Gavrel, Herrade Igersheim, Jean-François Laslier, and Isabelle Lebon (June 2014). “Who’s Favored by Evaluative Voting? An Experiment Conducted During the 2012 French Presidential Election”. In: *Electoral Studies* 34, pp. 131–145. DOI: [10.1016/j.electstud.2013.11.003](https://doi.org/10.1016/j.electstud.2013.11.003). URL: <https://hal.science/hal-00803024>.
- (Dec. 2018). “How voters use grade scales in evaluative voting”. In: *European Journal of Political Economy* 55, pp. 14–28. DOI: [10.1016/j.ejpoleco.2017.09.006](https://doi.org/10.1016/j.ejpoleco.2017.09.006). URL: <https://shs.hal.science/halshs-01618039>.
- Baujard, Antoinette, Herrade Igersheim, and Isabelle Lebon (May 2021). “Some regrettable grading scale effects under different versions of evaluative voting”. In: *Social Choice and Welfare* 56.4, pp. 803–834. DOI: [10.1007/s00355-020-01300-z](https://doi.org/10.1007/s00355-020-01300-z). URL: <https://hal.science/hal-03095898>.
- Baumeister, Dorothea, Sylvain Bouveret, Jérôme Lang, Nhan-Tam Nguyen, Trung Thanh Nguyen, Jörg Rothe, and Abdallah Saffidine (May 2017). “Positional scoring-based allocation of indivisible goods”. In: *Autonomous Agents and Multi-Agent Systems* 31.3, pp. 628–655. ISSN: 1573-7454. DOI: [10.1007/s10458-016-9340-x](https://doi.org/10.1007/s10458-016-9340-x). URL: <http://recherche.noiraudes.net/resources/papers/JAAMAS16.pdf>.
- Bei, Xiaohui, Youming Qiao, and Shengyu Zhang (2017). “Networked Fairness in Cake Cutting”. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 3632–3638.

- Besson, Nicolas, Sylvain Bouveret, Nadia Brauner, and Nicolas Brulard (Mar. 2024). *Traveling salesman games: semicore allocations for collaborative transportation in short food supply chains*. PREPRINT (Version 1) available at Research Square. URL: <https://doi.org/10.21203/rs.3.rs-4164447/v1>.
- Beviá, Carmen (1998). “Fair Allocation in a General Model with Indivisible Goods”. In: *Review of Economic Design* 3, pp. 195–213.
- Beynier, Aurélie, Sylvain Bouveret, Michel Lemaître, Nicolas Maudet, Simon Rey, and Parham Shams (May 2019). “Efficiency, Sequenceability and Deal-Optimality in Fair Division of Indivisible Goods”. In: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS’19)*. Montreal, Canada: IFAAMAS, 9 pages. URL: <http://recherche.noiraudes.net/en/cycle-deals.php>.
- Blanch, Renaud and Sylvain Bouveret (May 2018). *Simulations de systèmes de vote avec proportionnelle pour les législatives*. Tech. rep. Laboratoire d’Informatique de Grenoble. URL: <http://recherche.noiraudes.net/resources/2018-05-28-rapport.pdf>.
- Blanchard, Marianne, Milan Bouchet-Valat, Damien Cartron, Jérôme Greffion, and Julien Gros (Sept. 2022). “Concerned yet polluting: A survey on French research personnel and climate change”. In: *PLOS Climate* 1.9, pp. 1–20. DOI: [10.1371/journal.pclm.0000070](https://doi.org/10.1371/journal.pclm.0000070). URL: <https://doi.org/10.1371/journal.pclm.0000070>.
- Bouveret, Sylvain (2017). “Social Choice on the Web”. In: *Trends in Computational Social Choice*. Ed. by Ulle Endriss. AI Access. Chap. 20, pp. 387–402. URL: <http://research.illc.uva.nl/COST-IC1205/BookDocs/Chapters/TrendsCOMSOC-20.pdf>.
- Bouveret, Sylvain, Renaud Blanch, et al. (July 2018). *Voter Autrement 2017 - Online Experiment*. Dataset and companion article on Zenodo. DOI: [10.5281/zenodo.1199545](https://doi.org/10.5281/zenodo.1199545). URL: <https://doi.org/10.5281/zenodo.1199545>.
- (Nov. 2019). *Voter Autrement 2017 for the French Presidential Election – The data of the In Situ Experiment*. Dataset and companion article on Zenodo. DOI: [10.5281/zenodo.3548574](https://doi.org/10.5281/zenodo.3548574). URL: <https://doi.org/10.5281/zenodo.3548574>.
- Bouveret, Sylvain, Katarína Cechlárová, and Julien Lesca (June 2019). “Chore Division on a Graph”. In: *Autonomous Agents and Multi-Agent Systems* 33.5, pp. 540–563. URL: <http://recherche.noiraudes.net/resources/papers/JAAMAS19.pdf>.
- Bouveret, Sylvain, Yann Chevaleyre, François Durand, and Jérôme Lang (Aug. 2017). “Voting by Sequential Elimination with few Voters”. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-17)*. Melbourne, Australia. URL: <http://recherche.noiraudes.net/resources/papers/IJCAI17-sequences.pdf>.
- Bouveret, Sylvain, Yann Chevaleyre, and Nicolas Maudet (2016). “Fair Allocation of Indivisible Goods”. In: *Handbook of Computational Social Choice*. Ed. by Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. Cambridge University Press. Chap. 12. URL: <http://recherche.noiraudes.net/en/handbook.php>.
- Bouveret, Sylvain, Hugo Gilbert, Jérôme Lang, and Guillaume Méroué (2023). *Thou Shalt not Pick all Items if Thou are First: of Strategyproof and Fair Picking Sequences*. arXiv: [2301.06086](https://arxiv.org/abs/2301.06086) [cs.GT]. URL: <https://arxiv.org/abs/2301.06086>.
- Bouveret, Sylvain and Jérôme Lang (July 2011). “A general elicitation-free protocol for allocating indivisible goods”. In: *Proceedings of the 22st International Joint Conference on Artificial Intelligence (IJCAI’11)*. Barcelona, Spain. URL: <http://recherche.noiraudes.net/resources/papers/IJCAI11.pdf>.

- Bouveret, Sylvain and Jérôme Lang (Aug. 2014). “Manipulating picking sequences”. In: *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI’14)*. Prague, Czech Republic: IOS Press. URL: <http://recherche.noiraudes.net/resources/papers/ECAI14.pdf>.
- Bouveret, Sylvain and Michel Lemaître (May 2014). “Characterizing Conflicts in Fair Division of Indivisible Goods Using a Scale of Criteria”. In: *Proceedings of the 13th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS’14)*. Paris, France: ACM. URL: <http://recherche.noiraudes.net/resources/papers/AAMAS14.pdf>.
- (2015). “Characterizing conflicts in fair division of indivisible goods using a scale of criteria”. In: *Autonomous Agents and Multi-Agent Systems* 30.2, pp. 259–290. ISSN: 1573-7454. DOI: [10.1007/s10458-015-9287-3](https://doi.org/10.1007/s10458-015-9287-3). URL: <http://recherche.noiraudes.net/resources/papers/JAAMAS15.pdf>.
- (June 2016). “Efficiency and Sequenceability in Fair Division of Indivisible Goods with Additive Preferences”. In: *Proceedings of the Sixth International Workshop on Computational Social Choice (COMSOC’16)*. Toulouse, France. URL: <https://www.irit.fr/COMSOC-2016/proceedings/BouveretLemaitreCOMSOC2016.pdf>.
- Brainard, William C and Herbert E Scarf (2005). “How to compute equilibrium prices in 1891”. In: *American Journal of Economics and Sociology* 64.1, pp. 57–83.
- Brams, Steven J, Paul H Edelman, and Peter C Fishburn (2003). “Fair division of indivisible items”. In: *Theory and Decision* 55, pp. 147–180.
- Brams, Steven J, Marc D Kilgour, and Christian Klamler (2012). “The undercut procedure: an algorithm for the envy-free division of indivisible items”. In: *Social Choice and Welfare* 39.2-3, pp. 615–631.
- Brams, Steven J and Daniel King (2005). “Efficient Fair Division—Help the Worst off or Avoid Envy?” In: *Rationality and Society* 17.4, pp. 387–421.
- Brams, Steven J and Alan D Taylor (1996). *Fair Division — From Cake-cutting to Dispute Resolution*. Cambridge Univ. Press.
- (2000). *The Win-win Solution. Guaranteeing Fair Shares to Everybody*. W. W. Norton & Company.
- Brandt, Felix, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia, eds. (2015). *Handbook of Computational Social Choice*. Cambridge University Press.
- Brânzei, Simina, Hadi Hosseini, and Peter Bro Miltersen (2015). “Characterization and Computation of Equilibria for Indivisible Goods”. In: *Algorithmic Game Theory*. Ed. by Martin Hoefer. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 244–255.
- Brânzei, Simina and Fedor Sandomirskiy (2019). *Algorithms for Competitive Division of Chores*. arXiv: [1907.01766 \[cs.GT\]](https://arxiv.org/abs/1907.01766).
- Budish, Eric (2011). “The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes”. In: *Journal of Political Economy* 119.6.
- Budish, Eric and Estelle Cantillon (Aug. 2012). “The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard”. In: *American Economic Review* 102.5, pp. 2237–71.
- Caragiannis, Ioannis, Jugal Garg, Nidhi Rathi, Eklavya Sharma, Giovanna Varricchio, et al. (2023). “New Fairness Concepts for Allocating Indivisible Items.” In: *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI-23)*, pp. 2554–2562.
- Caragiannis, Ioannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang (2016). “The Unreasonable Fairness of Maximum Nash Welfare”. In: *Proceedings of the 2016 ACM Conference on Economics and Computation*. EC ’16. Maastricht, The Netherlands: ACM, pp. 305–322. ISBN: 978-1-4503-3936-0.
- Carlisle, Ketih and Rebecca L Gruby (2019). “Polycentric Systems of Governance: A Theoretical Model for the Commons”. In: *Policy Studies Journal* 47.4.

- Chaudhury, Bhaskar Ray, Jugal Garg, and Kurt Mehlhorn (2024). “EFX exists for three agents”. In: *Journal of the ACM* 71.1, pp. 1–27.
- Chevalere, Yann, Paul E Dunne, et al. (2006). “Issues in Multiagent Resource Allocation”. In: *Informatica* 30, pp. 3–31.
- Chevalere, Yann, Ulle Endriss, and Nicolas Maudet (2017). “Distributed fair allocation of indivisible goods”. In: *Artificial Intelligence* 242, pp. 1–22.
- Cohendet, Marie-Anne, Jérôme Lang, Jean-François Laslier, Thierry Pech, and Frédéric Sawicki (Mar. 2018). *Une “dose de proportionnelle” : pourquoi ? comment ? laquelle ?* Tech. rep. Terra Nova.
- Condorcet, Nicolas de (1785). *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. Paris: Imprimerie Royale.
- Conitzer, Vincent, Rupert Freeman, and Nisarg Shah (2017). “Fair Public Decision Making”. In: *Proceedings of the 2017 ACM Conference on Economics and Computation*. EC’17. Cambridge, Massachusetts, USA: Association for Computing Machinery, pp. 629–646. ISBN: 9781450345279. DOI: [10.1145/3033274.3085125](https://doi.org/10.1145/3033274.3085125).
- Darmann, Andreas, Julia Grundner, and Christian Klamler (2017). “Election outcomes under different ways to announce preferences: an analysis of the 2015 parliament election in the Austrian federal state of Styria”. In: *Public Choice* 173, pp. 201–216.
- de Keijzer, Bart, Sylvain Bouveret, Tomas Klos, and Yingqian Zhang (Oct. 2009). “On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences”. In: *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT’09)*. Lecture Notes in Artificial Intelligence. Venice, Italy: Springer Verlag. URL: <http://recherche.noiraudes.net/resources/papers/ADT09.pdf>.
- Demko, Stephen and Theodore P Hill (1988). “Equitable Distribution of Indivisible Items”. In: *Mathematical Social Sciences* 16, pp. 145–158.
- Díaz, S. et al., eds. (2019). *Summary for policymakers of the global assessment report on biodiversity and ecosystem services of the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services*. IPBES.
- Dugast, César and Alexia Soyeux (June 2019). *Faire sa part ? Pouvoir et Responsabilité des Individus, des Entreprises, et de l’État face à l’Urgence Climatique*. Tech. rep. <https://www.carbone4.com/publication-faire-sa-part>. Carbone 4.
- Endriss, Ulle, Nicolas Maudet, Fariba Sadri, and Francesca Toni (2006). “Negotiating Socially Optimal Allocations of Resources”. In: *Journal of Artificial Intelligence Research* 25, pp. 315–348.
- Feige, Uriel, Ariel Sapir, and Laliv Tauber (2021). “A tight negative example for MMS fair allocations”. In: *International Conference on Web and Internet Economics (WINE)*. Springer, pp. 355–372.
- Fisher, Irving (1892). *Mathematical Investigations in the Theory of Value and Prices, and Appreciation and Interest*. Augustus M. Kelley, Publishers.
- Foley, Duncan Karl (1967). “Resource Allocation and the Public Sector”. In: *Yale Econ. Essays* 7.1.
- Forte, Andrea, Vanessa Larco, and Amy Bruckman (2009). “Decentralization in Wikipedia Governance”. In: *Journal of Management Information Systems* 26.1, pp. 49–72.
- Fressoz, Jean-Baptiste (2024). *Sans transition. Une nouvelle histoire de l’énergie*. Ecocène. Seuil.
- Gardner, Martin (1978). *Aha! Insight*. Freeman.
- Garg, Jugal, Peter McGlaughlin, and Setareh Taki (2019). “Approximating maximin share allocations”. In: *2nd Symposium on Simplicity in Algorithms (SOSA 2019)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik.
- Garg, Jugal and Setareh Taki (2021). “An improved approximation algorithm for maximin shares”. In: *Artificial Intelligence* 300, p. 103547.

- Ghods, Mohammad, MohammadTaghi HajiAghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami (2018). “Fair allocation of indivisible goods: Improvements and generalizations”. In: *Proceedings of the 2018 ACM Conference on Economics and Computation*, pp. 539–556.
- Gorantla, Pranay, Kunal Marwaha, and Santhoshini Velusamy (2023). “Fair allocation of a multiset of indivisible items”. In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, pp. 304–331.
- Herreiner, Dorothea and Clemens Puppe (2002). “A Simple Procedure for Finding Equitable Allocations of Indivisible Goods”. In: *Social Choice and Welfare* 19.
- Institut National de l’Information Géographique et Forestière (2015). *GEOFLA*. <http://professionnels.ign.fr/geofla>. Accessed on July 2nd, 2015.
- (2016). *Découpage des cantons pour les élections départementales de mars 2015*. <https://www.data.gouv.fr/fr/datasets/decoupage-des-cantons-pour-les-elections-departementales-de-mars-2015/>. Accessed on April 4th, 2018.
- IPCC (2021). *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Vol. In Press. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press. DOI: [10.1017/9781009157896](https://doi.org/10.1017/9781009157896).
- Kalinowski, Thomas, Nina Narodytska, and Toby Walsh (Aug. 2013). “A Social Welfare Optimal Sequential Allocation Procedure”. In: *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI-13)*. Ed. by Francesca Rossi. Beijing, China: IJCAI/AAAI, pp. 227–233.
- Kalinowski, Thomas, Nina Narodytska, Toby Walsh, and Lirong Xia (July 2013). “Strategic Behavior when Allocating Indivisible Goods Sequentially”. In: *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI-12)*. Bellevue, WA: AAAI Press, pp. 452–458.
- Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ashesh Rambachan (2018). “Algorithmic fairness”. In: *Aea papers and proceedings*. Vol. 108. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 22–27.
- Kohler, David A and R Chandrasekaran (1971). “A class of sequential games”. In: *Operations Research* 19.2, pp. 270–277.
- Kurokawa, David, Ariel D Procaccia, and Junxing Wang (2016). “When can the maximin share guarantee be guaranteed?” In: *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*. AAAI’16. Phoenix, Arizona: AAAI Press, pp. 523–529.
- (2018). “Fair Enough: Guaranteeing Approximate Maximin Shares”. In: *J. ACM* 65.2. ISSN: 0004-5411. DOI: [10.1145/3140756](https://doi.org/10.1145/3140756). URL: <https://doi.org/10.1145/3140756>.
- Lang, Jérôme (2004). “Logical preference representation and combinatorial vote”. In: *Annals of Mathematics and Artificial Intelligence* 42, pp. 37–71.
- Laslier, Jean-François (2011). “And the loser is... plurality voting”. In: *Electoral systems: paradoxes, assumptions, and procedures*. Springer, pp. 327–351.
- Lemaître, Michel, Gérard Verfaillie, and Nicolas Bataille (Aug. 1999). “Exploiting a Common Property Resource under a Fairness Constraint: a Case Study”. In: *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*. Ed. by Thomas Dean. Stockholm, Sweden: Morgan Kaufmann, pp. 206–211.
- Lipton, Richard J, Evangelos Markakis, Elchanan Mossel, and Amin Saberi (2004). “On approximately fair allocations of indivisible goods”. In: *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pp. 125–131.
- Loosemore, John and Victor J Hanby (1971). “The theoretical limits of maximum distortion: Some analytic expressions for electoral systems”. In: *British Journal of Political Science* 1.4, pp. 467–477.

- Markakis, Evangelos and Christos-Alexandros Psomas (Dec. 2011). “On worst-case allocations in the presence of indivisible goods”. In: *International Workshop on Internet and Network Economics (WINE)*. Springer, Singapore, pp. 278–289.
- Ministère de l’Enseignement Supérieur, de la Recherche et de l’Innovation (Apr. 2021). *La situation des femmes universitaires dans l’enseignement supérieur en 2020*. <https://www.enseignementsup-recherche.gouv.fr/fr/la-situation-des-femmes-universitaires-dans-l-enseignement-superieur-en-2020-8231>.
- Ministère de l’Intérieur (2017a). *Elections législatives des 11 et 18 juin 2017 - Résultats du 1er tour*. <https://www.data.gouv.fr/en/datasets/elections-legislatives-des-11-et-18-juin-2017-resultats-du-1er-tour/>. Accessed on April 4th, 2018.
- (2017b). *Elections législatives des 11 et 18 juin 2017 - Résultats du 2nd tour*. <https://www.data.gouv.fr/en/datasets/elections-legislatives-des-11-et-18-juin-2017-resultats-du-2nd-tour/>. Accessed on April 4th, 2018.
- Moulin, Hervé (1990). “Fair Division under Joint Ownership: Recent Results and Open Problems”. In: *Social Choice and Welfare* 7, pp. 149–170.
- (2003). *Fair Division and Collective Welfare*. MIT Press.
- Nguyen, Trung Thanh and Jörg Rothe (2023). “Complexity Results and Exact Algorithms for Fair Division of Indivisible Items: A Survey”. In: *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI-23)*, pp. 6732–6740.
- Ogryczak, Włodzimierz and Tomasz Śliwiński (2003). “On Solving Linear Programs with the Ordered Weighted Averaging Objective”. In: *European Journal of Operational Research* 148, pp. 80–91.
- Ostrom, Elinor (1990). *Governing the Commons - The Evolution of Institutions for Collective Action*. Cambridge University Press.
- Othman, Abraham, Tuomas Sandholm, and Eric Budish (2010). “Finding approximate competitive equilibria: efficient and fair course allocation”. In: *Proceedings of AAMAS’10*. Vol. 10. Toronto, Canada, pp. 873–880.
- Pápai, Szilvia (2000). “Original papers: Strategyproof multiple assignment using quotas”. In: *Review of Economic Design* 5.1, pp. 91–105.
- (July 2001). “Strategyproof and Nonbossy Multiple Assignments”. In: *Journal of Public Economic Theory* 3.3, pp. 257–271.
- Parijs, Philippe Van (1997). *Real Freedom for All: What (if Anything) Can Justify Capitalism?* OUP Catalogue 9780198293576. Oxford University Press.
- Perry, Jonathan and Robert C Powers (2010). “Anonymity, monotonicity, and quota pair systems”. In: *Mathematical Social Sciences* 60.1, pp. 57–60.
- Plaut, Benjamin and Tim Roughgarden (2018). “Almost Envy-Freeness with General Valuations”. In: *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA ’18. New Orleans, Louisiana: Society for Industrial and Applied Mathematics, pp. 2584–2603. ISBN: 9781611975031.
- (2020). “Almost envy-freeness with general valuations”. In: *SIAM Journal on Discrete Mathematics* 34.2, pp. 1039–1068.
- Pörtner, H.-O. et al. (2022). *Climate Change 2022: Impacts, Adaptation and Vulnerability*. Technical Summary. Cambridge, UK and New York, USA: Cambridge University Press, pp. 37–118. ISBN: 9781009325844.
- Potters, Jos AM, Imma J Curiel, and Stef H Tijs (1992). “Traveling salesman games”. In: *Mathematical programming* 53.2, pp. 199–211. ISSN: 0025-5610.

- Procaccia, Ariel D and Junxing Wang (2014). “Fair Enough: Guaranteeing Approximate Maximin Shares”. In: *Proceedings of the 15th ACM Conference on Economics and Computation (EC’14)*.
- Rassenti, Stephen J, Vernon L Smith, and Robert L Bulfin (1982). “A combinatorial auction mechanism for airport time slot allocation”. In: *The Bell Journal of Economics*, pp. 402–417.
- Raworth, Kate (2017). *Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist*. London: Random House.
- Richardson, Katherine et al. (2023). “Earth beyond six of nine planetary boundaries”. In: *Science Advances* 9.37. DOI: [10.1126/sciadv.adh2458](https://doi.org/10.1126/sciadv.adh2458). eprint: <https://www.science.org/doi/pdf/10.1126/sciadv.adh2458>. URL: <https://www.science.org/doi/abs/10.1126/sciadv.adh2458>.
- Rockström, Johan et al. (2009). “A safe operating space for humanity”. In: *Nature* 461, pp. 472–475.
- Rothe, Jörg, ed. (2015). *Economics and computation*. Springer.
- Sandholm, Tuomas W (1998). “Contract Types for Satisficing Task Allocation: I. Theoretical Results”. In: *Proceedings of the AAAI Spring Symposium: Satisficing Models*. Ed. by Sandip Sen. Menlo Park, California: AAAI Press, pp. 68–75.
- Shams, Parham, Aurélie Beynier, Sylvain Bouveret, and Nicolas Maudet (Aug. 2020). “Fair in the Eyes of Others”. In: *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI’20)*. Santiago de Compostela, Spain: IOS Press. URL: <https://arxiv.org/abs/1911.11053>.
- (Nov. 2021). “Minimizing and balancing envy among agents using Ordered Weighted Average”. In: *Proceedings of the 7th International Conference on Algorithmic Decision Theory (ADT’21)*. Toulouse, France: Springer Verlag. URL: <http://recherche.noiraudes.net/resources/papers/ADT21.pdf>.
- (2022). “Fair in the Eyes of Others”. In: *Journal of Artificial Intelligence Research (JAIR)* 75, pp. 913–951. URL: <https://jair.org/index.php/jair/article/view/13778/26864>.
- Shapley, Lloyd S (1953). “A Value for n-Person Games”. In: *Contributions to the Theory of Games 2*, pp. 307–317. DOI: [10.1515/9781400881970-018](https://doi.org/10.1515/9781400881970-018).
- Steffen, Will, Wendy Broadgate, Lisa Deutsch, Owen Gaffney, and Cornelia Ludwig (2015). “The trajectory of the Anthropocene: The Great Acceleration”. In: *The Anthropocene Review* 2 (1), pp. 81–98. DOI: <https://doi.org/10.1177/2053019614564785>.
- Steffen, Will, Katherine Richardson, et al. (2015). “Planetary boundaries: Guiding human development on a changing planet”. In: *Science* 347.6223.
- Steinhaus, Hugo (1948). “The Problem of Fair Division”. In: *Econometrica* 16.1, pp. 101–104.
- Stromquist, Walter (2008). “Envy-free cake divisions cannot be found by finite protocols”. In: *the electronic journal of combinatorics* 15.1.
- Suksompong, Warut (2025). “Weighted fair division of indivisible items: A review”. In: *Information Processing Letters* 187, p. 106519. ISSN: 0020-0190. DOI: <https://doi.org/10.1016/j.ipl.2024.106519>. URL: <https://www.sciencedirect.com/science/article/pii/S0020019024000498>.
- Thomson, William (2015). “Introduction to the Theory of Fair Allocation”. In: *Handbook of Computational Social Choice*. Ed. by Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. Cambridge University Press. Chap. 11.
- Tinbergen, Jan (1953). *Redelijke Inkomensverdeling*. Haarlem: N. V. DeGulden Pers.
- Tomlinson, Bill, Juliet Norton, Eric P.S Baumer, Marcel Pufal, and Barath Raghavan (2015). “Self-Obviating Systems and their Application to Sustainability”. In: *iConference 2015*.
- Varian, Hal R (1974). “Equity, Envy and Efficiency”. In: *Journal of Economic Theory* 9, pp. 63–91.

- Vazirani, Vijay V (2007). “Combinatorial algorithms for market equilibria”. In: *Algorithmic Game Theory*. Cambridge University Press. Chap. 5, pp. 103–134.
- Yager, Ronald R (1988). “On Ordered Weighted Averaging Aggregation Operators in Multicriteria Decision Making”. In: *IEEE Transactions on Systems, Man, and Cybernetics* 18, pp. 183–190.
- Young, H Peyton (1994a). “Chapter 34 Cost allocation”. In: *Handbook of Game Theory with Economic Applications*. Vol. 2. Elsevier, pp. 1193–1235. DOI: [10.1016/S1574-0005\(05\)80066-9](https://doi.org/10.1016/S1574-0005(05)80066-9). URL: <https://www.sciencedirect.com/science/article/pii/S1574000505800669>.
- (1994b). *Equity in Theory and Practice*. Princeton University Press.
- Zwicker, William S (2015). “Introduction to the Theory of Voting”. In: *Handbook of Computational Social Choice*. Ed. by Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. Cambridge University Press. Chap. 2.