



HAL
open science

Procedures based on Exchanges and new Relaxations of Envy-Freeness in Fair Division of Indivisible Goods

Parham Shams

► **To cite this version:**

Parham Shams. Procedures based on Exchanges and new Relaxations of Envy-Freeness in Fair Division of Indivisible Goods. Artificial Intelligence [cs.AI]. Sorbonne Université, 2023. English. NNT : 2023SORUS477 . tel-04531057

HAL Id: tel-04531057

<https://theses.hal.science/tel-04531057>

Submitted on 3 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



PhD thesis in Computer Science

Procedures based on Exchanges and new Relaxations of Envy-Freeness in Fair Division of Indivisible Goods

Publicly presented on December 21st 2023 by
PARHAM SHAMS

Jury :

Supervisors :

Aurélié BEYNIER	Associate Professor, LIP6, Sorbonne Université
Sylvain BOUVERET	Associate Professor, LIG, Université Grenoble Alpes
Nicolas MAUDET	Professor, LIP6, Sorbonne Université

Reviewers :

Gauthier PICARD	Senior Research Scientist, CNRS, ONERA, Université de Toulouse
Agnieszka RUSINOWSKA	Senior Research Scientist, Centre d'Economie de la Sorbonne, CNRS, Université Paris 1 Panthéon Sorbonne

Examiners :

Bruno ESCOFFIER	Professor, LIP6, Sorbonne Université
Laurent GOURVES	Senior Research Scientist, CNRS, LAMSADE, Université Paris Dauphine

Remerciements

Je tiens tout d'abord à remercier Gauthier Picard et Agnieszka Rusinowska d'avoir accepté de rapporter ma thèse. De même, je remercie Bruno Escoffier et Laurent Gourvès d'avoir accepté le rôle d'examineur mais aussi celui de membre de mon comité de suivi tout au long de cette thèse. Nos échanges ainsi que leurs retours et conseils ont été précieux.

Ensuite, je tiens évidemment à remercier mes encadrants Aurélie Beynier, Sylvain Bouveret et Nicolas Maudet qui ont accepté d'encadrer ma thèse durant ces années. Tout d'abord sur l'aspect académique ils ont toujours été de bon conseil et m'ont toujours guidé dans la bonne direction tout en me laissant une liberté intellectuelle totale que je value plus que tout et je ne saurais les en remercier assez pour cela. Sur le plan humain, nos échanges étaient toujours très agréables et la bonne humeur ambiante rendait nos discussions plaisantes. Les réunions étaient donc non seulement des moments de travail mais aussi des moments de partage sur divers sujets. J'ai donc beaucoup appris à leur contact sur beaucoup d'aspects et je leur en suis pour cela très reconnaissant.

Dans mon parcours académique je tiens également à remercier Patrice Perny qui a accepté d'encadrer mon stage de M2 et m'a ensuite aiguillé vers Nicolas Maudet pour la thèse; ainsi sans lui cette thèse n'aurait pas été possible. Il a de plus toujours été à l'écoute et de bon conseil durant la thèse. De même, je remercie Nicolas Bredèche qui m'a offert l'opportunité d'un premier stage dans le monde de la Recherche en M1.

Je voudrais évidemment aussi remercier les personnes avec qui j'ai partagé le bureau. En commençant évidemment par Nadjat qui nous y a chaleureusement accueilli. Elle a tenu pour moi une place très spéciale durant ces années et a su rendre cette expérience d'autant plus enrichissante. Elle a été une des raisons principales pour lesquelles le temps passé était si agréable et qui font que je suis parvenu au terme de cette thèse. Ensuite Hugo avec qui j'étais en Master puis en thèse et qui est devenu un collègue avec qui j'ai eu de très longues discussions sur tous les sujets possibles et imaginables allant de la recherche à la politique, la philosophie, le cinéma, la littérature ou encore aux blagues (que devaient supporter Nadjat et Alexandre). Il a été d'une grande aide et est un des facteurs du bon temps que j'ai pu passer en thèse. Et enfin Alexandre qui nous a fait partager sa passion et son talent pour la musique. Enfin, je tiens aussi à remercier Fares, Jordan, Manuel et Valentin avec qui j'ai eu le plaisir de partager le bureau pendant 2 mois. Ce fut bref mais intense en échanges et je n'ai aucun doute qu'ils mèneront à bien leur thèse.

De manière plus générale je tiens à remercier toutes les personnes des équipes SMA, Décision et RO avec qui j'ai pu interagir et échanger durant toutes ces années et qui ont rendu tous ces moments, ces repas et ces pauses divertissants, enrichissants et épanouissants. Merci donc à Olivier, Thibaut, Bruno, Marvin, Gaspard, Anne-Elisabeth, Nawal, Pierre-Henri, Fanny, Evripidis, Cassandre, Clara, Kostas, David (Saulpic), David (Wu),

Nicolas (De Buffala), Anja, Magdaléna, Martin, Nathanaël, Santiago, Adèle, Cécile, Emmanuel, Pierre, Safia. C'était un plaisir de les côtoyer durant ces années.

Je tiens aussi à remercier le directeur de l'école doctorale Habib Mehrez qui a été à l'écoute, très compréhensif et qui a beaucoup oeuvré pour que je soutienne ma thèse.

Des remerciements s'imposent également envers tous mes amis qui m'ont beaucoup aidé à me changer les idées, me reposer et décompresser durant ces années de thèse. Tout d'abord mes amis de Licence et Master Alexandre, Jaime, Leslie et Tanguy avec qui j'ai entre autres eu le plaisir de m'échapper de nombreux "labyrinthes". Une pensée spéciale également à mes amis de prépa que j'ai toujours plaisir à retrouver : Manuella, Marie, Mohamed, Pierrick, Romain, Tanis et Virgile. Et enfin mes amis de lycée : Arthur, Inès, Jean-Loup, Kévin, Laure, Marie, Mathilde, Maya et Tikal. J'ai été très heureux de compter sur leur soutien tout au long de cette thèse et sur leur présence lors de ma soutenance.

Enfin, je tiens évidemment à remercier mes parents et mon frère qui ont été d'une aide précieuse par l'organisation du pot de thèse. Plus sérieusement, ils ont toujours été d'une grande aide sur tous les plans et leurs conseils m'ont toujours été précieux. Entre autres, j'ai pu grandir grâce à eux dans un environnement chaleureux et intellectuellement propice à la réflexion, la remise en question permanente, la curiosité scientifique et l'ouverture d'esprit. Ce sont ces aspects qui font que j'en suis là aujourd'hui, à écrire ces mots et je ne les en remercierai jamais assez.

Contents

Notations	9
Introduction	11
1 State Of the Art	19
1.1 Introduction	20
1.2 The fair division problem	20
1.2.1 Multiagent allocation of indivisible resources	21
1.2.2 Preferences	21
1.2.3 Domain Restriction	23
1.2.4 Preference cultures	25
1.3 Desirable criteria and properties	25
1.3.1 Classical fairness measures	25
1.3.2 Relaxations of fairness measures	30
1.3.3 Efficiency	31
1.4 Computing fair and efficient allocations	32
1.4.1 Desirable properties of a procedure	33
1.4.2 Centralized procedures	33
1.4.3 Decentralized protocols	36
1.5 Thesis purpose	40
1.6 Organization of the thesis	40
2 Efficiency, Sequenceability and Deal-Optimality	41
2.1 Introduction	42
2.2 Definitions	42
2.3 Sequenceable allocations	42
2.4 Pareto-optimality	43
2.5 Cycle deals-optimality	45
2.6 Restricted Domains	47
2.7 Envy-Freeness and CEEI	50
2.8 Experiments	51
2.9 Conclusion	52
3 Swap Dynamics in Single-Peaked Housing Markets	53
3.1 Introduction	53
3.2 Centralized allocation procedures for housing markets	55

3.2.1	Crawler algorithm	55
3.3	Swap dynamics: a family of procedures based on swap-deals	57
3.4	Properties of swap dynamics	60
3.4.1	Strategy-proofness	61
3.4.2	Pareto-optimality of swap dynamics	62
3.4.3	Maximality of the single-peaked domain	65
3.5	The “price of” swap dynamics	68
3.6	Experimental study	71
3.6.1	Experimental protocol	71
3.6.2	Analysis of the results	73
3.7	Conclusion	79
4	Balancing envy among agents	81
4.1	Introduction	82
4.2	MinOWA Envy	82
4.2.1	Fair OWA	82
4.2.2	Linearization of OWA minimization	84
4.3	Link with other fairness measures	85
4.3.1	Warm-up: $n = 2$	86
4.3.2	General case ($n \geq 3$)	87
4.4	Experimental results	88
4.4.1	Spliddit instances	89
4.4.2	Synthetic instances	90
4.5	Conclusion	90
5	Fair in the Eyes of Others	93
5.1	Introduction	93
5.2	K -approval envy	95
5.3	Some properties of K -app envy	98
5.4	K -approval non-proportionality	103
5.5	Properties of (K -app non-prop)-freeness	104
5.6	Link between approval envy-freeness and approval non-proportionality	107
5.7	Computation	109
5.7.1	A MILP formulation for K -app envy	111
5.7.2	A MILP formulation for K -app non-prop	112
5.8	House Allocation setting	113
5.9	Experimental results	117
5.9.1	Spliddit instances	117
5.9.2	Uniformly distributed preferences: general setting	118
5.9.3	Uniformly distributed preferences: house allocation problems	120
5.9.4	Correlated preferences	122
5.10	Conclusion	123
	Conclusion and future works	125
	A Determining if a solution is Pareto Optimal	145

B	MILP characterizing the EF1	147
C	MILP characterizing the EFX	149
D	Fair in the eyes of others	151

Notations

Notations of Chapter 1:

$\mathcal{N} = \{a_1, \dots, a_n\}$	set of n agents
$\mathcal{O} = \{o_1, \dots, o_m\}$	set of m objects
n	number of agents
m	number of objects
I	add-MARA instance
$u(i, j)$	utility that a_i has for o_j
\succ_i	linear order of preferences of a_i
L	preference profile with n linear orders
π_i	share of agent a_i in allocation π
\mathcal{I}	set of all add-MARA instances
$\mathcal{A}(I)$	set of all allocations for instance I
PROP	Proportionality
MMS	Max Min Fair Share
mMS	min Max Fair Share
CEEI	Competitive Equilibrium with Equal Incomes
EF	Envy-free
PO	Pareto-Optimal

Notations of Chapter 2:

μ	sequence of transfers
$\mathcal{S}(I)$	the set of possible sequences for the instance I
$s(I)$	defined by the fact that $(\vec{\sigma}, \pi) \in s(I)$ if and only if π can be generated by sequence $\vec{\sigma}$

Notations of Chapter 3:

\succ_i	linear order of preferences of a_i
L	preference profile with n linear orders
ark	average rank
mrk	minimum rank
$top(a_i)$	top object of agent a_i
$C_k(\pi)$	set of all the improving deals of size at most k that can be applied from π

Notations of Chapter 4:

\mathcal{I}^p	Set of add-MARA instances with only two agents
\mathcal{I}^b	Set of add-MARA instances with binary utilities
OWA	Ordered Weighted Average
MNW	Max Nash Welfare

Notations of Chapter 5:

(K-app) EF	K approval envy-free
(K-app) non-prop	K approval non-proportionality

Introduction

One of the biggest issues humanity has to face is the one of climate change. In order to tackle this issue, the latest Intergovernmental Panel on Climate Change (IPCC) report [IPCC, 2023] states that "Equity, inclusion, just transitions, broad and meaningful participation of all relevant actors in decision making at all scales enable deeper societal ambitions for accelerated mitigation, and climate action more broadly, and build social trust, support transformative changes and an equitable sharing of benefits and burdens (high confidence). Equity remains a central element in the UN climate regime, notwithstanding shifts in differentiation between states over time and challenges in assessing fair shares.". Not only do we see that there is a question of decision making with all (relevant) actors but also the issue of equity which remains central and the need to assess fair shares. These all are concepts that we will tackle in this thesis, in particular the notion of fairness for which there is not one and only one formal definition. Moreover, the part referring to decision making with actors sends back to the field that encapsulates the work presented in this thesis.

Computational Social Choice (often referred to as Comsoc) is the field that is at the intersection of *Social Choice* and *Computer Science*.

The Social Choice aspect is about creating and analysing methods for *Collective Decision Making*. Indeed, when you have a situation or a problem in which a lot of agents are involved (agents can refer to people or robots for instance), anyone of them having preferences over alternatives then, some procedure is needed to make a decision that satisfies them as best as possible. For example, in voting theory the agents are the voters and their preferences can be expressed via their votes. When we say that we want a decision that satisfies the agents, we can easily see that the meaning is not very clear. First, we can wonder which agents we want to satisfy. On the one hand, it could seem acceptable not to care about one or two unhappy agents if the others are very happy, this is what we usually call classical utilitarianism. On the other hand, we could be very careful about every agent and for example be looking for a solution (the elected candidate in our voting setting) for which the least happy voter is the happiest, this is usually called egalitarianism. It is obvious that we can navigate between those two extremes and that the value of a solution depends on how you aggregate the happiness of the agents. Moreover, we can wonder what happiness means for an agent. The satisfaction of an agent is expressed through preferences over alternatives. Usually, we can consider ordinal ones where the agents simply express the fact that they prefer some candidate to another or some group of candidates to another group of candidates. The other type of preferences is cardinal and refers to the agents having some numerical function associating the candidates to some utility. For example, imagine a voting setting with two voters (the agents) and two candidates (the alternatives). If the preferences are ordinal both voters can express

the fact that they prefer the first candidate over the second. However, if the preferences are cardinal one voter could express her preference by associating utility of 1 to the first candidate and 0 to the second while the other voter has utility 50 associated to the first candidate and 1 to the second. This raises the issue of interpersonal comparisons of utilities. Moreover, even before comparing utilities between agents, it might also be difficult to assess the meaning of numbers when they are supposed to reflect an agent's preference: for instance, why using 50, and not 100 or 1000? And finally, we can also see that if the voters have to choose two candidates instead of only one arises the problem of how the preferences add up.

Social Choice is about designing and studying procedures that return solutions respecting some desirable properties and/or in a similar way preventing the solutions to have undesired properties. It can also deal with the properties of the procedures themselves. For example, in our voting problem, it would be extremely natural to look for a procedure that always elects a candidate c_1 which is better ranked by a majority of voters considering successive duels between c_1 and the other candidates (it is called a Condorcet winner) [Condorcet, 1785]. Another very natural and common way to elect a candidate would be to associate scores to the voters' preferences and electing the candidate with the maximum sum of scores. A very popular one is the Borda score, axiomatized by Young [1974], expressing the number of candidates someone beats in the preferences of a voter [Baumeister et al., 2014]. Therefore, Social Choice is about finding a compromise between good and undesired properties about the solutions but also about the procedure itself. Finally, it is also very common to just study and analyse the existence of solutions that mix several desirable properties while potentially avoiding some bad ones. However, if such a procedure is easily subject to manipulability (meaning that voters would have an incentive to lie about their preferences and easy in the sense that the complexity of doing so is reasonable enough) then this procedure may not be that desirable. This aspect is within the scope of Computer Science.

Indeed, as we have talked about the Social Choice part of Computational Social Choice, let us now transition to the Computer Science part. In fact, Social Choice theory does not really care about the complexity of the procedures. Indeed, when the number of agents and/or the number of alternatives is too high, some procedures can become inapplicable because even a computer would take centuries to run them.

Hence, Computer Science (Algorithmic Theory and Artificial Intelligence) becomes useful as they take into account the complexity of procedures. For example, in our voting setting, even if a procedure is manipulable we could conclude that it is not that bad because for a voter to know how to manipulate his ballot such that she gets a better outcome is computationally hard. This notion of complexity of manipulation has been introduced by Bartholdi et al. [1989]. But this is not the only thing Computer Science deals with in this context.

Even if we will not list them all, a big aspect of it is the algorithmic and optimization part. Indeed, as discussed before we have to design procedures and algorithms that can return the desired solutions while it is also very useful to look at tools in order to optimize them. Note that there are a lot of other areas of Computer Science such as automatic proving in Social Choice [Geist and Peters, 2017; Ciná and Endriss, 2016; Grandi and Endriss, 2013]. Preference Learning is also a field that has attracted a lot of attention in recent years [Fürnkranz and Hüllermeier, 2010]. It refers to the induction of preference

models from empirical data. The goal is to predict the preferences of an agent or a group of agents.

So by combining these fields you obtain *Computational Social Choice* [Brandt et al., 2016]. This field has been gaining importance over the last years as illustrated by the many dedicated conferences.

As we have already discussed, voting theory is one of the main topics of *Computational Social Choice* but another one in which we will be interested throughout this thesis is *Resource Allocation* and more specifically *Fair Division of Indivisible Goods*. The problem is about fairly and/or efficiently allocating a set of m resources to n agents. First note that the resources can either be objects, goods, tasks or chores.

In the case where the objects are divisible, the problem is called the *cake cutting problem* [Steinhaus, 1948; Brams and Taylor, 1996]. For instance, a common problem is the land division problem where you have a piece of land of some size and you want to divide it between n agents but you also want a solution (a division of the land between the agents) such that you minimize the number of non-connex lands for each agent or you want the pieces of lands to have specific geometric shapes [Segal-Halevi et al., 2020]. Obviously, you want to achieve that by also having other desirable properties on the solution and on the procedure as discussed before.

When the objects are indivisible we face a problem of *Fair Division of Indivisible Goods*. There are m indivisible objects that you want to efficiently and fairly share between n agents. Efficiency in this context often refers to completeness and Pareto-Optimality (PO). Completeness states that every solution should be such that all the objects are allocated whereas Pareto-Optimality of a solution means that there is no other solution such that all the agents are not worse-off and at least one agent is strictly better-off.

While efficiency is commonly defined by these two concepts in this context it is much more complex for fairness as there are a lot of ways to express the fairness of a solution. Under cardinal preferences there are several aggregators that can express the fairness of a solution such that the Max Nash Welfare [Varian, 1974] which consists in the solution maximizing the product of the utilities of the agents, or the max fair Ordered Weighted Average (OWA) [Yager, 1998] which consists in the allocation maximizing some ordered weighted sum of the utilities of the agents.

Another widely studied fairness criterion is envy-freeness (EF) [Tinbergen, 1953; Foley, 1967]. An allocation is said to be envy-free if all the agents are happy with their own share/bundle, in the sense that there is no agent that prefers the bundle of another agent to her own bundle. This is a very strong property for an allocation. Indeed, in the indivisible setting such an allocation does not always exist. Think of the basic situation where two objects have to be shared between two people but both of them prefer the same object. Obviously, the one that gets this item will be envied by the other. This is why in the literature, a lot of relaxations of envy-freeness have been proposed. One of the most studied relaxations of envy-freeness is envy-freeness up to one good (EF1) [Lipton et al., 2004; Budish, 2011] in which we relax envy-freeness by allowing the agents to "forget" one object from the bundle of another agent. If by forgetting one object from the other agent's bundle she becomes not envious then, we will say that she is envy-free up to one good towards that agent. Another recently heavily studied notion is envy-freeness up to any good (EFX) [Caragiannis et al., 2016] which also consists in forgetting one item but this time it can be any item from the other agent's bundle and hence, in the worst

case it can be the least valued item. In another manner, we can relax envy-freeness by minimizing some degree of envy which could for example be the sum of the envies of the agents. These are orthogonal approaches and we will deal with both of them in this work.

In the same way, proportionality (PROP) is also a very natural fairness notion reached when all the agents have at least one n th of the utility of the whole set of objects [Steinhaus, 1948]. While it is a weaker notion than EF (as if an EF allocation exists then it is also proportional while the converse does not hold) we can still not guarantee the existence of a PROP allocation and hence relaxations such as PROP1 have been proposed and studied [Conitzer et al., 2017].

In light of what Computational Social Choice is about, we aim to build and analyze procedures that satisfy some combination of properties. For example, Max Nash Welfare satisfies Pareto-Optimality and EF1 [Caragiannis et al., 2016] so we can wonder whether we can define a procedure returning PO and EFX allocations. Another preliminary issue consists in determining if such an allocation always exists (note that in this case we do not even know if an EFX allocation exists). Even if we have existing procedures returning such allocations we could look for better procedures in terms of time complexity or strategy-proofness. Another track is about proposing new fairness measures or relaxations and find if there are links between them and other existing fairness measures.

Finally, an important question about the procedure is if it is centralized or decentralized. A centralized procedure involves a coordinator which asks the agents for their preferences. The coordinator then computes the outcome which is communicated to the agents. On the other hand, in decentralized procedures, the agents directly interact with each others to obtain the outcome in a distributed way. Such protocols are more often used in real life and answer to the need of simplicity, anonymity and confidence issues as agents do not have to communicate their preferences to some central authority. In this thesis we will in particular study procedures based on exchange sequences and picking sequences.

As we have seen what Computer Social Choice deals with and what it encompasses we will now give some real life problems where it is useful. Indeed, we have obviously already seen that the notions of collective decision making and fairness are prevailing topics nowadays at least via the IPCC example. But there are many others.

First, one that is widely discussed nowadays is the matching of the students to colleges for their studies which is currently done by the Parcoursup platform. Indeed, after finishing High School the students have to be matched to universities and schools. However, each university and school has a fixed number of seats and they also have preferences over the students. In the same way, each student can go to at most one facility and also have preferences over them. Hence, and even if the number of seats is sufficient for all the students, there will almost certainly be conflicts between students and also between schools. This is why there is a need to design an algorithm that returns a matching that is fair and efficient. As both parties have preferences on each other, one way to guarantee fairness for both students and schools is stability (we do not want a couple student/school not to be matched together while they prefer each other from what they get). But we could also want to look for a matching for which the lowest rank of the school a student gets is the highest possible. On the other hand, efficiency will obviously at least be to guarantee that there remains no seat while a student has no school because it would be

a shame to waste a seat. Besides, this raises other questions of mechanism design, for instance we want the algorithm that would do such a matching to be strategy-proof. It means that we want the students and the schools not to be tempted to lie about their preferences. A very similar problem called the roommates problem consists in a group of people that each has preferences over the others and that we want to associate in pairs in a stable way. Here it means that we want a partition so that we cannot find two people such that they prefer each other over their partner in the current partition [Irving, 1985]. Also note that these problems historically come from the Hospital Residents problem (in which you want to allocate residents to the hospitals that each have a certain acceptance quota) and the Stable Marriage Problem (where the goal is to form n couples from n men and n women that have preferences on each other). These two problems have been addressed by Gale and Shapley [1962] giving birth to the eponym algorithm that guarantees stability of the returned marriage/allocation.

Another very interesting real-life application is about resource allocation for constellations of satellites in order to observe Earth called the Orbit Slot Allocation Problem. As space programs and sending satellites in space represent very expansive projects, they are usually funded by several countries or companies. Hence when it comes to using these satellites that are very often requested and usually it implies conflicts between those demands, there has then to be a choice to know which requests are going to be prioritized. The goal would then be to fairly respond to these demands but also in an efficient way. Fairly in that case would be for example to take into account the importance of the request but also the one of the requestor as the level of funding is not necessarily the same. Efficiency aims here not to sub-exploit the satellites as it would be a shame not to use them while they are available. This problem has been widely presented and studied in [Bensana et al., 1999; Verfaillie and Lemaître, 2001; Fargier et al., 2004; Bianchessi et al., 2007; Picard, 2022; Roussel et al., 2023].

The House Allocation Problem (HAP) is also very widely studied [Abraham et al., 2005; Gan et al., 2019; Beynier et al., 2019b]. In particular a variant of this problem has recently gained importance in which are added diversity constraints [Gross-Humbert et al., 2022; Gross-Humbert et al., 2023]. It also consists in a matching where you want to allocate one and only one house to each agent. When considering diversity constraints we identify different groups and try to fairly allocate houses to them in order to mix groups in the neighbourhoods.

Finally, we will end up with more daily-life problems that can be helped by Comsoc and fair division of indivisible goods. The course allocation and scheduling problem is a prominent example, where courses have to be fairly allocated to students [Budish, 2011; Othman et al., 2010; Pascual et al., 2018]. Contrary to the university allocation problem, in this setting only students are assumed to have preferences over courses (what is usually called “one-sided” problem). Divorce is another classical situation where Comsoc can help two members of a couple splitting their belongings between them [Brams and Taylor, 1996]. In this situation, objects like furniture or houses are not divisible and that is where the field of fair division of indivisible goods comes in play. In a similar context of a household, a couple may find it challenging to fairly divide chores between themselves. This matter is addressed by Igarashi and Yokoyama [2023], who design an application based on fair division of chores.

These few examples of real-life problems show how important and modern this field is

and why the work we will display in this manuscript can be applied in addition to being theoretically interesting. Note that we obviously did not mention all the real-life problems linked to that field such as all the voting problems, scheduling and also bidding problems but only a few of them to show how our work can be anchored in our society.

In this thesis, we are looking at the fair division of indivisible goods and try to find fair and efficient allocations. However, fairness has not only one but many definitions. One that has received a lot of attention is EF but is not guaranteed to exist hence relaxations have been proposed and studied recently. This is why we concentrate on new relaxations of envy-freeness but also proportionality. Another focus is about designing, finding and studying ways to obtain such allocations that have desirable properties called procedures. We will mainly try to go towards more decentralized ones that do not rely on some central authority to run and make the agents participate in building the final allocation.

We will now talk about the positioning of our different chapters throughout the thesis.

While Chapter 1 will formally introduce the model that we will be using throughout this thesis, it will also anchor the subject within existing literature.

Chapter 2 is mostly centered around a very common procedure called sequences of sincere choices or picking sequences. More precisely we will study the notion of sequenceable allocations which are the ones that can be obtained through the aforementioned procedure. We will also investigate distributed procedures based on local exchanges. This chapter mainly studies connections between these procedures and other existing fairness and efficiency criteria such as envy-freeness and Pareto-Optimality among others.

In Chapter 3, we will focus on a specific setting called housing market in which each agent has initially one object and cannot get more or less objects. Moreover, we also consider the domain restriction of single-peakedness. In this setting, a procedure called Crawler satisfies the same three axioms as the Top Trading Cycle procedure (TTC) does in the general setting. However, both procedures can involve long trading cycles so we study how we could decompose them via the shortest possible cycles: swap-deals (cycles of size 2) in addition to studying them in our specific setting.

Chapter 4 is about proposing a new relaxation of envy-freeness whose idea comes from the domain of optimization. Indeed, a way to relax envy-freeness is by considering some degree of envy. For example, we may decide to sum the degree of envy which occurs among all the pairs of agents (which is clearly a relaxation of envy-freeness). However, by minimizing the sum of the envies we could have an allocation such that one agent is very envious while the others are not. Hence we encounter a problem of fairness: if there has to be some envy in our society, then a legitimate objective could be to share this envy as fairly as possible among agents which we try to achieve by computing a fair Ordered Weighted Average (OWA) of the envies.

Chapter 5 is also about a relaxation of envy-freeness but that comes from a more collective/voting setting. The idea comes from the work by Parijs [1997] and Guibet Lafaye [2006] about unanimous envy that is described as the worst case of envy. An agent is unanimously envious if she envies an agent and that everyone agrees that she should be envious of this agent. This is quite a change of approach as the notion of preference is intrinsically subjective but we break that by introducing objectivity into it and discuss the various reasons why it is legitimate according to us. From that notion we propose a scale of approval envy corresponding to the maximum number of agents that agree on the envies of the agents and we try to minimize that number. In this chapter, after

formally defining this notion we investigate its links with existing fairness measures and other relaxations of EF. Besides, we introduce the same approval relaxation notion for proportionality (PROP) and in addition to doing the same work as for approval envy we investigate the relation between our two approval notions. This comparison is powered by the fact that our approval envy notion implies the proportionality one but the converse does not hold.

Finally, we conclude by summarizing our main contributions and proposing ideas for future works.

Chapter 1

State Of the Art

Summary of the chapter

In this chapter we will mainly formally introduce the fair division problem, then we will formally see what desirable criteria and properties we are looking for in the solutions of our problem. We will also see how we can reach such desirable solutions by introducing some procedures and what we can be expecting from them.

Finally, we will see how this thesis is positioned on these fields and we will sum up how this document is organized and the contributions and topics of our different chapters.

1.1 Introduction

This thesis focuses on multiagent fair resource allocation which consists in fairly and efficiently allocating a set of resources to a set of agents taking into account the agents' preferences about these resources. Fair division is an ubiquitous problem in multiagent systems or economics [Steinhaus, 1948; Moulin, 2003; Young, 1994], with applications ranging from allocation of schools, courses or rooms to students [Abraham et al., 2005; Othman et al., 2010], to division of goods in inheritance or divorce settlement [Brams and Taylor, 1996].

1.2 The fair division problem

After listing in the Introduction many real life problems that fall in the spectrum of fair division of indivisible goods, we use a simpler example in order to introduce the fair division problem and its different aspects. It is a classical example inspired by Budish [2011]:

Example 1.1. *Imagine that two agents want to share a diamond and two rocks between them. Both agents heavily prefer the diamond over the rocks such as they also prefer to have the diamond instead of the two rocks. We obviously see that the agent that will get the diamond (even if the other agent gets two rocks) will be better off. A solution could be to throw away the diamond and only allocate the rocks as it would lead to a very desirable solution (of one rock to each agent) in terms of fairness but it would really be a big waste (we will see the notion of efficiency and completeness later on). Moreover, the problem would also be very easy if we could divide the diamond in two parts but it is not possible here.*

There are two main fields in fair allocation: fair allocation of indivisible goods and fair allocation of divisible goods. Divisible resources are such that you can give some portion of it to one agent and some other portion to another or other agent(s). Obviously, indivisible resources refer to the ones that cannot be split between two agents. The divisible case allows some allocations that exhibit desirable properties that we could not obtain if items are indivisible. Indeed, in Example 1.1 dividing the diamond in half would allow us to reach a solution such that none of the agents envy the other one. As we will see throughout this chapter and this thesis there is not one and only one way to define what fairness and efficiency are. Moreover, the resources can be chores or tasks which the agents do not want to be allocated while it can also be goods.

Note that we have focused in our work on the fair division problem of indivisible goods. However we can see that some results and methods from the fair division problem of divisible goods can be of great help and use in the indivisible setting.

This section will shortly introduce and talk about fair division of indivisible goods from its origins to its applications. A formal introduction to the model will be given. As mentioned in the beginning of this section, the fair and efficient allocation is made regarding the preferences of the agents towards the resources hence we thoroughly introduce the different kinds of preferences that can be encountered in the literature.

1.2.1 Multiagent allocation of indivisible resources

Let us first formalize our problem. We consider MultiAgent Resource Allocation problems (MARA) where we aim at fairly and efficiently dividing a set of indivisible goods (also called items or objects) among a set of agents. A MARA instance I is defined as a finite set of *objects* $\mathcal{O} = \{o_1, \dots, o_m\}$, a finite set of *agents* $\mathcal{N} = \{a_1, \dots, a_n\}$ and a preference profile \mathcal{P} representing the interest of each agent $a_i \in \mathcal{N}$ towards the objects $o_j \in \mathcal{O}$.

A *solution* of the problem is called an allocation π which is a function that maps each agent a_i to a subset π_i of the objects \mathcal{O} . In this thesis, π_i will denote the set of objects (the share) held by agent a_i . An allocation is such that $\forall a_i, \forall a_j$ with $i \neq j : \pi_i \cap \pi_j = \emptyset$ (a given object cannot be allocated to more than one agent). We will denote by $\mathcal{A}(I)$ the set of allocations for I .

We also formalize the House Allocation Problem (HAP) that we will consider in this thesis which consists in the same problem but with some additional constraints. First, $n = m$ and we have an initial allocation denoted by π^0 associating each agent with an object. Moreover a solution of the problem is such that each agent obtains one and only one object (also called house hence the name of the problem).

1.2.2 Preferences

As discussed at the beginning of this section, the preferences of the agents are of paramount importance as they are the key to evaluate the satisfaction of the agents. We first introduce the two main types of preferences that we can encounter: ordinal or cardinal preferences. Finally, we will see that preferences can exhibit some structure that can be useful in designing a good procedure. It can also have an impact on the complexity of a procedure.

Ordinal preferences

Under ordinal preferences, agents are assumed to express their preferences over the resources through orders. More precisely, we will assume that each agent a_i is equipped with a binary relation \succ_i , where $\pi \succ_i \pi'$ denotes the fact that a_i strictly prefers bundle π to bundle π' . Moreover, we will say that a_i weakly prefers π to π' (denoted by $\pi \succeq_i \pi'$) if and only if $\pi \succ_i \pi'$ or $\pi = \pi'$.

Finally, a preference profile $L = \langle \succ_i \mid a_i \in \mathcal{N} \rangle$ is a vector of n orders, one for each agent.

Notice that expressing preferences over the whole set of bundles is not feasible in practice as there exists an exponential number of bundles [Bouveret et al., 2010] (keep in mind that for m objects you have to express 2^m preferences in order to relate the full preferences of an agent). To overcome this issue, compact preference representations have been investigated in the literature such as graphical representation or logical languages [Chevaleyre et al., 2008b].

In this thesis, we will also consider the particular setting where individual preferences are expressed as linear orders over single objects (and not over bundles). Notice that in that case, lifting preferences on objects to a preference relation over bundles can be challenging as discussed by Barberà et al. [2004].

Let us now introduce some further notation in the specific case where the agents' preferences are given as linear orders on objects. In that case we use $top(\succ)$ to denote its top-ranked resource: $\forall o \in \mathcal{O} \setminus \{top(\succ)\}, top(\succ) \succ o$. Similarly, $snd(\succ)$ refers to the second most preferred resource in \succ . With a slight abuse of notation we will write $top(a_i)$ and $snd(a_i)$ to refer to $top(\succ_i)$ and $snd(\succ_i)$. When it is not clear from the context we will subscript these notations to specify the resource set. For instance $top_R(a_i)$ is the most preferred resource for agent a_i among resources in $R \subseteq \mathcal{O}$.

Given a resource $r \in \mathcal{O}$ and an agent $a_i \in \mathcal{N}$, we use $rank_{a_i}(r)$ to refer to the rank of r in \succ_i . We have then $rank_{a_i}(top(a_i)) = m$, $rank_{a_i}(snd(a_i)) = m - 1$, etc...

With slight notation abuse, we will also denote by L the vector of linear orders over objects, when the agents preferences are represented using this preference model.

Example 1.2. *Let us consider the following instance with 2 agents and 3 objects which is an echo of the Example 1.1. Let us assume that o_1 represents the diamond while o_2 and o_3 represent two rocks.*

$$\begin{aligned} a_1 : o_1 \succ_1 o_2 \succ_1 o_3 \\ a_2 : o_1 \succ_2 o_3 \succ_2 o_2 \end{aligned}$$

Here, $top(a_1) = top(a_2) = o_1$, while $snd(a_1) = o_2$ and $snd(a_2) = o_3$. Under this preference representation, it is not clear how to lift preferences over bundles of objects: for instance, given only \succ_1 , it is not clear how $\{o_1\}$ and $\{o_2, o_3\}$ compare with each other.

Cardinal preferences

On the other hand, cardinal preferences are such that the preferences of an agent a_i over bundles of objects are defined by a *utility function* $u_i : 2^{\mathcal{O}} \rightarrow \mathbb{R}$ measuring her satisfaction $u_i(\pi_i)$ when she obtains share π_i . Note however that we will stay in \mathbb{Q}^+ in this thesis as we will not consider chores (positive utilities). This will also allow for a more efficient use of linear programming.

In this thesis we make the assumption that utility functions are additive i.e., the utility of an agent a_i for a share π_i is defined as the sum of the utilities over the objects forming π_i :

$$u_i(\pi_i) \stackrel{def}{=} \sum_{o_k \in \pi_i} u(i, k),$$

where $u(i, k)$ is the utility given by agent a_i to object o_k . This assumption is commonly considered in MARA [Lipton et al., 2004; Procaccia and Wang, 2014; Dickerson et al., 2014; Caragiannis et al., 2016] as additive utility functions provide a compact but yet expressive way to represent the preferences of the agents. MARA instances with additive utility functions are called add-MARA instances for short.

Note that k -additive utility functions extend standard additive functions to represent synergies between sets of resources of size at most k [Grabisch, 1997; Chevaleyre et al., 2008a].

Moreover, in this thesis, we will only consider MARA instances with *commensurable* preferences, that is, such that:

$$\exists K \in \mathbb{N} \text{ s.t } \forall i \in \llbracket 1, n \rrbracket, \sum_{j=1}^m u(i, j) = K$$

Example 1.3. Let us consider the following add-MARA instance with 2 agents and 3 objects which is an echo of the Example 1.1. Let us assume that o_1 represents the diamond while o_2 and o_3 represent two rocks.

	o_1	o_2	o_3
a_1	97	2	1
a_2	97	1	2

Note that contrary to the Example 1.2 with ordinal preferences there is no need to assume anything more than the additivity of the utilities to conclude that both agents prefer the diamond over the two rocks.

We say that the agents' preferences are *strict on objects* if, for each agent and each pair of objects, the agent necessarily strictly prefers one object over the other. Similarly, we say that the agents' preferences are *strict on shares* if, for each agent a_i and each pair of shares $\pi \neq \pi'$, we have $u_i(\pi) \neq u_i(\pi')$. Note that strict preferences on shares entail strict preferences on objects; the converse is false.

1.2.3 Domain Restriction

However, in real life preference profiles often exhibit some underlying structure. This is why it is common to consider some structures of preferences. We present here some of them that are often encountered in the literature. They can also be referred to as domain restrictions.

Single-peakedness

A widely studied one is the single-peakedness. Under single-peaked preferences, the agents are assumed to share a common axis \triangleleft over the resources and individual rankings are defined with respect to this axis.

This domain of preferences has been introduced by Black [1948] and Arrow [1951]. Initially motivated in voting contexts, it is now a well studied domain of preferences [Moulin, 1988; Elkind et al., 2017]. Numerous works have explored single-peaked preferences in the context of fair division. Sprumont [1991] for instance characterized the *uniform allocation rule*, the unique strategy-proof, efficient and anonymous allocation procedure with single-peaked preferences and divisible objects. Kasajima [2013] investigated probabilistic allocation of indivisible objects with single-peaked preferences. More recently, Hougard et al. [2014] extended this research area to indivisible objects and considered the problem of assigning agents to a line under single-peaked preferences. On the empirical side, it is not clear whether single-peaked preferences can be observed in real-life scenarios, as mentioned by Puppe [2018]. On the one hand, Spector [2000]; DeMarzo et al. [2003]; List et al. [2013] argue that interactions between the agents lead towards single-peaked preferences. On the other hand, Egan [2014] shows that political preferences might be double-peaked (and not single-peaked) when it comes to some polarizing topics.

Definition 1.1. Let \mathcal{O} be a set of resources and \triangleleft an axis (i.e. a linear order) over \mathcal{O} . We say that a linear order \succ is *single-peaked with respect to \triangleleft* if we have:

$$\forall (o_1, o_2) \in \mathcal{O}^2 \text{ s.t. } \left. \begin{array}{l} o_2 \triangleleft o_1 \triangleleft \text{top}(\succ), \\ \text{or, } \text{top}(\succ) \triangleleft o_1 \triangleleft o_2 \end{array} \right\} \Rightarrow o_1 \succ o_2.$$

In words, \succ is single-peaked with respect to \triangleleft if \succ is decreasing on both the left and the right sides of $top(\succ)$, where left and right are defined by \triangleleft .

For a set of resources \mathcal{O} , we denote by $\mathcal{L}_{\mathcal{O}}$ the set of all linear orders over \mathcal{O} . Any subset $D \subseteq \mathcal{L}_{\mathcal{O}}$ is then called a preference domain. For a given axis \triangleleft , we call $\mathcal{SP}_{\triangleleft}$ the set of all the linear orders on resources single-peaked with respect to \triangleleft :

$$\mathcal{SP}_{\triangleleft} = \{\succ \in \mathcal{L}_{\mathcal{O}} \mid \succ \text{ is single-peaked w.r.t. } \triangleleft\}.$$

A preference domain D is called single-peaked if and only if there exists an axis \triangleleft such that $D \subseteq \mathcal{SP}_{\triangleleft}$. An instance I is said to be single-peaked if it is defined over a single-peaked preference domain.

Example 1.4. Consider the following four linear orders defined over 3 objects.

$$\begin{aligned} \succ_1: & o_1 \succ_1 o_2 \succ_1 o_3 \\ \succ_2: & o_3 \succ_2 o_2 \succ_2 o_1 \\ \succ_3: & o_2 \succ_3 o_1 \succ_3 o_3 \\ \succ_4: & o_2 \succ_4 o_3 \succ_4 o_1 \end{aligned}$$

One can check that these linear orders represent a single-peaked preference profile with respect to \triangleleft defined as: $o_1 \triangleleft o_2 \triangleleft o_3$. In fact these orders exactly correspond to $\mathcal{SP}_{\triangleleft}$.

Observation 1.1. Let D be a preference domain single-peaked over \triangleleft . For a subset of objects $\mathcal{O}' \subseteq \mathcal{O}$, the domain D' defined as the restriction of D to \mathcal{O}' is a single-peaked domain over \triangleleft' , the restriction of \triangleleft to \mathcal{O}' .

Ballester and Haeringer [2011] provided a characterization of single-peaked domains. In particular, they gave a necessary condition for a domain to be single-peaked: it should be worst-restricted [Sen, 1966].

Definition 1.2. An instance $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$ is worst-restricted if for any triplet of objects $(o_x, o_y, o_z) \in \mathcal{O}^3$, one of them is never ranked last in the restriction of L to these three objects.

Proposition 1.1 (Ballester and Haeringer, 2011). If an instance is single-peaked then it is worst-restricted.

Let us illustrate the single-peaked domain with a simple example.

Example 1.5. Consider the following three linear orders defined over 3 objects that are not single-peaked:

$$\begin{aligned} \succ_1: & o_1 \succ_1 o_2 \succ_1 o_3 \\ \succ_2: & o_3 \succ_2 o_1 \succ_2 o_2 \\ \succ_3: & o_2 \succ_3 o_3 \succ_3 o_1 \end{aligned}$$

It can be checked that there is no linear order \triangleleft over which these preferences are single-peaked. Indeed, they are not worst-restricted: every object of the triplet (o_1, o_2, o_3) is ranked last at least once, which violates Proposition 1.1.

Binary preferences

Another common domain restriction, derived from approval voting, is the binary preferences in which agents have only two possible valuations the objects [Halpern et al., 2020]. We say that an instance has binary preferences if for every agent a_i and every object o_k , the value of a_i for o_k is either 0 or 1, i.e. $u(i, k) \in \{0, 1\}$. This domain restriction corresponds to settings where the agents just tell whether they like or dislike each object [Aleksandrov et al., 2015a]. We will use this domain restriction in some complex problems to derive some intuitions on the more global scale.

1.2.4 Preference cultures

When drawing experiments to test and validate the theory, one may have to automatically generate preference profiles. In social choice, the probability distribution used to generate preference profiles is called the culture. Impartial culture refers to the case where the probability distribution is uniform [Garman and Kamien, 1968].

However, impartial culture may be considered as unrealistic since in real-world setting agents may have similar or dependent preferences [Deemen, 2014]. In strict ordinal settings, a classical way to capture correlated preferences is to use Mallows distributions [Mallows, 1957] allowing us to measure the impact of the similarity of the preferences between agents. This notion is conveyed by a parameter ϕ called dispersion: when $\phi = 0$ all the agents have the same preferences while the bigger ϕ gets the more we tend towards uniformly distributed preferences. For cardinal preferences, we can use a generalization of the Mallows distribution to cardinal preferences presented by Durand et al. [2016] based on Von Mises–Fisher distributions. Similarly to the dispersion parameter in Mallows distributions, the similarity between the preferences of the agents is tuned by the *concentration* parameter: when it is zero agents' preferences are uniformly distributed, whereas when it is infinite agents have the same preferences.

1.3 Desirable criteria and properties

Now that we have seen how the preferences of the agents can be expressed and that we have seen what the fair division problem of indivisible goods is about, we can look at what are the good properties that we can expect a solution (an allocation of the objects to the agents) to have.

1.3.1 Classical fairness measures

We will first see that fairness does not mean one and only one thing but can be expressed by a lot of criteria, fairness measures, some of which we describe here. A fairness measure is a score (that can be boolean) that estimates how fair a solution is.

Max min fair Share (MMS): I cut, you choose

Max min fair share is inherited from the notion of "I cut you choose" in the problem of cake cutting in the fair division problem of divisible goods. It has been built around the idea of determining what is the maximum amount an agent can get if she divides (cuts)

the objects in n bundles but chooses last her bundle. This notion was introduced by Budish [2011] building on concepts introduced by Moulin [1990].

Definition 1.3. For a MARA instance I we define the max-min share (MMS) of agent a_i as follows:

$$u_i^{MMS} = \max_{\pi \in \mathcal{A}(I)} \min_{j \in \mathcal{N}} u_i(\pi_j)$$

Besides, an allocation π is said max-min fair share if every agent gets at least her max-min share; formally: $\forall i, u_i(\pi_i) \geq u_i^{MMS}$.

Example 1.6. Let us consider the following add-MARA instance with 3 agents and 6 objects:

	o_1	o_2	o_3	o_4	o_5	o_6
a_1	1	2	3	1	3	2
a_2	2	0	7	2	1	0
a_3	0	3	5	0	1	3

We can compute the max-min fair shares of each agent introduced in Definition 1.3:

$$\begin{aligned} u_1^{MMS} &= 4 \\ u_2^{MMS} &= 2 \\ u_3^{MMS} &= 3 \end{aligned}$$

From that, we can easily see that the squared allocation is MMS as:

$$\begin{aligned} u_1^{MMS} &= 4 \leq 6 = 3 + 3 = u_1(\pi_1) \\ u_2^{MMS} &= 2 \leq 2 = 2 + 0 = u_2(\pi_2) \\ u_3^{MMS} &= 3 \leq 3 = 0 + 3 = u_3(\pi_3) \end{aligned}$$

Complexity The problem of finding an MMS allocation is NP-Hard [Bouveret and Lemaître, 2016a].

Proportionality

Proportionality is also a very natural notion reached when all the agents have at least $\frac{1}{n}$ of the utility of the whole set of objects [Steinhaus, 1948].

Definition 1.4. For a MARA instance I we define the proportional share (Prop) of agent a_i as follows:

$$Prop_i = u_i^{PROP} = \frac{u_i(\mathcal{O})}{n}$$

Besides, we say that an allocation π is proportional if every agent gets at least her proportional share; formally: $u_i(\pi_i) \geq u_i^{PROP}$.

Example 1.7. Let us consider the same instance as the one introduced in Example 1.6:

	o_1	o_2	o_3	o_4	o_5	o_6
a_1	1	2	3	1	3	2
a_2	2	0	7	2	1	0
a_3	0	3	5	0	1	3

We can compute the proportional shares of each agent introduced in Definition 1.4:

$$\begin{aligned} u_1^{PROP} &= \frac{1+2+3+1+3+2}{3} = 4 \\ u_2^{PROP} &= \frac{2+0+7+2+1+0}{3} = 4 \\ u_3^{PROP} &= \frac{0+3+5+0+1+3}{3} = 4 \end{aligned}$$

From that, we can easily see that the squared allocation is proportional as:

$$\begin{aligned} u_1^{PROP} &= 4 \leq 5 = 3 + 2 = u_1(\pi_1) \\ u_2^{PROP} &= 4 \leq 4 = 2 + 2 = u_2(\pi_2) \\ u_3^{PROP} &= 4 \leq 4 = 3 + 1 = u_3(\pi_3) \end{aligned}$$

Complexity The problem of finding a Proportional allocation is NP-Complete [Bouveret and Lemaître, 2016a].

min Max Share (mMS): You cut, I choose

min Max fair Share is inherited from the notion of "You cut, I choose" in the problem of cake cutting in the fair division problem of divisible goods:

Definition 1.5. For a MARA instance I we define the min-max share (mMS) of agent a_i as follows:

$$u_i^{mMS} = \min_{\pi \in \mathcal{A}(I)} \max_{j \in \mathcal{N}} u_i(\pi_j)$$

Besides, we say that an allocation π is min-max fair share if every agent gets at least her min-max share; formally: $u_i(\pi_i) \geq u_i^{mMS}$

Example 1.8. Let us consider the same instance than the one introduced in Example 1.6:

	o_1	o_2	o_3	o_4	o_5	o_6
a_1	1	2	3	1	3	2
a_2	2	0	7	2	1	0
a_3	0	3	5	0	1	3

By computing the min max fair shares we can see the squared allocation is mMS:

$$\begin{aligned} u_1^{mMS} &= 4 \leq 4 \\ u_2^{mMS} &= 7 \leq 7 \\ u_3^{mMS} &= 5 \leq 6 \end{aligned}$$

Complexity The problem of finding an mMS allocation is NP-Hard [Bouveret and Lemaître, 2016a].

Envy-freeness (EF)

Envy-freeness (EF), is one of the prominent notions studied in fair division [Foley, 1967; Brams and Fishburn, 2002; Lipton et al., 2004; de Keijzer et al., 2009; Segal-Halevi and Suksompong, 2019]. An allocation of items among a set of agents is said to be envy-free if no agent prefers the share of another agent to her own share.

Definition 1.6. An allocation π of an add-MARA instance is called *envy-free* if no agent strictly prefers the bundle of another agent over hers. Formally:

$$\forall a_i \in \mathcal{N}, \nexists a_j \in \mathcal{N} \quad u_i(\pi_i) < u_i(\pi_j)$$

Envy-freeness is a boolean criterion that can be refined to measure the degree of pairwise envy.

Definition 1.7. We define the notion of *pairwise envy* of an allocation π $pe(i, j, \pi)$ from the point of view of agent i between her and agent j as the maximum between 0 and the difference between the utility for agent i if she had agent j 's bundle and her actual utility in the allocation π . It can be read as how much agent i envies agent j 's bundle. Formally:

$$pe(i, j, \pi) \stackrel{\text{def}}{=} \max\{0, u_i(\pi_j) - u_i(\pi_i)\}$$

Definition 1.8. From that notion of pairwise envy, we derive the notion of *envy* $e_i(\pi)$ of an agent i as follows:

$$e_i(\pi) \stackrel{\text{def}}{=} \max_{j \in \mathcal{N}} pe(i, j, \pi)$$

Definition 1.9. Finally, we define $\vec{e}(\pi)$ as the vector of every agent's envy of an allocation π :

$$\vec{e}(\pi) \stackrel{\text{def}}{=} (e_1(\pi), \dots, e_n(\pi))$$

Note that an allocation π is envy-free if and only if $\vec{e}(\pi) = (0, \dots, 0)$.

Definition 1.10 (Degree of envy of the society). The *degree of envy* $de(\pi)$ of the society for an allocation π is defined as follows:

$$de(\pi) \stackrel{\text{def}}{=} \sum_{a_i \in \mathcal{N}} \sum_{a_j \in \mathcal{N}} pe(i, j, \pi).$$

Note that an allocation π is envy-free if and only if $de(\pi) = 0$. Moreover, keep in mind that the degree of envy of the society can be defined differently as discussed by Chevalyere et al. [2017]. Indeed, the aggregation of the envies of the agents can be done via a maximum. More than that, even pairwise envy could be defined as possibly negative if we want a negative envy (measuring how much the agent is not envious and so happy) to compensate for a positive envy. We will ourselves propose another definition of the degree of envy of the society in Chapter 4.

Example 1.9. Let us consider the same instance as the one introduced in Example 1.6:

	o_1	o_2	o_3	o_4	o_5	o_6
a_1	1	2	3	1	3	2
a_2	2	0	7	2	1	0
a_3	0	3	5	0	1	3

We can first see that the squared allocation is not envy-free. Indeed, a_1 envies a_3 as $u_1(\{o_4, o_5\}) = 4 < 5 = u_1(\{o_1, o_2, o_6\})$. However, from the latter allocation, by simply giving o_1 to a_1 we get this new allocation:

	o_1	o_2	o_3	o_4	o_5	o_6
a_1	1	2	3	1	3	2
a_2	2	0	7	2	1	0
a_3	0	3	5	0	1	3

This new allocation is envy-free.

Complexity It is known that the problem of finding an envy-free allocation is NP-Complete [Lipton et al., 2004].

Competitive Equilibrium with Equal Incomes (CEEI)

The notion of *competitive equilibrium* is an old and well-known concept in economics [Walras, 1874; Fisher, 1892]. If equal incomes are imposed among the stakeholders, this concept becomes the *competitive equilibrium from equal incomes* [Moulin, 2003], yielding a very strong fairness concept that has been recently explored both in artificial intelligence and in economics [Othman et al., 2010; Budish, 2011; Bouveret and Lemaître, 2016a].

Definition 1.11. Let $I = (\mathcal{N}, \mathcal{O}, w)$ be an add-MARA instance, π an allocation, and $\vec{p} \in [0, 1]^m$ a vector of prices. A pair (π, \vec{p}) is said to form a competitive equilibrium with equal incomes (CEEI) if

$$\forall a_i \in \mathcal{N} : \pi_i \in \operatorname{argmax}_{\pi \subseteq \mathcal{O}} \left\{ u_i(\pi) : \sum_{o_k \in \pi} p_k \leq 1 \right\}$$

In other words, π_i is one of the maximal shares that a_i can buy with a budget of 1, given that the price of each object o_k is p_k .

We will say that allocation π is a CEEI if there exists a vector \vec{p} such that (π, \vec{p}) forms a CEEI.

Example 1.10. Let us consider the same instance as the one introduced in Example 1.6 where we also consider a vector of prices \vec{p} :

		o_1	o_2	o_3	o_4	o_5	o_6
	p_j	0.2	0.5	1	0.2	0.6	0.5
1	a_1	1	2	3	1	3	2
1	a_2	2	0	7	2	1	0
1	a_3	0	3	5	0	1	3

The squared allocation is CEEI: Competitive Equilibrium with Equal Incomes with prices p_j .

Complexity The CEEI test is known to be NP-Hard [Brânzei et al., 2015], and, to the best of our knowledge, no practical method has been described before.

Link between fairness measures

Note that it has also been shown [Bouveret and Lemaître, 2016a] that when preferences are additive $\text{CEEI} \implies \text{EF} \implies \text{mMS} \implies \text{PROP} \implies \text{MMS}$, where $FM_1 \implies FM_2$ means that an allocation satisfying FM_1 necessarily satisfies FM_2 and that we can find instances for which an allocation satisfies FM_2 but none that satisfy FM_1 .

1.3.2 Relaxations of fairness measures

In recent literature a lot of relaxations of fairness notions have been introduced and studied. Indeed, when it is not guaranteed to find an allocation satisfying a given fairness criterion it seems interesting to study relaxations of it. Note that none of the classical fairness measures introduced in Section 1.3.1 are guaranteed to exist. Indeed, for every fairness introduced in Section 1.3.1 we can find an instance for which no allocation satisfies the said fairness measure. It can be easily observed for Proportionality (which implies no guarantee of existence of stronger). However, for MMS the problem remained open for a while but it was settled by Procaccia and Wang [2014] that an MMS allocation may not exist either. These measures are hence good candidates to investigate relaxations.

Envy-freeness relaxations

For example, to cope with the possible non-existence of an envy-free allocation, another approach is to alleviate the requirements of the fairness notion. Several relaxations of envy-freeness have been proposed such as envy-freeness up to one good (EF1) [Lipton et al., 2004; Budish, 2011] or envy-freeness up to any good (EFX) [Caragiannis et al., 2016]. An allocation is said to be envy-free up to one good (resp. up to any good) if no agent a_i envies the share π_j of another agent a_j after removing from π_j *one* (resp. *any*) item.

We give the formal definitions of these two relaxations here:

Definition 1.12. *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and π be an allocation of I . π is envy-free up to one good EF1 if and only if $\forall a_i, a_j \in \mathcal{N}$, either $u_i(\pi_i) \geq u_i(\pi_j)$ or $\exists o_k \in \pi_j$ such that $u_i(\pi_i) \geq u_i(\pi_j \setminus \{o_k\})$.*

Definition 1.13. *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and π be an allocation of I . π is envy-free up to any (strictly positively valued) good EFX if and only if $\forall a_i, a_j \in \mathcal{N}$, either $u_i(\pi_i) \geq u_i(\pi_j)$ or $\forall o_k \in \pi_j$ for which $u(i, k) > 0$, $u_i(\pi_i) \geq u_i(\pi_j \setminus \{o_k\})$.*

An even more demanding notion called EFX_0 [Plaut and Roughgarden, 2018; Kyropoulou et al., 2020] differs on the fact that an agent can forget any object even the ones valued to 0:

Definition 1.14. *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and π be an allocation of I . π is envy-free up to any good if and only if $\forall a_i, a_j \in \mathcal{N}$, either $u_i(\pi_i) \geq u_i(\pi_j)$ or $\forall o_k \in \pi_j$, $u_i(\pi_i) \geq u_i(\pi_j \setminus \{o_k\})$.*

Clearly, we have $EF \implies EFX_0 \implies EFX \implies EF1$.

Existence for EF1 is guaranteed, but this is still to the best of our knowledge an open question for EFX in the general case. However, the existence guarantee of an EFX solution has been proven for few agents (at most 3 agents) and specific utility functions. For instance it has been proved that an EFX allocation always exists for instances with identical valuations and for instances involving two agents with general and possibly distinct valuations [Plaut and Roughgarden, 2018], as well as for three agents with additive valuations [Chaudhury et al., 2020]. When the objects have only two possible valuations and with additive preferences, Amanatidis et al. [2020] proved that any allocation maximizing

the Nash Social Welfare is EFX. This result provides a polynomial algorithm for computing EFX allocations in the two-agent setting. More recently, Bu et al. [2023] proved that an EFX allocation always exists under binary preferences even if not additive. An important effort has also been dedicated to EFX approximations such as the notion of α -EFX [Plaut and Roughgarden, 2018; Feldman et al., 2023] or partial EFX allocations where some objects remain unallocated [Berger et al., 2022].

Amanatidis et al. [2018] studied four fairness notions – Envy-freeness up to one good (EF1), envy-freeness up to any good (EFX), maximin share fairness (MMS), and pairwise maximin share fairness (PMMS) – and investigated the relations between these notions and their relaxations. Although PMMS is a stronger notion than EFX, Amanatidis et al. [2018] proved that both notions provide the same worst-case guarantee for MMS. In the same vein, they showed that EFX and EF1 both provide similar approximation for PMMS.

Proportionality relaxations

Another example of relaxations that can be encountered in the literature is the one of proportionality. Even though proportionality is a less demanding fairness criterion than envy-freeness [Bouveret and Lemaître, 2016a], the existence of a proportional allocation is not guaranteed. For that reason, relaxations of this notion such as proportionality up to one item (PROP1) has been proposed [Conitzer et al., 2017]. An allocation satisfies PROP1 if every agent gets at least her proportional share when one item is added to her current bundle. For example, Conitzer et al. [2017] proved the existence of PROP1 allocations for a public decision setting where a decision has to be made on several public issues. Each issue has several possible alternatives and each agent has a utility for each alternative. The decision problem consists in choosing one alternative for each issue. Considering PROP1 allocations, Barman and Krishnamurthy [2019] presented a strongly polynomial-time algorithm to find such allocations for positive utilities. Note that some of these papers deal with fair division with chores [Brânzei and Sandomirskiy, 2019] or with mixed utilities [Aziz et al., 2019, 2020]. It can be noticed that there is a similar kind of link between EF1 and PROP1 as the one that exists between EF and proportionality. Namely, any EF1 allocation is also PROP1 (which we could write $EF1 \implies PROP1$ for short).

Another example of an approximation of a fairness measure is the one of CEAI (Competitive Equilibrium with Almost equal Incomes) which consists in relaxing the fact that every object has to have the same price for every agent [Budish, 2011; Segal-Halevi, 2020].

1.3.3 Efficiency

There exists an extensive literature investigating how to define efficiency requirements for an allocation (see Chevaleyre et al., 2006; Thomson, 2016 for some surveys).

Completeness

The first and simplest way to express efficiency is through completeness. We do not want any item to be wasted meaning we want all the objects from \mathcal{O} to be allocated:

$$\bigcup_{a_i \in \mathcal{N}} \pi_i = \mathcal{O}$$

It would indeed seem very inefficient to throw away one or multiple objects in order to achieve fairness. Indeed, if we take the example of two agents that have to share two rocks and one diamond, we could easily say that we do not allocate the diamond and give one rock to each agent. However, one can sense that it is not an acceptable solution and that is why we assume completeness throughout this document.

Pareto efficiency

Pareto efficiency is considered and known as the gold standard in terms of efficiency.

Definition 1.15 (Pareto-Optimality). *Let π and π' be two allocations. We say that π' Pareto-dominates π if and only if:*

$$\forall a_i \in \mathcal{N}, \pi'_i \succeq_i \pi_i, \text{ and } \exists a_j \in \mathcal{N}, \pi'_j \succ_j \pi_j.$$

An allocation π is then said to be Pareto-Optimal (PO) if there is no other allocation π' that Pareto-dominates π .

Example 1.11. *Let us go back to this example with 2 agents and 3 objects :*

	o_1	o_2	o_3
a_1	$\boxed{97}^*$	2^*	$\boxed{1}$
a_2	97	$\boxed{1}$	2^*

We can see that the squared allocation is Pareto dominated by the starred one (which is a PO allocation). Moreover, note that to go from the squared allocation to the starred one we make a Pareto improving exchange between two agents. This will be thoroughly studied and used through this thesis. We can also notice that if we give o_1 to a_2 in the starred allocation it is still a Pareto Optimal allocation.

We can also see how Pareto-Optimality similarly works for the HAP.

Example 1.12. *Let us consider the following instance with 5 agents and 5 resources. The preferences presented below are single-peaked with respect to $o_1 \triangleleft o_2 \triangleleft o_3 \triangleleft o_4 \triangleleft o_5$. The initial allocation $\pi^0 = \langle o_1, o_2, o_3, o_4, o_5 \rangle$ is represented by the underlined resources.*

$$\begin{aligned} a_1 : o_3 \succ_1 o_4 \succ_1 o_5 \succ_1 \boxed{o_2} \succ_1 \underline{o_1} \\ a_2 : o_3 \succ_2 \boxed{o_4} \succ_2 o_5 \succ_2 \underline{o_2} \succ_2 o_1 \\ a_3 : o_4 \succ_3 \boxed{o_5} \succ_3 \underline{o_3} \succ_3 o_2 \succ_3 o_1 \\ a_4 : \boxed{o_3} \succ_4 \underline{o_4} \succ_4 o_5 \succ_4 o_2 \succ_4 o_1 \\ a_5 : \boxed{o_1} \succ_5 o_2 \succ_5 o_3 \succ_5 o_4 \succ_5 \underline{o_5} \end{aligned}$$

The allocation π^0 is not Pareto-optimal as it is Pareto-dominated by the squared allocation $\boxed{\pi} = \langle o_2, o_4, o_5, o_3, o_1 \rangle$.

1.4 Computing fair and efficient allocations

We have seen several desirable properties about the solution of our problems. However, we have not yet seen the major issue consisting in finding allocation procedures that return an allocation satisfying a mixture of desirable properties. An allocation procedure M is a mapping taking as input an instance I and returning an allocation $M(I)$.

1.4.1 Desirable properties of a procedure

As we have seen in Section 1.3 there are desirable properties we aim an allocation to have (such as Pareto-Optimality) but there are also for the procedures. As discussed in Schmidlein and Endriss [2023] in the context of voting, several axioms have been identified in the literature such as anonymity (changing the names of the agents should not change the outcome of the vote) and neutrality (changing the names of the objects should not change the outcome of the vote). We will detail here two other desirable properties of allocation procedures that are relevant in our context and throughout this work: individual rationality and strategy-proofness.

Individual rationality

In the House Allocation Problem (HAP), given that the agents initially own the objects in our setting, another very common efficiency requirement is that of *individual rationality*. It stipulates that no agent should be worse-off in the final allocation.

Definition 1.16 (Individual rationality). *A procedure M is said to be individually rational if for every instance $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$ and every agent $a_i \in \mathcal{N}$, we have:*

$$M(I) \succeq_i \pi_{a_i}^0$$

Strategy-proofness

We conclude by introducing another desirable property that is *strategy-proofness*. Informally, an allocation procedure is strategy-proof if no agent could get a strictly better outcome by lying instead of revealing her true preferences.

Definition 1.17. *Let $L = \langle \succ_j \rangle_{a_j \in \mathcal{N}}$ be a preference profile. For a given agent $a_i \in \mathcal{N}$, an i -variant $L_{-i} = \langle \succ'_j \rangle_{a_j \in \mathcal{N}}$ of L is a preference profile such that:*

$$\forall a_j \in \mathcal{N} \setminus \{a_i\}, \succ_j = \succ'_j \text{ and } \succ_i \neq \succ'_i.$$

An allocation procedure M is said to be strategy-proof if for every instance $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$, every agent $a_i \in \mathcal{N}$, every i -variant L_{-i} of L , we have:

$$M(\langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle) \succeq_i M(\langle \mathcal{N}, \mathcal{O}, L_{-i}, \pi^0 \rangle).$$

1.4.2 Centralized procedures

In centralized procedures we often have a central coordinator that refers to some central authority that runs the procedure. It basically runs the algorithm by communicating with agents, by harvesting their preferences or the choices they make during the algorithm runs. The central coordinator may be a real human being but it can just also be a computer or a robot playing this role.

Optimization of a collective utility function

There are a lot of known collective utility functions that we can seek to optimize in order to express fairness in our problem. For example, the maximization of Max Nash Welfare [Caragiannis et al., 2016] but also the maximization of the Ordered Weighted Average of the utilities [Perny and Spanjaard, 2003]. They are indeed centralized procedures as we can solve them via solving a Mixed Integer Linear Program.

We can also define collective functions if we deal with ordinal preferences and while we have seen that efficiency can be evaluated by achieving Pareto-Optimality, efficiency can also be evaluated via social welfare measures such as the *average rank* (ark) of the objects held by the agents, defined in the context of the House Allocation Problem as:

$$ark(\pi) = \frac{1}{n} \cdot \sum_{a_i \in \mathcal{N}} rank_{a_i}(\pi_{a_i}).$$

Maximizing the average rank is of course equivalent to maximizing the sum of ranks, and can also be interpreted as the utilitarian social welfare, under the assumption that agents have Borda utilities.

However, maximizing the average rank or searching for Pareto-optimal solutions may not be satisfactory as it can lead to particularly unfair allocations. For this reason, many fairness criteria have been introduced. We will focus here on maximizing the *minimum rank* (mrk) of the objects held by the agents, defined in the context of the House Allocation Problem as:

$$mrk(\pi) = \min_{a_i \in \mathcal{N}} rank_{a_i}(\pi_{a_i}).$$

Again, if we were to interpret the rank as a cardinal utility function, the minimum rank would be equivalent to the egalitarian welfare. Maximizing the minimum rank follows Rawls' principle of maximizing the welfare of the worst-off [Rawls, 1971]. It has been introduced by Pazner and Schmeidler [1978] and is now a very common rule in fair division [Thomson, 1983; Sprumont, 1996; Nguyen et al., 2014].

Example 1.13. *Let us go back to Example 1.12 to see how we would compute these ranks.*

We have $ark(\overline{\pi}) = \frac{2+4+4+5+5}{5} = 4$ and $mrk(\overline{\pi}) = \min(2, 4, 4, 5, 5) = 2$.

Top Trading Cycle

Let us present the following procedure in the context of the House Allocation Problem and when preferences on objects are distinct. The Top Trading Cycle (TTC) algorithm [Shapley and Scarf, 1974] takes as input an instance $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$ and proceeds as follows. The algorithm maintains a set of available agents N and a set of available objects O where initially $O = \mathcal{O}$ and $N = \mathcal{N}$. At each step of the algorithm, a directed bipartite graph $G = \langle V, E \rangle$, with $V = N \cup O$, is defined. The nodes of G represent the agents in N and the objects in O , and the set of edges E is such that:

- there is a directed edge (a_i, o_i) between a_i and o_i if and only if $o_i = top_O(a_i)$ i.e., agents are linked to their preferred object in O ,
- there is a directed edge (o_i, a_i) between o_i and a_i if and only if $o_i = \pi_{a_i}^0$ i.e., objects are linked to their owner in π^0 .

Algorithm 1.1: TTC algorithm

Input: An instance $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$

Output: An allocation π

```
1  $\pi \leftarrow$  empty allocation
2  $O \leftarrow \mathcal{O}$ : list of objects
3  $N \leftarrow \mathcal{N}$ : list of agents
4 while  $N \neq \emptyset$  do
5    $E \leftarrow \emptyset$ 
6    $V \leftarrow N \cup O$ 
7   for each agent  $a \in N$  do
8     | Add a directed edge in  $E$  between  $a$  and  $top_O(a)$ 
9   end
10  for each object  $o \in O$  do
11    | Add a directed edge in  $E$  between  $o$  and her owner in  $\pi^0$ 
12  end
13  Select a cycle  $\mu$  from the graph  $G = \langle V, E \rangle$ 
14   $do(\mu, \pi)$ 
15 end
16 return  $\pi$ 
```

Note that there always exists at least one cycle in G and that cycles correspond to improving deals (even in the worst case where only one agent is involved and hence the deal is weakly improving as she keeps her preferred object). The cycle-deals constructed can be of size 1 if an agent already owns her top object in O . Note that the size of a cycle-deal refers here to the number of agents involved so the size of the cycle in the TTC graph would be the double. The TTC algorithm selects one of the cycles μ in G and performs a reallocation of the objects following the cycle. The agents and objects involved in μ are then removed from N and O . A new graph G' is computed with the remaining agents and objects. The process is iterated on the new graph G' and $\pi^0[\mu]$ until an empty graph has been reached. Note that the outcome of TTC is unique (regardless of the possibly different choices of cycles to implement during the process).

A formal description of TTC procedure is given in Algorithm 1.1. Note that we make use of the sub-procedure $do(\mu, \pi)$ which simply implements the cycle-deal μ (reallocating the objects to the agents following the arcs of the cycle) and adds the resulting allocation (restricted to the agents involved in the deal) in the allocation π . It then removes the agents and the objects involved in the deal μ from the lists of available agents and objects, N and O respectively.

Example 1.14. Let us consider the instance defined in Example 1.12 with $\pi^0 = \langle o_1, o_2, o_3, o_4, o_5 \rangle$ being the underlined allocation:

$$\begin{aligned}
a_1 &: o_3 \succ_1 o_4 \succ_1 o_5 \succ_1 o_2 \succ_1 \underline{o_1} \\
a_2 &: o_3 \succ_2 o_4 \succ_2 o_5 \succ_2 \underline{o_2} \succ_2 o_1 \\
a_3 &: o_4 \succ_3 o_5 \succ_3 \underline{o_3} \succ_3 o_2 \succ_3 o_1 \\
a_4 &: o_3 \succ_4 \underline{o_4} \succ_4 o_5 \succ_4 o_2 \succ_4 o_1 \\
a_5 &: o_1 \succ_5 o_2 \succ_5 o_3 \succ_5 o_4 \succ_5 \underline{o_5}
\end{aligned}$$

The first graph used by the TTC algorithm is shown in Figure 1.1(a).

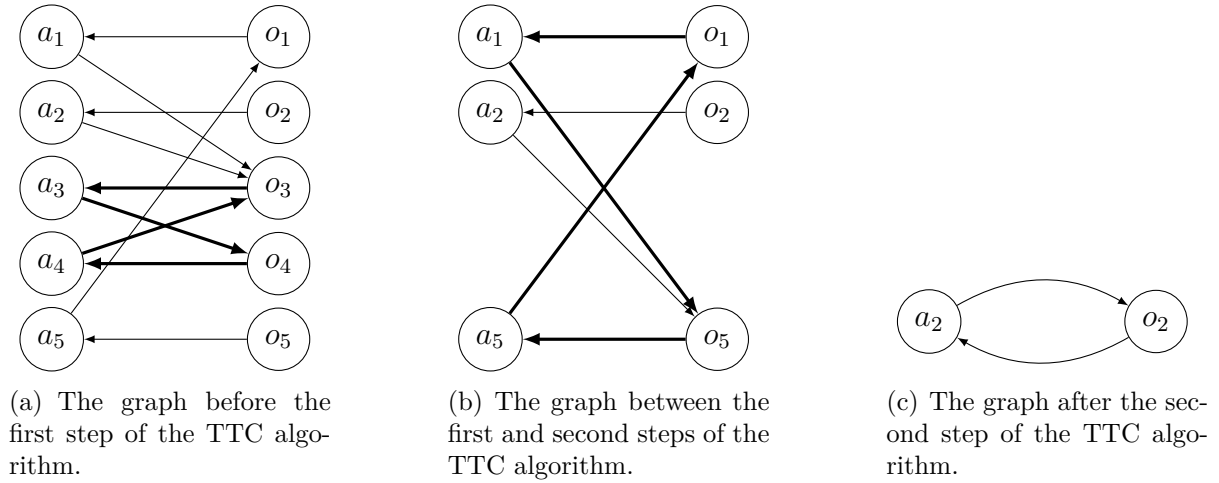


Figure 1.1: Evolution of the reallocation graph during the execution of the TTC algorithm on Example 1.14.

The first cycle-deal that can be applied is then $\mu_1 = \langle a_3, a_4 \rangle$ (highlighted in thick lines in Figure 1.1(a)) which means that a_3 gives her object to a_4 and vice-versa. Which gives us allocation $\pi^1 = \langle o_1, o_2, o_4, o_3, o_5 \rangle$ with the resulting graph given in Figure 1.1(b).

This means that the next cycle-deal is $\mu_2 = \langle a_1, a_5 \rangle$ (highlighted in thick lines in Figure 1.1(b)) so a_1 and a_5 swap their objects giving us $\pi^2 = \langle o_5, o_2, o_4, o_3, o_1 \rangle$. The next TTC graph is given in Figure 1.1(c).

Finally, a_2 remains with her initial object o_2 which corresponds to the cycle $\mu_3 = \langle a_2 \rangle$. Algorithm is finished as there are no agents left. Hence TTC returns $\pi^{TTC} = \langle o_5, o_2, o_4, o_3, o_1 \rangle$.

TTC algorithm which can be used without any domain restriction, is well known to satisfy the three main desirable properties of an allocation procedure: *Pareto-Optimality*, *individual rationality* and *strategy-proofness*.

1.4.3 Decentralized protocols

Even if for the sake of simplicity we have chosen to present centralized and distributed procedures as two distinct groups, it would actually be more accurate to consider that there is a continuum between the different kinds of protocols. Indeed, the most centralized

procedures rely on a central authority to which the agents just give their preferences and wait for the outcome. We move further towards decentralization by involving the agents in the procedure even in the presence of a central authority and we go towards full decentralization by getting rid of the central coordinator.

The use of a central coordinator induces a weak point in the system: the coordinator causes a bottleneck whose default often leads to the failure of the whole allocation process. The system may also be inherently distributed and the use of a central coordinator may not be possible because of limitations in the communication infrastructure or because of privacy requirements. In addition, centralized procedures may be perceived as less fair by the agents who are not allowed to actively take part in the allocation process [Van den Bos et al., 1997; Leventhal, 1980; Thibaut and Walker, 1975]. While TTC is presented as a centralized procedure (*i.e.* a central authority computes each trading cycle that should be implemented, and dictates it to the agents¹), distributed procedures [Sandholm, 1998; Chevaleyre et al., 2017] take a different perspective. Agents autonomously negotiate over the objects and locally agree on deals, and the outcome of the object allocation problem results from the sequence of such local deals. Such dynamics may achieve interesting efficiency and fairness results [Chevaleyre et al., 2007; Endriss et al., 2006]. However, long cycles still pose real challenges for these procedures [Rosenschein and Zlotkin, 1994] as they involve distributed coordination among numerous agents.

Resource exchanges

As stated in the beginning of this section it would be interesting to design procedures that rely more heavily on the agents rather than on a central authority. And the ones relying on resource exchanges seem to be good candidates as the agents negotiate and make deals themselves with one or more agents by exchanging one or several objects in order to be more satisfied.

As discussed in the introduction of this manuscript and this section and motivated by Damamme et al. [2015] resource exchanges are a very powerful tool to implement decentralized procedures and even if they may induce a high efficiency loss. Indeed, that loss is not that punishing under the assumption of individual rationality and that in single-peaked domains PO can be guaranteed.

Improving deals

A deal is a vector of agents, usually denoted by $\mu = \langle a_1, \dots, a_k \rangle$, where $a_i \in \mathcal{N}, \forall i \in \{1, \dots, k\}$. It represents an exchange where agent a_i gives her object to agent a_{i+1} for each $i \in \{1, \dots, k-1\}$ and agent a_k gives her object to agent a_1 . With a slight abuse of notations, given a deal $\mu = \langle a_1, \dots, a_k \rangle$, agent a_{k+1} will refer to agent a_1 and agent a_0 (*i.e.* the agent before a_1) to agent a_k . For the particular case of deals involving only two agents, *i.e.* $k = 2$, we will talk about *swap-deals*. Pareto-optimality can be thought as a reallocation of objects among agents using improving *deals* [Shapley and Scarf, 1974; Sandholm, 1998], *Trading cycles* or *cycle deals* constitute a sub-class of deals, which is classical and used, *e.g.*, by Varian [1974, page 79] and Lipton et al. [2004, Lemma 2.2]

¹In principle, it is possible to distribute the execution of central procedures, by letting all agents broadcast, compute locally their own version of the central algorithm, and coordinate. This might however induce a significant cost.

in the context of envy-freeness. Trying to link efficiency concepts with various notions of deals is thus a natural idea.

Definition 1.18. Let $\langle \mathcal{N}, \mathcal{O}, u \rangle$ be an add-MARA instance and π be an allocation of this instance. A (N, M) -cycle deal of π is a sequence of transfers of items $\mu = \langle (\mu_1, \mathcal{O}_1), \dots, (\mu_N, \mathcal{O}_N) \rangle$, where, for each $j \in \{1, \dots, N\}$, μ_j denotes the j^{th} agent involved in the sequence and $\mu_j \in \mathcal{N}$, $\mathcal{O}_j \subseteq \pi_{\mu_j}$, and $|\mathcal{O}_j| \leq M$. The allocation $\pi[\mu]$ resulting from the application of μ to π is defined as follows:

- $\pi[\mu]_{\mu_j} = \pi_{\mu_j} \setminus \mathcal{O}_j \cup \mathcal{O}_{j-1}$ for $j \in \{2, \dots, N\}$;
- $\pi[\mu]_{\mu_1} = \pi_{\mu_1} \setminus \mathcal{O}_1 \cup \mathcal{O}_N$;
- $\pi[\mu]_i = \pi_i$ if $a_i \notin \{\mu_1, \dots, \mu_N\}$.

A cycle deal $\langle (\mu_1, \mathcal{O}_1), \dots, (\mu_N, \mathcal{O}_N) \rangle$ will be written

$$\mu_1 \xrightarrow{\mathcal{O}_1} \mu_2 \dots \mu_{N-1} \xrightarrow{\mathcal{O}_{N-1}} \mu_N \xrightarrow{\mathcal{O}_N} \mu_1.$$

In other words, in a cycle deal (we omit N and M when they are not necessary to understand the context), each agent gives a subset of at most M items from her share to the next agent in the sequence and receives in return a subset from the previous agent. $(N, 1)$ -cycle deals will be denoted by N -cycle deals. 2-cycle deals will be called *swap*-deals. Among these cycle deals, some are more interesting: those where each agent improves her utility by trading objects. More formally, a deal d will be called *weakly improving* if $u_i(\pi[\leftarrow d]_i) \geq u_i(\pi_i) \forall i \in \mathcal{N}$ with at least one of these inequalities being strict, and *strictly improving* if all these inequalities are strict.

Intuitively, if it is possible to improve an allocation by applying an improving cycle deal, then it means that this allocation is inefficient. Reallocating the items according to the deal will make everyone better-off. It is thus natural to derive a concept of efficiency from this notion of cycle-deal.

Definition 1.19. An allocation is said to be $>$ - (N, M) -cycle optimal (resp. \geq - (N, M) -cycle optimal) if it does not admit any strictly (resp. weakly) improving (K, M) -cycle deal for any $K \leq N$.

Picking sequences

Following the discussion made above about centralized versus decentralized protocols, we notice that even if there is a central coordinator the agents communicate way fewer preferences than with TTC for example. Besides, they are part of the procedure that is running. This is why we consider it is, to a certain level, a decentralized procedure. This very simple protocol works as follows. A central authority chooses a sequence of agents before the protocol starts, having as many agents as the number of objects (some agents may appear several times in the sequence). Then, each agent appearing in the sequence is asked to choose in turn one object among those that remain. For instance, according to the sequence $\langle a_1, a_2, a_2, a_1 \rangle$, agent a_1 will choose first, then agent a_2 will pick two objects in a row, and agent a_1 will take the last object. This protocol, used in a lot of everyday situations, has been studied for the first time by Kohler and Chandrasekaran

[1971]. Later, Brams and Taylor [2000] have studied a particular version of this protocol, namely alternating sequences, in which the sequence of agents is restricted to a balanced ($\langle 1, 2, 2, 1, \dots \rangle$) or strict ($\langle a_1, a_2, a_1, a_2, \dots \rangle$) alternation of agents. Bouveret and Lang [2011] have further formalized this protocol, whose properties (especially related to game theoretic aspects) have later been characterized by Kalinowski et al. [2013a,b]. Finally, Aziz et al. [2015b] have studied the complexity of problems related to finding whether a particular assignment (or bundle) is achievable by a particular class of picking sequences. Picking sequences have also been considered by Brams and King [2005], that focus on a situation where the agents have ordinal preferences. They make an interesting link between this protocol and Pareto-optimality, showing, among others, that picking sequences always result in a Pareto-optimal allocation, but also that every Pareto-optimal allocation can be obtained by a sequence of sincere choices.

Definition 1.20. *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance. A picking sequence or sequence of sincere choices (or simply sequence when the context is clear) is a vector of \mathcal{N}^m . We will denote by $\mathcal{S}(I)$ the set of possible sequences for the instance I .*

Let $\vec{\sigma} \in \mathcal{S}(I)$ be a sequence of agents and let σ_t be the t^{th} agent of the sequence. $\vec{\sigma}$ is said to *generate* allocation π if and only if π can be obtained as a possible result of the non-deterministic (the algorithm contains an instruction **choose** splitting the control flow into several branches, building all the allocations generated by $\vec{\sigma}$) Algorithm 1.2 on input I and $\vec{\sigma}$.

Algorithm 1.2: Execution of a sequence

Input: an instance $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ and a sequence $\vec{\sigma} \in \mathcal{S}(I)$

Output: an allocation $\pi \in \mathcal{A}(I)$

- 1 $\pi \leftarrow$ empty allocation (such that $\forall i \in \mathcal{N} : \pi_i = \emptyset$);
 - 2 $\mathcal{O}_1 \leftarrow \mathcal{O}$;
 - 3 **for** t **from** 1 **to** m **do**
 - 4 $a_i \leftarrow \sigma_t$;
 - 5 **choose** object $o_t \in \text{best}(\mathcal{O}_t, i)$;
 - 6 $\pi_i \leftarrow \pi_i \cup \{o_t\}$;
 - 7 $\mathcal{O}_{t+1} \leftarrow \mathcal{O}_t \setminus \{o_t\}$
 - 8 **end**
-

Definition 1.21. *An allocation π is said to be sequenceable if there exists a sequence $\vec{\sigma}$ that generates π , and non-sequenceable otherwise. For a given instance I , we will write $s(I)$ the binary relation defined by $(\vec{\sigma}, \pi) \in s(I)$ if and only if π can be generated by $\vec{\sigma}$.*

Example 1.15. *Let I be the following instance:*

	o_1	o_2	o_3
a_1	8	2	1
a_2	5	1	5

For instance, sequence $\langle a_2, a_1, a_2 \rangle$ generates two possible allocations: $\langle o_1, o_2 o_3 \rangle$ and $\langle o_2, o_1 o_3 \rangle$, depending on whether agent 2 chooses object o_1 or o_3 that she both prefers. Allocation $\langle o_1 o_2, o_3 \rangle$ can be generated by three sequences. Allocations $\langle o_1 o_3, o_2 \rangle$ and $\langle o_3, o_1 o_2 \rangle$ are non-sequenceable.

For any instance I , $|\mathcal{S}(I)| = |\mathcal{A}(I)| = n^m$. Also note that the number of objects allocated to an agent by a sequence is the number of times the agent appears in the sequence.

1.5 Thesis purpose

As we have seen throughout this chapter the problem of fair division of indivisible goods consists in finding fair and/or efficient allocations. However, we do not know what fairness really means as there are a lot of ways to express that notion. Among those fairness notions that we have presented, the one of envy-freeness is very appealing but there are cases where it cannot be found. Therefore literature has proposed relaxations of envy-freeness and hence we propose two new relaxations of envy-freeness that we present and thoroughly study.

Moreover, in order to obtain such allocations we need procedures that can be either centralized or decentralized each having their strengths and weaknesses. In particular, centralized procedures present problems of confidence, simplicity and elicitation of preferences. This is why we look at procedures that lead to a certain degree of decentralization and allow agents to be involved in the construction of the allocation: exchange sequences and picking sequences.

Finally, we have chosen to stay in the context of additive preferences meaning that the utility of a bundle of objects is simply the sum of the utilities of the objects. Moreover, we have sometimes considered domain restrictions such as single-peakedness or binary utilities to see if procedures have desirable properties on those domain preferences.

1.6 Organization of the thesis

We have seen throughout this chapter a very large display of definitions of fairness and a very large set of problems in which fairness and efficiency is studied. From that we show how the thesis is built.

Through Chapters 2 and 3 we will mainly focus on decentralized procedures, their design and their properties in different contexts.

We will then come back in Chapters 4 and 5 in a more centralized setting by proposing new approximations of envy-freeness and proportionality. Chapter 4 will take inspiration from classical multi-objective optimization to build a new approximation while Chapter 5 will be about a more collective way of thinking about envy-freeness and proportionality which are classically personal visions.

Chapter 2

Efficiency, Sequenceability and Deal-Optimality

Summary of the chapter

In fair division of indivisible goods, using sequences of sincere choices (or picking sequences) is a natural way to allocate the objects. The idea is as follows: at each stage, a designated agent picks one object among those that remain. Another intuitive way to obtain an allocation is to give objects to agents in the first place, and to let agents exchange them as long as such “deals” are beneficial. This chapter investigates these procedures, when agents have additive preferences over objects, and unveils surprising connections between them, and with other efficiency and fairness notions. In particular, we show that an allocation is sequenceable if and only if it is optimal for a certain type of deals, namely cycle deals involving a single object. Furthermore, any Pareto-optimal allocation is sequenceable, but not the converse. Regarding fairness, we show that an allocation can be envy-free and non-sequenceable, but that every competitive equilibrium with equal incomes is sequenceable. To complete the picture, we show how some domain restrictions may affect the relations between these notions. Finally, we experimentally explore the links between the scales of efficiency and fairness.

This chapter is based in full on a paper published in the Autonomous Agents Multi Agent Systems conference [Beynier et al., 2019a]. It is a joint work with the co-authors of the papers.

2.1 Introduction

In this chapter, we focus on a particular allocation protocol: *sequences of sincere choices* (also known as *picking sequences*). We analyze the links between sequences, certain types of deals among agents, and some efficiency and fairness properties, in a more general model in which the agents have numerical additive preferences on the objects. Our main contributions are the following. We give a formalization of the link between allocations and sequences of sincere choices, with the notion of *sequenceability*, and highlighting a simple characterization of the sequenceability of an allocation (Section 2.3). Then, we show that in this slightly more general framework than the one by Brams and King [2005], Pareto-optimality and sequenceability are not equivalent anymore (Section 2.4). By unveiling the connection between sequenceability and cycle deals among agents (Section 2.5), we obtain a rich “scale of efficiency” that allows us to characterize the degree of efficiency of a given allocation. Interestingly, some domain restrictions have significant effects on this hierarchy (Section 2.6). We also highlight (Section 2.7) a link between sequenceability and another important economical concept: the competitive equilibrium from equal income (CEEI). Another contribution is the experimental exploration of the links between the scale of efficiency and fairness properties (Section 2.8).

2.2 Definitions

We will first give a definition that will be useful throughout this chapter for the notion of sequenceability.

Definition 2.1. *Given an agent a_i and a set of objects \mathcal{O}' , we will write $\text{best}(\mathcal{O}', i) = \text{argmax}_{o_k \in \mathcal{O}'} u(i, k)$ the objects from \mathcal{O}' having the highest weight for a_i (they will be called top objects of a_i).*

A (sub-)allocation $\pi^{|\mathcal{O}'}$ is said frustrating if no agent receives one of her top objects in $\pi^{|\mathcal{O}'}$ (formally: $\text{best}(\mathcal{O}', i) \cap \pi_i^{|\mathcal{O}'} = \emptyset$ for each agent a_i), and non-frustrating otherwise.

In the following, we will consider a particular way of allocating objects to agents: sequences of sincere choices as presented in Algorithm 1.2. Informally the agents are asked in turn, according to a predefined sequence, to choose and pick a top object among the remaining ones. This is closely linked to the notion of sequenceable allocation described in Definition 1.21.

The notion of frustrating allocation and sequenceability were already implicitly present in the work by Brams and King [2005], and sequenceability has been extensively studied by Aziz et al. [2015b] with a focus on sub-classes of sequences (*e.g.* alternating sequences). However, a fundamental difference is that in our setting, the preferences might be non strict on objects, which entails that the same sequence can yield different allocations (in the worst case, an exponential number), as Example 1.15 shows.

2.3 Sequenceable allocations

We have seen in Example 1.15 that some allocations are non-sequenceable. We will now formalize this and give a precise characterization of sequenceable allocations. That is, we

will try to identify under which conditions an allocation is achievable by the execution of a sequence of sincere choices. We first start by noticing that in every sequenceable allocation, the first agent of the sequence gets a top object, so every frustrating allocation is non-sequenceable. However, being non-frustrating is not a sufficient condition for an allocation to be sequenceable, as the following example shows:

Example 2.1. *Consider this instance:*

	o_1	o_2	o_3	o_4
a_1	9	8	2	1
a_2	2	5	1	4

In the squared allocation $\pi = \langle o_1o_4, o_2o_3 \rangle$, each agent receives her top object. However, after o_1 and o_2 have been allocated (they must be allocated first by all sequences generating π), the dotted sub-allocation remains. This sub-allocation is obviously non-sequenceable because it is frustrating. Hence π is not sequenceable either.

It was shown by Bouveret and Lemaître [2016b] that this property of containing a frustrating sub-allocation exactly characterizes the set of non-sequenceable allocations:

Proposition 2.1 (Bouveret and Lemaître, 2016b). *Let $I = \langle \mathcal{N}, \mathcal{O}, u \rangle$ be an instance and π be an allocation of this instance. The two following statements are equivalent:*

- (A) π is sequenceable.
- (B) No sub-allocation of π is frustrating (in every sub-allocation, at least one agent receives a top object).

Besides characterizing a sequenceable allocation, the proof of Proposition 2.1 gives a practical way of checking if an allocation is sequenceable, and, if it is the case, of computing a sequence that generates this allocation.

Proposition 2.2 (Bouveret and Lemaître, 2016b). *Let $I = \langle \mathcal{N}, \mathcal{O}, u \rangle$ be an instance and π be an allocation of this instance. We can decide in time $O(n \times m^2)$ if π is sequenceable.*

The proof is based on the execution of Algorithm 2.1. This algorithm is similar in spirit to the one proposed by Brams and King [2005] but is more general because (i) it can involve non-strict preferences on objects, and (ii) it can conclude with non-sequenceability.

2.4 Pareto-optimality

When an allocation is generated from a sequence, in some sense, a weak form of efficiency is applied to build the allocation: each successive (picking) choice is “locally” optimal. This raises a natural question: is every sequenceable allocation Pareto-optimal?

This question has already been extensively discussed independently by Aziz et al. [2016b] and Bouveret and Lemaître [2016b]. We complete the discussion here to give more insights about the implications of the previous results in our framework.

Brams and King [2005, Proposition 1] prove the equivalence between sequenceability and Pareto-optimality. However, they have a different notion of Pareto-optimality, because the agents’ preferences are given as linear orders over *objects*. To be able to compare

Algorithm 2.1: Sequencing an allocation

Input: $I = \langle \mathcal{N}, \mathcal{O}, u \rangle$ and $\pi \in \mathcal{A}(I)$
Output: a sequence $\vec{\sigma}$ generating π or **NonSeq**

- 1 $(\vec{\sigma}, \mathcal{O}') \leftarrow (\langle \rangle, \mathcal{O})$;
- 2 **for** t **from** 1 **to** m **do**
- 3 **if** $\exists i$ *such that* $\text{best}(\mathcal{O}', i) \cap \pi_i \neq \emptyset$ **then**
- 4 Append a_i to $\vec{\sigma}$;
- 5 let $o_k \in \text{best}(\mathcal{O}', i) \cap \pi_i$;
- 6 $\mathcal{O}' \leftarrow \mathcal{O}' \setminus \{o_k\}$;
- 7 **end**
- 8 **else return NonSeq** ;
- 9 **end**
- 10 **return** $\vec{\sigma}$;

bundles, these preferences are lifted on subsets using the *responsive set extension* \succ_{RS} . This extension leaves many bundles incomparable and leads to define *possible* and *necessary* Pareto-optimality. Brams and King’s notion is possible Pareto-optimality. Aziz et al. show that, given a linear order \succ on objects and two bundles π and π' , $\pi \succ_{RS} \pi'$ if and only if $u(\pi) > u(\pi')$ for *all* additive utility functions u compatible with \succ (that is, such that $u(o_k) > u(o_l)$ if and only if $o_k \succ o_l$). This characterization of responsive dominance yields the following reinterpretation of Brams and King’s result: an allocation π is sequenceable if and only if for each other allocation π' , there is a set u_1, \dots, u_n of additive utility functions, respectively compatible with \succ_1, \dots, \succ_n such that $u_i(\pi_i) > u_i(\pi'_i)$ for at least one agent a_i .

As noted by Bouveret and Lemaître [2016b], the latter notion of Pareto-optimality is not very demanding, because it only requires to find one utility function that works. They show that under our stronger notion, the equivalence between sequenceability and Pareto-optimality no longer holds.¹

Example 2.2 (Bouveret and Lemaître, 2016b). *Let us consider the following instance:*

	o_1	o_2	o_3
a_1	5	4	2
a_2	8	2	1

The sequence $\langle a_1, a_2, a_2 \rangle$ generates allocation $\pi = \langle o_1, o_2 o_3 \rangle$ giving utilities $\langle 5, 3 \rangle$. π is then sequenceable but it is dominated by $\pi' = \langle o_2 o_3, o_1 \rangle$, giving utilities $\langle 6, 8 \rangle$ (and generated by $\langle a_2, a_1, a_1 \rangle$). Observe that, under ordinal linear preferences, π' would not dominate π , but they would be incomparable.

The last example shows that a sequence of sincere choices does not necessarily generate a Pareto-optimal allocation. What about the converse? We can see, as a trivial corollary of the reinterpretation of Brams and King’s result in our terminology, that the answer is positive *if the preferences are strict on shares*. The following result is more general, because it holds even without this assumption:

¹Actually, since it is known [de Keijzer et al., 2009; Aziz et al., 2016a] that testing Pareto-optimality with additive preferences is coNP-complete, and that testing sequenceability is in P (Proposition 2.2), they cannot be equivalent unless $P = \text{coNP}$.

Proposition 2.3 (Aziz et al., 2016b; Bouveret and Lemaître, 2016b). *Every Pareto-optimal allocation is sequenceable.*

As already noticed by Aziz et al. [2016b], the proof follows from an adaptation of Brams and King’s Proposition 1 (necessity part of the proof) [Brams and King, 2005]. However, we find useful to give the proof, because it is more general than the previous one, and will be reused in subsequent results of this chapter. Before giving this proof, we illustrate it on a concrete example from Bouveret and Lemaître [2016b].

Example 2.3 (Bouveret and Lemaître, 2016b). *Let us consider the following instance:*

	o_1	o_2	o_3	o_4	o_5
a_1	$\boxed{12}^*$	15	11^*	$\boxed{7}$	2
a_2	2	12	$\boxed{7}$	15^*	$\boxed{11}^*$
a_3	15	$\boxed{20}^*$	9	2	1

The squared allocation π is not sequenceable: indeed, every sequence that could generate it should start with $\langle a_3, a_1, \dots \rangle$, leaving the frustrating dotted sub-allocation \vec{p} .

Let us consider a_1 for instance. Since the suballocation is frustrating, she does not receive o_3 (which is her top object), but a_2 does. This latter agent, however, does not get her top object, o_4 , because a_1 receives it. Obviously, if a_1 gives o_4 to a_2 and receives o_3 in return, we have built a cycle in which each agent gives a regular object and receives a top one. Doing this, we have built the starred allocation π' strictly dominating π .

Corollary 2.1 (Bouveret and Lemaître, 2016b). *No frustrating allocation can be Pareto-optimal (equivalently, in every Pareto-optimal allocation, at least one agent receives a top object).*

Proposition 2.3 implies that there exists, for a given instance, three classes of allocations: (1) non-sequenceable (therefore non Pareto-optimal) allocations, (2) sequenceable but non Pareto-optimal allocations, and (3) Pareto-optimal (hence sequenceable) allocations. These three classes define a “scale of efficiency” that can be used to characterize the allocations. What is interesting and new here is the intermediate level. We will see that this scale can be further refined.

2.5 Cycle deals-optimality

We have seen in Section 1.4.3 that Pareto-optimality can be thought as a reallocation of objects among agents using improving *deals* so trying to link efficiency concepts with various notions of deals strikes as a natural idea.

We begin with easy observations. First, \geq -cycle optimality implies $>$ -cycle optimality, and these two notions become equivalent when the preferences are strict on shares (like in Chapter 3). Moreover, restricting the size of the cycles and the size of the bundles exchanged yield less possible deals and hence lead to weaker optimality notions.

Note that for $N' \leq N$ and $M' \leq M$ (at least one of these inequalities being strict), $>$ -(N, M)-cycle-optimality and \geq -(N', M')-cycle-optimality are incomparable. These observations show that cycle-deal optimality notions form a (non-linear) hierarchy of efficiency

concepts of diverse strengths. The natural question is whether they can be related to sequenceability and Pareto-optimality. Obviously, Pareto-optimality implies both $>$ -cycle-optimality and \geq -cycle-optimality. An easy adaptation of the proof of Proposition 2.3 leads to the following stronger result:

Proposition 2.4. *An allocation π is sequenceable if and only if it is $>$ - n -cycle optimal (with $n = |\mathcal{N}|$).*

Proof. Let π be a non-sequenceable allocation. Then by Proposition 2.1, there is at least one frustrating sub-allocation in π . Using the same line of arguments as in the proof of Proposition 2.3 we can build a strictly improving k -cycle. Hence π is not $>$ -cycle optimal. Conversely, suppose that π admits a strictly improving k -cycle deal. Then obviously this cycle yields a sub-allocation that is frustrating. \square

The scale of efficiency introduced in Section 2.4 can then be refined with a hierarchy of $>$ -cycle optimality notions below sequenceable allocations: Pareto-optimal \Rightarrow sequenceable $\Leftrightarrow >$ - n -cycle optimal $\Rightarrow >$ - $(n - 1)$ -cycle optimal $\Rightarrow \dots \Rightarrow >$ -swap optimal.

As for \geq -cycle optimality, it forms a parallel hierarchy between Pareto-optimal and non-sequenceable allocations. Note that sequenceability does not involve any \geq - n -cycle-optimality. Thus, as soon as $k < n$, \geq - k -cycle optimality and sequenceability become incomparable notions.

For instance, for 3 agents, there exist allocations which are \geq -swap optimal but not sequenceable and the other way around:

	o_1	o_2	o_3	o_4
a_1	2	1*	2	3*
a_2	3*	3	1	2
a_3	1	2	3*	1

Here the squared allocation is \geq -swap optimal, but not sequenceable: there exists a strictly improving 3-cycle. At the same time, the starred allocation is sequenceable (by $\langle a_2, a_3, a_1, a_1 \rangle$), but not even \geq -swap optimal, since a_1 and a_2 may agree on a weakly improving deal.

The observations previously made in this section suggest that, in some situations, the most complex cycle deals could be required to reach Pareto-optimal allocations. This is indeed the case—we now make this claim more precise. Observe that to be involved in a weakly improving cycle deal, each agent must pass at least one item, thus for a (n, k) -cycle deal, we have that $k \leq m - (n - 1)$ (i.e. the “largest bundle” circulating in this cycle deal can be at most $m - n + 1$). The following generic example shows that it may be necessary to implement a $(n, m - n + 1)$ -cycle to reach a Pareto-optimal allocation.

	α_1	α_2	\dots	α_{n-1}	β_1	\dots	β_{m-n+1}
1	1	0	0	0	$1/(m - n + 1)$	$1/(m - n + 1)$	$1/(m - n + 1)$
2	1	1	0	0	0	0	0
3	0	1	1	0	0	0	0
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
n	0	0	0	2	$1/(m - n + 1)$	$1/(m - n + 1)$	$1/(m - n + 1)$

Initially, every agent a_i for $i = 1, \dots, n - 1$ holds item α_i , while agent n holds $\beta_1, \dots, \beta_{m-n+1}$. Hence everyone enjoys utility 1. This allocation is dominated by the

allocation where each agent a_i for $i = 2, \dots, n$ holds α_{i-1} while a_1 holds $\beta_1, \dots, \beta_{m-n+1}$. In this allocation, the utilities of agents are instead $\langle 1, 1, \dots, 2 \rangle$. But to obtain α_{n-1} , a_n must get it from a_{n-1} who would only release it if she gets α_{i-2} , etc. In the end, a_1 will only release α_1 if she gets the full bundle $\beta_1, \dots, \beta_{m-n+1}$. Overall there are n agents involved in the deal, exchanging up to $m - n + 1$ items. By construction, it is easy to check that no simpler cycle deal (either in terms of number of items or number of agents) would allow to reach this allocation. Furthermore, there is clearly no other allocation Pareto-dominating the initial allocation.

However, it is important to note that cycle-deals may not be sufficient to reach Pareto-optimal outcomes when there are more items than agents.

Example 2.4. Consider the following example:

	o_1	o_2	o_3	o_4	o_5	o_6
a_1	$\boxed{3}$	$\boxed{6}$	6^*	0	6	4^*
a_2	2^*	0	$\boxed{6}$	$\boxed{3}$	7^*	0
a_3	0	5^*	0	4^*	$\boxed{6}$	$\boxed{3}$

Note that in the squared allocation, all agents enjoy the same utility, $\langle 9, 9, 9 \rangle$, and that it is Pareto-dominated by the starred allocation which induces the vector of utilities $\langle 10, 9, 9 \rangle$. We leave it to the reader to check that no swap deal, nor 3-cycle, would be weakly improving. In fact, the only way to reach the starred allocation from this initial allocation would require to implement simultaneously two $(3, 1)$ -cycle deals ($a_1 \xrightarrow{o_1} a_2 \xrightarrow{o_3} a_3 \xrightarrow{o_6} a_1$ and $a_3 \xrightarrow{o_5} a_2 \xrightarrow{o_4} a_1 \xrightarrow{o_2} a_3$).

Finally, a corollary of Propositions 2.2 and 2.4 is that checking whether an allocation is $>n$ -cycle optimal can be made in polynomial time (by checking whether it is sequenceable).

More generally, we can observe that checking whether an allocation is (k, k') -cycle optimal can be done by iterating over all k -uples of agents², and for each one iterating over all possible transfers involving less than k' objects. In total, there are $k! \binom{n}{k}$ k -uples of agents (which is upper-bounded by n^{k+1}). For each k -uple, there are at most $\left(\sum_{k''=0}^{k'} \binom{m}{k''}\right)^k$ possible transfers, which is again upper-bounded by $(1+m)^{kk'}$. Hence, in total, checking whether an allocation is (k, k') -cycle optimal can be done in time $O(n^{k+1} \times (1+m)^{kk'})$. This is polynomial in n and m if both k and k' are bounded (as for swap deals).

2.6 Restricted Domains

We now study the impact of several preference restrictions on the hierarchy of efficiency notions introduced in Section 2.5.

²We do not need to also run through all cycles of strictly less than k agents: such a cycle can be simulated just by appending at the end some agents whose role is just to pass the objects they receive to the next agent.

Strict preferences on objects Bouveret and Lemaître [2016b] studied the setting of strict preferences on objects.

When the preferences are strict on objects, then obviously every sequence generates exactly one allocation. They show that the converse is also true:

Proposition 2.5 (Bouveret and Lemaître, 2016b). *Preferences are strict on objects iff $s(I)$ is a mapping from $\mathcal{S}(I)$ to $\mathcal{A}(I)$.*

Same order preferences We say that the agents have *same order preferences* [Bouveret and Lemaître, 2016a] if there is a permutation $\eta : \mathcal{O} \mapsto \mathcal{O}$ such that for each agent a_i and each pair of objects o_k and o_l , if $o_{k'} = \eta(o_k) < \eta(o_l) = o_{l'}$ then $u(i, k') \geq u(i, l')$. Bouveret and Lemaître [2016b] also studied that setting.

Proposition 2.6 (Bouveret and Lemaître, 2016b). *All the allocations of an instance with same order preferences are sequenceable (and actually cycle-deal optimal). Conversely, if all the allocations of an instance are sequenceable, then this instance has same order preferences.*

Let us now characterize the instances for which $s(I)$ is a one-to-one correspondence.

Proposition 2.7 (Bouveret and Lemaître, 2016b). *For a given instance I , the following two statements are equivalent.*

- (A) *Preferences are strict on objects and in the same order.*
- (B) *The relation $s(I)$ is a one-to-one correspondence.*

The proof is a consequence of Propositions 2.5 and 2.6.

Single-peaked preferences Interestingly, when preferences are single-peaked, the hierarchy of n -cycle optimality collapses at the second level. A similar result was proven by Damamme et al. [2015]. We reformulate it more generally here:

Proposition 2.8. *If the preferences are single-peaked and additive, then an allocation π is $\geq n$ -cycle optimal iff it is swap-optimal.*

Proof. First, note that $\geq n$ -cycle optimality trivially implies swap-stability. Let us now show the converse.

Let us consider an allocation π that is ≥ 2 -cycle optimal but not $\geq n$ -cycle optimal. Since π is not $\geq n$ -cycle optimal, there exists at least one improving deal in π . Let μ be the smallest improving cycle-deal, i.e., which involves the smallest number of agents. Assume without loss of generality that $\mu = \langle a_1, \dots, a_k \rangle$, with $2 < k \leq n$.

Since μ is an improving deal, every agent in μ is happy to exchange with the agent coming before her in μ :

$$\pi_{a_{i-1}} \succ_i \pi_{a_i}, \forall a_i \in \mu. \quad (2.1)$$

Moreover, as there exists no improving swap-deal in π (π is ≥ 2 -cycle optimal) agents do not want to exchange with the agent coming after them in μ :

$$\pi_{a_i} \succ_i \pi_{a_{i+1}}, \forall a_i \in \mu. \quad (2.2)$$

Indeed, because of (2.1), an improving swap-deal would otherwise exist in π .

We now show by induction on the size of μ , denoted by k , that such an improving deal cannot exist.

If $k = 3$, consider without loss of generality that $\mu = \langle a_1, a_2, a_3 \rangle$. From (2.1) and (2.2), we obtain the following profile where underlined objects indicate the initial allocation:

$$\begin{aligned} a_1 : \pi_{a_3} \succ_1 \underline{\pi_{a_1}} \succ_1 \pi_{a_2}, \\ a_2 : \pi_{a_1} \succ_2 \underline{\pi_{a_2}} \succ_2 \pi_{a_3}, \\ a_3 : \pi_{a_2} \succ_3 \underline{\pi_{a_3}} \succ_3 \pi_{a_1}. \end{aligned}$$

The triplet of objects $\langle \pi_{a_1}, \pi_{a_2}, \pi_{a_3} \rangle$ is thus a witness of the violation of the worst-restrictedness condition for a profile to be single-peaked (Proposition 1.1). Indeed, all the three objects are ranked last by an agent when we restrict our attention to these objects. The contradiction is thus established.

Suppose now that π is \geq -($k - 1$)-cycle optimal. We will show that no improving deal of size k exists in π . From the induction hypothesis, we get that:

$$\pi_{a_i} \succ_i \pi_{a_j}, \forall a_i, a_j \in \mu, a_j \notin \{a_{i-1}, a_i\}. \quad (2.3)$$

Indeed there would otherwise exist two agents a_l and a'_l , that are not next to one another in μ , such that $\pi_{a'_l} \succ_l \pi_{a_l}$. It would then have been possible to “cut” μ between those two agents so that a_l receives $\pi_{a'_l}$. The new cycle-deal obtained would also have been improving and then an improving deal of size strictly smaller than k would exist.

Because the profile is single-peaked, it is also worst-restricted (Proposition 1.1) and there exist thus at most two objects ranked last by an agent among the ones appearing in μ . Call π_{a_w} one such object holds by agent a_w , and consider the triplet of objects $O = \langle \pi_{a_{w-1}}, \pi_{a_w}, \pi_{a_{w+1}} \rangle$. From (2.1), (2.2) and (2.3) we get:

$$\begin{aligned} a_w : \pi_{a_{w-1}} \succ_w \underline{\pi_{a_w}} \succ_w \pi_{a_{w+1}}, \\ a_{w+1} : \pi_{a_w} \succ_{w+1} \underline{\pi_{a_{w+1}}} \succ_{w+1} \pi_{a_{w-1}}. \end{aligned}$$

Hence when restricting preferences to O , for every object in O , there exists an agent ranking it last among the objects in O . This violates the worst-restrictedness condition for the single-peaked profile and sets the contradiction.

Overall we have proved that no improving deal exists in π which entails that π is \geq - n -cycle optimal. \square

Together with Proposition 2.4, Proposition 2.8 gives another interpretation of sequenceability in this domain:

Corollary 2.2. *If preferences are single-peaked (and additive), then an allocation π is sequenceable if and only if it is swap-optimal.*

Proposition 1 by Damamme et al. [2015] is much stronger than our Corollary 2.2, as it shows that swap-optimality is actually equivalent to Pareto-efficiency *when each agent receives a single resource*. Unfortunately, in our context where each agent can receive several items, this is no longer the case, as the following example shows:

Example 2.5. Consider this instance, single-peaked with respect to $1 \triangleright \dots \triangleright 6$:

	o_1	o_2	o_3	o_4	o_5	o_6
a_1	$\boxed{1}^*$	$\boxed{2}$	3	4	5	6^*
a_2	1	3^*	$\boxed{4}$	$\boxed{5}$	6^*	2
a_3	1	2	4^*	5^*	$\boxed{6}$	$\boxed{3}$

The squared allocation is swap-optimal, but Pareto-dominated by the starred allocation.

2.7 Envy-Freeness and CEEI

The use of sequences of sincere choices can also be motivated by the search for a *fair* allocation protocol. Here, we will focus on two fairness properties and analyze their link with sequenceability.

As Bouveret and Lemaître [2016a] and Brânzei et al. [2015] have shown, with additive preferences, every CEEI allocation is envy-free. Bouveret and Lemaître [2016b] investigated whether an envy-free or CEEI allocation is necessarily sequenceable. For envy-freeness, they answered in the negative.

Proposition 2.9 (Bouveret and Lemaître, 2016b). *There exist non-sequenceable envy-free allocations, even if the agents' preferences are strict on shares.*

Proof. A counterexample with strict preferences on shares is given in Example 2.3 above, for which we can check that the circled allocation π is envy-free and non-sequenceable. \square

Concerning CEEI, it is already well known that any CEEI allocation is Pareto-optimal (hence sequenceable) if the preferences are strict on shares [Bouveret and Lemaître, 2016a]. This is also a consequence of the First Welfare Theorem introduced by Babaioff et al. [2017] for indivisible goods.

However, surprisingly, this result does not hold anymore if the preferences are not strict on shares, as the following example shows:

	o_1	o_2	o_3	o_4
a_1	$\boxed{2}^*$	3^*	3	$\boxed{2}$
a_2	2	3	$\boxed{4}^*$	1
a_3	0	$\boxed{4}$	2	4^*

The squared allocation is CEEI (with prices 0.5, 1, 1, 0.5) but is ordinally necessary (hence also additively) dominated by the starred allocation. In spite of this negative result, a certain level of efficiency for CEEI allocations can still be guaranteed:

Proposition 2.10 (Bouveret and Lemaître, 2016b). *Every CEEI allocation is sequenceable.*

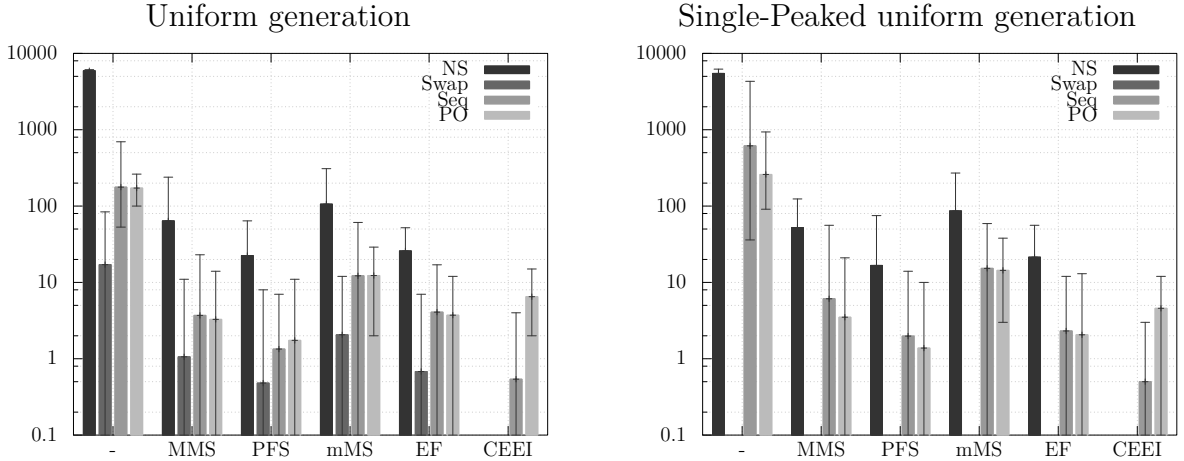


Figure 2.1: Distribution of the number of allocations by pair of (efficiency, fairness) criteria.

2.8 Experiments

We have exhibited in Sections 2.4 and 2.5 a “hierarchy of allocation efficiency” made of several steps: Pareto-optimal (PO), sequenceable (Seq), {cycle-deal-optimal}, non-sequenceable (NS). A natural question is to know, for a given instance, which proportion of allocations are located at each level of the scale. We give a first answer by experimentally studying the distribution of allocations between the different levels. For cycle-deal optimality, we focus on the simplest type of deals, namely, $>$ -swap-deals. We thus have a linear scale of efficiency concepts, from the strongest to the weakest: PO \rightarrow Seq \rightarrow Swap \rightarrow NS. We also analyze the relation between efficiency and various notions of fairness by linking this latter scale with the 6-level scale of fairness introduced by Bouveret and Lemaître [2016a]: CEEI \rightarrow Envy-Freeness (EF) \rightarrow min-max share (mMS) \rightarrow proportionality (PROP) \rightarrow max-min share (MMS) \rightarrow NS. Our experimental protocol is the following. We generate 50 add-MARA instances involving 3 agents-8 objects, using two different models. For both models, a set of weights are uniformly drawn in the interval $[[0, 100]]$ and the instances are then normalized. For the second model, these weights are reordered afterwards to make the preferences single-peaked. For each instance, we generate all 6561 allocations, and identify for each of them the *highest* level of fairness and efficiency satisfied. The average number of allocations with min-max interval is plotted as a box for each level on a logarithmic scale in Figure 2.1.

Note that some fairness and efficiency tests require to solve NP-hard or coNP-hard problems (MMS, mMS, and PO tests). These tests are delegated to an external ILP solver. This is especially interesting for the CEEI test which is known to be NP-hard [Brânzei et al., 2015], and for which, to the best of our knowledge, no practical method had been described before. The implementation is available as a fully documented and tested Free Python library.³

We note several interesting facts. First, a majority of allocations do not have any efficiency nor fairness property (first black bar on the left). Second, the distribution of allocations on the efficiency scale seems to be related to the fairness criteria: a higher proportion of swap-optimal or sequenceable allocations are found among envy-free allocations than among allocations that do not satisfy any fairness property, and for CEEI allocations,

³Available at: <https://gricad-gitlab.univ-grenoble-alpes.fr/bouveres/fairdiv>.

there are even more Pareto-optimal allocations than just sequenceable ones. Lastly, the absence of vertical bar for swap-optimality in the experiments concerning single-peaked preferences confirms the results of Corollary 2.2: in this context, no allocation can be swap-optimal but not Sequenceable; hence, all the allocations that are swap-optimal are contained in the bars concerning sequenceable or Pareto-optimal allocations. Similarly, the absence of bars for swap-optimality and NS (non-sequenceable) in both graphs for the CEEI case confirms the result of Proposition 2.10.

2.9 Conclusion

In this chapter, we have shown that picking sequences and cycle-deals can be reinterpreted to form a rich hierarchy of efficiency concepts. Many interesting questions remain open, such as the complexity of computing cycle-deals or the link between efficiency concepts and social welfare. One could also think of further extending the efficiency hierarchy by studying restrictions on possible sequences (*e.g.* alternating) or extending the types of deals to non-cyclic ones.

Chapter 3

Swap Dynamics in Single-Peaked Housing Markets

Summary of the chapter

This chapter focuses on the problem of fairly and efficiently allocating objects to agents. We consider a specific setting, usually referred to as a *housing market*, where each agent must receive exactly one object (and initially owns one). In this framework, in the domain of linear preferences, the Top Trading Cycle (TTC) algorithm is the only procedure satisfying Pareto-optimality, individual rationality and strategy-proofness. Under the restriction of single-peaked preferences, we will investigate another procedure called Crawler that enjoys the same properties. These two centralized procedures might however involve long trading cycles. In this chapter we focus instead on procedures involving the shortest cycles: bilateral swap-deals. In such swap dynamics, the agents perform pairwise mutually improving deals until reaching a swap-stable allocation (no improving swap-deal is possible). We prove that in the single-peaked domain every swap-stable allocation is Pareto-optimal, showing the efficiency of the swap dynamics. In fact, this domain turns out to be maximal when it comes to guaranteeing this property. Besides, both the outcome of TTC and Crawler can always be reached by sequences of swaps. However, some Pareto-optimal allocations are not reachable through improving swap-deals. We further analyze the outcome of swap dynamics through social welfare notions, in our context the average or minimum rank of the objects obtained by agents in the final allocation. We start by providing a worst-case analysis of these procedures. Finally, we present an extensive experimental study in which different versions of swap dynamics are compared to other existing allocation procedures. We show that they exhibit good results on average in this domain, under different cultures for generating synthetic data.

This chapter is based in full on a paper published at the *Autonomous Agents and MultiAgent Systems* journal in 2021 [Beynier et al., 2021]. It is a joint work with the co-authors of the papers.

3.1 Introduction

This chapter studies the problem of allocating fairly and efficiently a set of indivisible objects to a set of agents. We consider the specific setting of *housing market* by moreover

assuming that the preferences are *single-peaked*. Under this assumption, there exists a common ordering of the objects such that the further you go from your most preferred object—either going left or right according to the ordering—the less preferred the objects are.

The design of object allocation procedures is primarily guided by the the properties we want the final allocation, or the procedure itself, to satisfy. In single-peaked housing markets, the two prominent allocation procedures are *Top Trading Cycle* (TTC) [Shapley and Scarf, 1974] and *Crawler* [Bade, 2019]. They both satisfy a set of key desirable properties: Pareto-optimality, strategy-proofness and individual rationality [Roth, 1982; Shapley and Scarf, 1974; Ma, 1994; Bade, 2019]. In fact, *Crawler* has an additional advantage over TTC in terms of strategy-proofness, since it can be implemented in obviously dominant strategies [Bade, 2019].

However, both TTC and *Crawler* require a significant amount of global coordination. Indeed, they are both based on *trading cycles*—cyclic exchanges of objects between several agents—which can potentially involve many agents. Such long cycles may not be acceptable or even feasible in practice, for instance because of the risk of failure they induce. As an example, kidney exchange programs usually restrict the size of the cycles to two or three [Roth et al., 2005] because of time constraints.

Limiting the length of these cycles is thus a relevant agenda to reduce the coordination complexity and make procedures more robust. It is in particular a prerequisite to the development of more *decentralized* approaches that rely on the agents autonomously performing simple deals, without being hampered by a prohibitive coordination and communication cost.

In this chapter, we pursue this line of research and focus on the simplest possible local deals in the context of housing markets: *swap-deals* [Damamme et al., 2015]. Under a *swap dynamics*, agents meet each other, in a pairwise fashion, and exchange their objects if they both benefit from it. The process iterates until a stable state, an equilibrium, is reached. Under this approach, once a trading cycle is selected, very little coordination is required. Nevertheless, selecting *which* swap-deal (with which agent) to perform may still require significant prior coordination, depending on the heuristic used. Hence this approach is best described as a family of dynamics, with different degrees of decentralization depending on the heuristics used to select the improving deals to be implemented. We shall study several of them in this chapter.

We first establish that a large class of swap dynamics is vulnerable to manipulation, unveiling a tension between efficiency and strategy-proofness for such decentralized procedures. On the positive side, we demonstrate that in housing markets with single-peaked preferences, every allocation that is stable with respect to swap-deals is Pareto-optimal. We also prove that the single-peaked domain is maximal in that respect: in other words, any larger domain would fail to offer such a guarantee. Moreover, even though some Pareto-optimal allocations may not be reached by a sequence of improving swap-deals, we show that the allocations returned by TTC and *Crawler* are both reachable: there exist sequences of swap-deals simulating these procedures, such that (potentially) long trading cycles are not necessary any longer. We further investigate how swap dynamics behave with respect to social welfare, taken as the average or minimum rank of the object obtained by the agents. After a worst-case analysis, we explore experimentally the influence of different heuristics used to select deals. These experiments highlight that swap

dynamics perform particularly well with respect to these objectives.

Section 3.2 describes the centralized allocation procedure Crawler discussed in this chapter, while swap dynamics are introduced in Section 3.3. A formal analysis of these procedures is presented in Section 3.4. Section 3.5 offers a comparison between swap dynamics outcomes and social welfare optimization, in a “price of anarchy” perspective. The experimental analysis is presented in Section 3.6. Section 3.7 concludes.

We recall that in this chapter we deal with the House Allocation Problem with single-peaked preferences.

3.2 Centralized allocation procedures for housing markets

This section deals with two centralized allocation procedures that will be used as references in the chapter. The first one is the TTC algorithm that has already been introduced (see Section 1.4.2) and which can be used without any domain restriction. We recall that it satisfies the three main desirable properties of an allocation procedure: *Pareto-optimality*, *individual rationality* and *strategy-proofness*. The second one, Crawler [Bade, 2019], is specially designed for single-peaked domains. It satisfies the same properties as TTC. Both of these procedures are based on the notion of deals introduced in Section 1.4.3.

3.2.1 Crawler algorithm

In Crawler algorithm, agents are initially ordered along the single-peaked axis according to the object they hold. The first agent is then the one holding the object on the left end side of the axis. As in TTC, the list of available objects is denoted by O and is ordered according to the single-peaked axis. A list of available agents N such that the i^{th} agent of the list is the one who holds the i^{th} object in O is also maintained.

Considering agents sequentially from the first in N to the last in N , Crawler checks for each agent a_i where her top object, $top_R(a_i)$, is on the axis.¹

- If $top_R(a_i)$ is on her right, the algorithm moves to the next agent.
- If a_i holds her top object $top_R(a_i)$, then $top_R(a_i)$ is allocated to a_i . Agent a_i and the object $top_R(a_i)$ are removed from N and O . The algorithm starts again from the agent on the left end of the axis.
- If $top_R(a_i)$ is on the left of a_i , the agent is allocated her top object $top_R(a_i)$. We denote by t^* the index of $top_R(a_i)$ and t the index of a_i (we have $t^* < t$). Then, every agent between t^* and $t - 1$ receives the object held by the agent on her right (the objects “crawl” towards left). Once again, a_i and $top_R(a_i)$ are removed from N and O and the algorithm restarts from the first agent.

Once all the objects have been allocated, the algorithm terminates.

A formal description of Crawler procedure is given in Algorithm 3.1. Note that we make use of the sub-procedure $pick(a_{t^*}, o, N, O, \pi)$ which simply assigns the object o to

¹Note that the algorithm can equivalently be executed from the last agent to the first one.

Algorithm 3.1: Crawler algorithm

Input: An instance $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$ single-peaked with respect to \triangleleft

Output: An allocation π

```
1  $\pi \leftarrow$  empty allocation
2  $O \leftarrow \mathcal{O}$ : list of objects sorted accordingly to  $\triangleleft$ 
3  $N \leftarrow \mathcal{N}$ : list of agents such that the  $i^{\text{th}}$  agent is the one who initially holds the
    $i^{\text{th}}$  object in  $O$ 
4 while  $N \neq \emptyset$  do
5    $t^* \leftarrow |N|$ 
6   for  $i = 0$  to  $|N| - 1$  do
7     if  $o_i \succ_i o_{i+1}$  then
8        $t^* \leftarrow i$ 
9       Break
10    end
11  end
12   $o \leftarrow \text{top}_R(a_{t^*})$ 
13   $\text{pick}(a_{t^*}, o, N, O, \pi)$ 
14 end
15 return  $\pi$ 
```

the given agent a_{t^*} in the allocation π , and then removes the agent and the object from the lists of available agents and objects, N and O respectively. Since the list of objects is ordered following the single-peaked axis and the i^{th} agent in N corresponds to the owner of the i^{th} object in O , the removal of o and a_{t^*} is in fact equivalent to assigning r to agent a_{t^*} and crawling the objects from right to left.

Let us illustrate the execution of Crawler:

Example 3.1. *On the instance of Example 1.12, agent a_4 is the first agent whose top object is not on her right, she thus receives her top object o_3 . The second step matches agent a_3 to o_4 . On the third step, agents a_1 and a_2 both have their top objects (among the remaining objects) on the right but the last agent a_5 has hers on her left. a_5 is then matched to o_1 . Object o_2 crawls to agent a_1 and object o_5 crawls to agent a_2 . On the fourth step, a_2 picks her current object o_5 . Finally, a_1 is assigned object o_2 . The allocation returned by Crawler is $\pi^C = \langle o_2, o_5, o_4, o_3, o_1 \rangle$. At each step i of the procedure, an improving cycle-deal μ_i is applied (with the last agent in the cycle picking her top object among the remaining ones): $\mu_1 = \langle a_3, a_4 \rangle$, $\mu_2 = \langle a_3 \rangle$, $\mu_3 = \langle a_1, a_2, a_5 \rangle$, $\mu_4 = \langle a_2 \rangle$, $\mu_5 = \langle a_1 \rangle$. One can observe that on this example the allocation returned by Crawler is not the same as the one returned by TTC. However, both procedures lead to the same minimum rank mrk and average rank ark .*

Interestingly, a variant of this procedure allows to check Pareto-optimality in single-peaked domains in linear time [Beynier et al., 2020].

3.3 Swap dynamics: a family of procedures based on swap-deals

We will now focus on dynamics based on local exchanges between the agents as presented in Section 1.4.3. We however reintroduce them as they may be simplified as we are in the context of the House Allocation Problem.

Interestingly, it can be observed that in housing markets, any reallocation (permutation of objects) can be implemented as a collection of disjoint cycle-deals [Shapley and Scarf, 1974]. This notion of deal is thus sufficient in this context. Let us first introduce some additional definitions.

Let us first introduce the concept of improving deals in the context of House Allocation Problem as each agent has one and only one item.

Definition 3.1. *Let π be an allocation, and $\mu = \langle a_1, \dots, a_k \rangle$ a deal involving $k \geq 1$ agents. The allocation $\pi[\mu]$ obtained by applying the deal μ to π is defined by:*

$$\begin{cases} \pi[\mu]_{a_i} = \pi_{a_{i-1}} & \text{if } i \in \{1, \dots, k\}, \\ \pi[\mu]_{a_i} = \pi_{a_i} & \text{otherwise.} \end{cases}$$

A deal is said to be improving if $\pi[\mu] \succ_i \pi$ for every agent a_i involved in μ .

Observe that for a deal μ of length 1, we have $\pi[\mu] = \pi$. This trivial case consists of an agent giving her object to herself. In particular, when an agent already holds her top object, she should also hold it in the final allocation because of individual rationality.

It is also straightforward to see that a procedure applying only improving deals trivially satisfies individual rationality.

Given an allocation π , we denote by $C_k(\pi)$, $k \geq 2$, the set of all the improving deals of size at most k that can be applied from π :

$$C_k(\pi) = \{\mu \mid \mu \text{ is an improving deal and } |\mu| \leq k\}.$$

When investigating procedures based on improving exchanges, we will try to reach allocations that are stable with respect to some deals.

Definition 3.2. *For a given $k \in \{1, \dots, n\}$, an allocation π is said to be stable with respect to C_k if $C_k(\pi) = \emptyset$.*

We notice that this is equivalent to an allocation being \succ -(N, M)-cycle (or \geq -(N, M)-cycle as preferences are strict here) optimal with $M = 1$ and $N = n$ from Definition 1.19. It is obvious from this definition that if an allocation is stable with respect to C_k for a given k , it is also stable with respect to any $k' < k$. Moreover, since we are considering housing markets, it can be observed that an allocation is Pareto-optimal if and only if it is stable with respect to C_n .

Dynamics based on local exchanges start from an initial allocation and let the agents negotiate improving cycle-deals involving at most k agents until reaching an allocation stable with respect to C_k .

Following this process, to each $k \in \{1, \dots, n\}$ corresponds a family of allocation procedures based on C_k . Indeed, at each step, the improving deal to be implemented can

Algorithm 3.2: Cycle-deals dynamics

Input: An instance $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$,
a maximum size k for the deals,
a selection heuristic σ
Output: An allocation π stable with respect to C_k

```
1  $\pi \leftarrow \pi^0$ 
2 while  $\pi$  is not stable with respect to  $C_k$  do
3    $\mu \leftarrow \sigma(\mathcal{N}, k, \dots)$ 
4   if  $\mu$  is successful then
5      $do(\mu, \pi)$ 
6   end
7 end
8 return  $\pi$ 
```

be selected in many different ways. We call *selection heuristic* a mapping giving the deal to implement at a given step of the procedure. A dynamic based on local exchanges will thus be defined by a given k and a specific selection heuristic. Depending on the selection heuristic, it may be the case that the selected deal is not an improving one. Such a deal could then be refused by the agents. We will say that a deal is *successful* if all agents involved in it agree to exchange. The sequence of proposed deals, together with the fact that they were successful or not will be called a *history*.

A general formal description of cycle-deals dynamics is given in Algorithm 3.2. The $do(\mu, \pi)$ method is similar to the one described in Algorithm 1.1 and implements the deal μ (as chosen by the selection heuristic σ) on the allocation π . Note that the selection heuristic itself takes as input at least the set of agents \mathcal{N} and the maximum size k of the cycle-deals (but it may be more informed and take more parameters, as we shall discuss later in this section). Unless stated otherwise we shall assume from now on that $k = 2$ and omit this parameter, as we will mainly be interested in deals of size 2. We will also use *swap-stability* to refer to stability with respect to C_2 .

We say that the procedure has *reached* the allocation obtained upon termination. Termination occurs when, for every possible swap-deal $\langle a_x, a_y \rangle$, there exists a latest unsuccessful proposal, such that there was subsequently no successful swap-deal involving a_x nor a_y . This guarantees in particular that the process cannot end as long as there remains an improving deal that has not yet been proposed. To ensure termination, the selection heuristic should not prevent some deals from happening. We will thus require heuristics to satisfy a property of full coverage as defined below.

Definition 3.3. A selection heuristic σ has minimal (resp. full) coverage if for any $a_x, a_y \in \mathcal{N}^2$, there exists at least one round in the sequence (resp. after the latest successful swap-deal involving a_x or a_y if there is one), when the swap-deal $\langle a_x, a_y \rangle$ is proposed.

\mathcal{M}_2 will denote the family of swap dynamics defined with respect to a selection heuristic with full coverage. It is useful to make a further distinction between different types of selection heuristics, depending on the information they take as input. In particular, heuristics may require preferential information (e.g. which objects agents would be happy to swap their current object with), or on the contrary, be solely based on observable information (e.g. the history of deals). We will mostly focus on the latter in what follows.

A history-based selection heuristic σ is a function taking as input the set of agents \mathcal{N} , the history of deals h , and returning a deal $\sigma(\mathcal{N}, h)$ to be proposed. Swap dynamics which rely on history-based selection heuristics will be called *history-based swap dynamics*.

A swap dynamics equipped with a given history-based selection heuristic σ on the instance $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$, proceeds iteratively as follows. The history h is initially empty. At each round corresponding to an allocation π and a history h , the allocation is updated to π' defined as:

$$\pi' = \begin{cases} \pi[\sigma(\mathcal{N}, h)] & \text{if the agents in } \sigma(\mathcal{N}, h) \text{ agree on swapping} \\ \pi & \text{otherwise.} \end{cases}$$

The deal $\sigma(\mathcal{N}, h)$ together with its “success status” are then added to the history h .

We now give some examples of selection heuristics which will be studied in this chapter. We start with *round-robin* heuristics, which simply specify a predefined way to order the different pairs of agents (i.e. possible swaps), and repeat it until termination. There are several ways to proceed, we give here two prominent examples:

- *Round-Robin over the Agents* (M_2 -RRA): agents are ordered and paired following their name ($\langle a_1, \dots, a_n \rangle$) in a round-robin fashion. The first agent a_1 is paired with each other agent, the second agent a_2 is paired with each other agent a_j with $j > 2$ and so on. The agents then go over possible deals by iterating over the following sequence:

$$(a_1, a_2), (a_1, a_3), \dots, (a_1, a_n), (a_2, a_3), \dots, (a_{n-1}, a_n), (a_1, a_2), \dots$$

- *Round-Robin over Pairs of agents* (M_2 -RRP): agents are ordered and paired following their name ($\langle a_1, \dots, a_n \rangle$) in a round-robin fashion. In this case, the first agent is paired with the second agent, the second agent is paired with the third... Hence, the agents go over possible deals by iterating over the following sequence:

$$(a_1, a_2), (a_3, a_4), \dots, (a_{n-1}, a_n), (a_1, a_3), (a_2, a_4), \dots, (a_1, a_n), \dots$$

There is an obvious bias in the way deals are selected, it is thus natural to introduce some stochasticity in the process. One way to do this is as follows:

- *Randomized Round-Robin over deals* (M_2 -RRR): the heuristic picks uniformly at random a permutation among all the possible deals. This permutation defines a round-robin order in which the deals are considered. The agents then go over possible deals by iterating over this permutation.

The three previous heuristics all guarantee by construction that all possible swap-deals were proposed before a swap-deal gets proposed again. The following natural heuristic do not have such guarantee.

- *Uniform* (M_2 -U): a pair of agents (i.e. a swap-deal) (a_i, a_j) is selected uniformly at random among all possible pairs.

- *Random matchings* (M_2 -RM): this heuristic proceeds in succession of matching steps where at each step, every agent is matched to a unique other agent and all the resulting swap-deals are proposed to the agents. This simulates a natural market where agents are paired randomly and simultaneously try to perform bilateral deals.

Note that these selection heuristics do not all have the same degree of decentralization. In particular, M_2 -U can be easily executed in a fully decentralized way. Round-robin heuristics require a central entity to broadcast the sequence of pairs. Then, the agents can meet in a distributed way. Finally, M_2 -RM can be implemented in a distributed way using a protocol ensuring that an agent can not encounter several agents simultaneously. Each agent a_i then selects uniformly at random an agent a_j to encounter and contact her. If agent a_j is already engaged in another encounter, a_i selects another agent.

We conclude the section by a straightforward observation: different selection heuristics can lead to different outcomes given the same initial allocation. Let us illustrate this with a simple example.

Example 3.2. *Consider the instance described in Example 1.12:*

$$\begin{aligned}
a_1 &: o_3 \succ_1 o_4 \succ_1 o_5 \succ_1 o_2 \succ_1 \underline{o_1} \\
a_2 &: o_3 \succ_2 o_4 \succ_2 o_5 \succ_2 \underline{o_2} \succ_2 o_1 \\
a_3 &: o_4 \succ_3 o_5 \succ_3 \underline{o_3} \succ_3 o_2 \succ_3 o_1 \\
a_4 &: o_3 \succ_4 \underline{o_4} \succ_4 o_5 \succ_4 o_2 \succ_4 o_1 \\
a_5 &: o_1 \succ_5 o_2 \succ_5 o_3 \succ_5 o_4 \succ_5 \underline{o_5}
\end{aligned}$$

We showed that *Crawler* and *TTC* return different allocations. In fact, these allocations can be reached by a sequence of swap-deals. Observe first that the cycle-deals applied by *TTC* are at most of length 2 (Example 1.14). The allocation $\pi^{TTC} = \langle o_5, o_2, o_4, o_3, o_1 \rangle$ is obtained through two swap-deals: $\langle a_3, a_4 \rangle$ and $\langle a_1, a_5 \rangle$. The allocation $\pi^C = \langle o_2, o_5, o_4, o_3, o_1 \rangle$ returned by *Crawler* also is reachable by swap-deals. It is obtained by applying the following sequence: $\langle a_3, a_4 \rangle$, $\langle a_2, a_5 \rangle$ and $\langle a_5, a_1 \rangle$. Since these two allocations are Pareto-optimal, they are swap-stable. This shows that the way swap-deals are selected may affect the final stable allocation. Notice that these are not the only two swap-stable allocations reachable from π^0 of the instance: $\langle o_3, o_2, o_5, o_4, o_1 \rangle$ and $\langle o_2, o_3, o_5, o_4, o_1 \rangle$ are the two other ones.

This example suggests that the allocations returned by *Crawler* and by *TTC* can both be reached via swap-deals — a point we will make formal in the next section. More generally, as *TTC* and *Crawler* both provide desirable guarantees, on Pareto-optimality, individual rationality and strategy-proofness, it is natural to study whether swap dynamics enjoy similar properties.

3.4 Properties of swap dynamics

We now investigate properties of swap dynamics. Recall first the easy observation made in Section 3.2: swap dynamics, because they rely on improving deals, are individually

rational. But what about strategy-proofness and Pareto-optimality? We first discuss strategy-proofness, showing that such procedures are in general subject to manipulation. We next prove that any allocation stable with respect to swap-deals is Pareto-optimal and that both the allocation returned by TTC and Crawler can be reached via swap-deals. Finally, we show that the single-peaked domain is a maximal domain when it comes to guaranteeing Pareto-optimality.

3.4.1 Strategy-proofness

Although strategy-proofness is usually defined for centralized procedures, the question is still relevant for swap dynamics. The basic definitions of single-peakedness have been introduced in Section 1.2.3. Note that based on these definitions, it is straightforward that restricting a single-peaked preference domain preserves its single-peakedness.

In a swap dynamic, an agent has the opportunity to behave strategically only when she is asked to accept or reject a proposed deal. As usual, we assume here that a potential manipulator has full knowledge of the preferences of the other agents and is aware of the fact that the selection heuristics has minimal coverage. In other words, the manipulator only needs to know whether she will ever have the opportunity to swap with some other agent or not. A swap dynamic will be called *strategy-proof* when, for every instance, at no point during the procedure an agent can be better off by accepting a non-improving deal or by refusing an improving one. Otherwise, the procedure is *manipulable*.

Proposition 3.1. *Any history-based swap dynamic with minimal coverage is manipulable.*

Proof. We are given an arbitrary history of deals h starting with (a_x, a_y) , as produced by the selection heuristic. We are now going to show that we can build an instance $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$ such that agent a_x will have an incentive to accept the first swap-deal (a_x, a_y) , even though this is not rational. The instance involves a_x, a_y, a_z , as well as $n - 3$ other dummy agents, and assumes preferences to be single-peaked with respect to $o_1 \triangleleft o_2 \triangleleft o_3 \triangleleft o_{d_1} \triangleleft \dots \triangleleft o_{d_{n-3}}$.

$$\begin{aligned}
a_x &: o_2 \succ_1 \underline{o_3} \succ_1 o_1 \succ_1 o_{d_1} \dots \\
a_y &: o_2 \succ_2 o_3 \succ_2 \underline{o_1} \succ_2 o_{d_1} \dots \\
a_z &: o_1 \succ_3 \underline{o_2} \succ_3 o_3 \succ_3 o_{d_1} \dots \\
d_1 &: \underline{o_{d_1}} \succ \dots \\
&\vdots \\
d_{n-3} &: \underline{o_{d_{n-3}}} \succ \dots
\end{aligned}$$

Observe first that all the dummy agents have their top object from the start, hence they will not be involved in any deal. Now, consider the situation when a_x is given the opportunity to deal with a_y . If a_x is truthful (as all the other agents), no improving swap-deal involving agent a_x can occur, and she will end up with object o_3 . Now, suppose instead that a_x strategically accepts $\langle a_x, a_y \rangle$. This deal is improving for a_y so if a_x agrees on it, it will be implemented. In that case, a_y can no longer exchange with a_z , hence the only improving swap-deal left is $\langle a_x, a_z \rangle$. Because the selection heuristic has minimal

coverage, this opportunity will occur at some point in the future. In the final allocation, agent a_x would then hold her top object o_2 (obtained from the swap-deal $\langle a_x, a_z \rangle$). \square

Now, suppose that for a given selection heuristic, there exists a swap-deal $\langle a_x, a_y \rangle$ that will never be proposed. Consider then the instance where all the agents have their top object, except for a_x and a_y who would be happy to perform a swap-deal—but will never get a chance to. As no other swap-deal is possible, the outcome is certainly not Pareto-optimal. This leads to the following observation.

Observation 3.1. *Consider a history-based swap dynamic M_2 defined with respect to a selection heuristic σ . If M_2 is Pareto-optimal, then σ has minimal coverage.*

Together with Proposition 3.1, this tells us that there is a fundamental tension between strategy-proofness and Pareto-optimality for history-based swap-deal procedure.

Proposition 3.2. *No history-based swap dynamic can be both Pareto-optimal and strategy-proof.*

The reason why other types of dynamics may not be concerned by this result is that they can potentially condition the selection of deals to the preferences.

3.4.2 Pareto-optimality of swap dynamics

We have already shown in Chapter 2 (Proposition 2.8) that swap-stability is equivalent to Pareto-optimality in this setting, hence any allocation reached by swap dynamics is Pareto-optimal.

This theorem states that the C_k -stability hierarchy collapses at the C_2 level in single-peaked housing markets. Remember that every Pareto-optimal allocation is stable with respect to C_n . This result provides then a new characterization of Pareto-optimality in our setting.

Corollary 3.1. *In a single-peaked housing market, an allocation π is Pareto-optimal if and only if it is stable with respect to C_2 .*

Stating this result in terms of stability with respect to C_n and not just Pareto-optimality gives us more flexibility. Indeed, in the more general setting where there are more objects to allocate than the number of agents, the result of Proposition 2.8 from Chapter 2 still holds [Beynier et al., 2019a] but it is no longer the case that Pareto-optimality implies C_n -stability.

As we have proven that the allocation reached by swap-deals is Pareto-optimal, a natural question is then whether *every* allocation that Pareto-dominates the initial allocation can be reached by swap-deals. It is not the case.

Proposition 3.3. *There exists an instance $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$ for which there is an allocation π that Pareto-dominates π^0 and that can not be reached by a sequence of improving swap-deals.*

Proof. Let us consider the following instance with three agents where the initial allocation $\pi^0 = \langle \underline{o_3}, o_2, o_1 \rangle$ is the one underlined.

$$\begin{array}{l} a_1 : \quad o_1 \quad \gamma_1 \quad \boxed{o_2} \quad \gamma_1 \quad \underline{o_3} \\ a_2 : \quad \boxed{o_1} \quad \gamma_2 \quad \underline{o_2} \quad \gamma_2 \quad o_3 \\ a_3 : \quad o_2 \quad \gamma_3 \quad \boxed{o_3} \quad \gamma_3 \quad \underline{o_1} \end{array}$$

The squared allocation $\boxed{\pi} = \langle o_2, o_1, o_3 \rangle$ is Pareto-optimal. However from π^0 only two deals are possible: $\mu_1 = \langle a_1, a_3 \rangle$ that reaches allocation $\pi' = \langle o_1, o_2, o_3 \rangle$, or $\mu_2 = \langle a_2, a_3 \rangle$ that leads to $\pi'' = \langle o_3, o_1, o_2 \rangle$. No sequence of improving swap-deals can thus reach π . \square

It is however interesting to note that both the allocation returned by TTC and by Crawler can always be reached through improving swap-deals.

Proposition 3.4. *Let $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$ be an instance and let π^{TTC} be the allocation returned by the TTC algorithm on I . Then π^{TTC} is reachable by swap-deals from π^0 .*

Proof. We show that any cycle-deal applied by TTC can be implemented as a sequence of improving swap-deals. For a given step of the algorithm, we consider \mathcal{O} the set of objects and \mathcal{N} the set of agents remaining at this step and π the current partial allocation.

Let μ be the next cycle-deal to be applied, and let o_μ and N_μ be respectively the set of objects and the set of agents involved in μ . We must prove that the allocation $\pi[\mu]$ can be reached by a sequence of swap-deals. When $|o_\mu| = 2$, μ is actually a swap-deal and the claim is trivially true.

Assume then that $|o_\mu| > 2$, let us consider the instance I' obtained from I by restricting \mathcal{O} to o_μ and \mathcal{N} to N_μ . In this restricted instance, the allocation π^* where each agent receives her top object is feasible. Indeed, it is the allocation obtained by applying μ (by definition of μ in TTC). This allocation is trivially Pareto-optimal as every agent has her top objects and it thus Pareto-dominates every other allocation. It is then the only Pareto-optimal allocation in I' . Remember from Observation 1.1 that I' is single-peaked. Then, by virtue of Proposition 2.8 from Chapter 2, this implies that π^* must be reachable by improving swap-deals in I' . Overall, π_μ is then reachable by improving swap-deals in I .

The same argument can be stated for any step of the algorithm. By concatenating the sequences of improving swap-deals for every step of the algorithm, we then obtain a sequence of improving swap-deals leading to π^{TTC} from π^0 . \square

Proposition 3.5. *Let $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$ be an instance and let π^C be the allocation returned by Crawler. Then π^C is reachable by swap-deals from π^0 .*

Proof. We show that every cycle-deal applied by Crawler can be implemented as a sequence of improving swap-deals. For clarity reasons and without loss of generality, we assume that each agent a_j currently holds object o_j .

For a given iteration of the algorithm, consider agent a_i who picks object o_k currently held by agent a_k . From the definition of the procedure, a_k is on the left of a_i (with respect to the single-peaked axis) and a_k has already been considered at this step before considering a_i . In fact, all the agents between a_k (included) and a_i (excluded) on the left of the single-peaked axis have already been considered at the current iteration before reaching a_i . Moreover, all these agents have passed their turn because their peak is on

their right. In other words, each agent a_j between a_k (included) and a_i (excluded) prefers the object held by the agent a_{j+1} on her right, that is:

$$o_{j+1} \succ_j o_j, \forall j \in \{k, \dots, i-1\}.$$

Let $\mu = \langle a_i, a_{i-1}, \dots, a_{k+1}, a_k \rangle$ be the cycle-deal implemented by Crawler at the current step. In this deal, a_k gives her object o_k to a_i and all the other agents of the deal give their object to the next agent in the sequence which is the agent on their left with respect to the single-peaked axis. The decomposition of μ into a sequence of swap-deals consists in using agent a_i as a hub for the swap-deals. Agent a_i first exchanges with a_{i-1} then, a_i exchanges with a_{i-2} and so on until a_i performs a swap-deal with a_k . At the end, a_i holds o_k and each other agent a_j in μ_i holds the object initially held by a_{j+1} . The sequence of swap-deals is thus equivalent to μ_i .

We now show that all these swap-deals are improving. In the first deal $\langle a_i, a_{i-1} \rangle$, a_{i-1} receives the object o_i held by a_i that she prefers to her current object o_{i-1} (as shown previously $o_i \succ_{i-1} o_{i-1}$). Simultaneously, agent a_i receives the object o_{i-1} held by a_{i-1} that she prefers to her current object since her peak is on the left of a_{i-1} (it is held by a_k), that is:

$$o_j \succ_i o_{j+1}, \forall j \in \{k, \dots, i-1\}.$$

The first deal is thus mutually improving. Concerning the next swap-deal $\langle a_i, a_j \rangle$ with $j \in \{k, \dots, i-2\}$, a_i exchanges o_{j+1} that she obtained from her previous swap-deal, against o_j held by a_j . Agent a_j receives the object o_{j+1} that she prefers to her current object o_j (as shown previously $o_{j+1} \succ_j o_j$). a_i receives the object o_j that she prefers to o_{j+1} (since $\forall j \in \{k, \dots, i-1\}, o_j \succ_i o_{j+1}$). All the swap-deals are thus mutually improving. \square

These two results show that both the outcome of TTC and Crawler could be implemented as sequences of mutually beneficial swap-deals in single-peaked domains. This attenuates a bit a critical feature of these procedures. However, even though these long trading cycles can be broken down into swap-deals, they would still need to be pre-computed, and carefully coordinated.

The proof of Proposition 3.5 shows how to compute the set of swap-deals that achieves Crawler outcome: for each cycle-deals $\mu = \langle a_i, a_{i-1}, \dots, a_{k+1}, a_k \rangle$ implemented, a_i is used as a hub to decompose the deal into mutually beneficial swap-deals. On the contrary, the proof of Proposition 3.4 does not provide explicitly the corresponding sequence of deals. We show now how to compute the set of swap-deals that achieves the TTC outcome using Crawler. Note that for each cycle-deal selected by TTC, we can build a sub-instance of the initial instance and apply Crawler on this sub-instance to obtain the set of swap-deals that implements the cycle-deal. Since Crawler is Pareto-optimal, the following observation is straightforward:

Observation 3.2. *In a single-peaked housing market, if the allocation where each agent has her top object is feasible (i.e. each agent has a different top object), then the allocation is returned by Crawler.*

We can now compute the set of swap-deals that achieves the TTC outcome using Crawler. For each deal μ implemented by TTC:

- If $|\mu| \leq 2$, it is already a swap-deal and we are done.

- If $|\mu| > 2$, the deal has to be decomposed into a set of swap-deals. Let I' be the sub-instance restricted to the agents and the objects involved in μ . This instance is guaranteed to remain single-peaked (see Observation 3.1). By definition of TTC, in I' , each agent obtains her top object. By Observation 3.2, the Crawler outcome will be the same as the outcome of μ when applying Crawler to I' . The swap-deal decomposition of μ is then obtained by applying Crawler on I' and decomposing each cycle-deal μ' of Crawler as explained above.

Again, it should be kept in mind that in practice, a coordination mechanism would have to ensure that the agents indeed execute the correct sequence of swap-deals (and thus the desired allocation would be obtained).

Figure 3.1 concludes this section by summarizing the different findings related the properties of swap dynamics.

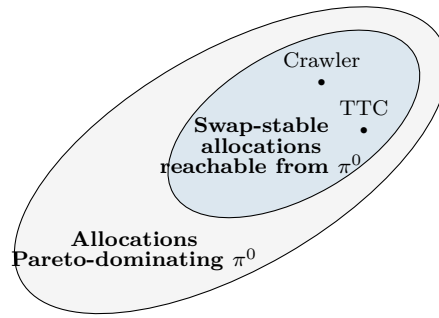


Figure 3.1: Summary of the results of Proposition 2.8 from Chapter 2 and Propositions 3.3, 3.4 and 3.5 for a given initial allocation π^0 . Remember from Example 1.12 that TTC and Crawler can return different allocations.

3.4.3 Maximality of the single-peaked domain

We now show that the single-peaked domain is maximal for the Pareto-optimality of swap dynamics: For every preference domain D such that $\mathcal{SP}_{\triangleleft}$ is a proper subset of D (for a given linear order \triangleleft) there exists an instance such that no swap dynamics can reach a Pareto-optimal allocation.

Before going through the proof, let us first start with a simple example. Consider the following profile and let \triangleleft be the order $o_1 \triangleleft o_2 \triangleleft o_3$.

$$\begin{aligned}
 a_1 &: o_1 \succ_1 \underline{o_2} \succ_1 o_3 \\
 a_2 &: o_3 \succ_2 \underline{o_1} \succ_2 o_2 \\
 a_3 &: o_2 \succ_3 \underline{o_3} \succ_3 o_1
 \end{aligned}$$

It is clear that the profile is not single-peaked with respect to \triangleleft (it actually is not for any order) as it is clearly not worst-restricted. Note that the preferences of agents a_1 and a_3 are single-peaked with respect to \triangleleft but not those of agent a_2 .

Take the underlined objects to form the initial allocation. It has been constructed as follows. Agents a_1 and a_3 receive their second best object (respectively o_2 and o_3). For agent a_2 , the violation of single-peakedness comes from her ranking of objects o_1 and o_2 . She is thus allocated her most preferred object among these two, i.e. o_1 . It can easily

be checked that this allocation is stable with respect to C_2 but is Pareto-dominated by $\langle o_1, o_3, o_2 \rangle$.

This construction is generalized to an arbitrary number of agents below.

Theorem 3.1. *Let \mathcal{O} be a set of objects and \triangleleft a linear order over \mathcal{O} . For every preference domain D such that $\mathcal{SP}_{\triangleleft} \subsetneq D$, there exists an instance $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$ defined over D such that swap dynamics on I do not reach a Pareto-optimal allocation.*

Proof. Let us construct an instance $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$ defined over D such that swap dynamics on I do not reach a Pareto-optimal allocation.

Without loss of generality and for the ease of the reader, let us assume that \triangleleft is such that $o_1 \triangleleft o_2 \triangleleft \dots \triangleleft o_n$.

Since $\mathcal{SP}_{\triangleleft} \subsetneq D$, there exists a linear order $\succ^* \in D$ that is not single-peaked with respect to \triangleleft . There exist then two objects $(o_{s-1}, o_s) \in \mathcal{O}^2$ such that:

$$\left. \begin{array}{l} o_s \triangleleft o_{s-1} \triangleleft \text{top}(\succ^*), \\ \text{or, } \text{top}(\succ^*) \triangleleft o_{s-1} \triangleleft o_s \end{array} \right\} \text{ and } o_s \succ_i^* o_{s-1}.$$

Without loss of generality, let us assume that $\text{top}(\succ^*) \triangleleft o_{s-1} \triangleleft o_s$. We will use index t to refer to the top ranked object in \succ^* , i.e., $o_t = \text{top}(\succ^*)$.

Moreover, as $\mathcal{SP}_{\triangleleft} \subset D$, for every object $o_i \in \mathcal{O}$ there exist in D two linear orders \succ_i^1 and \succ_i^2 that are single-peaked with respect to \triangleleft such that:

$$\text{top}(\succ_i^1) = \text{top}(\succ_i^2) = o_i \text{ and } \text{snd}(\succ_i^2) = o_{i-1}.$$

Thus, in the preference order \succ_i^1 the top object is o_i and the other objects are not constrained. For \succ_i^2 , the top object is o_i and the second best must be o_{i-1} .

We introduce the preference profile $L = \{\succ_i \mid a_i \in \mathcal{N}\}$ defined as follows:

$$\begin{aligned} \gamma_1 &= \gamma^*, \\ \gamma_i &= \gamma_{i-1}^1, \quad \forall a_i \in \mathcal{N}, i \in \{2, \dots, t\}, \\ \gamma_i &= \gamma_i^2, \quad \forall a_i \in \mathcal{N}, i \in \{t+1, \dots, s\}, \\ \gamma_i &= \gamma_i^1, \quad \forall a_i \in \mathcal{N}, i \in \{s+1, \dots, n\}. \end{aligned}$$

The initial allocation π^0 is then defined as:

$$\begin{aligned} \pi_{a_1}^0 &= o_s, \\ \pi_{a_i}^0 &= \text{top}(a_i) = o_{i-1}, \quad \forall a_i \in \mathcal{N}, i \in \{2, \dots, t\}, \\ \pi_{a_i}^0 &= \text{snd}(a_i) = o_{i-1}, \quad \forall a_i \in \mathcal{N}, i \in \{t+1, \dots, s\}, \\ \pi_{a_i}^0 &= \text{top}(a_i) = o_i, \quad \forall a_i \in \mathcal{N}, i \in \{s+1, \dots, n\}. \end{aligned}$$

To get a better understanding of the instance $I = \langle \mathcal{N}, \mathcal{O}, L, \pi^0 \rangle$ constructed in this proof, Figure 3.2 presents the preference profile L and the initial allocation π^0 .

We claim that π^0 is stable with respect to C_2 but not Pareto-optimal. Allocation π^0 is clearly not Pareto-optimal as the allocation in which every agent receives her top object is feasible—because no two agents have the same top object—and this allocation clearly Pareto-dominates π^0 . Let us now show that $C_2(\pi^0) = \emptyset$.

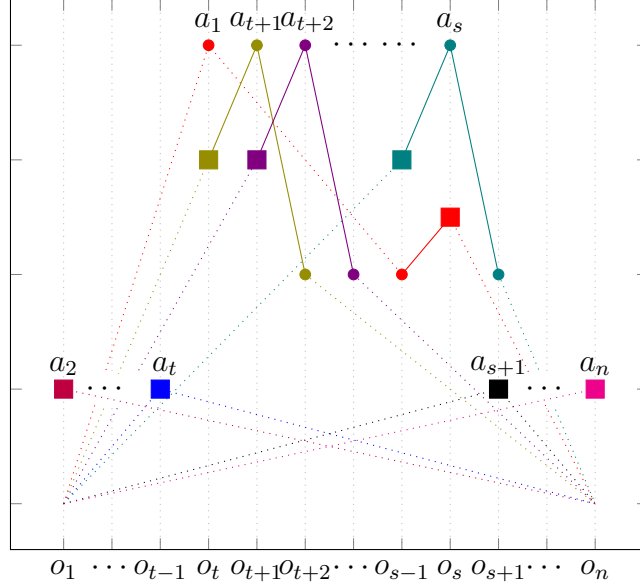


Figure 3.2: The instance constructed in the proof of Theorem 3.1. The dotted lines represent the trend of the utilities, dots fixed points and squares the object allocated to each agent.

First observe that every agent $a_i, i \in \{2, \dots, t\} \cup \{s+1, \dots, n\}$ owns her top object, hence can not be involved in an improving swap-deal.

Consider now agent $a_i, i \in \{t+1, \dots, s-1\}$. She owns her second most preferred object which she would only trade against her top object that is owned by agent a_{i+1} . However, agent a_{i+1} is not interested in π_{a_i} , hence no improving swap-deal is possible.

Finally, let us consider agent a_s who owns her second most preferred object o_{s-1} and whose top object o_s is owned by agent a_1 . By the hypothesis that a_1 's preferences are not single-peaked, we have $o_s \succ_1 o_{s-1}$. Once again, there is no improving swap-deal involving agent a_s .

Overall, we have proved that $C_2(\pi^0) = \emptyset$, hence any swap dynamic returns π^0 on I . As π^0 is Pareto-dominated, this concludes the proof. \square

This result is particularly interesting since it shows that the single-peaked domain captures in a “tight” way the domain under which swap dynamics are Pareto-optimal (in the vein of similar results obtained by Chevaleyre et al. [2010] in different settings).

Note that this is not a characterization result: there are some domains that are not single-peaked but for which swap dynamics return Pareto-optimal allocations.

Example 3.3. *Let us consider the following preference profile:*

$$\begin{aligned} a_1 &: o_1 \succ_1 o_2 \succ_1 o_3 \\ a_2 &: o_1 \succ_2 o_3 \succ_2 o_2 \\ a_3 &: o_3 \succ_3 o_2 \succ_3 o_1 \end{aligned}$$

This profile is not single-peaked over any linear order \triangleleft : the triplet $\langle o_1, o_2, o_3 \rangle$ is a witness of the violation of the worst-restrictedness condition (Proposition 1.1), however we can show that for every initial allocation, swap dynamics return a Pareto-optimal allocation. Let us consider the different initial allocations:

1. $\pi^0 = \langle o_1, o_2, o_3 \rangle$: $C_2(\pi^0) = \emptyset$ and π^0 is Pareto-optimal.

2. $\pi^0 = \langle o_1, o_3, o_2 \rangle$: $C_2(\pi^0) = \emptyset$ and π^0 is Pareto-optimal.
3. $\pi^0 = \langle o_2, o_1, o_3 \rangle$: $C_2(\pi^0) = \emptyset$ and π^0 is Pareto-optimal.
4. $\pi^0 = \langle o_2, o_3, o_1 \rangle$: two swap-deals are possible:
 - (a) $\langle a_1, a_3 \rangle$ which leads to case 2.
 - (b) $\langle a_2, a_3 \rangle$ which leads to case 3.
5. $\pi^0 = \langle o_3, o_1, o_2 \rangle$: one swap-deal is possible: $\langle a_1, a_3 \rangle$ which leads to case 3.
6. $\pi^0 = \langle o_3, o_2, o_1 \rangle$: three swap-deals are possible:
 - (a) $\langle a_1, a_3 \rangle$ which leads to case 1.
 - (b) $\langle a_1, a_2 \rangle$ which leads to case 4.
 - (c) $\langle a_2, a_3 \rangle$ which leads to case 5.

This example does not contradict Theorem 3.1: Swap dynamics are Pareto-optimal on the domain $D = \{\succ_1, \succ_2, \succ_3\}$ because D does not include every single-peaked linear order over \triangleleft . For instance take \triangleleft to be $o_1 \triangleleft o_2 \triangleleft o_3$, extend D by adding the two following linear orders: $o_2 \succ o_1 \succ o_3$ and $o_2 \succ o_3 \succ o_1$ so that $\mathcal{SP}_{\triangleleft} \subset D'$ and Theorem 3.1 will apply.

3.5 The “price of” swap dynamics

It is natural to now ask to what extent swap dynamics induce a cost in terms of social welfare. In that perspective we will first discuss the “standard” *price of anarchy (PoA)* [Koutsoupias and Papadimitriou, 1999; Anshelevich et al., 2013], that is, the (worst-case, over all instances) ratio between the worst stable outcome and the social welfare optimum. More formally, for an allocation procedure M :

$$PoA_v(M) = \max_{I \in \mathcal{I}_D} \frac{\max_{\pi \in M(I)} v(\pi)}{\min_{\pi \in M(I)} v(\pi)}, \text{ with } v \in \{ark, mrk\}.$$

with *ark* the average rank and *mrk* the minimum rank as introduced in Chapter 1.

But, as we have discussed, swap dynamics constitute a family of procedures. In that context, a perhaps more relevant metric is the ratio between the best and worst stable outcome which can be obtained by such procedures. In particular, this would tell us the price to pay for not being guided by a central planner (in the selection of deals) towards maximizing our social welfare notion. We call this notion the *price of dynamics (PoD)*, and for a family of dynamics \mathcal{M} , we define it as follows:

$$PoD_v(\mathcal{M}) = \max_{M, M' \in \mathcal{M}} \frac{\max_{\pi \in M'(I)} v(\pi)}{\min_{\pi \in M(I)} v(\pi)}, \text{ with } v \in \{ark, mrk\}.$$

Note that this notion is not specific to swap dynamics and can be applied to any family of dynamics. Also, observe that an upper bound on the *PoA* also applies to the *PoD*.

These definitions are parametrized by the social welfare notion considered. In terms of average rank, we recall that Damamme et al. [2015] established that the price of anarchy

is 2 for swap dynamics *in the general domain*. While the upper bound remains valid in our restricted domain, the instance exhibited to show this bound to be tight in [Damamme et al., 2015] violates single-peakedness. We now show that this result still holds under our domain restriction.

Proposition 3.6. *For any $M \in \mathcal{M}_2$, $PoA_{ark}(M)$ is 2 in the single-peaked domain.*

Proof. Let us consider the following single-peaked instance involving n agents, and assume without loss of generality n to be odd. Take for now the shaded allocation as the initial allocation.²

$$\begin{array}{l}
a_1 : \boxed{o_{n-1}} \succ o_n \succ o_{n-2} \succ o_{n-3} \cdots o_3 \succ \underline{o_2} \succ o_1 \\
a_2 : \boxed{o_1} \succ \boxed{o_2} \succ \underline{o_3} \succ o_4 \cdots o_{n-2} \succ o_{n-1} \succ o_n \\
a_3 : \boxed{o_2} \succ o_1 \succ \boxed{o_3} \succ \underline{o_4} \cdots o_{n-2} \succ o_{n-1} \succ o_n \\
\vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots \\
a_{n-2} : \boxed{o_{n-3}} \succ o_{n-4} \succ o_{n-5} \succ o_{n-4} \cdots \boxed{o_{n-2}} \succ \underline{o_{n-1}} \succ o_n \\
a_{n-1} : \boxed{o_{n-2}} \succ o_{n-3} \succ o_{n-4} \succ o_{n-5} \cdots o_{n-1} \succ \boxed{o_n} \succ \underline{o_1} \\
a_n : \boxed{o_1} \succ o_2 \succ o_3 \succ o_4 \cdots o_{n-2} \succ o_{n-1} \succ \boxed{o_n}
\end{array}$$

Let π be the shaded allocation, and π^* be the squared allocation (a_1 holds o_{n-1} in both cases). Observe that both π and π^* are Pareto-optimal allocations. For π^* this is obvious as only a_n does not hold her top object. For π , notice that a_2 would only wish to swap for o_1 , which is held by an agent who ranks it first. Then a_3 would only wish to swap for o_1 or o_2 , and so on until a_{n-2} . Finally, a_{n-1} may swap with anyone but a_n , but no one wants to swap with her.

We see that $ark(\pi^*) = [(n-1) \cdot n + 1]/n$, while $ark(\pi) = [n \cdot (n+1)/2 + n - 1]/n$, thus this instance shows that (asymptotically) the PoA is at least 2. In Damamme et al. [2015] it is shown that 2 is an upper bound for the PoA of any swap-deal procedure. We thus conclude that the PoA is 2 here, as in the general domain. \square

Following Damamme et al. [2015], we can also make some further observations. In fact, as both allocations are Pareto-optimal, this instance shows that the PoA of any procedure satisfying individual rationality must be at least 2. Indeed, taking the shaded allocation π as the initial allocation, any individually rational procedure would output π . Furthermore, as any Pareto-optimal allocation is stable for swap-deals, our upper bound remains valid.

Observation 3.3. *For any individually rational allocation procedure M , $PoA_{ark}(M)$ is 2 in the single-peaked domain. This holds in particular for TTC and Crawler.*

To address the price of dynamics, we must now make use of a different initial allocation, and show that both allocations are reachable by sequences of swaps. This would demonstrate that this gap could also only be due to the selection heuristics.

Proposition 3.7. *For the family of swap dynamics \mathcal{M}_2 , $PoD_{ark}(\mathcal{M}_2)$ is 2 in the single-peaked domain.*

²The underlined allocation will only be relevant in forthcoming proofs.

Proof. Consider again the instance presented in the proof of Proposition 3.7. Take now the initial allocation to be $\pi^0 = \langle \underline{o_2}, o_3, o_4, \dots, o_{n-1}, o_1, o_n \rangle$, the underlined allocation. We first show that the shaded allocation π can be reached from π^0 via improving swap-deals. First the deal $\langle a_n, a_{n-1} \rangle$ is implemented. Then a_1 acts as a hub, and the following sequence is implemented $\langle a_1, a_2 \rangle \dots \langle a_1, a_{n-2} \rangle$.

Now, the squared allocation π^* can also be reached with swap-deals. Agent a_1 will act as a hub for odd agents, while a_{n-1} will do the same for even agents. In other words, the following sequence is implemented: $\langle a_1, a_3 \rangle, \langle a_1, a_5 \rangle \dots \langle a_1, a_{n-2} \rangle$ (note that a_n gets her final object from the start), followed by the sequence of swaps $\langle a_{n-1}, a_2 \rangle, \langle a_{n-1}, a_4 \rangle \dots \langle a_{n-1}, a_{n-3} \rangle$.

The result follows then from the computations of the average rank presented in Proposition 3.7. □

Finally, we turn our attention to the minimum rank counterparts of the same notions. Note that n is certainly an upper bound in that case, as this is highest possible ratio between two allocations.

Proposition 3.8. *For any $M \in \mathcal{M}_2$, $PoA_{mrk}(M)$ is $\Theta(n)$ in the single-peaked domain.*

Proof. Let us consider the following single-peaked instance involving n agents. Take the shaded allocation as the initial allocation.³ Let π be the shaded allocation, and π^* be the squared allocation.

$$\begin{array}{cccccccccccc}
a_1 : & \boxed{o_1} & \succ & o_2 & \succ & o_3 & \succ & o_4 & \cdots & \underline{o_{n-1}} & \succ & o_n \\
a_2 : & \boxed{o_2} & \succ & \boxed{o_3} & \succ & \underline{o_1} & \succ & o_4 & \cdots & o_{n-1} & \succ & o_n \\
a_3 : & \boxed{o_3} & \succ & \boxed{o_4} & \succ & \underline{o_2} & \succ & o_1 & \cdots & o_{n-1} & \succ & o_n \\
a_4 : & \boxed{o_4} & \succ & \boxed{o_5} & \succ & \underline{o_3} & \succ & o_2 & \cdots & o_{n-1} & \succ & o_n \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
a_n : & \boxed{o_2} & \succ & o_1 & \succ & o_3 & \succ & o_4 & \cdots & o_{n-1} & \succ & \boxed{o_n}
\end{array}$$

It is easy to see that $mrk(\pi^*) = n - 1$, while $mrk(\pi) = 1$, thus PoA is $\Omega(n)$. Clearly, the PoA cannot be worse, thus it is $\Theta(n)$. □

Proposition 3.9. *$PoD_{mrk}(\mathcal{M}_2)$ is $\Theta(n)$ in the single-peaked domain.*

Proof. Let us consider the single-peaked instance of Proposition 3.8 and let us assume that the initial allocation is now $\pi^0 = \langle o_{n-1}, o_1, o_2, \dots, o_{n-2}, o_n \rangle$, the underlined allocation.

From this initial allocation π^0 , it is possible to reach the shaded allocation $\pi = \langle o_1, o_2, o_3, \dots, o_{n-1}, o_n \rangle$ with improving swap deals. This allocation is reached by performing $n - 2$ swap-deals $\langle a_1, a_i \rangle, \forall i \in \{n - 1, \dots, 2\}$ starting with $\langle a_1, a_{n-1} \rangle$ and finishing with $\langle a_1, a_2 \rangle$.

From π^0 , it is also possible to reach $\pi^* = \langle o_1, o_3, o_4, \dots, o_n, o_2 \rangle$ (the squared allocation) which is both Pareto-optimal and optimal in terms of minimum rank. If n is even, we perform the following sequence of couples of swap-deals $\langle a_{i-1}, a_n \rangle$ and $\langle a_{i-2}, a_1 \rangle \forall i \in \{n, n - 2, \dots, 6, 4\}$. If n is odd, we perform the same sequence of couples of swap-deals but $\forall i \in \{n, n - 2, \dots, 7, 5\}$ and add two final swap-deals $\langle a_2, a_n \rangle$ and $\langle a_1, a_n \rangle$.

³Again, the underlined allocation will only be relevant in a later proof.

The result follows from the previous computations of the minimum rank in these two allocations. □

These different theoretical results remain a worst-case analysis. To complete the picture, it would be valuable to know how well these swap dynamics do in practice. This is what we investigate in the next section.

3.6 Experimental study

Our objective in this section is to get some insights about the empirical behaviour of swap dynamics. Our study includes swap dynamics based on the different history-based selection heuristics introduced in Section 3.3, as well as two other selection heuristics added for comparison: one which takes as input more than the mere history of deals (M_2 -PW), and another one which relaxes the constraint of using swaps only (M_3 -U). For completeness, we also compare their performance to that of TTC and Crawler (keeping in mind though that these two centralized procedures are not specifically designed to optimize our social welfare measures).

We now give the full detail of our protocol by specifying (i) how deals are selected in the swap dynamics, (ii) how preferences are generated, as well as (iii) the full specification of the parameters used in the experiments.

3.6.1 Experimental protocol

Selection heuristics for swap dynamics

Concerning swap-deal procedures, our study includes the history-based selection heuristics already introduced in Section 3.3: M_2 -RRA, M_2 -RRP, M_2 -RRR, M_2 -U and M_2 -RM. We also include for comparison a preference-based heuristics (M_2 -PW) and a procedure allowing cycles involving up to 3 agents (M_3). More precisely:

- *Priority to the worst-off agent* (M_2 -PW): agents are ordered considering the rank of the object they own from the one with the lowest rank to the one with the highest rank. Agents are then paired in a round-robin fashion like M_2 -RRA does following this ordering. Note that M_2 -PW is more demanding than history-based heuristics as it requires some information about the agent's current rank which can only be collected via a central entity.
- *Uniform up to three agents* (M_3 -U): a deal is selected uniformly at random among all possible deals involving 2 or 3 agents. If the deal is rational for all the agents involved, it is implemented.

Generation of single-peaked preferences

Different methods can be envisioned to generate single-peaked preferences. We consider impartial culture for single-peaked domain (IC-SP) and uniform peak for the single-peaked domain (UP-SP).

Single-peaked preferences under impartial culture (IC-SP) are drawn using the method proposed by Walsh [2015]. Given an axis, single-peaked preferences are built recursively from the end (i.e. the worst object of the agent) to the top object. At each iteration, the next object in the preference order is randomly selected between the two extremes of the axis. The selected object is then removed from the axis and so on until the axis is empty.

In the uniform peak culture (UP-SP), presented by Conitzer [2009], preferences are constructed by first picking uniformly at random a object to be the peak. The second-highest ranked object is chosen with equal probability from the two adjacent alternatives, and so on until a full order is obtained.

As already mentioned by Walsh [2015], the probabilities of the preference orders significantly differ from one method to another. Under IC-SP, each single-peaked preference order has a uniform probability $1/2^{n-1}$ to be selected. On the contrary, under UP-SP, probabilities over preference orders are not uniform. In fact, the peak is uniformly drawn (with probability $1/n$) and single-peaked preferences are then built from this peak. Since there is only one preference order with its peak at one end of the axis, these orders are more likely to be drawn than preference orders with a peak in the middle of the axis, for instance.

Let us consider the case where $n = 5$ and the axis is $o_1 \triangleleft o_2 \triangleleft o_3 \triangleleft o_4 \triangleleft o_5$. Preference orders $o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5$ and $o_5 \succ o_4 \succ o_3 \succ o_2 \succ o_1$ both have a probability $1/5$ to be generated whereas the probability to generate some preference order with the peak o_3 is $1/5$, and there are 6 orders with this peak. Under IC-SP, the probability to generate the preference order $o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5$ is $1/16$.

Closer inspection reveals that the frequency of each rank for a given object is more evenly distributed with UP-SP than with IC-SP. For instance, under IC-SP, objects on the end-sides of the single-peaked axis have a high probability to be ranked among the top objects of an agent but they have a very low probability to be ranked in the second half of the preference linear order. Under UP-SP, these objects have a high probability to be selected as the top object of an agent but they also have a more uniform probability distribution among all other possible ranks. A object on the end-side of the single-peaked axis is thus more likely to obtain a low ranking under UP-SP.

This suggests that profiles under UP-SP are more diverse than under IC-SP. To quantify this more precisely, we have computed two diversity indices suggested by [Hashemi and Endriss, 2014]: the sum of Kendall’s tau and the sum of Spearman distances.⁴ The results confirm that UP-SP generates more diverse profiles on average compared to IC-SP. Take for instance the case of seven objects ($n = 7$). We generated 10000 preference profiles under both UP-SP and IC-SP. With respect to the Kendall’s tau index, profiles under UP-SP achieved an average diversity score of 75% (normalized over the empirically observed maximum) while this value reaches only 48% in the case of IC-SP. With the Spearman distance, the numbers are still in favour of UC-SP: 85% versus 64%.

⁴The Kendall’s tau distance between two preference order \succ_i and $\succ_{i'}$ over \mathcal{O} is the number of pairs $(r, r') \in \mathcal{O}^2$ such that \succ_i and $\succ_{i'}$ do not rank r and r' in the same order. The Spearman distance between two preference order \succ_i and $\succ_{i'}$ over \mathcal{O} is defined as $\sum_{r \in \mathcal{O}} |\text{rank}_{a_i}(r) - \text{rank}_{a_{i'}}(r)|$. In both cases, the diversity index of a profile L is the sum of the distance between every pair of preference orders $(\succ_i, \succ_{i'})$ of L .

Experiments conducted

We conducted two types of experiments:

- (i) A study of the average efficiency and fairness of swap dynamics. The *efficiency* (respectively *fairness*) ratio is defined as the ratio between the average (respectively minimum) rank realized by the procedure and the optimal average (respectively min) rank achievable for the instance (disregarding the individual rationality constraint). The optimal values are obtained using matching techniques [Garfinkel, 1971]. These investigations can be viewed as the empirical counterpart of the worst-case analysis of Section 3.5.

To isolate the influence of the individual rationality (IR) constraints alone, we also include the rules “max *ark* IR” and “max *mrk* IR” which respectively returns the allocation maximizing the average and the minimum ranks under the constraint of being individually rational. We also study the performance of swap dynamics with respect to TTC and Crawler.

- (ii) A study of the number of swaps performed by swap dynamics with respect to the number of cycles performed by TTC and Crawler. Of course this comparison needs to acknowledge that the size of cycles may be very different in both approaches. We thus also investigate the size of the cycles computed by TTC and Crawler.

Parameters of the experiments

For every experiment, we considered several number of objects, varying from $n = 2$ to $n = 60$ and we randomly generated 1000 instances in each case. For a given instance, the initial allocation is selected uniformly at random among all possible allocations. For every instance, we ran the different procedures and report the results for a fixed number of objects, averaged over the 1000 instances.

3.6.2 Analysis of the results

Efficiency and fairness of swap dynamics

Comparison among swap dynamics. Figure 3.3 presents the efficiency and fairness ratio for each swap dynamic for preferences generated both under the IC-SP model (left side) and under the UP-SP model (right side). Results for max *mrk* IR and max *ark* IR are also presented for comparison purposes.

Regarding the minimum rank—Figures 3.3(a) and 3.3(b)— M_2 -PW reaches significantly better allocations than other dynamics. This is expected: this heuristic favours deals between agents holding low ranked objects and thus tends to improve the satisfaction of the poorest agents. On the contrary, round-robin heuristics with their fixed ordering tend to always favour the same agents and thus often leads to lower minimum rank. However, notice that the randomized round-robin version (M_2 -RRR) slightly corrects this.

Interestingly, the performance of M_2 -PW is highly sensitive to the culture considered. This seems to be due to the way preferences are generated. Indeed, the results for max *mrk* IR follow similar patterns, indicating that IC-SP (Figure 3.3(a)) is much more constrained

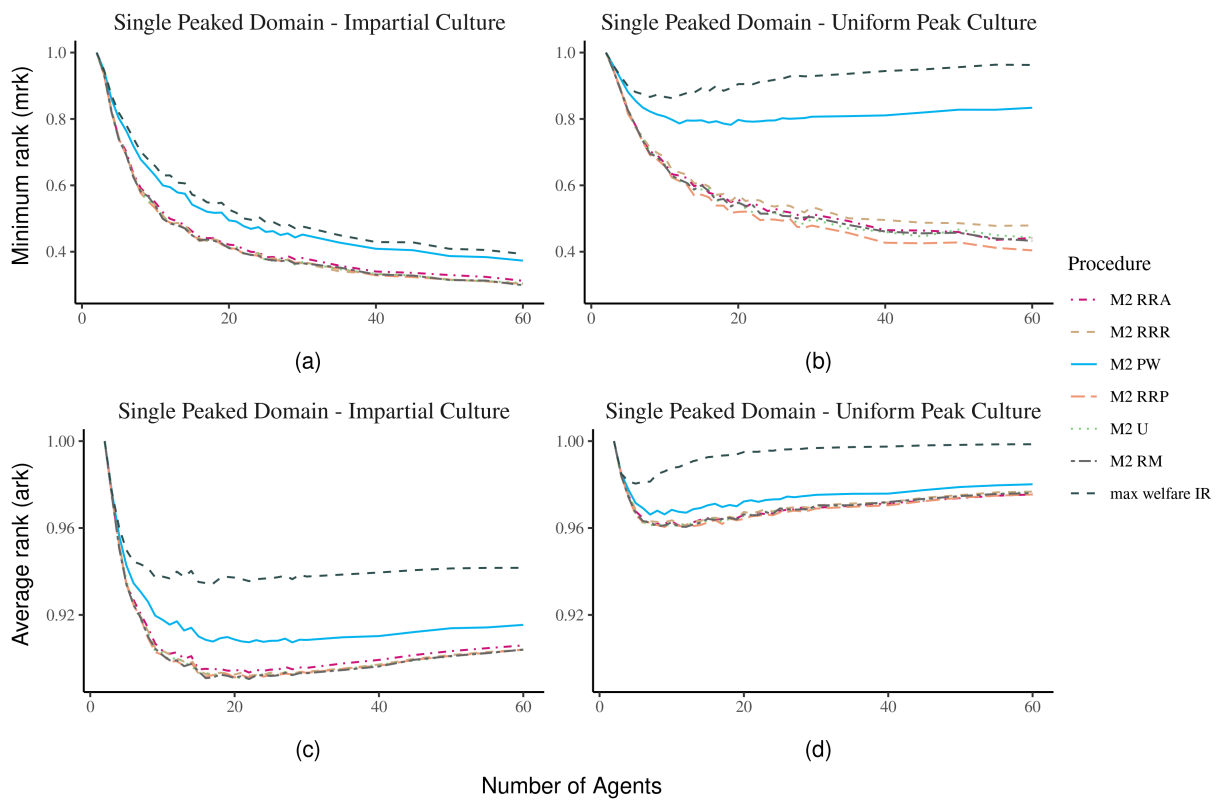


Figure 3.3: Average efficiency and fairness ratios for each heuristic of the swap-deal procedure and for each preference culture. For comparison baseline, the average maximum welfare achievable under individual rationality ($\max mrk$ IR for the upper part and $\max ark$ IR for the lower part) is also presented. The number of agents varies from 2 to 60.

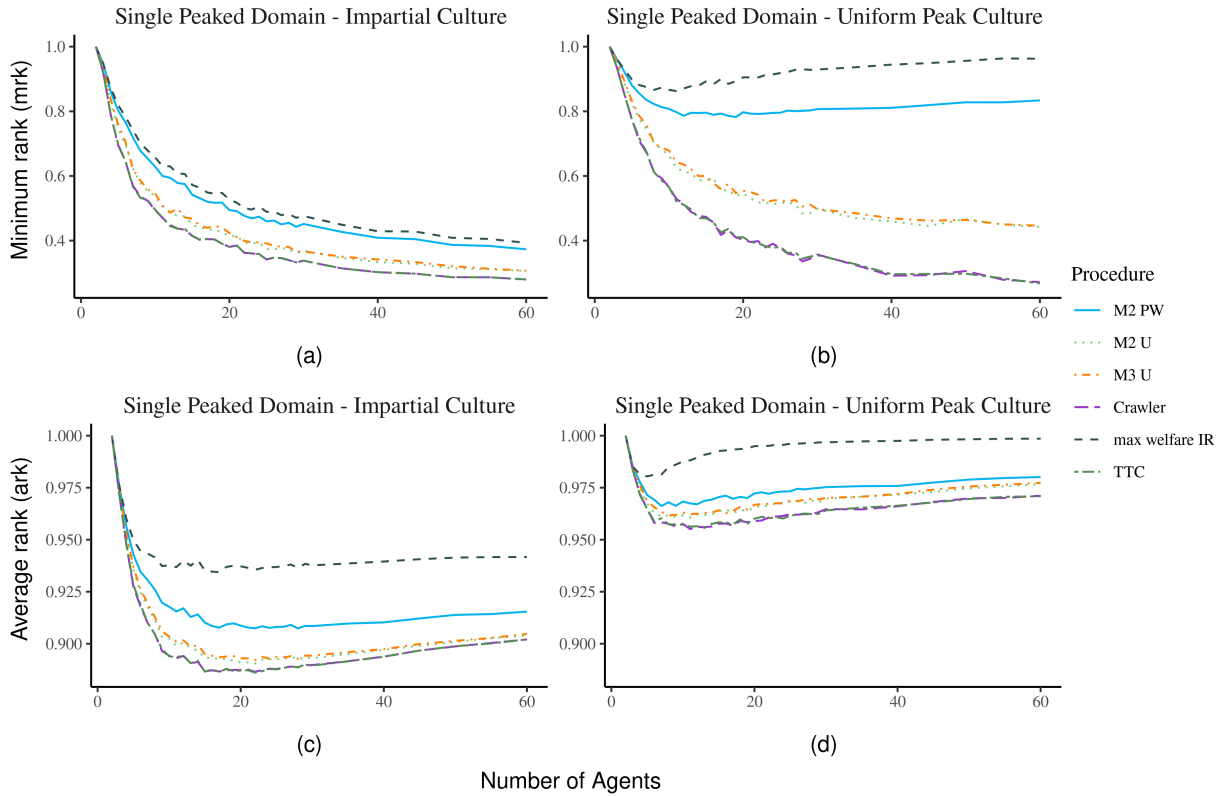


Figure 3.4: Average efficiency and fairness ratios for each procedure and for each preference domain. Curves for “max welfare IR” correspond to max mrk IR for the upper part of the figure and to max ark IR for the lower part. The number of agents varies from 2 to 60.

by the initial allocation than UP-SP (Figure 3.3(b)) is. Still, under UP-SP, M_2 -PW returns close to optimal allocations, while other dynamics remain with a fairness ratio of about 50%. Following this observation, as the allocation returned by M_2 -PW can be taken as a witness of a good allocation reachable by swap dynamics, we see that the empirical minimum rank price of dynamics for our studied family of dynamics is much higher in UP-SP than in IC-SP.

Regarding the average rank of the outcomes—Figures 3.3(c) and 3.3(d)—all the heuristics obtain very good results (above 90% under IC-SP (Figure 3.3(c)) and above 96% under UP-SP (Figure 3.3(d))). It can be observed that all heuristics give very similar values. Recall that our price of anarchy results informed us that in principle, swap dynamics can return allocation with an efficiency ratio as low as 50%. Given these high performances, it is not surprising that very little room is left to observe difference among selection heuristics. Still, it is noteworthy that M_2 -PW provides, again, the best results (while this selection heuristics doesn’t seem designed to optimize this measure of welfare at first sight). Note that our remark about the relative performance of P2-PW under the two cultures does not hold here: this suggests that only a few agents are left with low-ranked objects with our swap dynamics, even under UP-SP.

Comparison with TTC and Crawler. We now compare swap dynamics to other procedures. For the sake of readability, we only keep M_2 -PW (which offered the best performance) and M_2 -U (the most decentralized selection heuristics, requiring hardly any

communication between the agents). Results are presented in Figure 3.4.

Regarding the minimum rank—Figures 3.4(a) and 3.4(b)—a first observation is that M_2 -U performs significantly better than TTC and Crawler algorithm (which give similar results). The poor performances of these procedures, not designed for this purpose, are not surprising: they typically leave some agents with very little opportunities to be involved in a cycle-deal. For instance, by selecting top-trading cycles, TTC can remove all the agents involved once a deal is implemented, thus significantly limiting the range of possible deals for the remaining agents. On the other hand, as discussed, M_2 -PW favours low ranked agents and gives more opportunities to these agents to exchange their initial objects.

Regarding the average rank of the outcomes obtained by the different procedures—Figures 3.4(c) and 3.4(d)—swap dynamics slightly outperform TTC and Crawler under both cultures.

Finally, it can also be observed that M_3 -U shows no significant difference with M_2 -U under the same heuristics (uniform selection of the exchanges), either for the average or minimum rank: slightly increasing the size of the deals leads raises complex coordination issues with no evidence of improvements in terms of welfare.

These experiments promote the relevance of swap dynamics: besides being simple to implement, they also provide very good results both in terms of average and minimum rank, close to the optimal when individual rationality is enforced. For the minimum rank and under UP-SP, we observed that our history-based swap dynamics may incur a significant cost, even though better allocations would be reachable by swaps (this was shown by comparison to a more specifically designed, not history-based, selection heuristic).

Length of swap dynamics

Figure 3.5 represents the mean number of swap-deals performed (solid lines) when varying the size of the instances. For the sake of readability, we only show some selection heuristics (the other ones performed a number of swaps similar to either M_2 -RRA or M_2 -U). Dotted lines represent the highest and the lowest numbers of swaps registered for an instance of a given size (averaged over 1000 randomly generated instances).

It can be observed that the different selection heuristics are close in terms of the average number of swaps. However, the number of swaps performed under UP-SP is significantly higher than under IC-SP. This phenomenon is related to the method used to generate single-peaked preferences. According to the discussion of Section 3.6.1, UP-SP offers a greater diversity of profiles, which leads to more opportunities of exchanges. However, even though we noted in previous experiments that the results in terms of welfare were better in UP-SP than in IC-SP, note that the proportion is of a different order here: there are on average three times more swap-deals performed under UP-SP.

The higher diversity of profiles under UP-SP can also explain the higher variance in the number of swaps. Typically, when agents have completely opposite preference orders, the number of deals can greatly vary depending on which agents are encountered. This is much more likely to happen under UP-SP.

Comparison with TTC and Crawler. The number of deals performed by swap dynamics can be compared with the number of deals induced by TTC or Crawler. To do so we introduce the number of unit exchanges performed by a procedure. It corresponds

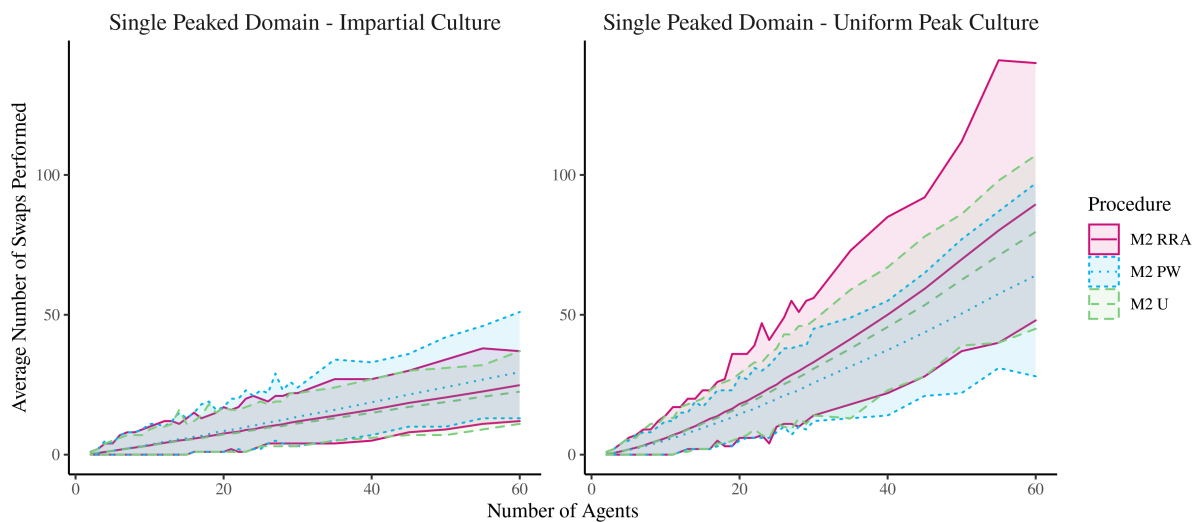


Figure 3.5: Average number of swap-deals performed, the filled area represents the range, from the minimum to the maximum. The number of agents varies from 2 to 60.

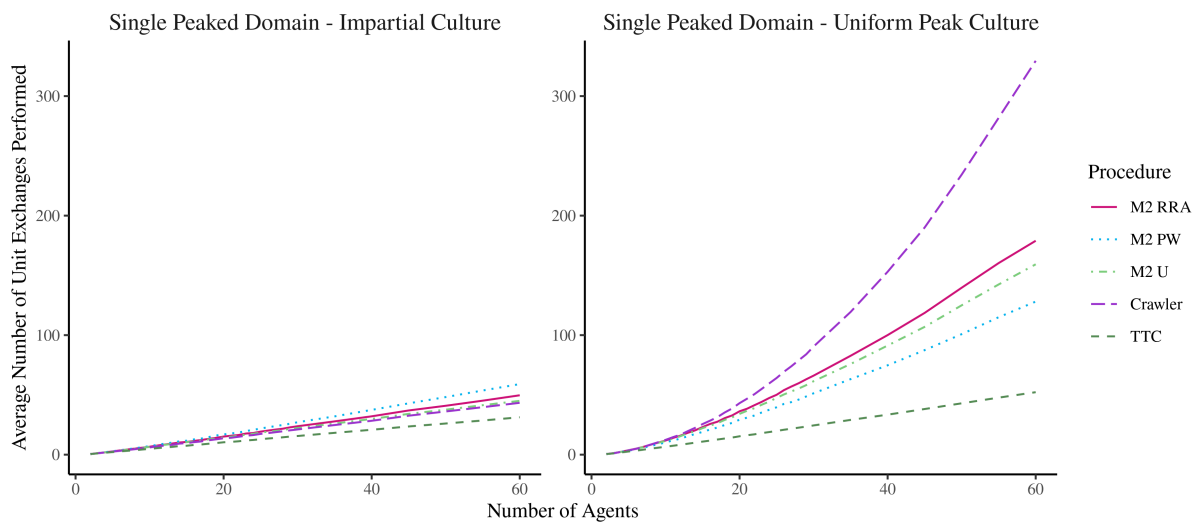


Figure 3.6: Average number of unit exchanges performed by some swap dynamics, TTC and Crawler. A cycle-deal of length k counts for k unit exchanges. The number of agents varies from 2 to 60.

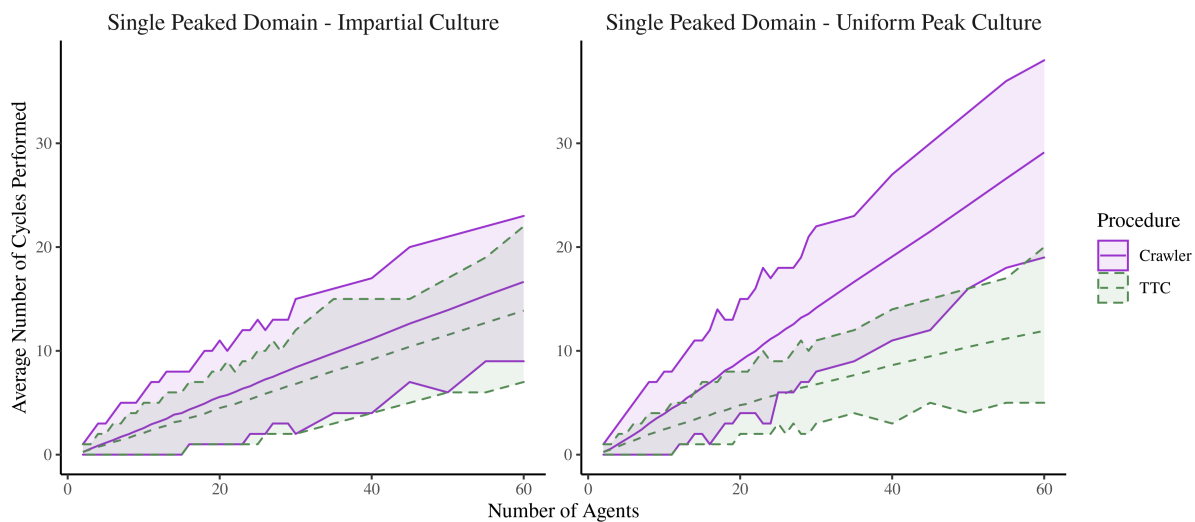


Figure 3.7: Average number of cycles performed by TTC and Crawler, the filled area represents the range, from the minimum to the maximum. The number of agents varies from 2 to 60.

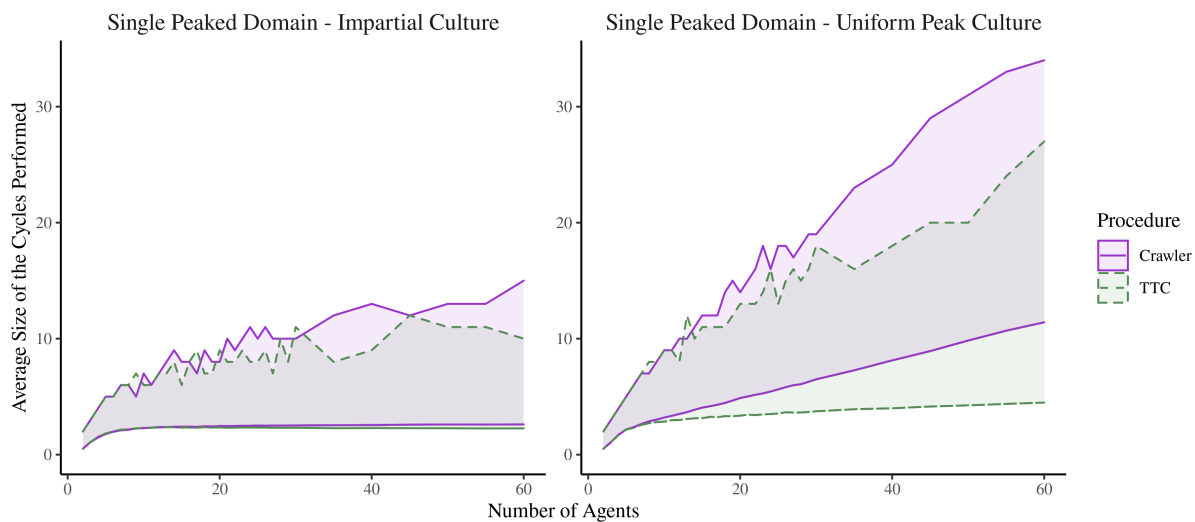


Figure 3.8: Average size of the cycles performed by TTC and Crawler, the filled area represents the maximum. The minimum, not shown, is constantly equal to zero. The number of agents varies from 2 to 60.

	TTC IC-SP	TTC UP-SP	Crawler IC-SP	Crawler UP-SP
β_0	4.64	4.73	4.27	2.86
β_1	0.15	0.36	0.20	0.56
O^2	0.7872	0.9032	0.883	0.9862

Table 3.1: Linear regression of the maximum size of the deals over the number of agents. β_0 and β_1 are the coefficients of the regression: $maxSize = \beta_0 + \beta_1 * n$ where n is the number of agents, O^2 the coefficient of determination. The p-value is omitted as it is not meaningful for simulations (see, e.g., [Lee et al., 2015]).

to the sum of the length of the cycle-deals applied. For a swap dynamics, it is twice the number of swap-deals performed for instance. This measure allows us to compare all the procedures even though they perform cycle-deals of different sizes. As shown in Figure 3.6, the procedure can be clustered in three distinct classes. The swap dynamics can be grouped together, offering a middle way between TTC which performs particularly well and Crawler which performs particularly badly in terms of number of unit exchanges.

The difference between Crawler and TTC can be explained by the fact that the former performs much more deals on average than TTC. As shown in Figure 3.7, this can be observed for any number of agents and the gap increases as the number of agents increases. The size of the deals is also larger when implementing Crawler (Figure 3.8). Both of these facts naturally lead to a higher number of unit exchanges.

The larger number of deals and the larger sizes of deals performed by Crawler are related to the fact that agents are ordered with respect to the object they initially hold and with respect to the order of these objects on the single-peaked axis. Based on this order, Crawler only considers deals $\mu = \langle a_1, \dots, a_k \rangle$ such that a_i and a_{i+1} (with $i \in \{1, 2, \dots, k-1\}$) are owners of adjacent objects on the axis. Hence, TTC allows for considering a larger range of cycle-deals than Crawler.

Putting aside the comparison between Crawler, and TTC, the size of the deals to implement can be extremely large for both procedures as depicted in Figure 3.8. A cycle-deal may involve more than half (respectively 35%) of the agents under UP-SP and a fifth (respectively 15%) of the agents under IC-SP for Crawler (respectively TTC). The linear regressions explaining these values are presented in Table 3.1. Overall, even though TTC performs less unit exchanges than our swap dynamics, it can still need to perform extremely large cycle-deals.

These experiments show that the number of swap-deals can be significant, especially under UP-SP where long sequences of slightly improving steps are more likely to occur. On the other hand, we show empirically that TTC and Crawler are indeed prone to require the implementation of cycle-deals of large size. This illustrates the trade-off which occurs between the coordination requirements and the length of the procedure.

3.7 Conclusion

This chapter studied the property of swap dynamics for the allocation of indivisible objects in the restricted setting of single-peaked housing markets. The basic principle of the procedure is to let agents perform pairwise improving exchanges. We showed in particu-

lar that the single-peaked domain happens to be maximal for guaranteeing convergence to Pareto-optimal outcomes with such dynamics. We also showed that the outcomes of TTC and of Crawler are reachable by swap-deal sequences. To refine our analysis, we have studied two further notions: the average rank and the minimum rank of the objects obtained by the agents. None of the procedures discussed in this chapter are specifically designed for optimizing these ranks, even though these notions capture very natural criteria of efficiency (for the average rank) and fairness (for the minimum rank). It thus seems important to study how these allocation procedures behave on that respect. To complement worst-case theoretical bounds on the loss of social welfare induced by swap dynamics, we ran experiments which demonstrated that they actually provide good results in practice.

Our focus on swap-deals is motivated by their minimal coordination requirements. In the context of this chapter, our experimental results suggest that there is little gain to expect when allowing deals involving three agents. A complementary relaxation is to permit deals possibly not improving for some of the agents involved (but at least a majority). This may unveil interesting connections with other notions: Kondratev and Nesterov [2019] recently established for instance that an allocation is popular in a housing market if and only if no (majority) improving exchanges between three agents exist.

To go further with the experiments, it would be interesting to use real data. Our attempt to use data from Preflib [Mattei and Walsh, 2013] was not successful as there is no dataset that is single-peaked when there are more than 5 agents. Getting such preferences would be an interesting way to confirm our results. Regarding the model itself, Bade [2019] extended Crawler to single-peaked domains with indifferences. Whether our results with swap dynamics could be similarly generalized is an avenue for future research. Overall, this chapter raises the exciting issue of giving a characterization of rules that are efficient, individually rational and strategy-proof for the single-peaked domain. Such a characterization would also provide more insights about the type of selection heuristic for which the swap dynamics are strategy-proof. It would also be interesting to tackle the characterization of the swap dynamics efficiency, that is, identifying the domain on which the swap dynamics are Pareto-optimal. The maximality result is a significant step in this sense, it would be nice to complete the picture.

Chapter 4

Balancing envy among agents

Summary of the chapter

In this chapter, we mainly focus on a new relaxation of envy-freeness motivated by the fact that if an envy-free allocation does not exist then we can focus on some sort of aggregation of the envies of the agents and could focus on how envy is spread among agents. Indeed, as it is common to consider in multi-objective problems, we would rather like envy to be fairly spread among agents. But as you must know by now, fairness means a lot of things so here by fairness we mean fair Ordered Weighted Average (OWA).

This chapter is based in full on a paper published at the Algorithmic Decision Theory conference [Shams et al., 2021].

4.1 Introduction

As we have already seen, envy-freeness is an attractive criterion: the fact that each agent is better off with her own share than with any other share is a guarantee of social stability. Besides, it does not rely on any interpersonal comparability. Unfortunately, envy-freeness is also a demanding notion as soon as we require all goods to be allocated, and it is well-known that in many situations, no such allocation exists (consider for instance the situation where the number of items to allocate is strictly less than the number of agents at stake). Hence several relaxations of envy-freeness have been studied in recent years. Two orthogonal approaches have been considered as discussed in Section 1.3.2.

In this chapter, we elaborate on this idea of minimizing the degree of envy but seek to offer a broader perspective. More precisely, we explore the possibility of finding allocations where envy is “fairly balanced” amongst agents. For that purpose, we start from the notion of individual degree of envy and use a *fair* Ordered Weighted Average operator (by “fair”, we mean an OWA where weights are non-increasing.) to aggregate these individual envies into a collective one, that we try to minimize. This family of operators contains both the egalitarian and utilitarian operators mentioned previously. But doing so also sometimes allows us to draw results valid for the whole family of fair operators. Along our way, we shall for instance see that no algorithm fairly minimizing envy can be guaranteed to return an envy-free allocation up to any good, even though such allocation does exist. More generally, we provide several insights regarding the behaviour of such fair minimizing operators, comparing their outcomes with alternative approaches, either analytically or experimentally. Technically, this is made possible through to the use of linearization techniques which alleviate the burden of computing these outcomes.

The remainder of this chapter is as follows. We first formally introduce our fairness minOWA envy criterion (Section 4.2) and we show that OWA minimization problems can be formulated as linear programs. We then investigate the link between minimizing the OWA of the envy vector and other fairness notions (Section 4.3). We thus study fairness guarantees of the minOWA solutions. Finally, we present some experimental results investigating the fairness of min OWA solutions (Section 4.4).

4.2 MinOWA Envy

Our approach elaborates on minimizing the degree of envy of the agents while balancing the envy among the agents as suggested by Lipton et al. [2004]. The general idea would be to look for allocations that minimize this vector of envy in some sense: the lower this vector is, the less envious the agents are. This corresponds to a multiobjective optimization problem where each component of the envy vector is a different objective to minimize.

4.2.1 Fair OWA

There are different ways to tackle this minimization problem, each approach conveying a different definition of minimization. Our approach, guided by the egalitarian notion of fairness [Rawls, 1971], is to ensure that, while being as low as possible, the envy is

also distributed as equally as possible amongst agents. To this end, we use a prominent aggregation operator that can convey fairness requirements: order weighted averages.

Ordered Weighted Averages (OWA) have been introduced by Yager [1988] with the idea to build a family of aggregators that can weight the importance of objectives (or agents) according to their relative utilities, instead of their identities. In this way, we can explicitly choose to favour the poorest (or richest) agents, or to concentrate the importance of the criterion on the middle-class agents. Formally, the OWA operator is defined as follows:

Definition 4.1. Let $\vec{\alpha} = \langle \alpha_1, \dots, \alpha_n \rangle$ be a vector of weights. In the context of minimization, the ordered weighted average parameterized by $\vec{\alpha}$ is the function $owa^{\vec{\alpha}} : x \mapsto \sum_{i=1}^n \alpha_i \times x_i^\downarrow$, where x^\downarrow denotes a permutation of x such that $x_1^\downarrow \geq x_2^\downarrow \geq \dots \geq x_n^\downarrow$.

Amongst all OWA, only those giving more weight to the unhappiest agents can be considered fair in the egalitarian sense. This property can be formalized as follows. Let x be a vector such that $x_j \geq x_i$ (a_i is better off than a_j) and let ε be such that $0 \leq \varepsilon \leq 2(x_j - x_i)$. Then, for any non-increasing vector $\vec{\alpha}$: $owa^{\vec{\alpha}}(x) \leq owa^{\vec{\alpha}}(\langle x_1, \dots, x_i + \varepsilon, \dots, x_j - \varepsilon, \dots, x_n \rangle)$.

In other words, such an OWA favours any transfer of wealth from a happier agent to an unhappier agent. Such a transfer is called a *Pigou-Dalton* transfer, and the OWA with non-increasing weight vectors $\vec{\alpha}$ are called *fair OWA*. Moreover, we have considered without loss of generality in this chapter that the weight vector sums to 1 so we will make no difference between weights $\langle 1, 1, 1 \rangle$ and $\langle \frac{1}{3}, \frac{1}{3}, \frac{1}{3} \rangle$. Note that fair OWA is also referred to as *Generalized Gini Index* [Weymark, 1981] in the literature. In matching problems [Lesca et al., 2019] and multiagent allocation problems [Heinen et al., 2015], fair OWA has been applied to the utility vector so as to maximize a global utility function while reducing inequalities. However, we can note that maximizing the OWA of the utility vector does not necessarily return an EF allocation even when such an allocation exists:

Example 4.1. Consider this add-MARA instance with 3 agents and 4 objects:

	o_1	o_2	o_3	o_4
a_1	1	2*	3	4*
a_2	2	2	5*	1
a_3	4*	0	4	2

The squared allocation is the allocation that maximizes the value of the OWA of the utility vector with weight $\langle 1, 0, 0 \rangle$. We can easily notice that this allocation is not envy free as a_1 envies a_2 . Moreover, the star allocation is obviously an EF one. Note that in the context of maximization, a fair OWA is also defined with non-increasing weights but by sorting the components by increasing value.

Since our motivation is to return an EF allocation when there is one and otherwise minimize the envy while equally distributing it between the agents, we propose to minimize the fair OWAs of the envy vector.

Definition 4.2. Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and $\vec{\alpha}$ be a non-increasing vector. An allocation $\hat{\pi}$ is an $\vec{\alpha}$ -minOWA Envy allocation if:

$$\hat{\pi} \in \operatorname{argmin}_{\pi \in \mathcal{A}(I)} (owa^{\vec{\alpha}}(\vec{e}(\pi))).$$

It is important to note that a major advantage of this solution is that it always exists as it is the result of an optimization process. Moreover, this optimization problem can be modeled as an Integer Linear Program, which will give a way to compute optimal allocations. Keep also in mind that there can be several allocations with the same OWA envy value.

Let us now see some helpful properties of fair OWA. Note that we will consider here that we are in a minimization context.

Definition 4.3. By denoting v_k^\downarrow the k^{th} biggest component of a given vector v , the Lorenz vector \vec{L} of v is defined as $\vec{L}(v) = \langle v_1^\downarrow, v_1^\downarrow + v_2^\downarrow, \dots, \sum_{i=1}^n v_i^\downarrow \rangle$.

Definition 4.4. Let x and y be two vectors of the same size and x_i (respectively y_i) be the i^{th} component of x (respectively y). We say that x Pareto dominates y iff for every component $x_i \leq y_i$ and there is one component $x_j < y_j$ and x strongly Pareto dominates y iff for every component $x_i < y_i$.

Definition 4.5. We say x (strongly) Lorenz dominates y iff $\vec{L}(x)$ (strongly) Pareto dominates $\vec{L}(y)$.

Theorem 4.1. Perny and Spanjaard 2003

If x Lorenz dominates y then for any non-increasing weight $\vec{\alpha} : owa^{\vec{\alpha}}(x) \leq owa^{\vec{\alpha}}(y)$. Similarly if x strongly Lorenz dominates y then for any non-increasing weight $\vec{\alpha} : owa^{\vec{\alpha}}(x) < owa^{\vec{\alpha}}(y)$.

This helpful property is shown in [Perny and Spanjaard, 2003]. As (strong) Pareto dominance implies (strong) Lorenz dominance, the same theorem holds with (strong) Pareto dominance.

4.2.2 Linearization of OWA minimization

By using a linearization introduced by Ogryczak and Śliwiński [2003] we can model our problem of minimizing the OWA of the envy vector as a linear program. This linearization smartly uses the definition of OWA with its Lorenz components. Moreover we consider decreasing OWA weights (fair OWA) so $\alpha_1 \geq \alpha_2 \dots \geq \alpha_n$ and we denote by $\vec{\alpha} = \langle \alpha_1 - \alpha_2, \alpha_2 - \alpha_3, \dots, \alpha_n \rangle$. The OWA of a vector can be written with its Lorenz components:

$$\begin{aligned} owa(\vec{e})(\pi) &= \sum_{k=1}^n \alpha_k e_k^\downarrow \\ &= \sum_{k=1}^n \alpha'_k L_k^\downarrow(\vec{e}) \end{aligned}$$

It is known that $L_k^\downarrow(\vec{e})$ can be defined as the result of a LP [Ogryczak and Tamir, 2003]:

$$\begin{aligned} L_k^\downarrow(\vec{e}) &= \max \sum_{i=1}^n c_i^k e_i \\ \text{s.t. } \left\{ \begin{array}{l} \sum_{i=1}^m c_i^k = k \quad *r_k \\ c_i^k \in [0, 1] \quad \forall i \in \llbracket 1, n \rrbracket \quad *b_i^k \end{array} \right. \end{aligned}$$

The difficulty here is that we cannot directly inject this LP in the previous definition of $owa(\vec{e})(\pi)$: this LP is defined as a maximization problem, whereas our goal is to minimize $owa(\vec{e})(\pi)$. To circumvent this problem, we will use the dual version of this LP (with the dual variables shown in gray in the LP above):

$$\begin{aligned} L_k^\downarrow(\vec{e}) = \min \quad & kr_k - \sum_{i=1}^n b_i^k \\ \text{s.t.} \quad & \begin{cases} r_k + b_i^k \geq e_i & \forall i \in \llbracket 1, n \rrbracket \\ b_i^k \geq 0 & \forall i \in \llbracket 1, n \rrbracket \end{cases} \end{aligned}$$

Now that we have a way express the OWA as a linear program, we need to encode the part related to the fair division problem and the computation of the envy vector. For that, we introduce a set of $n \times m$ Boolean variables z_i^j : z_i^j is 1 iff o_j is allocated to a_i . The preemption constraint and completeness can be easily expressed as the constraint set $\sum_{i=1}^n z_i^j = 1 \forall j \in \llbracket 1, m \rrbracket$ (just stating that each object should be allocated to exactly one agent). The pairwise envy between a_i and a_h is simply computed as $\sum_{j=1}^m w(a_i, o_j)(z_h^j - z_i^j)$. The envy e_i of an agent a_i is just defined to be greater than or equal to pairwise envies from a_i to a_h for all h . Since we are in a minimization problem, e_i will exactly be equal to the maximal pairwise envy in the returned solution.

Putting things together, we obtain the following MIP:

$$\begin{aligned} \min \quad & owa(\vec{e})(\pi) = \min \sum_{k=1}^n \alpha'_k (kr_k + \sum_{i=1}^n b_i^k) \\ \text{s.t.} \quad & \begin{cases} r_k + b_i^k \geq e_i & \forall i, k \in \llbracket 1, n \rrbracket \\ e_i \geq \sum_{j=1}^m u(i, j)(z_h^j - z_i^j) & \forall i, h \in \llbracket 1, n \rrbracket \\ \sum_{i=1}^n z_i^j = 1 & \forall j \in \llbracket 1, m \rrbracket \\ z_i^j \in \{0, 1\} \quad \forall j \in \llbracket 1, m \rrbracket & \forall i \in \llbracket 1, n \rrbracket \\ b_i^k \geq 0, \quad e_i \geq 0 & \forall i, k \in \llbracket 1, n \rrbracket \end{cases} \end{aligned}$$

4.3 Link with other fairness measures

We focus here on the possible links between the min OWA allocation and other fairness measures. We recall that if an envy-free allocation exists, it will be returned by the min OWA optimization. For any instance I , we denote by $\text{PROP}(I)$ the set of allocations satisfying $\text{PROP} \in \{EF1, EFX, EFX_0, MMS\}$. We also denote by $\vec{\alpha}$ -min OWA(I) the set of all min OWA optimal allocation for the specific weight vector $\vec{\alpha}$, and by \forall -min OWA(I) the set of $\vec{\alpha}$ -min OWA, for *all* (fair) weight vectors $\vec{\alpha}$.

Moreover, in the following, we recall that we denote by \mathcal{I} the set of all add-MARA instances. Furthermore, different domain restrictions will be of interest: we denote by \mathcal{I}^p the set of add-MARA instances involving only two agents (pairwise instances), and by \mathcal{I}^b the set of add-MARA instances where agents have binary utilities.

4.3.1 Warm-up: $n = 2$

In the special case where the allocation problem involves only two agents, we highlight strong connections between min OWA allocations and other fairness measures (MMS, EF1 and EFX).

Proposition 4.1. $\forall I \in \mathcal{I}^p : \forall\text{-min OWA}(I) \subseteq \text{MMS}(I) \subseteq \text{EFX}(I)$

Proof. For add-MARA instances where an envy-free allocation exists, our proof is straightforward as min OWA returns the EF allocation. It is thus also MMS, EF1 and EFX.

We now focus on add-MARA instances for which there is no EF allocation. In the presence of only 2 agents any min OWA allocation π is such that only one of the two agents is envious. Indeed, if no agent is envious then it means the add-MARA instance has an envy-free allocation (which is a contradiction). Similarly, if both agents are envious it means there is an envy-free allocation (which is again a contradiction) as the agents would just have to exchange their bundles to obtain that allocation. Consequently, the sorted envy vector will be of the form $(e, 0)$. Suppose for the sake of contradiction that such an allocation is not MMS. The agent that is envy-free (let us say w.l.o.g it is a_2) obviously has her max-min share. So, under the assumption that the allocation is not MMS, a_1 does not have her max-min share. It means that there is an allocation π' such that $\min(u_1(\pi'_1), u_1(\pi'_2)) > u_1(\pi_1)$ and a_2 is still not envious (if a_2 is envious in π' , just swap her share with a_1 's). Obviously, a_1 's pairwise envy for a_2 has decreased in π' compared to that of π , and a_2 's envy is still 0. This contradicts the fact that π is the optimal min-OWA envy allocation. Finally, it is known [Caragiannis et al., 2016] in the two-agents setting that MMS implies EFX, which completes the proof. \square

However, even though an MMS allocation is EFX, this does not hold for EFX_0 even for 2 agents as we can see in Example 4.2.

Example 4.2. Consider this add-MARA instance with 2 agents and 3 objects:

	o_1	o_2	o_3
a_1	1	0	2
a_2	0	1	2

It is easy to see that the squared allocation is MMS as the max-min share of each agent is 1. Moreover, we can see that this allocation is EFX (a_1 can forget o_3) whereas it is not EFX_0 (because a_1 has to forget o_2 which does not make here becoming envy-free).

However, we show that we can very easily build an EFX_0 allocation from an arbitrary min OWA envy one.

Proposition 4.2. For any instance $I \in \mathcal{I}^p$ and for any weight vector $\vec{\alpha}$: $\vec{\alpha}\text{-minOWA}(I) \cap \text{EFX}_0(I) \neq \emptyset$. Furthermore, it can be obtained from an arbitrary $\vec{\alpha}\text{-min-OWA}$ envy optimal allocation in linear time.

Proof. Let us call π an arbitrary min OWA allocation. If π is envy-free then it is obviously EFX_0 and the proof concludes. Note that envy-freeness is checked in $O(1)$ as we just have to check the values of both variables e_1 and e_2 . Otherwise, it means that one and exactly one agent is envious, by using a same argument as in the proof of Proposition 4.1. W.l.o.g.

we consider a_1 is the envious agent. We start from π and transfer to a_1 all the objects that she values with utility zero. The resulting allocation is called π' . We show that π' is EFX_0 . a_1 still envies a_2 in π' but is EFX by Proposition 4.1. By transferring all zero-valued objects to her share, she becomes EFX_0 in π' . Now consider a_2 . If a_2 envies a_1 in π' then by swapping their bundles, we can obtain an envy-free allocation. This contradicts the fact that π is min-OWA envy optimal. Hence, a_2 still does not envy a_1 in π' , and thus is also EFX_0 obviously. Since in π' a_2 is still envy-free and the pairwise envy from a_1 to a_2 has not changed, π' is still min-OWA envy optimal. The complexity is linear in the number of objects since we have to implement the transfer of zero-valued objects to a_1 's bundle. \square

On Example 4.2, this means that a_1 should receive o_2 . This adjustment is inefficient: by construction, it returns an allocation which is Pareto-dominated by the original min OWA envy optimal allocation. Intuitively, it can be seen as the price to pay to get EFX_0 : by assigning those items that the agent does not value to her, the mechanism offers the strongest possible fairness guarantees.

4.3.2 General case: $n \geq 3$

We now turn to more general settings involving at least 3 agents. Since an EF1 allocation is guaranteed to exist, we more specifically focus on the relation between min OWA and EF1. Unfortunately, we notice that in the general case these two sets can be disjoint i.e., there are instances for which no allocation is both EF1 and min-OWA, for any weight vector:

Proposition 4.3. $\exists I \in \mathcal{I} : \text{EF1}(I) \cap \forall\text{-min OWA}(I) = \emptyset$

Proof. Let us consider the add-MARA instance with 4 agents and 5 objects:

	o_1	o_2	o_3	o_4	o_5
a_1	20	2	2	2	4
a_2	20	2	2	2	4
a_3	13	1	1	1	14
a_4	0	0	0	0	30

In order to prove the proposition we will show that the squared allocation is the only min OWA envy allocation (for any given weight vector) and that it is (obviously) not EF1. First note that as a_1 and a_2 have similar preferences the allocation derived from the squared allocation where we swap the bundles of these agents will be the same in terms of Lorenz envy vector. The squared allocation has a vector of envy $\vec{e} = \langle 0, 14, 14, 0 \rangle$ and $L(\vec{e}) = \langle 14, 28, 28, 28 \rangle$. First consider the allocations in which a_4 does not possess o_5 . We have $\vec{e}_1 = \langle e_1, e_2, e_3, 30 \rangle$ and $L(\vec{e}_1) = \langle 30, L_2, L_3, L_4 \rangle$ with L_2, L_3, L_4 being greater than or equal to 30. \vec{e}_1 is thus strongly Lorenz dominated by \vec{e} . Let us now consider the other possible allocations (in which a_4 possesses o_5): if a_3 has o_1 instead of a_1 then $\vec{e}_2 = \langle 20, 14, 1, 0 \rangle$ and $L(\vec{e}_2) = \langle 20, 34, 35, 35 \rangle$. \vec{e}_2 is thus strongly Lorenz dominated by \vec{e} . Finally, we focus on allocations in which a_3 has one to three items from the set of objects $\{o_1, o_2, o_3\}$. If a_3 has one of these items we have $\vec{e}_3 = \langle 0, 16, 13, 0 \rangle$ and $L(\vec{e}_3) = \langle 16, 29, 29, 29 \rangle$. If a_3 has two of these items we have $\vec{e}_4 = \langle 0, 18, 12, 0 \rangle$ and

$L(\vec{e}_4) = \langle 18, 30, 30, 30 \rangle$. Finally if a_3 has all these items we have $\vec{e}_5 = \langle 0, 20, 11, 0 \rangle$ and $L(\vec{e}_5) = \langle 20, 31, 31, 31 \rangle$. All \vec{e}_3 , \vec{e}_4 and \vec{e}_5 are strongly Lorenz dominated by \vec{e} . As we know that minimizing fair OWA of a vector is consistent with the Lorenz dominance (see Theorem 4.1), it means that if a solution strongly Lorenz dominates another, then its fair OWA value will be strictly lower (in a minimization problem such as ours) for any non-creasing weight. We can then conclude that the squared allocation is indeed the only min OWA envy one and it is not EF1. \square

However, a significant number of experiments actually suggest that for almost any instance, some EF1 allocation is also min-OWA, either for the weight vector $\langle 1, 0, \dots, 0 \rangle$, or for the weight vector $\langle 1, 1, \dots, 1 \rangle$. Moreover, we have a positive result in the restricted domain where agents have binary utilities.

Proposition 4.4. $\forall I \in \mathcal{I}^b : \text{EFX}_0(I) \cap \forall\text{-min OWA}(I) \neq \emptyset$

Proof. First note that if the instance is EF then the min OWA envy allocation will be EF and thus EF1 and the proof concludes. Hence we will consider instances that are not EF. As we consider binary utilities, we know thanks to [Chaudhury et al., 2020] that an EFX_0 allocation always exists. We can easily notice that any such allocation is such that the envy of each agent is at most 1. Hence, as with the weight vector $\langle 1, 0, \dots, 0 \rangle$ the OWA envy value of an EFX_0 allocation is 1 (as we supposed no EF allocation exists), it is the minimum OWA envy value possible. It can thus be returned by minimizing the OWA envy value. \square

4.4 Experimental results

We drew some experiments to compare the performances of the allocations obtained by min OWA envy with the Maximization of Nash Welfare. More precisely we implemented the linearization described in [Caragiannis et al., 2016] that returns an allocation approximating MNW but closely enough to keep interesting properties such as EF1 and Pareto Optimality. As we have seen through this chapter the range of possibilities offered by the fact that OWA is parameterized is interesting. We will see how three different weights $\vec{\alpha}_1 = \langle 1, 0, \dots, 0 \rangle$, $\vec{\alpha}_2 = \langle \frac{1}{2}, \frac{1}{4}, \dots, \frac{1}{2^n} \rangle$ and $\vec{\alpha}_3 = \langle 1, 1, \dots, 1 \rangle$ compare to each other. $\vec{\alpha}_1$ and $\vec{\alpha}_3$ correspond to respectively minimize the max envy and the sum of the envies. $\vec{\alpha}_2$ is somewhere in the middle of those two extrema with a strictly decreasing weight vector.

All the tests presented in this section have been run on an Intel(R) Core(TM) i7-2600K CPU with 16GB of RAM and using the Gurobi solver to solve Mixed Integer Programs. We have tested our methods on two types of instances: Spliddit instances [Goldman and Procaccia, 2015] and synthetic instances under uniformly distributed commensurable preferences (that is, for each agent a_i and object o_j , utilities are drawn i.i.d. following the uniform distribution on some interval $[x, y]$ and such that the utilities of each agent sums to $5m$).

We evaluate the performances of the OWA envy minimization outcome for both types of instances through the following criterion: EF, EFX_0 , EFX, EF1 and Pareto dominance. Tables 4.1 and 4.2 present the percentage of min OWA envy outcomes that satisfy each criterion. We also study how the vector of weights of the OWA influences the characteristics of the outcomes. The computation time (in seconds) of each approach is

Table 4.1: Performances for minimizing the OWA envy (with weights $\vec{\alpha}_1, \vec{\alpha}_2, \vec{\alpha}_3$) or maximizing the Nash Welfare on Spliddit instances

	$\vec{\alpha}_1$	$\vec{\alpha}_2$	$\vec{\alpha}_3$	MNW
%EF	65.4	65.4	65.4	57.2
%EFX ₀	90.0	93.0	92.7	90.9
%EFX	98.5	99.4	99.0	94.9
%EF1	99.4	99.8	99.3	100
%Pareto	77.1	78.7	79.2	100
%EF+PO	45.7	45.6	46.0	57.2
time(s)	$3.5 \cdot 10^{-3}$	$5.7 \cdot 10^{-7}$	$6.9 \cdot 10^{-7}$	$1.1 \cdot 10^{-6}$

also mentioned. We recall the strong connections between the 4 first fairness notions as $EF \implies EFX_0 \implies EFX \implies EF1$. As it can be checked in Tables 4.1 and 4.2, the percentage of EF allocations should always be lower than or equal to the number of EFX_0 ones which should be lower than or equal to the number of EFX allocations and so on.

The methodology to check whether an allocation is Pareto-Optimal or not thanks to a particular MILP is described in Appendix A.

4.4.1 Spliddit instances

Our first set of experiments has been performed on real-world data from the fair division website Spliddit [Goldman and Procaccia, 2015]. There is a total of 3535 instances from 2 agents to 15 agents and up to 93 items. Note that 1849 of these instances involve 3 agents and 6 objects. By running the MIPs minimizing the OWA envy with the three different weights' vectors described above with a timeout of 1 minute (after this duration the best current solution, if it exists, is returned) we were able to solve all the instances to optimal. The results of these experiments are presented in Table 4.1. The first three columns respectively correspond to the results of minimization of the OWA envy with respectively $\vec{\alpha}_1, \vec{\alpha}_2$ and $\vec{\alpha}_3$, while the fourth column presents the results of the optimization of MNW.

Minimizing the OWA envy provably returns an EF allocation if there exists one. Hence, among the Spliddit instances 65.4 % are envy-free. Note that only 57.2 % of the allocations returned by MNW are EF which means that for around 8.2 % of the Spliddit instances, an envy-free allocation exists but MNW failed to return it. Moreover, without any surprise as Pareto optimality (PO) of the MNW allocations is guaranteed, minimizing OWA envy returns fewer PO allocations than MNW. However, around slightly less than 80% of the min OWA envy allocations are Pareto optimal. It is guaranteed that MNW returns an EF1 solution. However, we can observe, for every weight, that more than 99% of the allocations returned by min OWA envy are EF1. This balances the negative result in Proposition 4.3. Moreover, it can be very interestingly observed that the percentage of EFX_0 is greater for $\vec{\alpha}_2$ and $\vec{\alpha}_3$ than for MNW. The same holds for the percentage EFX but for the 3 weights' vectors and by a more noticeable margin of around 5%. However, MNW performs slightly better than min OWA when we consider EF alongside with Pareto. Finally, we can see that all the optimization programs run very quickly in average with a slightly bigger time for $\vec{\alpha}_1$.

4.4.2 Synthetic instances

For each couple $(|\mathcal{N}|, |\mathcal{O}|)$ from $(3, 4)$ to $(10, 12)$, we generated 100 synthetic add-MARA instances with uniformly distributed preferences. We then ran the four optimization methods described above on the generated instances. We considered such couples of values in order to produce settings where few EF allocations exist as suggested in [Dickerson et al., 2014]. Although it is interesting to consider EF instances to compare with MNW, minimizing OWA envy is even more relevant when no EF allocation exists. Due to lack of space, Table 4.2 presents the results for only 4 couples (n, m) but similar trends can be observed for the other couples of values. As witnessed for the Spliddit instances, MNW often fails to return an EF allocation even when there exists one. As shown in Table 4.2, the number of EF allocations missed by MNW can be quite important as shown by the gap between the percentage of EF allocations returned by min OWA envy and the percentage for MNW. This is exemplified in Table 4.2 for 2 agents and 5 agents where the gap is respectively of 16% and 31%. Even more significantly, it turns out that min OWA outperforms MNW when we consider EF together with Pareto. Once again and in an even stronger way than for the Spliddit instances, these results heavily balance the result of Proposition 4.3: in practice the allocations returned by the min OWA envy were always EF1. Concerning EFX_0 and EFX we also obtained very positive results. Indeed, min OWA envy returns around 10% more EFX_0 and EFX instances than MNW. Note that we confirm Proposition 4.1 as we have 100% of EFX allocations when $n = 2$. Note that we did not adjust the allocation returned by the min OWA optimization to break ties as discussed in the proof of Proposition 4.2. Thus, we get 97% of EFX_0 but this percentage could be even higher. However, these positive results about EF , EFX_0 and EFX come with a price on efficiency as we can see that PO is not guaranteed and the percentage gets lower as the number of agents increases but is still above 60% for $\vec{\alpha}_2$ and $\vec{\alpha}_3$. This highlights the inherent compromise and tension between efficiency and fairness. Besides, as it was the case for the Spliddit instances we can see that the computation is overall quite fast. We can notice that the MNW computation never surpasses 0.02 seconds whereas for 10 agents, min OWA envy optimization is slightly faster than a second for $\vec{\alpha}_1$ and $\vec{\alpha}_3$ and around 2 seconds for $\vec{\alpha}_2$. Finally, we can see that the three different weights considered here lead to quite similar performances. We can globally notice more encouraging results for $\vec{\alpha}_3$ except for EFX . However, keep in mind that the advantage of using a parameterized function is its rich expressiveness so we could see our method as a combination of the results of the 3 weights.

4.5 Conclusion

In this chapter, we introduced a new fairness concept following the idea of minimizing envy. More particularly, we used an OWA to express fairness in the distribution of envy between agents. This generalizes several approaches using various definitions of degree of envies, which can be captured by adequate weight vector. In practice, we put a special focus on the egalitarian variant (minimizing the highest envy), the utilitarian variant (minimizing the sum of envies), and the compromise consisting of using the fair vector of decreasing weights. After implementing a MIP to compute min OWA allocations, we unveil several connections between the min OWA allocation and other famous fairness

Table 4.2: Performances for minimizing the OWA envy (with weights $\vec{\alpha}_1, \vec{\alpha}_2, \vec{\alpha}_3$) or maximizing the Nash Welfare on synthetics instances (as a function of the number of agents and objects (n, m) ($\epsilon \leq 10^{-3}$)).

	(2,3)				(5,7)				(8,10)				(10,12)			
	$\vec{\alpha}_1$	$\vec{\alpha}_2$	$\vec{\alpha}_3$	MNW	$\vec{\alpha}_1$	$\vec{\alpha}_2$	$\vec{\alpha}_3$	MNW	$\vec{\alpha}_1$	$\vec{\alpha}_2$	$\vec{\alpha}_3$	MNW	$\vec{\alpha}_1$	$\vec{\alpha}_2$	$\vec{\alpha}_3$	MNW
%EF	74	74	74	58	48	48	48	17	10	10	10	1	1	1	1	0
%EFX ₀	97	97	97	88	96	96	96	88	88	86	88	78	72	82	83	80
%EFX	100	100	100	92	97	97	98	91	98	96	93	85	87	95	92	84
%EF1	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100
%Pareto	100	100	100	100	73	76	72	100	64	66	67	100	51	62	64	100
%EF+PO	74	74	74	58	33	34	32	17	5	6	5	1	1	1	1	0
time(s)	ϵ	ϵ	ϵ	ϵ	0.02	0.02	0.02	0.01	0.1	0.4	0.1	0.04	0.5	2.5	0.7	0.07

measures.

In particular, we compare our approach with the alternative relaxations consisting of seeking “envy-freeness up to some/any good”. Some of our conclusions show that these approaches correspond to very different perspectives: we show in particular that no algorithm minimizing a fair OWA can ever guarantee to return an EF1 (and thus nor EFX) allocation. This is however balanced by the fact that it never occurred in our experiments. Indeed, even in the very few cases for which the min OWA allocation was not EF1 we easily found a weight for which it was the case. This raises the question of choosing the appropriate weight vector for example by eliciting it. We left that question open for now. Indeed, we also ran some experiments to test the performances of our method and compared it with other allocation protocols. The results are extremely encouraging. Our min OWA approaches do very well (in particular regarding the likelihood to return an EFX allocation, which may be somewhat paradoxical given our previous remarks) in terms of fairness, both on real Spliddit instances and randomly generated ones. In comparison, Nash social welfare –despite its guarantee to return an EF1 allocation– is dominated on that respect, as well as on the likelihood to return an EF and Pareto optimal allocation.

For future work, an interesting idea would be to look at the MILPs characterizing the EFX and EF1 allocations. Indeed, by solving such MILPs we can obtain EFX or EF1 allocations but what would be interesting is to use such MILPs to try and add Pareto-Optimality (for instance) by using the Lagrangian relaxation method for Integer Programming [Geoffrion, 1974; Fisher, 1981] which consists in relaxing an optimization problem by moving some constraints as a penalty in the objective function. Note that those MILPs have been presented by Kaczmarczyk [2021] and our modelization of the EF1 and EFX are respectively thoroughly presented in Appendices B and C. We can also, via Lagrangian relaxation, relax the constraints on EFX or EF1 if we want to mix it with other desirable properties.

Moreover, such MILP could also be used to study these properties through the prism of polyhedral approach such as presented by Edmonds [1965]. It consists in studying the polytope of a problem and allows to obtain min-max relaxations or polynomial-time algorithms.

Chapter 5

Fair in the Eyes of Others

Summary of the chapter

In this chapter, we mainly focus on a new relaxation of envy-freeness motivated by the fact that if an agent is envious then we would like to avoid, if possible, the worst situation in which all the agents agree on the fact that this agent should be envious. From this notion we also derive a similar notion for proportionality. This chapter is fully based on a conference paper [Shams et al., 2020] and on a longer journal version of the same paper [Shams et al., 2022].

5.1 Introduction

As we have seen in the state of the art and in the previous section, literature has proposed a lot of relaxations of envy-freeness, such as for instance envy-freeness up to one good, EF1 [Budish, 2011], or envy-freeness up to any good, EFX [Gourvès et al., 2014; Caragiannis et al., 2016]. When agents interact among a social network, local notions of fairness have been investigated where an agent can only make comparisons with her neighbors in the network. In the divisible setting, local envy-freeness and local proportionality have been studied for instance by Abebe et al. [2017]; Bei et al. [2017]. In the indivisible setting, similar notions have been studied by Aziz et al. [2018]; Chevaleyre et al. [2017], while complexity issues related to local envy-freeness have been investigated in oriented graphs [Bredereck et al., 2018] and non-oriented graphs in house allocation problems [Beynier et al., 2019b].

In this chapter, we propose a slightly different relaxation of envy. To illustrate the notion we introduce, consider a given instance where no envy-free allocation exists. Now suppose that in this instance there exist two allocations π and π' that make a single agent (say, a_i) envious of some other agent a_j (for simplicity). Furthermore, assume that in allocation π , no agent but a_i thinks that a_j 's bundle is better than a_i 's, while in allocation π' all the other agents concur with the assessment that a_i envies a_j . According to Pajouh [1997], π' exhibits *unanimous envy*, and it seems difficult to justify that π' should be returned in place of π . Inspired by this notion, we define in this chapter the notion of K -approval envy, as a way to introduce a continuum between envy-freeness and unanimous envy. As may be clear from the name, the idea is simply to ask agents to express their own view about envy relations expressed by other agents. The objective will thus be to seek

allocations minimizing social support for the expressed envy relations i.e., minimizing the number of agents K approving the envy. Of course, this approach may be controversial: after all, the notion of preference is inherently subjective. Introducing this flavour of objectivity may lead to undesirable consequences. At the extreme, one may simply replace individual preferences by some unanimous “mean” profile, thus profoundly changing the very nature of the notion. We believe though that there are several justifications to this approach:

- First, note that we only seek the approval of other agents in the case the agent herself explicitly expresses envy: absence of envy thus remains completely subjective. While a symmetrical treatment may also be justifiable in some situations, there is an obvious reason which motivates us to start with the proposed definition, namely the fact that the notion would no longer be a relaxation of envy-freeness.
- Secondly, all other things being equal, we believe that an allocation minimizing K is socially more desirable. We do not necessarily regard this notion as a compelling choice, but we think this can enrich the picture of fallback allocations when no envy-free allocation exists, as other relaxations do [Amanatidis et al., 2018]. In particular, in repeated settings, the fact that agents perceive outcomes as globally fair (not only for themselves, but also for others) may be important as an incentive for participation.
- Finally, one further motivation of our work is that our approach can be seen as providing guidance regarding agents and more specifically agents’ preferences, in order to progress towards envy-freeness by helping them revise their utilities for example. In particular, if we envision systems integrating deliberative phases in the collective decision-making process, our model could be used to set the agenda of such deliberations. If a vast majority of agents contradict an agent on her envy towards another agent, it may indicate for instance that she lacks information regarding the actual value of (some items of) her share. Initiating a discussion might help to solve such “objectively unjustified” envies when they occur.

While envy-freeness is a widely studied notion in fair division of indivisible goods, another prominent notion in the literature is *proportionality*. This notion is based on the proportional share: the proportional share of an agent is equal to the n^{th} of the utility this agent gives to the whole set of objects (with n the number of agents). An allocation is proportional if and only if each agent receives at least her proportional share. Note that there are strong links between proportionality and envy-freeness, namely, any envy-free allocation is also proportional, whereas, on the contrary, there are instances for which a proportional allocation exists but no envy-free one [Bouveret and Lemaître, 2016a]. As is the case for envy-freeness, a proportional allocation is not guaranteed to exist. As a result, there has been a lot of work in recent literature about a relaxation of proportionality called proportionality up to one item (PROP1) [Conitzer et al., 2017]. In the same spirit as EF1, PROP1 requires each agent to get her proportional share by obtaining the object of some other agent that she values the most. In light of these remarks and given the strong relationships between envy-freeness and proportionality, we also explore an approval version of this latter notion.

The remainder of this chapter is as follows. Our notion of K -approval envy is presented in Section 5.2. Some properties of this notion are then studied in Section 5.3: it is shown in particular, that if the hypothetical situation of allocation π described at the beginning of the introduction occurs, then an EF allocation must also exist. We also show that our notion inherits from the complexity of related problems. After introducing the approval notion around proportionality in Section 5.4, some properties are put forward in Section 5.5 and some links between our two approval notions are studied in Section 5.6. As we did for approval envy, we show that the problem inherits from the hardness of the classical notion. This hardness results motivates the MIP formulations that we detail in Section 5.7. We next turn to the *house allocation* setting and we show that if each agent exactly holds a single item, then we can define an efficient algorithm returning an allocation minimizing the value of K for both our approval notions. One caveat of our notions is that (unlike other relaxations) it is not guaranteed to exist, as intuitively observed in the case of unanimous envy and unanimous non-proportionality. In Section 5.9, we provide empirical evidences showing that allocations with reasonable values of K exist under synthetic cultures as well as in real datasets.

5.2 K -approval envy

The notion of envy being inherently subjective, it might be the case that an agent envies another agent, but that she has no reason to do so from the point of view of the other agents. The difficulty here is to define the notion of objectivity, since no ground-truth can properly serve as a basis of this definition. In her book, Guibet Lafaye [2006] recalls the notion of *unanimous envy*, that was initially discussed by Parijs [1997], and that can be defined as follows: an agent a_i unanimously envies another agent a_j , if all the agents think that a_i indeed envies a_j . Here, unanimity is used as a proxy for objectivity.

As we can easily imagine, looking for allocations that are free of unanimous envy will be too weak to be interesting: as soon as one agent disagrees with the fact that a_i envies a_j , this potential envy will not be taken into account. Here, we propose an intermediate notion between envy-freeness and (unanimous envy)-freeness:

Definition 5.1 (K -approval envy). *Let π be an allocation, a_i, a_j be two different agents, and $1 \leq K \leq n$ be an integer. We say that a_i K -approval envies (K -app envies for short) a_j if there is a subset \mathcal{N}_K of K agents including a_i such that:*

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

In other words, at least $K - 1$ agents amongst $\mathcal{N} \setminus \{a_i\}$ agree with a_i on the fact that she should actually envy agent a_j .

Example 5.1. *Let us consider the following add-MARA instance with 3 agents and 6 objects:*

	o_1	o_2	o_3	o_4	o_5	o_6
a_1	0	3	3	1	3	2
a_2	2	0	7	2	1	0
a_3	0	3	5	0	1	3

Note that there is no envy-free allocation for this instance. In the squared allocation, a_1 is not envious, a_2 envies a_3 and a_3 envies a_1 . Concerning the envy of a_2 towards a_3 , a_1 disagrees with a_2 being envious of a_3 whereas agent a_3 agrees. Hence, agent a_2 2-app envies agent a_3 . Concerning the envy of a_3 towards a_1 , agent a_1 agrees with a_3 being envious of a_1 whereas agent a_2 does not. Hence, a_3 2-app envies a_1 .

Note that in the definition, as soon as a_i does not envy a_j , then, a_i does not K -app envy a_j , no matter what the value of K is or how many agents think that a_i should actually envy a_j . Doing so, we ensure that our approval notion is a relaxation of envy-freeness.

Let us start with an easy observation:

Observation 5.1. *Given an allocation π of an add-MARA instance, if a_i K -app envies a_j in π , then a_i $(K-1)$ -app envies a_j in π .*

Moreover, if a_i n -app envies a_j , we will say that a_i *unanimously* envies a_j . Finally, we can observe that a_i 1-app envies a_j if and only if a_i envies a_j .

We can naturally derive from Definition 5.1 the counterpart of envy-freeness:

Definition 5.2 ($(K$ -approval envy)-free allocation). *An allocation π is said to be $(K$ -app envy)-free if and only if a_i does not K -app envy a_j in π for all pairs of agents (a_i, a_j) .*

Definition 5.3 ($(K$ -approval envy)-free instance). *An add-MARA instance I is said to be $(K$ -app envy)-free if and only if it accepts a $(K$ -app envy)-free allocation.*

Example 5.2. *Going back to Example 5.1, the squared allocation is (3-app envy)-free so the instance is (3-app envy)-free.*

A threshold of special interest is obviously $\lfloor n/2 \rfloor + 1$, since it requires a strict majority to approve the envy under inspection. A Strict Majority approval envy-free (SM-app-EF) allocation is a $(K$ -app envy)-free allocation such that $K \leq \lfloor n/2 \rfloor$, translating the fact that every time envy occurs, there is a strict majority of agents that do not agree with that envy.

Going further, it is important to notice that $(K$ -app envy)-freeness is not guaranteed to exist. Indeed, for all number of agents n and all number of objects m , there exist instances for which no $(K$ -app envy)-free allocation exists, no matter what K is. Suppose for instance that all the agents rank the same object (say o_1) first, and that for all a_i , $u(i, 1) > \sum_{k=2}^m u(i, k)$. Then obviously, everyone agrees that all the agents envy the one that will receive o_1 . Such instances will be called *unanimous envy instances*:

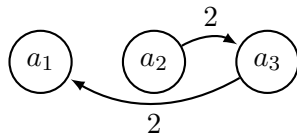
Definition 5.4 (Unanimous envy instance). *An add-MARA instance I is said to exhibit unanimous envy if I is not $(K$ -app envy)-free for any value of K .*

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 it 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)”.

A useful representation, for a given allocation, is the induced envy graph [Lipton et al., 2004]: vertices are agents, and there is a directed edge from a_i to a_j if and only if a_i envies a_j . An allocation is envy-free if and only if the envy graph has no arc. In our context, we can define a weighted notion of the envy graph.

Definition 5.5 (Weighted envy graph). The weighted envy graph of an allocation π is defined as the weighted graph (\mathcal{N}, E) where nodes are agents, such that there is an edge $(a_i, a_j) \in E$ if a_i envies a_j , with the weight $w(a_i, a_j)$ corresponding to the number of agents (including a_i) approving this pairwise envy in π .

Example 5.3. The induced weighted envy graph of Example 5.1 is as follows:



Our notion of K -approval envy can be interpreted as a vote on envy, that works as follows. For each pair of agents (a_i, a_j) , if a_i declares to envy a_j , we ask the rest of the agents to vote on whether they think that a_i indeed envies a_j . Then, a voting procedure is used to determine whether a_i envies a_j according to the society of agents. Several voting procedures can be used. However, since there are only two candidates (yes / no), the most reasonable voting rules are based on quotas: a_i envies a_j if and only if there is a minimum quota of agents that think so.¹ This makes a connection with a related work of Segal-Halevi and Suksompong [2019] which uses voting to decide upon envy-freeness, but in the context of fair division of resources *jointly* owned by *groups* of agents.

Finally, we want to emphasize that our notion of K -approval envy is based on *pairwise* envy. Namely, if agent a_i envies a_j , we will try to evaluate how many other agents think that this envy is justified. Another possibility² would be, for each envious agent a_i , to evaluate how many other agents think that a_i has indeed reasons to be envious, no matter which agent a_i envies. The difference is subtle. To illustrate this, suppose that a_i envies a_j and another agent a_k disagrees with this particular envy, but thinks that a_i has indeed reasons to envy $a_l (\neq a_j)$. With the first notion (our notion), a_k 's opinion will be discarded, whereas in the second one, it will be counted.

In practice, we believe that this alternative notion of approval envy will be much less discriminating than ours. The intuition can be explained as follows. Suppose that for some allocation π there is a bundle π_i that is not the top one for any agent. In that case, not only the agent a_i receiving π_i will envy another agent (since π_i is not a_i 's top bundle), but every other agent will also agree that a_i should be envious. Hence, π will exhibit unanimous envy. Now suppose in the contrary that every bundle of π is the top one for some agent. If preferences are strict on bundles, then by the pigeon hole principle, the top bundle of each agent has to be a different bundle of π , meaning that there is an envy-free allocation in that case. Said otherwise, if preferences are strict on bundles, an instance can only either be envy-free or unanimous envy. The only edge case happens when agents can have several tied top bundles, which does not happen very often in practice.

The experiments we run on the Spliddit instances (see Section 5.9) tend to confirm this intuition. This is why we decided not to investigate further this alternative notion of approval envy.

¹More precisely, these rules exactly characterize the set of anonymous and monotonic voting rules [Perry and Powers, 2010].

²We warmly thank an anonymous reviewer for pointing this alternative notion to us.

5.3 Some properties of K -app envy

There are natural relations between the properties of (K -app envy)-freeness for different values of K . The following observation is a direct consequence of Observation 5.1.

Observation 5.2. *Let π be an allocation, and $K \leq N$ be an integer. If π is (K -app envy)-free, then π is also $((K+1)$ -app envy)-free.*

However, the converse does not hold. More precisely, the following proposition shows that the implication stated in Observation 5.2 is strict.

Proposition 5.1. *Let π be an allocation, and $3 \leq K \leq n$ be an integer. If π is (K -app envy)-free, π is not necessarily $((K-1)$ -app envy)-free.*

Proof. Let $h \in \{2, \dots, n-1\}$ be an integer, and let us consider the instance with n agents and n objects defined as follows:

- $u(1, 1) = 1 - (n - 1)\varepsilon$;
- $u(i, 1) = u(i, i) = \frac{1-(n-2)\varepsilon}{2}$ for $i \in \{2, \dots, h-1\}$;
- $u(i, i) = 1 - (n - 1)\varepsilon$ for $i \in \{h, n-1\}$;
- $u(n, 1) = \frac{2}{n+1}$ and $u(n, j) = \frac{1}{n+1}$ for $j > 1$;

and $u(i, j) = \varepsilon$ for other pairs with $\varepsilon < \frac{1}{n+1}$.

This construction is illustrated in the general case in Figure 5.1. Moreover, one instance with $n = 4$ agents, $m = 4$ objects and $h = 3$ is shown in Example 5.4.

	o_1	o_2	o_3	\dots	o_h	\dots	o_n
a_1	$1 - (n - 1)\varepsilon$	ε	ε	\dots	ε	\dots	ε
a_2	$\frac{1-(n-2)\varepsilon}{2}$	$\frac{1-(n-2)\varepsilon}{2}$	ε	\dots	ε	\dots	ε
a_3	$\frac{1-(n-2)\varepsilon}{2}$	ε	$\frac{1-(n-2)\varepsilon}{2}$	\dots	ε	\dots	ε
\vdots							
a_h	ε	ε	ε	\dots	$1 - (n - 1)\varepsilon$	\dots	ε
\vdots							
a_n	$\frac{2}{n+1}$	$\frac{1}{n+1}$	$\frac{1}{n+1}$	\dots	$\frac{1}{n+1}$	\dots	$\frac{1}{n+1}$

Figure 5.1: Instance used in the proof of Proposition 5.1

Consider the allocation π where each agent a_i gets item o_i . Obviously, the only envy in this allocation concerns a_n towards a_1 . Moreover, only a_1, \dots, a_{h-1} agree on this envy. Therefore, a_n h -app envies a_1 , but does not $(h+1)$ -app envy her. Moreover, π is $((h+1)$ -app envy)-free, but not $(h$ -app envy)-free. \square

Example 5.4. *In order to illustrate the previous proof, let us consider the following instance with 4 agents, 4 objects (and $h=3$) and the squared allocation π :*

	o_1	o_2	o_3	o_4
a_1	$1 - 3\varepsilon$	ε	ε	ε
a_2	$\frac{1}{2} - \varepsilon$	$\frac{1}{2} - \varepsilon$	ε	ε
a_3	ε	ε	$1 - 3\varepsilon$	ε
a_4	$\frac{2}{5}$	$\frac{1}{5}$	$\frac{1}{5}$	$\frac{1}{5}$

In this allocation, the only envy concerns a_4 towards a_1 . Moreover, only a_1 and a_2 agree with a_4 on her envy. Hence, π is (4 -app envy)-free but is obviously not (3 -app envy)-free as we can find 3 agents (a_1 , a_2 and a_4) agreeing on the envy of a_4 towards a_1 (in other words a_4 3 -app envies a_1).

Proposition 5.2. For any $K \geq 3$, there exist instances which are (K -app envy)-free but not $((K-1)$ -app envy)-free.

Proof. Consider the same instance as in Proposition 5.1. We have already shown that we have an allocation π that is $((h+1)$ -app envy)-free which means that the instance is $((h+1)$ -app envy)-free. We just have to show that there is no $(h$ -app envy)-free allocation in order to conclude. For that purpose, we first note that each agent has to get one and exactly one object. Indeed, if it is not the case at least one agent a_i will have no object and will thus be envious of any agent a_j that has an object. Moreover, as all agents value the empty bundle with utility 0 and every object is valued with a strictly positive utility, this envy will be unanimous. Hence, each agent has to get one and exactly one object in order to minimize the $(K$ -app envy)-freeness. Now consider objects o_j for $j \in \{h, n\}$. The agents a_j that receive an object o_j and that are envious will h -app envy the agent that received o_1 . Indeed, agents a_i for $i \in \{h, n-1\}$ value objects o_j with a utility higher than (or equal to) the one of o_1 (and thus do not approve the envy) while it is the opposite for the other agents who are exactly h hence the h -app envy. So if we want to avoid that envy, we have to give the objects o_j to agents so that they do not experience envy at all but it is not possible as such agents are agents a_p for $p \in \{h, n-1\}$. It means that we have $n-1-h+1$ agents that have to receive the $n-h+1$ objects which is obviously impossible. This means that we cannot avoid h -app envy which implies that no allocation can be $(h$ -app envy)-free. \square

Proposition 5.2 proves that the hierarchy of $(K$ -app) envy-free instances is strict for $K \geq 3$. Rather surprisingly, we will see that it is not the case for $K = 2$.

In order to show this result, we will resort to a tool that has been proved to be useful and powerful in many contexts dealing with envy [Biswas and Barman, 2018; Amanatidis et al., 2019; Beynier et al., 2019a]: the *bundle reallocation cycle technique*. This technique, originating from the seminal work of Lipton *et al.* 2004, consists in performing a cyclic reallocation of *bundles* so that every agent is strictly better in the new allocation. Thus, such a reallocation corresponds to a cycle in the opposite direction of the edges in the — weighted — envy graph introduced in Definition 5.5. It is known that performing a reallocation cycle decreases the degree of envy [Lipton et al., 2004]. Unfortunately, our first remark is that it does not necessarily decrease the level of K -app envy. Worse than that, it can actually increase it:

	o_1	o_2	o_3	o_4	\dots	o_{h+3}	o_{h+4}	\dots	o_n
a_1	$\boxed{\varepsilon}$	$\textcircled{2\varepsilon}$	$1 - 3\varepsilon$	0	\dots	0	0	\dots	0
a_2	$\textcircled{1 - \varepsilon}$	$\boxed{\varepsilon}$	0	0	\dots	0	0	\dots	0
a_3	0	0	$\boxed{1}$	0	\dots	0	0	\dots	0
a_4	ε	0	ε	$\boxed{\frac{1-2\varepsilon}{n-3}}$	\dots	$\frac{1-2\varepsilon}{n-3}$	$\frac{1-2\varepsilon}{n-3}$	\dots	$\frac{1-2\varepsilon}{n-3}$
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
a_{h+3}	ε	0	ε	$\frac{1-2\varepsilon}{n-3}$	\dots	$\boxed{\frac{1-2\varepsilon}{n-3}}$	$\frac{1-2\varepsilon}{n-3}$	\dots	$\frac{1-2\varepsilon}{n-3}$
a_{h+4}	0	ε	2ε	$\frac{1-3\varepsilon}{n-3}$	\dots	$\frac{1-3\varepsilon}{n-3}$	$\boxed{\frac{1-3\varepsilon}{n-3}}$	\dots	$\frac{1-3\varepsilon}{n-3}$
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
a_n	0	ε	2ε	$\frac{1-3\varepsilon}{n-3}$	\dots	$\frac{1-3\varepsilon}{n-3}$	$\frac{1-3\varepsilon}{n-3}$	\dots	$\boxed{\frac{1-3\varepsilon}{n-3}}$

Figure 5.2: The instance used in the proof of Proposition 5.3.

Proposition 5.3. *Let π be a $(K\text{-app envy})$ -free allocation, for $3 \leq K \leq n - 1$. After performing an improving bundle reallocation cycle (even between two agents), there can be an integer $K' > K$ such that the resulting allocation is $(K'\text{-app envy})$ -free (and not $(K\text{-app envy})$ -free).*

Proof. Let $h \in \{0, \dots, n - 4\}$ be an integer, and let us consider the instance with n agents and n objects defined by the following utility functions:

- a_1 : $u(1, 1) = \varepsilon$, $u(1, 2) = 2\varepsilon$, $u(1, 3) = 1 - 3\varepsilon$;
- a_2 : $u(2, 1) = 1 - \varepsilon$, $u(2, 2) = \varepsilon$;
- a_3 : $u(3, 3) = 1$;
- a_l for $l \in \{4, h + 3\}$: $u(l, 1) = u(l, 3) = \varepsilon$, $u(l, j) = \frac{1-2\varepsilon}{n-3}$ for $j \geq 4$;
- a_m for $m \in \{h + 4, n\}$: $u(m, 2) = \varepsilon$, $u(m, 3) = 2\varepsilon$, $u(m, i) = \frac{1-3\varepsilon}{n-3}$ for $i \geq 4$;

and $u(i, j) = 0$ for other pairs. We assume in this construction that $\varepsilon \leq \frac{1}{2n-1}$.

This construction is illustrated in Figure 5.2.

Consider the allocation π where each agent a_i gets item o_i (corresponding to the squared allocation in Figure 5.2). Obviously, in this allocation, there is no envy, except:

- a_1 envying a_2 (agents $a_{h+4} \dots a_n$ agree on that);
- a_1 envying a_3 (agents a_3 and $a_{h+4} \dots a_n$ agree on that);
- a_2 envying a_1 (agents $a_4 \dots a_{h+3}$ agree on that).

Hence the allocation is $((\max\{n-h, h+2\})\text{-app envy})$ -free. We now consider the allocation π' resulting from the improving bundle reallocation cycle between a_1 and a_2 (circled allocation in Figure 5.2). Observe that the only envy in π' is the one of a_1 towards a_3 , which is approved by everyone except a_2 . This allocation is thus $(n\text{-app envy})$ -free and not $((n-1)\text{-app envy})$ -free. If $h > 0$, then $\max\{n-h, h+2\} < n$, which proves the proposition. \square

Now consider a slight generalization of Lipton's cycles, *weakly improving cycles* (WIC), that correspond to a reallocation of bundles where all the agents in the cycle receive a bundle they like at least as much as the one they held, with one agent at least being strictly happier. Of course, our example of Proposition 5.3 still applies. On the other hand, this notion suffices to guarantee the decrease of the degree of envy (note that *identifying* the cycles themselves may not be easy any longer, but this is irrelevant for our purpose). The proof follows directly from the arguments of Lipton *et al.* 2004 (proof of Lemma 2).

Observation 5.3. *Let π be an allocation, and π' the allocation obtained after performing a weakly improving cycle. It holds that $de(\pi') < de(\pi)$.*

Proof. Let us consider an allocation π' obtained after performing a WIC on an allocation π . First note that the envies agents who are not involved in the WIC stay unchanged. By definition of a WIC, all the agents get at least as much in π' as they had in π . Thus basically $de(\pi') \leq de(\pi)$. Moreover, at least one agent gets a strictly better bundle so her envy strictly decreases. We finally get that $de(\pi') < de(\pi)$. \square

We now show that (2-app envy)-freeness exhibits a special behaviour: in contrast with Proposition 5.3, improving cycles (in fact, even weakly improving cycles) enjoy the property of preserving the (2-app envy)-freeness level of an allocation. We provide this result for *swaps* (cycles involving two agents only) as this is sufficient to establish our main result.

Lemma 5.1. *Let π be a (2-app envy)-free allocation that is not EF. There always exists a WIC (that we can identify) between two agents such that the resulting allocation is (K' -app envy)-free, with $K' \leq 2$.*

Proof. Let a_i be an envious agent (there is at least one). We identify the agent that a_i envies the most and call her a_j (if there are several agents that a_i envies the most, we can pick randomly one of them). As a_i envies a_j and a_j necessarily does not agree on this envy because otherwise it would contradict (2-app envy)-freeness of π , swapping the bundle of a_i and a_j is a WIC. Let us call π' this new allocation. We will now show that π' is a (K' -app envy)-free allocation with $K' \leq 2$.

In π' , all the agents except a_i and a_j have the same approval envy. Moreover, a_i is now EF in π' as she has received her preferred bundle. Suppose for contradiction that π' is (K' -app envy)-free with $K' > 2$. Then necessarily, this is due to a_j 2-app envying (at least) some other agent a_h (that can obviously not be a_i). For this to be the case, a_j has to envy a_h and another agent a_l has to approve this envy: (1) $u_j(\pi'_j) < u_j(\pi'_h)$, (2) $u_l(\pi'_j) < u_l(\pi'_h)$. However, as a_i envies a_j in π then (3) $u_i(\pi_i) < u_i(\pi_j)$ and as π is (2-app envy)-free and (3) holds, every agent a_l (except a_i of course) verifies (4) $u_l(\pi_i) \geq u_l(\pi_j)$.

Besides, π' is obtained after swapping the bundles of a_i and a_j in π so $\pi'_j = \pi_i$, $\pi'_i = \pi_j$ and $\pi'_h = \pi_h$; and from (2) we get: (5) $u_l(\pi_i) < u_l(\pi_h)$. By transitivity with (5) and (4), we obtain: (6) $u_l(\pi_j) < u_l(\pi_h)$. However, we know that a_j has the same utility in π and π' so $u_j(\pi'_j) = u_j(\pi_j)$. The latter combined with (1) (and the fact that $\pi'_h = \pi_h$) gives: (7) $u_j(\pi_j) < u_j(\pi_h)$. Finally, note that (6) and (7) translate the fact that a_j 2-app envies a_h in π which contradicts the fact that π is (2-app envy)-free. \square

Putting Lemma 5.1 and Observation 5.3 together allows us to prove that (2-app envy)-freeness is essentially a vacuous notion, in the sense that any instance enjoying an allocation with this property will have an EF allocation as well.

Proposition 5.4. *If an add-MARA instance is (2-app envy)-free then it is also envy-free.*

Proof. Take π as being an arbitrary (2-app envy)-free allocation. First note that if there is no envious agent in π then, by definition, π is envy-free and the proposition holds. We perform a WIC leading to π' that is still (2-app envy)-free (see Lemma 5.1). If π' is envy-free then we are done. Otherwise, from Observation 5.3 we know the degree of envy has strictly decreased and that the resulting allocation is still (2-app envy)-free by Lemma 5.1. Hence we can repeat this process until the current allocation is EF. The process is guaranteed to stop because the degree of envy of the society is bounded below by zero and the degree of envy of the society strictly decreases at each step until it reaches zero (which corresponds to an envy-free allocation). \square

Another consequence is that, for two agents, instances fall either in the envy-free or unanimous envy category:

Corollary 5.1. *Let I be an add-MARA instance with $n = 2$, if there is no envy-free allocation in I then I is a unanimous envy instance.*

Complexity We conclude this section with a few considerations on the computational complexity of the problems mentioned so far. First of all, as envy-freeness is (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.

One may also wonder how hard the problem is to determine whether a given instance exhibits unanimous envy or not, *i.e.* whether a (K -app envy)-free allocation exists for *some* value of K . For this question, instances where agents all have the same preferences provide insights.

Proposition 5.5. *For any add-MARA instance, if all the agents have the same preferences then the notions of (1-app envy)-freeness and (n -app envy)-freeness coincide.*

Proof. We already know from Observation 5.2 that (1-app envy)-freeness implies (n -app envy)-freeness for any add-MARA instance. So we just have to prove that if all the agents have the same preferences then (n -app envy)-freeness implies (1-app envy)-freeness. If an allocation π is (n -app envy)-free then it means that for any pair (a_i, a_j) of agents, a_i does not envy a_j or there is at least one agent a_h that disagrees on the envy of a_i towards a_j . Obviously, if for every pair of agents (a_i, a_j) we have a_i envy-free towards a_j then the allocation π is envy-free and the proof concludes. Besides, for every pair of envious/envied agents there is at least one agent disagreeing on the envy. But all the agents have the same preferences so it means that every agent should agree with each other. Hence, no envied agent can exist and we have (1-app envy)-freeness of allocation π . \square

From Proposition 5.5 we get that the problem of deciding the existence of unanimous envy is at least as hard as deciding the existence of an EF allocation when agents have similar preferences which is known to be NP-hard [Lipton et al., 2004]. As membership in NP is direct, we thus get as a corollary that:

Corollary 5.2. *Deciding whether an instance exhibits unanimous envy is NP-Complete.*

5.4 K -approval non-proportionality

As there is a clear hierarchy in the notions of fairness deriving from envy-freeness, it can be natural to consider how the different notions of this hierarchy would behave in an approval setting as we studied in Sections 5.2 and 5.3. Indeed, let us recall that it has been shown by Bouveret and Lemaître [2016a] that envy-freeness implies proportionality. Moreover, some relaxations of proportionality have been studied such as PROP1 in very recent works [Aziz et al., 2020; Barman and Krishnamurthy, 2019; Brânzei and Sandomirskiy, 2019; Conitzer et al., 2017]. This motivates us to investigate how we can derive an approval notion of proportionality.

In this section, we will introduce the approval version of proportionality. Observe first that our approval version of envy-freeness was based on a pairwise notion that we do not have in proportionality. This is why we slightly adapt the approval notion to this property.

Definition 5.6 (K -approval non-proportionality). *Let π be an allocation, a_i be an agent, and $1 \leq K \leq n$ be an integer. We say that π_i is K -approval non-proportional (K -app non-prop for short) in π if there is a subset \mathcal{N}_K of K agents including a_i such that:*

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

In other words, at least $K - 1$ agents amongst $\mathcal{N} \setminus \{a_i\}$ agree with a_i on the fact that she does not have her proportional share. We emphasize that we chose to focus on non-proportionality rather than on proportionality, to be consistent with our definition of K -app envy. The other related notions are defined accordingly as follows.

Definition 5.7 ($(K$ -approval non-proportional)-free allocation). *An allocation π is said to be $(K$ -app non-proportional)-free if and only if no π_i is K -app non-proportional.*

Once again, observe that the interpretation of this property is that an allocation is free of K -app non-prop: each agent a_i either thinks she receives a proportional share, or, if it is not the case, no more than $K - 2$ agents agree with a_i .

Definition 5.8 ($(K$ -approval non-proportional)-free instance). *An add-MARA instance I is said to be $(K$ -app non-proportional)-free if and only if it accepts a $(K$ -app non-proportional)-free allocation.*

Definition 5.9 (Unanimous non-proportional allocation). *An add-MARA allocation π is said to exhibit unanimous non-proportionality if π is not $(K$ -app non-proportional)-free for any value of K .*

Definition 5.10 (Unanimous non-proportional instance). *An add-MARA instance I is said to exhibit unanimous non-proportionality if I is not $(K$ -app non-proportional)-free for any value of K .*

Example 5.5. *Let us consider the add-MARA instance introduced in Example 5.1:*

	o_1	o_2	o_3	o_4	o_5	o_6
a_1	0	3*	3	$\boxed{1}$	3	$\boxed{2}^*$
a_2	$\boxed{2}^*$	$\boxed{0}$	7	2*	$\boxed{1}^*$	0
a_3	0	3	$\boxed{5}^*$	0	1	3

It is easy to notice that the proportional share of all the agents is the same and is worth 4. In the squared allocation, π_1 is not proportional as she values her bundle 3. Moreover a_2 and a_3 agree on the non-proportionality of her bundle. Hence π_1 is unanimous non-proportional. Besides π_2 is not proportional either. However neither a_1 nor a_3 agree on this non-proportionality as they value a_2 's bundle with respective utilities 6 and 4. So π_2 is 1-app non-prop. Finally, π_3 is proportional. Consequently, we can say that the squared allocation is unanimous non-prop because of a_1 . However, the instance itself is not unanimous non-prop as we can easily notice that the star allocation is proportional and hence (1-app non-prop)-free.

5.5 Properties of (K-app non-prop)-freeness

In this section, we will present some properties about the notion of (K -app non-prop)-freeness for different values of K . We will also present some complexity results.

There are natural relationships between the different notions of (K -app envy)-freeness, for different values of K . We start with an easy observation which is the counterpart for (K -app non-prop)-freeness of Observation 5.2:

Observation 5.4. *Given an allocation π of an add-MARA instance, if π_i is K -app non-proportional in π , then π_i is $(K-1)$ -app non-proportional in π .*

The following observation is a direct consequence of Observation 5.4.

Observation 5.5. *Let π be an allocation, and $K \leq N$ be an integer. If π is (K -app non-prop)-free, then π is also $((K+1)$ -app non-prop)-free.*

However, the converse does not hold. More precisely, the following proposition shows that the implication stated in Observation 5.5 is strict.

Proposition 5.6. *Let π be an allocation, and $3 \leq K \leq n$ be an integer. If π is (K -app non-prop)-free, π is not necessarily $((K-1)$ -app non-prop)-free.*

Proof. Let us consider the following instance with 3 agents and 3 objects and the squared allocation π . Recall that $Prop_i$ denotes the proportional share of a_i as stated in Definition 1.4:

	o_1	o_2	o_3
a_1	$Prop_1 + 1$	$Prop_1 + 1$	$Prop_1 - 2$
a_2	$Prop_2 - 1$	$Prop_2 - 1$	$Prop_2 + 2$
a_3	$Prop_3 - 1$	$Prop_3 - 1$	$Prop_3 + 2$

In this allocation, the only agent that does not hold her proportional share is a_2 . Moreover, we can easily see that a_3 agrees with this non-proportionality whereas a_1 does not. So a_2 experiences 2-app non-prop and thus π is a (3-app non-prop)-free allocation but not (2-app non-prop)-free. □

Proposition 5.7. *For any $K \geq 3$, there exists instances which are (K -app non-prop)-free but not $((K-1)$ -app non-prop)-free.*

Proof. Consider the same instance as in Proposition 5.6. We have already shown that we have an allocation π that is (3-app non-prop)-free which means that the instance is (3-app non-prop)-free. We just have to show that there is no (2-app non-prop)-free allocation in order to conclude. For that purpose, we first note that each agent has to get one and exactly one object. Indeed, if it is not the case at least one agent a_i will have no object and will thus not obtain her proportional share. Moreover, as all agents value the empty bundle with utility 0 this non-proportionality will be unanimous. Hence, each agent has to get one and exactly one object in order to minimize the (K -app non-prop)-freeness. Moreover, as a_2 and a_3 have the same preferences and only o_3 fulfils their proportional share then there is obviously no proportional allocation. Finally, this means that one of them will get either o_1 or o_2 , and the non-proportionality of their bundles will be approved by the other, leading to a 2-app non-prop. Thus there is no (2-app non-prop)-free allocation and so the instance is (3-app non-prop)-free and not (2-app non-prop)-free. \square

Proposition 5.7 proves that the hierarchy of K -app non-prop instances is strict for $K \geq 3$. As it was the case for the approval notion derived from envy-freeness we will see that it is not the case for $K = 2$ by show that (2-app non-prop)-freeness exhibits a special behaviour. For that, we start with a simple result.

Lemma 5.2. *Let π be an allocation. For each agent a_i , there is at least one bundle π_j such that $u_i(\pi_j) \geq Prop_i$.*

Proof. Let us consider for the sake of contradiction that there exists one allocation π in which an agent a_i cannot find any bundle that fulfils her proportional share. This means that every bundle is valued strictly less than $Prop_i = \frac{\sum_{j=0}^m u(i,j)}{n}$. By adding all the bundles (there are by definition n bundles in any allocation) we get that a_i values all the bundles strictly less than $n \times Prop_i = n \times \frac{\sum_{j=0}^m u(i,j)}{n} = \sum_{j=0}^m u(i,j)$ which is an obvious contradiction. \square

We now establish a result similar to Lemma 5.1:

Lemma 5.3. *Let π be a (2-app non-prop)-free allocation that is not proportional. There always exists a bundle exchange between two agents (swap), that is not necessarily improving, such that the resulting allocation is (K' -app non-prop)-free (with $K' \leq 2$) and such that the number of agents with a non-proportional bundle has strictly decreased.*

Proof. Let π be a (2-app non-prop)-free allocation that is not proportional. Let a_i be an agent whose π_i is non-proportional in π (there is at least one). According to Lemma 5.2, there is (at least) one share π_j such that $u_i(\pi_j) \geq Prop_i$. Let π' be the allocation resulting from swapping a_i 's and a_j 's bundles in π . In π' , all the agents except a_i and a_j have bundles with the same approval non-proportionality. Moreover, π'_i is now proportional in π' by definition of the swap we chose. Finally, π'_j is also proportional: suppose for contradiction that it is not the case. Then it would mean that $u_j(\pi'_j) < Prop_j$, which in turns implies $u_j(\pi_i) < Prop_j$. In other words, in π , π_i was not proportional and a_j agreed, which contradicts the fact that π was (2-app non-prop)-free. Hence, π'_j is proportional, and as a result, π' is still (2-app non-prop)-free, and the number of agents with a non-proportional bundle has increased by at least 1 (a_i is the new agent with a proportional bundle). \square

Putting together Lemma 5.2 and Lemma 5.3 allows us to prove that (2-app non-prop)-freeness is essentially a vacuous notion, in the same sense as it is for (2-app envy)-freeness (Proposition 5.4):

Proposition 5.8. *If an add-MARA instance is (2-app non-prop)-free then it is also proportional.*

Proof. Let π be an arbitrary (2-app non-prop)-free allocation. First note that if all the agents have proportional bundles in π then, by definition, π is proportional and the proposition holds. Otherwise, we perform a swap leading to π' that is still (2-app non-prop)-free (see Lemma 5.3). If π' is proportional then we are done. Otherwise, thanks to the second part of Lemma 5.3 we know the number of agents with a non-proportional bundle has strictly decreased. We can repeat this process until the current allocation is proportional. The process is guaranteed to stop because the number of agents with a non-proportional bundle is bounded below by zero and decreases at each step until it equals zero (which corresponds to a proportional allocation). \square

Another consequence is that, for two agents, instances are either proportional or unanimous non-proportional:

Corollary 5.3. *Let I be an add-MARA instance with $n = 2$, if there is no proportional allocation in I then I is an unanimous non-proportional instance.*

Proof. For any add-MARA instance involving exactly 2 agents, we can (by definition) only find (1-app non-prop)-free allocations or (2-app non-prop)-free allocations (as $1 \leq K \leq n$ for any add-MARA instance). By the contraposition of Proposition 5.8 we conclude the proof. \square

We also note that, as it was the case for K -app envy, performing a reallocation cycle can increase the level of K -app non-prop:

Proposition 5.9. *Let π be a (K -app non-prop)-free allocation, for $3 \leq K \leq n - 1$. After performing an improving bundle reallocation cycle (even between two agents), the resulting allocation may be (K' -app non-prop)-free (and not (K -app non-prop)-free) such that $K' > K$.*

Proof. Let us consider the following instance with 3 agents and 3 objects:

	o_1	o_2	o_3
a_1	$\overline{Prop_1 - 1}$	$\overline{Prop_1 - 2}$	$Prop_1 + 3$
a_2	$\underline{Prop_2}$	$\underline{Prop_2 + 3}$	$Prop_2 - 3$
a_3	$\underline{Prop_3 - 1}$	$Prop_3$	$\underline{Prop_3 + 1}$

First consider the squared allocation that is (2-app non-prop)-free as only a_1 does not hold her proportional share and that it is not approved by any other agent. Let us now consider the underlined allocation π that is the result of the improving bundle reallocation between a_1 and a_2 . We can see that only a_1 does not hold her proportional share and that this time a_3 approves it leading to a 2-app non-prop and thus a (3-app non-prop)-free allocation. \square

Complexity. We conclude this section with a few considerations on the computational complexity of the problems mentioned so far around the approval notion of proportionality. First of all, as proportionality is equivalent to (1-app non-prop)-freeness, the problem of finding the minimum K for which there exists a (K -app non-prop)-free allocation is at least as hard as determining whether a proportional allocation exists which is known to be NP-complete.

One may also wonder how hard the problem of determining whether a given instance exhibits unanimous non proportionality or not is, *i.e.* whether a (K -app non-prop)-free allocation exists for *some* value of K . For this question, as in Proposition 5.5, instances where agents all have the same preferences provide the answer.

Proposition 5.10. *For any add-MARA instance, if all the agents have the same preferences then the notions of (1-app non-prop)-freeness and (n -app non-prop)-freeness coincide.*

Proof. We already know from Observation 5.5 that (1-app non-prop)-freeness implies (n -app non-prop)-freeness for any add-MARA instance. So we just have to prove that if all the agents have the same preferences then (n -app non-prop)-freeness implies (1-app non-prop)-freeness. Let π be an (n -app non-prop)-free allocation. Then for any agent a_i , either $u_i(\pi_i) \geq Prop_i$, or there exists an agent a_j such that $u_j(\pi_i) \geq Prop_j$. Since all the agents have identical preferences, the last inequality reduces to $u_i(\pi_i) \geq Prop_i$, showing that a_i receives her proportional share. Hence in this case, π is proportional. \square

From Proposition 5.10 we get that the problem of deciding the existence of a unanimous non-proportional allocation is at least as hard as deciding the existence of a proportional allocation when agents have similar preferences which is known to be NP-hard (see for instance Bouveret and Lemaître, 2016a). As membership in NP is direct, we thus get as a corollary that:

Corollary 5.4. *Deciding whether an allocation exhibits unanimous non-prop is NP-Complete.*

5.6 Link between approval envy-freeness and approval non-proportionality

After having introduced some properties of approval non-proportionality, we will now investigate the relationships between this notion and approval envy-freeness introduced earlier.

We first recall that envy-freeness implies proportionality and that this implication is still valid for EF1 and PROP1. It is thus natural to wonder whether it is also the case for our approval notions. As we will see, the answer is negative.

Proposition 5.11. *A unanimous envy instance can be proportional.*

Proof. Let us consider the following generic add-MARA instance (here, $\varepsilon \leq 1/n$):

	o_1	o_2	\dots	o_{n-1}	o_n
a_1	$\boxed{1/n}$	$1/n - \varepsilon$	\dots	$1/n - \varepsilon$	$1/n + (n-2)\varepsilon$
a_2	$1/n - \varepsilon$	$\boxed{1/n}$	\dots	$1/n - \varepsilon$	$1/n + (n-2)\varepsilon$
\vdots					
a_{n-1}	$1/n - \varepsilon$	$1/n - \varepsilon$	\dots	$\boxed{1/n}$	$1/n + (n-2)\varepsilon$
a_n	$1/n - \varepsilon$	$1/n - \varepsilon$	\dots	$1/n$	$\boxed{1/n + (n-2)\varepsilon}$

In this instance, the squared allocation is proportional (and so (1-app non-prop)-free) whereas it is easy to see that the instance is a unanimous envy one as o_n is the top object of every agent. Hence the agent that gets o_n will be envied and this envy will be approved by everyone. \square

From this result, we can generalize the statement to any level of (K -app) envy and any level of (L -app non-prop)-freeness. First of all, from Observation 5.5, it is clear that the counter-example of Proposition 5.11 establishes that a unanimous envy instance can be (L -app non-prop)-free, for any $L \geq 1$. But note also that if (counterfactually) it was the case that proportionality (or indeed any level of (L -app non-prop)-freeness) would imply some level of (K -app)-envy freeness, then by invoking Observation 5.2 this would also imply (unanimous envy)-freeness, a contradiction with Proposition 5.11. Putting all these remarks together allows us to state the following result.

Corollary 5.5. *For any $K \geq 1$ and any $L \geq 1$, an allocation exhibiting (K -app) envy can be (L -app non-prop)-free.*

Since proportionality is a weaker notion than envy-freeness, the previous result may not come as a surprise. It seems much more likely to obtain a positive result in the other direction, that is, that some level of (K -app envy)-freeness actually implies some level of (K -app non-prop)-freeness. It turns out that this is not the case.

Proposition 5.12. *An instance that exhibits unanimous non-proportionality can be (3-app envy)-free.*

Proof. Let us consider the following add-MARA instance for which $Prop_i = \frac{1}{n}$ for all i :

	o_1	o_2	o_3	\dots	o_{n-1}	o_n
a_1	$\boxed{\varepsilon}$	ε	ε	\dots	ε	$1 - (n-1)\varepsilon$
a_2	ε	ε	ε	\dots	ε	$\boxed{1 - (n-1)\varepsilon}$
a_3	ε	ε	$\boxed{1 - (n-1)\varepsilon}$	\dots	ε	ε
\vdots						
a_n	ε	ε	ε	\dots	$\boxed{1 - (n-1)\varepsilon}$	ε

It is obvious that in any allocation the agent that gets o_1 will not get her proportional share and that this non-proportionality will be approved by everyone. However, the squared allocation is (3-app envy)-free since the only envy in this allocation is a_1 's towards a_2 , and only a_2 approves this envy. \square

Again, this allows us to state a more general result. First of all, it is direct from Observation 5.2 that the counter-example of Proposition 5.12 establishes that an unanimous non-proportional instance can be (K -app envy)-free, for any $K \geq 3$. But note also that if (counterfactually) it was the case that (3-app envy)-freeness (or indeed any level $L \geq 3$ of (L -app envy)-freeness) would imply some level of (K -app)-non-prop freeness, then by invoking Observation 5.5 this would also imply (unanimous non-prop)-freeness, a contradiction with Proposition 5.12. Putting all these remarks together allows us to state the following result.

Corollary 5.6. *For any $K \geq 3$ and any $L \geq 1$, a (K -app envy)-free instance can exhibit (L -app) non-proportionality.*

Note that this is the best we can do, since by Observation 5.2, Observation 5.5 and the well-known implication between envy-freeness and proportionality, we have an implication from (2-app envy)-freeness and any level of (L -app non-prop)-freeness.

Now in principle, and even if counter-intuitive at first sight, it could still be that exhibiting unanimous envy could imply proportionality; or that exhibiting unanimous non-proportionality could imply (3-app)-envy-free. The following result shows that both implications do not hold.

Proposition 5.13. *An instance can be at the same time unanimous envy and unanimous non-proportional.*

Proof. Let us consider the following instance with n agents and commensurable utilities ($Prop_i = \frac{1}{n}$ for all i and we assume that $\varepsilon < \frac{1}{n}$):

	o_1	o_2	\dots	o_n
a_1	ε	ε	\dots	$1 - (n-1)\varepsilon$
a_2	ε	ε	\dots	$1 - (n-1)\varepsilon$
\vdots				
a_n	ε	ε	\dots	$1 - (n-1)\varepsilon$

It is obvious to see that any agent getting an object different from o_n (say w.l.o.g o_1) will not be proportional and will envy the agent receiving o_n . Moreover, since all the agents have the same preferences, they will all agree with this non-proportionality and envy. \square

We have summed up the relations between approval envy notions and approval non-proportionality ones in Figure 5.3.

5.7 Computation

We have seen at the end of Section 5.3 (respectively Section 5.5) that the problem of determining, for a given instance I , the minimal value of K such that a (K -app envy)-free (respectively a (K -app non-prop)-free) allocation exists inherited from the high complexity of determining whether an envy-free (respectively a proportional) allocation exists.

To address this problem, we present in this section two Mixed Integer linear Programs that return, for a given add-MARA instance I , a (K -app envy)-free (respectively (K -app

[Bouveret and Lemaître, 2016a]

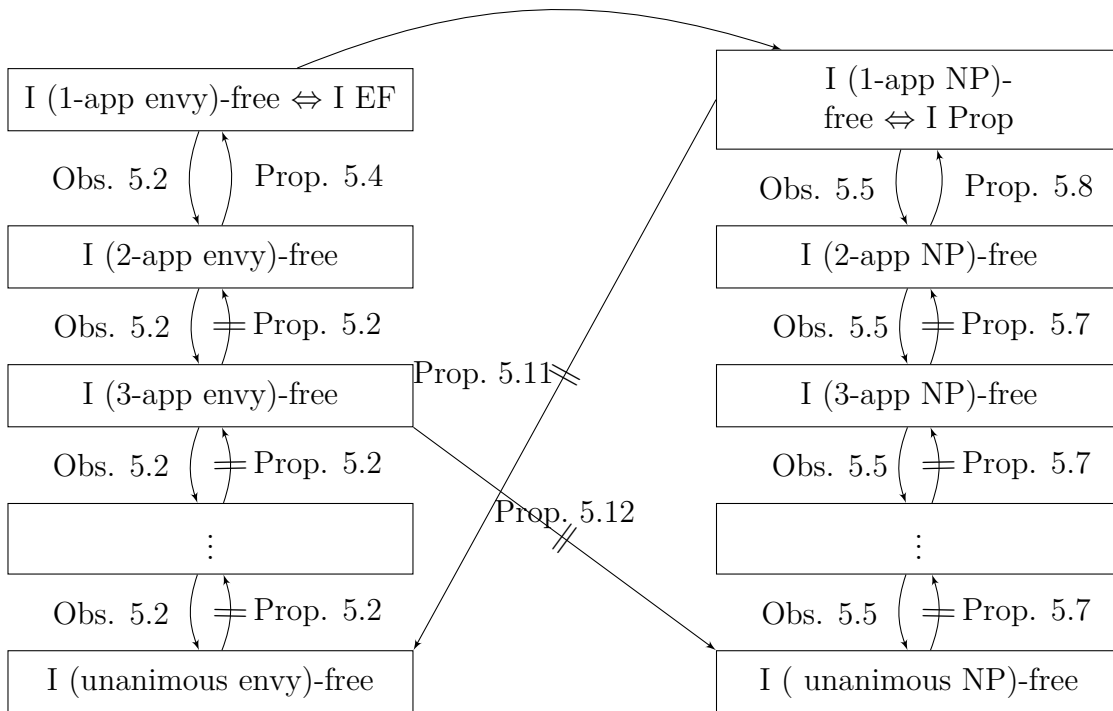


Figure 5.3: Hierarchy among instance properties. A simple edge denotes an implication relation. A striked out edge has been drawn when we have found a counter-example showing that this implication is not valid. Edges obtained by transitivity are not shown. All the remaining missing arcs are non-implication edges which can be obtained thanks to Corollaries 5.5 and 5.6.

non-prop)-free) allocation with the minimal K and no solution when I is an unanimous envy (respectively non-prop) instance. We will first introduce and thoroughly explain the MIP for K -app envy. Then, we will show how to adapt it to K -app non-prop.

In this section, we assume that all the utilities are integers. If they are not (recall that they are assumed to be in \mathbb{Q}^+) we can transform the instance at stake into a new one only involving integer utilities by multiplying them by the least common multiple of their denominators.

5.7.1 A MILP formulation for K -app envy

In this MILP, we use $n \times m$ Boolean variables \mathbf{z}_i^j (we use bold letters to denote variables) to encode an allocation: $\mathbf{z}_i^j = 1$ if and only if a_i gets item o_j . We also introduce n^3 Boolean variables \mathbf{e}_{kih} such that $\mathbf{e}_{kih} = 1$ if and only if a_i envies agent a_h according to a_k 's preferences. We also need to add n^2 Boolean variables \mathbf{x}_{ih} used to linearize the constraints on \mathbf{e}_{kih} . Finally, we use an integer variable \mathbf{K} corresponding to the K -app envy we seek to minimize.

We first need to write the constraints preventing an item from being allocated to several agents:

$$\sum_{i=1}^n \mathbf{z}_i^j = 1 \quad \forall j \in \llbracket 1, m \rrbracket \quad (5.1)$$

By adding these constraints we also guarantee completeness of the returned allocation (all the items have to be allocated to an agent).

Secondly, we have to write the constraints that link the variables \mathbf{e}_{kih} with the allocation variables \mathbf{z}_i^j :

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

As the utilities are integers, we can replace > 0 by ≥ 1 . In order to linearize the equivalence between the two constraints we introduce a number M that can be arbitrarily chosen such that $M > \max_{a_k \in \mathcal{N}} \sum_{j=1}^m u(k, j)$:

$$M\mathbf{e}_{kih} \geq \sum_{j=1}^m u(k, j)(\mathbf{z}_h^j - \mathbf{z}_i^j) \quad \forall k, i, h \in \llbracket 1, n \rrbracket \quad (5.2)$$

$$\sum_{j=1}^m u(k, j)(\mathbf{z}_h^j - \mathbf{z}_i^j) \geq 1 - M(1 - \mathbf{e}_{kih}) \quad \forall k, i, h \in \llbracket 1, n \rrbracket \quad (5.3)$$

Finally, we have to write the constraints that convey the fact that the allocation we look for is (K -app envy)-free:

$$\mathbf{e}_{iih} = 0 \vee \sum_{k=1}^n \mathbf{e}_{kih} \leq \mathbf{K} - 1 \quad \forall i, h \in \llbracket 1, n \rrbracket$$

Since $\mathbf{e}_{\mathbf{i}ih}$ are Boolean variables, we can replace $\mathbf{e}_{\mathbf{i}ih} = 0$ by $\mathbf{e}_{\mathbf{i}ih} \leq 0$. Now, these logical constraints are linearized as follows:

$$\mathbf{e}_{\mathbf{i}ih} \leq \mathbf{x}_{\mathbf{i}ih} \quad \forall i, h \in \llbracket 1, n \rrbracket \quad (5.4)$$

$$\sum_{k=1}^n \mathbf{e}_{\mathbf{k}ih} \leq \mathbf{K} - 1 + n(1 - \mathbf{x}_{\mathbf{i}ih}) \quad \forall i, h \in \llbracket 1, n \rrbracket \quad (5.5)$$

We can now put things together. Let I be an instance. Then, we will denote by $\mathcal{M}_1(I)$ the MIP defined as:

$$\begin{aligned} & \text{minimize} && \mathbf{K} \\ & \text{such that} && \mathbf{z}_i^j, \mathbf{e}_{\mathbf{k}ih}, \mathbf{x}_{\mathbf{i}ih} \in \{0, 1\} && \forall k, i, h \in \llbracket 1, n \rrbracket, j \in \llbracket 1, m \rrbracket \\ & && \mathbf{K} \in \llbracket 1, N \rrbracket \\ & && + \text{Constraints (5.1, 5.2, 5.3, 5.4, 5.5)} \end{aligned}$$

Proposition 5.14. *Let I be an add-MARA instance. Then, there is an optimal solution with $\mathbf{K} = L$ to $\mathcal{M}_1(I)$ if and only if I is an $(L\text{-app envy})$ -free instance and not an $((L-1)\text{-app envy})$ -free one. Moreover, $\mathcal{M}_1(I)$ does not admit any solution if and only if I is an unanimous envy instance.*

The proof of this result can be found in Appendix D.

5.7.2 A MILP formulation for K -app non-prop

In the previous subsection, we have introduced a Mixed Integer Linear Program that returns a $(K\text{-app envy})$ -free allocation with the minimal K and no solution when I is an unanimous envy (respectively non-prop) instance. We will now explain how to adapt it to K -app non-proportionality.

In this adapted MIP, we use the same Boolean variables \mathbf{z}_i^j . We also introduce n^2 Boolean variables $\mathbf{p}_{\mathbf{k}i}$ such that $\mathbf{p}_{\mathbf{k}i} = 1$ if and only if according to a_k 's preferences a_i 's bundle is worth strictly less than the proportional share of a_k . We also need to add n Boolean variables \mathbf{x}_i used to linearize the constraints on $\mathbf{p}_{\mathbf{k}i}$. Finally, we use an integer variable \mathbf{K} corresponding to the K -app non-proportionality we seek to minimize.

Recall that we assume in this section that the utilities are integers. We will further assume that $Prop_k = \sum_{j=1}^m u(k, j)/n$ is also an integer for each k . If it is not the case, they all the utilities can be multiplied by n without changing the result.

We first need Constraint (5.1) to ensure the correctness of the allocation.

Secondly, we have to write the constraints that link variables $\mathbf{p}_{\mathbf{k}i}$ with the allocation variables \mathbf{z}_i^j :

$$\sum_{j=1}^m u(k, j) \cdot \mathbf{z}_i^j < \frac{\sum_{j=1}^m u(k, j)}{n} (= Prop_k) \iff \mathbf{p}_{\mathbf{k}i} = 1 \quad \forall k, i \in \llbracket 1, n \rrbracket$$

As the utilities are integers, we can replace > 0 by ≥ 1 . In order to linearize the equivalence between the two constraints we introduce a number M that can be once again arbitrarily chosen such that $M > \max_{a_k \in \mathcal{N}} \sum_{j=1}^m u(k, j)$:

$$M\mathbf{p}_{ki} \geq Prop_k - \sum_{j=1}^m u(k, j)\mathbf{z}_i^j \quad \forall k, i \in \llbracket 1, n \rrbracket \quad (5.6)$$

$$Prop_k - \sum_{j=1}^m u(k, j)\mathbf{z}_i^j \geq 1 - M(1 - \mathbf{p}_{ki}) \quad \forall k, i \in \llbracket 1, n \rrbracket \quad (5.7)$$

Finally, we have to write the constraints that convey the fact that the allocation we look for is (K -app non-prop)-free:

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

Since \mathbf{p}_{ii} is a Boolean variable for each i , we can replace $\mathbf{p}_{ii} = 0$ by $\mathbf{p}_{ii} \leq 0$. Now, this logical constraint is linearized as follows:

$$\mathbf{p}_{ii} \leq \mathbf{x}_i \quad \forall i \in \llbracket 1, n \rrbracket \quad (5.8)$$

$$\sum_{k=1}^n \mathbf{p}_{ki} \leq \mathbf{K} - 1 + n(1 - \mathbf{x}_i) \quad \forall i \in \llbracket 1, n \rrbracket \quad (5.9)$$

We can now put things together. Let I be an instance. Then, we will denote by $\mathcal{M}_2(I)$ the MIP defined as:

$$\begin{aligned} & \text{minimize} && \mathbf{K} \\ & \text{such that} && \mathbf{z}_i^j, \mathbf{p}_{ki}, \mathbf{x}_i \in \{0, 1\} && \forall k, i \in \llbracket 1, n \rrbracket, j \in \llbracket 1, m \rrbracket \\ & && \mathbf{K} \in \llbracket 1, N \rrbracket \\ & && + \text{Constraints (5.1, 5.6, 5.7, 5.8, 5.9)} \end{aligned}$$

Proposition 5.15. *Let I be an instance. Then, there is an optimal solution with $\mathbf{K} = L$ to $\mathcal{M}_2(I)$ if and only if I is an (L -app non-prop)-free instance and not an $((L - 1)$ -app non-prop)-free one. Moreover, $\mathcal{M}_2(I)$ does not admit any solution if and only if I is an unanimous non-proportional instance.*

The proof of this result can be found in Appendix D.

5.8 House Allocation setting

We have seen in Corollaries 5.2 and 5.4 the problems of finding the minimal level K for which there exists a (K -app envy)-free or a (K -app non-prop)-free allocation are difficult in the general case. A natural way to tackle this difficulty is to look for particular restrictions where these problems can be solved efficiently. In this section, we will deal with the House Allocation setting.

As mentioned in Chapter 1, the House Allocation Problem (HAP for short) is a standard setting where there are exactly as many items as agents, and each agent receives exactly one item. This setting is relevant in many situations and has been extensively studied [Shapley and Scarf, 1974; Roth and Sotomayor, 1992; Abraham et al., 2005]. In

House Allocation Problems, computing an envy-free allocation and a proportional allocation reduces to the problem of finding a matching in a bipartite graph, which can be done in $O(n^3)$ [Gondran and Minoux, 1984]³. Indeed, an envy-free allocation exists if and only if all the agents get (one of) their top item(s) and a proportional allocation exists if and only if each agent a_i gets an item whose value is greater than $Prop_i$. It is therefore natural to wonder whether an allocation minimizing K -app envy or K -app non-proportionality could also be computed efficiently.

Our first observation hints in that direction. Indeed, characterizing unanimous envy becomes easy in house allocation problems.

Proposition 5.16. *Let I be an instance of HAP. I is an unanimous envy instance if and only if there exists at least one pair of items (o_i, o_j) such that all the agents unanimously strictly prefer o_i to o_j .*

Proof. (\Rightarrow) Suppose that for any pair of items (o_i, o_j) , there are two agents (a_k, a_l) such that $u(k, i) \geq u(k, j)$ and $u(l, i) \leq u(l, j)$. Let π be an allocation, and suppose w.l.o.g that $\pi_i = \{o_i\}$. Then for any pair of agents (a_i, a_j) , either (i) $u(i, i) \geq u(i, j)$, in which case a_i does not envy a_j , or (ii) $u(i, i) < u(i, j)$, in which case a_i envies a_j , but there is another agent a_k such that $u(k, i) \geq u(k, j)$. In the latter case, a_k disagrees with a_i 's envy towards a_j . Hence a_i does not unanimously envy a_j . Therefore I is not an unanimous envy instance.

(\Leftarrow) Suppose now that there is a pair of items (o_i, o_j) such that $u(k, i) > u(k, j)$ for all agents a_k . In any allocation one agent (say a_i) holds o_i while another agent (say a_j) holds o_j : a_j envies a_i and all the agents approve this envy. Therefore I is an unanimous envy instance. \square

Incidentally, we get as a corollary:

Corollary 5.7. *One can check in $O(n^3)$ whether an instance I of HAP is a unanimous envy instance or not.*

From this characterization we can also derive a result on the likelihood that unanimous envy exists when the utilities are uniformly distributed (that is, for each agent a_i and object o_j , utilities are drawn i.i.d. following the uniform distribution on some interval $[x, y]$). The interested reader can find this result in Appendix D.

We will now investigate the case of approval non-proportionality in the context of HAP. Interestingly, it is also possible to exactly characterize the set of unanimous non-prop instances.

Proposition 5.17. *Let I be an HAP instance. I is an unanimous non-prop instance if and only if there exists at least an item o_p such that $u(k, p) < Prop_k$ for all agents a_k .*

Proof. (\Rightarrow) Suppose no such item o_p exists. Let π be any allocation giving to each agent a_i an item (say o_i w.l.o.g). Then either $u(i, i) \geq Prop_i$, in which case a_i receives her proportional share, or $u(i, i) < Prop_i$, in which case there is another agent a_k such that

³Slightly better complexities have since been proposed, for more details check [Lovász and Plummer, 2009]

$u(k, i) > Prop_k$. a_k thus disagrees with π_i being non-proportional. Hence the instance is not unanimous non-prop.

(\Leftarrow) Now suppose that there is an item o_p such that $u(k, p) < Prop_k$ for all agents a_k . In any allocation one agent (say a_p) holds o_p . By definition, a_p does not get her proportional share, and all the agents agree with that. Therefore, the instance is unanimous non-prop. \square

As for approval envy-freeness, this result yields an efficient way of checking whether an instance is unanimous non-prop or not:

Corollary 5.8. *One can check in $O(n^2)$ whether an instance I of HAP is an unanimous non-prop instance or not.*

Like in the approval envy case, we can derive from this characterization an upper bound on the probability for an instance to be unanimous non-proportional (see Appendix D).

We will now show that finding an allocation minimizing (K -app envy)-freeness can be done in polynomial time. Before introducing the idea, we need an additional notation. For any pair of objects $(o_j, o_{j'})$, let $\#_{\prec}(o_j, o_{j'})$ denote the number of agents strictly preferring $o_{j'}$ to o_j . For any agent a_i and object o_j , we will also define $maxEnvy(i, j)$ as follows:

$$maxEnvy(i, j) = \max_{o_{j'} \text{ s.t. } u(i, j') > u(i, j)} \#_{\prec}(o_j, o_{j'})$$

In other words, $maxEnvy(i, j)$ denotes the maximal value of $\#_{\prec}(o_j, o_{j'})$ among the objects that are strictly preferred to o_j by a_i . As we can imagine, this will exactly be the value of the K -app envy experienced by a_i if she gets item o_j (note that if o_j is among a_i 's top objects, this value will be 0).

The key to the algorithm is to see that for a given K , determining whether a (K -app envy)-free allocation exists can be done in polynomial time by solving a matching problem. Namely, for each K , we build the following bipartite graph: $\mathcal{N} \cup \mathcal{O}$ is the set of nodes, and we add an edge $(a_i, o_j) \in \mathcal{N} \times \mathcal{O}$ if and only if $maxEnvy(i, j)$ is lower than or equal to K . We can observe that any perfect matching in this graph corresponds to a ($(K+1)$ -app envy)-free allocation. More precisely, if there exists a perfect matching, that means that the allocation π resulting from the perfect matching is ($(K+1)$ -app envy)-free but there could exist another allocation with lower (approval envy)-freeness. If there is no perfect matching, then there could exist a (h -app envy)-free allocation with $h > K + 1$. The only thing that remains to do is to run through all possible values of K , which can be done by dichotomous search between 0 and n . This is formalized in Algorithm 5.1.

Proposition 5.18. *For any HAP instance, we can find (one of) its optimal (K -app envy)-free allocations in $O(n^3 \log(n))$.*

Proof. First, the computation of the matrix $maxEnvy$ runs in $O(n^3)$. Indeed, to compute $maxEnvy(i, j)$ we first need to compute $\#_{\prec}(o_j, o_{j'})$ which already runs in $O(n^3)$ as we have to ask for each couple of objects (n^2 in total) the point of view of all the agents (n in total). From that, since

$$maxEnvy(i, j) = \max_{o_{j'} \text{ s.t. } u(i, j') > u(i, j)} \#_{\prec}(o_j, o_{j'})$$

Algorithm 5.1: Minimizing (K -app envy)-freeness in the HAP

input : $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ a HAP instance
output: Allocation π and its level minimizing the (K -app envy)-freeness or **None**
 if I is a unanimous envy instance

```

1  $maxEnvy \leftarrow \text{computeMaxEnvy}()$ ;
2  $res \leftarrow \text{None}$ ;
3  $low \leftarrow 0, high \leftarrow n$ ;
4 while  $low \leq high$  do
5    $K \leftarrow \lfloor (low + high)/2 \rfloor$ ;
6    $G \leftarrow \text{buildBipartiteGraph}(maxEnvy, K)$ ;
7    $\pi \leftarrow \text{perfectMatching}(G)$ ;
8   if  $\pi$  is not None then
9      $res \leftarrow \pi, K + 1$ ;
10     $high \leftarrow K - 1$ ;
11  end
12  else
13     $low \leftarrow K + 1$ ;
14  end
15 end
16 return  $res$ 
  
```

we can compute $maxEnvy(i, j)$ in $O(n)$. As there are n^2 different pairs (a_i, o_j) we have the final $O(n^3)$ complexity of computing $maxEnvy$.

Due to the dichotomous search, the algorithm needs to solve $\log(n)$ perfect matching problems, that can be solved in $O(n^3)$ [Gondran and Minoux, 1984]. The overall complexity of Algorithm 5.1 is thus $O(n^3 \log(n))$. \square

Following the same idea, we can propose an algorithm that returns an allocation minimizing (K -app non-prop)-freeness in polynomial time. For this case, we no longer need the matrix $maxEnvy$, but we have to replace it by some vector $\#nonProp$ that tells for each object o_j how many agents think this object is not worth their proportional share:

$$\#nonProp(j) = |\{a_i \text{ s.t. } u(i, j) < Prop_i\}|$$

In Algorithm 5.1 we then replace Line 5.1 by an instruction computing $\#nonProp$ for each o_j . Then, we replace the bipartite graph computed at Line 5 by the graph defined as follows: $\mathcal{N} \cup \mathcal{O}$ is still the set of nodes, and we add an edge $(a_i, o_j) \in \mathcal{N} \times \mathcal{O}$ if and only if $u(i, j) < Prop_i$ or $\#nonProp(j)$ is lower than or equal to K .

Proposition 5.19. *For any HAP instance, we can find (one of) its optimal K -app non-prop-free allocations in $O(n^3 \log(n))$.*

Proof. We know from the proof of Proposition 5.18 that the algorithm runs in at least $O(n^3 \log(n))$ due to the dichotomous search associated with the perfect matching problem resolutions. But the complexity could be worse because of the computation of $\#nonProp$ and the construction of the bipartite graph. To compute $\#nonProp$, it is enough for each object o_j to run through all the agents and count how many of them think o_j is not

worth their proportional share. This can be done in $O(n^2)$ steps, provided that we have pre-computed the values $Prop_i$ first (which can be done in $O(n)$ for each agent, that is, $O(n^2)$ in total). Computing the bipartite graph does not take longer than before, since we just have to check for each pair (a_i, o_j) whether $u(i, j) < Prop_i$ or $\#nonProp(j) \leq K$ (which can be made in constant time if the values $Prop_i$ and $\#nonProp(j)$ have been pre-computed). Thus in total, the adaptation of the algorithm does not cause any added complexity, so the global complexity is still $O(n^3 \log(n))$. \square

5.9 Experimental results

We conducted an experimental evaluation of our approval notions and solving methods. These experiments serve two purposes: (i) evaluate the behaviour of the MIPs we presented in Section 5.7 and of the polynomial algorithms described in Section 5.8 for the HAP setting, and (ii) observe the relevance of our two approval notions when varying the number of agents, of items, and the type of preferences. All the tests presented in this section have been run on an Intel(R) Core(TM) i7-2600K CPU with 16GB of RAM and using the Gurobi solver to solve the Mixed Integer Program.

We have tested our methods on three types of instances: Spliddit instances [Goldman and Procaccia, 2015], instances under uniformly distributed preferences and instances under an adaptation of Mallows distributions to cardinal utilities [Durand et al., 2016].

5.9.1 Spliddit instances

We have first experimented our MIPs on real-world data from the fair division website Spliddit [Goldman and Procaccia, 2015]. There is a total of 3535 instances from 2 agents to 15 agents and up to 96 items. Note that 1849 of these instances involve 3 agents and 6 objects. The program we ran for Spliddit instances proceeds as follows. It first checks whether the instance is HAP. If it is the case, it runs Algorithm 5.1 to compute the optimal level of approval envy. If this level is 1, it means that the instance is EF, and hence proportional [Bouveret and Lemaître, 2016a]. We stop there in this case. Otherwise, we run the adaptation of Algorithm 5.1 to compute the level of approval non-proportionality. If the instance is not HAP, we proceed the same way, replacing Algorithm 5.1 and its adaptation by MIPs \mathcal{M}_1 and \mathcal{M}_2 .

Approval envy Concerning approval envy, by setting a timeout of 1 minute, the program was able to solve all but 6 instances optimally. By extending the timeout to 10 minutes, we were able to solve 4 additional instances. We were however unable to solve the last 2 remaining instances optimally within 5 hours. Those instances respectively concern 6 agents and 15 objects, and 4 agents and 29 objects. However, by examining this latter instance, we could notice that all the agents had the same preferences. Running MIP \mathcal{M}_2 on this instance lead us to find an allocation that is proportional, meaning that this allocation is also envy-free in that case. Hence, in the end, only one instance still resists to our attempts. Among the 3534 instances that have been solved optimally, 63.8% admit an EF allocation, while 24.6% exhibit unanimous envy. Moreover, 29% of the 83 instances with more than 5 agents are Strict Majority-app EF (SM-app EF).

We have also implemented the alternative notion of (K -app envy)-freeness mentioned at the end of Section 5.2 and computed the optimal K for the 3469 easiest Spliddit instances (we removed those that timed out after 20 seconds). Among these instances, only 47 were found to be neither EF nor unanimous-envy, that is, about 1.4%, which confirms our intuition that this alternative notion is much less discriminating than the notion of K -app envy we use in this chapter.

Approval non-prop Concerning approval non-proportionality, all Spliddit instances have been solved optimally within 1 minute. 69.3% of the instances turn out to be proportional, while 25.4% exhibit unanimous non-prop. Note that since we know [Bouvet and Lemaître, 2016a] that envy-freeness implies proportionality, we knew from the previous experiments that the percentage of proportional instances would be greater than 63.8%. So we can notice that around 5.3% of the instances actually are neither proportional nor unanimous non-prop against the 11.7% we had for the approval envy notion.

5.9.2 Uniformly distributed preferences: general setting

We also ran tests on instances under uniformly distributed preferences, with n varying from 2 to 10 and m such that we produce appropriate settings to study our notions of approval envy-freeness and approval non-proportionality. Under Impartial Culture, all preference profiles are equally likely. It is a commonly studied in computational social choice [Black et al., 1958; Gehrlein and Fishburn, 1976] as a limit case, also providing an easy way to get syntactic instances without knowledge on preference characteristics from a particular concrete problem.

Approval envy We first studied the notion of approval envy and thus considered settings where few EF allocations exist [Dickerson et al., 2014]. More precisely we took m almost equal to n , for example 2 agents with 3 objects, 5 agents with 7 objects and 10 agents with 13 objects. As shown by Dickerson et al. [2014], the percentage of EF instances is tightly related to the ratio between the number of agents and the number of objects. The probability of EF instances is small when the number of objects is not much larger than the number of agents. For each couple (n, m) , Table 5.1 reports the percentage of envy-free instances obtained over 1000 randomly generated instances. It can be noticed that the number of EF instances decreases as the numbers of agents and objects increase. The worst-case in Table 5.1 is obtained for 9 agents and 11 objects where only 90 over 1000 instances are envy-free. For each couple (n, m) , we randomly picked 60 instances over the instances not EF that were randomly generated. Indeed, we wanted to investigate the behavior of our notion when no EF allocation exists (we know that if an EF allocation exists it will be returned by our methods). As we are in the general setting we solved the instances via the MIP \mathcal{M}_1 with a timeout of 10 minutes. Experimental results are depicted in Table 5.2.

The first three rows of Table 5.2 respectively report the percentage of instances that have been solved to optimal (a solution has been returned before the timeout), the percentage of unanimous envy instances and the percentage of Strict Majority-app-EF instances (SM-app-EF instances). The mean value of K/n gives a good insight on how many agents agree on the fairness notion (in Table 5.2, on the envy of an agent). Moreover, as it is a

(n, m)	(2,3)	(3,4)	(4,5)	(5,7)	(6,8)	(7,9)	(8,10)	(9,11)	(10,13)
% EF	86	59	42	58	36	29	11	9	14

Table 5.1: Percentage of envy-free instances as a function of the number of (agents,objects).

n	2	3	4	5	6	7	8	9	10
% optimal	100	100	100	100	100	68.3	1.7	1.7	0
% UE I	100	21.7	5	0	0	0	0	0	0
% SM-app EF I	0	0	0	50	50	75	40	33.3	6.7
mean(K/n)	-	1	0.85	0.72	0.61	0.57	0.59	0.63	0.66
time(s)	ε	0.008	0.04	0.21	1.97	21.29	50.09	56.16	-

Table 5.2: Performances of MIP \mathcal{M}_1 on randomly generated no EF instances

normalised measure it allows us to compare the level of approval non-proportionality and envy for instances with different number of agents. Finally, we store the mean computation time (in seconds) of the instances (solved to optimal).

First note that considering 2 agents is a special case as shown in Corollary 5.1. Indeed, as we have removed the EF instances, all the remaining instances are unanimous envy ones (denoted by - in the tables). Moreover, we observe that the percentage of SM-app-EF allocations is zero for up to 4 agents, which can be easily explained. Indeed, for 3 or 4 agents, being SM-app-EF means that there exists a (K -app envy)-free allocation with $K \leq 2$, which comes down (by Proposition 5.4) to say that there exists an envy-free allocation. Since all the EF instances have been removed, we cannot find an SM-app-EF allocation for $n \leq 4$.

Besides, without any surprise, the computation time rapidly increases while the percentage of instances solved to optimal (under a timeout of 10 minutes) starts decreasing for 7 agents.

Finally, positive results can be pinpointed. The percentage of unanimous envy instances is very low. This highlights the relevance of the K -approval envy-free notion. Indeed, in most instances, there exists allocation where we can find a subset of the agents supporting the absence of envy. Minimizing the number of agents approving the envy is thus relevant in almost all instances. Moreover, the experiments show that the percentage of SM-app-EF instances is higher than 30% except for 10 agents. Such instances are desirable as it means that the absence of envy is supported by more than half the agents: from the point of view of the social acceptance, it is thus possible to find an allocation where the fairness is supported by a majority of agents.

Approval non-prop First note that a proportional allocation is likely to exist as soon as $m \geq n$ [Suksompong, 2016]. As we do not want to be in the House Allocation setting yet, we considered instances for which $m = n + 1$. We have tested our MIP \mathcal{M}_2 described in Section 5.7.2 on such instances with a timeout of 10 minutes. For each couple (n, m) , we generated 10 000 instances.

The first four rows of Table 5.3 respectively represent the percentage of instances

n	2	3	4	5	6	7	8	9	10
% optimal	100	100	100	100	100	100	100	100	100
% proportional	86.94	83.77	98.41	99.98	99.76	99.95	100	100	100
% UNPI	100	79.2	69.8	100	70.8	100	-	-	-
mean(K/n)	-	1.0	0.96	-	0.95	-	-	-	-
time(s)	0.002	0.03	0.1	1.8	1.5	8.6	-	-	-

Table 5.3: Performances of MIP \mathcal{M}_2 on randomly generated instances.

that have been solved to optimal (a solution has been returned before the timeout), the percentage of proportional instances, the percentage of unanimous non-proportional instances (among the ones that are not proportional) and the mean value of K/n that gives a insight on how many agents agree on the non-proportionality of an agent. Finally, we store the mean computation time (in seconds) of the instances that are not proportional.

We can first notice that all the instances have been solved to optimal and the number of proportional instances remains very high even if we considered a favourable context with $m = n + 1$. Notably, for more than 8 agents, all the instances were proportional leaving no space for our relaxation to be useful.

Note that considering 2 agents is a special case as shown in Corollary 5.3. Indeed, as we do not consider the proportional instances, all the remaining instances are unanimous non-proportional ones. For more than 2 agents, we can see that the percentage of unanimous non-prop instances is almost 70% among non-proportional instances. Besides, we can notice that when it is relevant to look at the mean K/n metric, it tells us that the level of approval is very high. In light of these results, we could conclude that while proportionality is a much less demanding notion, it turns out that when it is not satisfied it is extremely often unanimously not satisfied.

5.9.3 Uniformly distributed preferences: house allocation problems

We have also tested our polynomial algorithms on HAP instances under uniformly distributed preferences. We have generated 20 instances for each number of agents from 5 to 100 agents (and objects) by steps of 5.

Approval envy Figure 5.4 shows the evolution of K/n as a function of the number of agents n (and hence also as a function of m as $n = m$) when minimizing the K -approval envy. First, note that we have only found 5 unanimous envy instances and all of them involved 5 agents. Indeed the probability of unanimous envy instance can be shown to quickly converge to 0 –see Proposition D.3 in Appendix. In HAP, agents are very likely to be envious as an agent envies someone as soon as she does not obtain her most preferred object. Let consider an agent a_j that holds o_j and that envies another agent a_k holding o_k . This envy is approved by all the agents that rank o_k over o_j . This envy is likely to be approved but it is also unlikely that all agents agree on this envy. In such contexts where the agents are likely to have mixed opinions, the K -approval envy-free notion and our related algorithm allow for computing allocations where the envy is supported by the

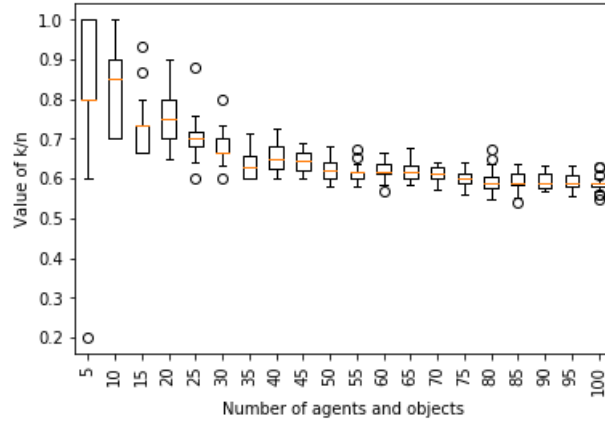


Figure 5.4: Optimal K/n (envy) in the HAP as a function of n

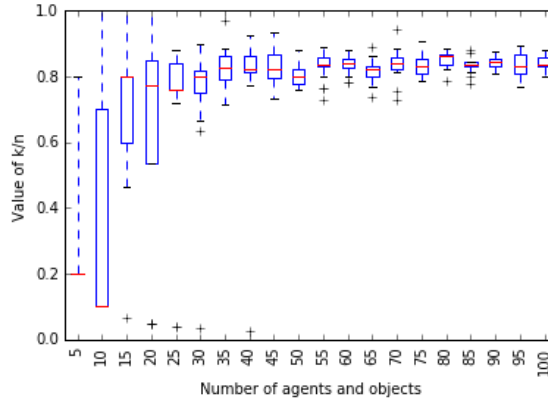


Figure 5.5: Optimal K/n (non-proportionality) in the HAP as a function of $n(=m)$

smallest subset of agents. As shown in Figure 5.4, even if the optimal K/n value is high for small problems, it slightly decreases as the size of the instances increases.

Note that the algorithm runs, without any surprise (in light of Proposition 5.18) much faster than our MIP \mathcal{M}_1 . Indeed, the mean runtime for 100 objects and agents is still around 2 seconds whereas we already observed that our MIP cannot solve easier problems within 10 minutes.

Approval non-prop We have also tested our polynomial algorithm to find an optimal K -approval non-proportional-free allocation. Although the algorithm was running very fast even for 100 agents and objects (confirming what we showed in Proposition 5.19), we almost only obtained proportional instances. We thus decided to test other instances: by using Borda utilities for each agent and randomly choosing one object per agent whose utility has been multiplied by the number of agents, we built a instances where only one object per agent fulfills the proportional share. We can see in Figure 5.5 that the value of K/n is stabilising around 0.8 meaning that around 80% of the agents agree with (at least) one agent’s non-proportionality.

5.9.4 Correlated preferences

As Impartial Culture may not reflect realistic preference profiles, we also generated instances where the preferences of the different agents may have similarities. In strict ordinal settings, a classical way to capture correlated preferences is to use Mallows distributions [Mallows, 1957] allowing us to measure the impact of the similarity of the preferences between agents. In these experiments, we used a generalization of the Mallows distribution to cardinal preferences based on Von Mises–Fisher distributions [Durand et al., 2016]. Like the dispersion parameter in Mallows distributions, the similarity between the preferences of the agents is tuned by a *concentration* parameter: when the concentration is zero the agents’ preferences are uniformly distributed, whereas when the concentration is infinite all the agents have the same preferences. The concentration can be viewed as the degree of conflicts among the resources. High concentration values lead to similar preferences among the agents for a given item.

We expected that the more similar the preferences between the agents are, the higher the degrees of K -app envy and non-proportionality would get and the more likely unanimous envy and non proportionality would occur. The results of our experiments both in the general setting and in HAP support this: the number of envy-free and proportional instances is decreasing along with the concentration value, and from a given threshold, all the instances exhibit unanimous envy and unanimous non-prop. We can see it for example through Figure 5.6.

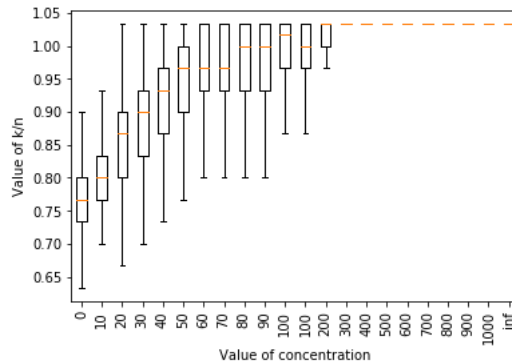


Figure 5.6: Optimal K/n (approval envy)-freeness in the HAP as a function of the value of concentration for $n = 30$ and $m = 30$

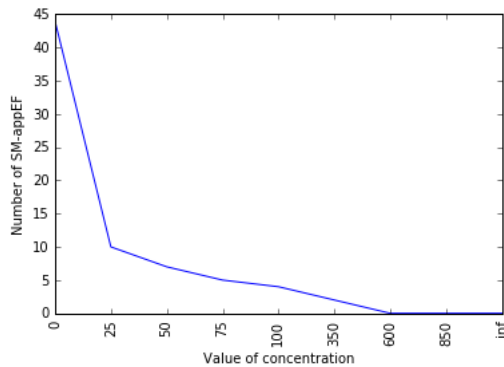


Figure 5.7: Number of SM-app EF instances as a function of the value of concentration for $n = 7$ and $m = 9$

Even though at the extreme (when all agents have the same preferences) all notions become unanimous, one may still wonder whether some degree of correlation among preferences may help to find large majorities of agents that contradict the envy of an agent. We thus studied how the number of SM-app EF instances varies as a function of the concentration. We considered instances involving 7 agents and 9 objects as we had previously noticed that under uniformly distributed preferences (which is equivalent to a value of concentration of 0), it was very likely to find SM-app EF instances. We then varied the concentration value. For each value, we generated 100 instances and counted the number of SM-app EF instances. As shown in Figure 5.7, the higher the concentration (and hence the more similar the preferences), the less SM-app EF instances are found, contradicting our hypothesis that correlation might make large majorities of agents contradicting envy more likely to occur.

5.10 Conclusion

In this chapter, we have introduced a new relaxation of envy-freeness and proportionality. These relaxations use a consensus notion, approval envy or non-proportionality, as a proxy for an idealized notion of envy between pairs of agents or proportionality of an agent. We have proposed algorithms to compute an allocation minimizing the approval envy or non-proportionality, and we have experimentally studied how these notions behave on real world data, as well as on instances with uniformly distributed or correlated preferences; more particularly in situations where no envy-free allocation exists and where no proportional allocation exists. We have shown that our notion of approval envy (less so for approval non-proportionality) strikes an interesting balance allowing to discriminate in practice among instances depending on the social support envy relations experience. In comparison, using consensus to determine whether a given agent should be envious or not in general proves to be of limited interest: except in rare cases, instances will either be envy-free or unanimous envy.

This work also opens up to a more general study of consensus-based notions of envy. One could for instance look for allocations that are judged envy-free by a given quota of agents. Restrictions of the approval notions such as an underlying social graph constraining the agents that can approve or disapprove –those agents you deem legitimate to express her view about a specific envy relation– could also be of interest for future work. Other domain restrictions, beyond house allocation, could be studied. For instance, the domain of binary additive preferences, with a cap on the number of items that an agent can like, may offer other tractable cases for our problem. Besides, the approval notions introduced in this article also call for a study of the manipulation that could arise from it. Indeed, asking the opinion of the agents gives birth to new ways of manipulating. More generally, an axiomatic study of the notions proposed here could nicely complement the results obtained.

Besides, it could also be interesting to propose extensions to the case where some items can be shared. Indeed, the approval concepts are a way to mix voting concepts with fair division and shared items is another way of building a continuum between voting and fair division. There may exist a potential link between both approaches. We leave the study of these notions for future work.

It could also be interesting to look deeper in other approval notions derived from the

notion of mMS and MMS as we did for PROP. If we consider that we have the mMS share and the MMS share of each agent then our guess is that it can be computed as done for the approval notion of PROP. However it can still be interesting to look for the links between those different measures as we did for the other two measures.

Finally, one could also think about using those approval notions to build a new kind of procedure based on picking sequences. The idea would be to have a picking order chosen by an approval notion of the agents and try to obtain an allocation with good approval or classical fairness and efficiency properties.

Conclusion and future works

We will first sum up the main contributions that we have made in this thesis.

After having formally introduced the model and having settled our problem regarding existing literature in Chapter 1, Chapter 2 focuses on the study of picking sequences and more particularly procedures around it and around cycle deals and their connections with Pareto-Optimality, envy-freeness but also CEEI in both the general setting and single-peaked domains. Following with procedures around cycle deals, we focus on swap-deals with single-peaked preferences for the House Allocation Problem in Chapter 3. We showed in particular that the single-peaked domain is maximal for guaranteeing convergence to Pareto-optimal outcomes with such dynamics. We also showed that the allocations returned by TTC and Crawler are reachable by swap-deal sequences. We also refined the analysis by studying two further notions: the average rank and the minimum rank of the resources obtained by the agents. None of the procedures discussed in this chapter aim to optimize these ranks, even though they respectively capture very natural criteria of efficiency and fairness. Hence we studied how these allocation procedures behave on that respect. Finally, to complement worst-case theoretical bounds on the loss of social welfare induced by swap dynamics, our experiments exposed that they actually behave well in practice.

After being focused on building and analyzing procedures that bring some level of decentralization we focus on envy-freeness and more particularly new relaxations of envy-freeness. Indeed, in Chapter 4, we proposed a way to relax envy-freeness by minimizing some degree of envy while balancing envy (when it cannot be avoided) as fairly as possible among agents thanks to the OWA aggregator. Then we looked at the connection between this new relaxation and other existing ones like EF1 and EFX and showed that theoretically, no algorithm minimizing a fair OWA can ever guarantee to return an EF1 (and thus nor EFX) allocation. However, by implementing our method we saw that there are very few cases for which the returned allocation was not EF1 and that for those few cases we could easily find OWA weights for which it was the case. Not only that but we also figured out that Max Nash Welfare is dominated by our method in regards of returning an EFX, EF or Pareto-Optimal allocations.

Finally, in Chapter 5 following the idea of relaxing envy-freeness we proposed two new relaxations of envy-freeness and proportionality based on the idea that when we cannot guarantee envy-freeness or proportionality, we at least do not want to return an allocation where everyone agrees on the envy or non-proportionality of an agent. We have seen a scale of approval envy and proportionality but also the links with other fairness measures and relaxations. We also show that our approval envy notion implies the one on proportionality while the converse does not hold as it is the case for the classical notions. Besides, implementing algorithms to compute allocations minimizing approval envy and

non-proportionality allow us to see how they behave and seem interesting in order to discriminate between instances when EF or PROP cannot be found.

Now that we have summed up the main contributions of our work let us also talk about what we could derive from it for future perspectives. Fair division is a very active and quickly expanding field. Our ambition here is not to propose an exhaustive overview of recent and emergent topics (for this we refer the reader to the recent surveys [Amanatidis et al., 2023]), but instead to discuss how the questions and results presented in this thesis could be expanded and connected to related notions.

Explainability and Interpretability

First, in Chapters 2 and 3 by using more decentralized procedures we go towards some explainability as these procedures are more interpretable. Indeed, the fact that agents are involved in the building of the solutions facilitates its explanation compared to black-box algorithms. But we could also try to use our work in Chapter 5 to bring a level of explainability as a procedure built around these approval notions could be considered explainable by being constructed by the agents themselves. Indeed, we could first explain the outcome with a contrastive explanation by highlighting that we chose the returned allocation rather than other one (that some agent would have preferred) because it minimises the approval envy or that its approval envy is lower than the one the agent was suggesting. We could also explain the procedure itself with a contrastive explanation by pinpointing that we chose this procedure over some other (than the agent asking for an explanation would have preferred) for some desirable property or axiom that it satisfies (such as individual rationality, strategy-proofness or anonymity among many) while the other procedure fails to. However, it should be kept in mind that axioms are not all equivalent in their simplicity, as illustrated for example by the distinction made between strategy-proofness and obvious strategy proofness, as mentioned in Chapter 3. This idea of building explanations based on axioms is discussed in the context of voting by Boixel and Endriss [2020]; Boixel et al. [2022]; Schmidtlein and Endriss [2023] and in approval-based sorting by [Belahcene et al., 2018]. Hence if some agent is unhappy about the outcome of the procedure we can explain it by rewinding it using the agents’ perspectives as it is through their vote (approval) that the outcome was built. This way of bringing explainability is discussed by Suryanarayana et al. [2022a]; Pozanco et al. [2022]; Suryanarayana et al. [2022b]; Zahedi et al. [2023] and differs from the XAI (eXplainable Artificial Intelligence) whose goal is merely to explain black-box algorithms.

Adaptation to other problems or fields

Another possible extension of our work would be to look beyond the setting studied in this thesis. For instance, a related but very different setting are two-sided markets, meaning that not only agents have preferences over resources but the converse is also true. Hence it is then interesting to try to look for fairness for both sides (“two-sided fairness”). While when resources are objects it might not be natural to consider fairness, for problems like allocating students to schedules/classes or in ranking problems it is the case [Do et al., 2021; Freeman et al., 2021].

When referring to other related problems, another path of research would be to try and apply the procedures or the new fairness relaxation measures that we have proposed

to problems which exhibit specific structure. For instance, this has already been done in the scheduling problem (in which you want to allocate tasks to one or several machines) by using classical voting rules such as Kemeny or Condorcet [Pascual et al., 2018]. In this particular problem, the agents would be the machines while the objects would be the jobs that we have to schedule on the multiple machines. We could for example look for envy-free allocations while trying to minimize the makespan. The difficulty resides in the fact that not only we have to look at the classical fairness measures of the fair division problem but also the constraints and objectives of a scheduling problem such as minimizing the makespan or the fact that some job has to start before some time and finish before some other time. Fair division under constraints has seen many recent developments [Suksompong, 2021].

In the same context, data science involves the processing of massive quantities of data requiring distributed file system and parallel programming which brings the problem of task allocation and load balancing. Baert et al. [2021] address the practical problem where the data required to execute a task are split between nodes that can themselves experience execution hazards or communication lags. Hence, in order to execute a task it is necessary to fetch data from several other nodes, resulting in an extra execution time of it. In this class of problems, MapReduce is the most significant distributed data processing models. However, several data skews in the MapReduce applications can induce an unbalanced workload during the reduce phase and an unfair allocation can arise during the reduce phase because of the heterogeneous performance of nodes. In order to face these problems, Dean and Ghemawat [2008] formalize this as a multiagent task allocation problem and propose a dynamic task reallocation process based on negotiations between agents. It would then be interesting to see if we would apply our decentralized procedures in this context of load balancing, specifically with our approval notion and our resource exchange setting as negotiations seem to be relevant specially as Beauprez et al. [2022] have already looked at swap-deals in this context.

More generally, it could be interesting to take an overview and walk outside of multi-agent decision making and try to see what our notions and procedures could mean in the multi-criteria decision field or decision under risk and uncertainty. Indeed, for example in the domain of decision under risk and uncertainty we have first-order (SD1) and second-order Stochastic Dominance (SD2) [Hadar and Russell, 1971]. In the multi-criteria decision domain, if you consider that the different criteria are commensurable and expressed on a common scale (a utility scale for example), then the pendant of SD2 is the weighted Lorenz Dominance [Gonzales and Perny, 2020]. Hence using tools from our field and adapting them to other fields could be interesting. Indeed, there are already collective utility functions that are used in several fields such as Max Nash Welfare or OWA of the utilities.

So in our case, would it make sense to say that a criterion envies another one or that a scenario envies another scenario? If it is the case then one could also wonder if we could also extend it to the relaxations we have studied. Our work could then help solve or bring a new perspective to other problems from other fields.

Moreover, we could also think of making a bridge between decision under uncertainty and fair division of indivisible goods in the context of online fair division [Aleksandrov et al., 2015b; Aleksandrov and Walsh, 2020]. They consider the problem of food banks that want to fairly (and efficiently) allocate food among different charities but without

knowing how much food and when they will get. Hence the modelization of this problem as an online one. By considering probabilities on the event of getting food (and how much) we would switch to decision under uncertainty and could use the tools from this field. It would then be interesting to see if we can for example achieve ex-ante and ex-post envy-freeness at the same time as discussed (in the classical way) by Freeman et al. [2020].

Group Fairness

Finally, another widely studied subject that is very close to what we tackled in this thesis is the fair division of indivisible goods but not among simple agents but among groups [Conitzer et al., 2019; Benabbou et al., 2019; Aziz and Rey, 2020; Gross-Humbert et al., 2022; Aziz et al., 2023; Scarlett et al., 2023; Gross-Humbert et al., 2023].

Hence it would be natural to try to use our work that is at the agent level and try to adapt it to the group level. For example, the notion of EF, EFX and EF1 are redefined in the context of groups and it would be interesting to see if our approval-based notions of envy make sense in this context. For instance, in House Allocation Problems with groups of agents, we could define in a straightforward way approval envy of a group if every group agrees on the fact that she should be envious. We could also use approval-based envy inside groups to decide whether a group is envious or not: for example a group is envious if a certain number of agents agree on the fact that an agent (or a certain number) is envious. This could give scales of approval fairness in this context as we did in ours. However note that a big challenge regarding this adaptation is the fact that the groups have different sizes. It is then necessary to think about how we could use OWA given the difference in the number of agents involved in each group. Moreover this number disparity implies a question about the approval notion. Is unanimous envy attained when all the agents agree on it or when all groups agree? Even envy between groups can be defined in several ways. For example, we could simply divide the utility of a group by its size but Gross-Humbert et al. [2023] propose a sharper way in which envy depends on a subset of objects of the bigger group.

Bibliography

- Abebe, R., Kleinberg, J., and Parkes, D. C. (2017). Fair division via social comparison. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS '17*, page 281–289, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems. (Cited page 93.)
- Abraham, D. J., Cechlárová, K., Manlove, D. F., and Mehlhorn, K. (2005). Pareto optimality in house allocation problems. In Fleischer, R. and Trippen, G., editors, *Algorithms and Computation: 15th International Symposiums*, pages 3–15. Springer Berlin Heidelberg, Berlin, Heidelberg. (Cited pages 15, 20, and 113.)
- Aleksandrov, M., Aziz, H., Gaspers, S., and Walsh, T. (2015a). Online fair division: Analysing a food bank problem. In *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI'15*, page 2540–2546. AAAI Press. (Cited page 25.)
- Aleksandrov, M., Aziz, H., Gaspers, S., and Walsh, T. (2015b). Online fair division: Analysing a food bank problem. In Yang, Q. and Wooldridge, M. J., editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 2540–2546. AAAI Press. (Cited page 127.)
- Aleksandrov, M. and Walsh, T. (2020). Online fair division: A survey. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 13557–13562. AAAI Press. (Cited page 127.)
- Amanatidis, G., Aziz, H., Birmpas, G., Filos-Ratsikas, A., Li, B., Moulin, H., Voudouris, A. A., and Wu, X. (2023). Fair division of indivisible goods: Recent progress and open questions. *Artificial Intelligence*, 322:103965. (Cited page 126.)
- Amanatidis, G., Birmpas, G., Filos-Ratsikas, A., Hollender, A., and Voudouris, A. A. (2020). Maximum nash welfare and other stories about EFX. In Bessiere, C., editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 24–30, Yokohama, Japan (relocated and postponed). ijcai.org. (Cited page 30.)
- Amanatidis, G., Birmpas, G., and Markakis, V. (2018). Comparing approximate relaxations of envy-freeness. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 42–48. International Joint Conferences on Artificial Intelligence Organization. (Cited pages 31 and 94.)

- Amanatidis, G., Ntokos, A., and Markakis, E. (2019). Multiple birds with one stone: Beating $1/2$ for EFX and GMMS via envy cycle elimination. ArXiv-1909.07650. (Cited page 99.)
- Anshelevich, E., Das, S., and Naamad, Y. (2013). Anarchy, stability, and utopia: creating better matchings. *Autonomous Agents and Multi-Agent Systems*, 26(1):120–140. (Cited page 68.)
- Arrow, K. J. (1951). *Social choice and individual values*. Wiley. (Cited page 23.)
- Aziz, H., Biró, P., Lang, J., Lesca, J., and Monnot, J. (2016a). Optimal reallocation under additive and ordinal preferences. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'16)*, pages 402–410, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems. (Cited page 44.)
- Aziz, H., Bouveret, S., Caragiannis, I., Giagkousi, I., and Lang, J. (2018). Knowledge, fairness, and social constraints. In *Proceedings of the 32nd AAAI conference on Artificial Intelligence (AAAI'18)*, New Orleans, Louisiana, USA. AAAI Press. (Cited page 93.)
- Aziz, H., Caragiannis, I., Igarashi, A., and Walsh, T. (2019). Fair allocation of indivisible goods and chores. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 53–59. International Joint Conferences on Artificial Intelligence Organization. (Cited page 31.)
- Aziz, H., Gaspers, S., Mackenzie, S., and Walsh, T. (2015a). Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92. (Cited page 44.)
- Aziz, H., Kalinowski, T., Walsh, T., and Xia, L. (2016b). Welfare of sequential allocation mechanisms for indivisible goods. In Kaminka, G. A., Fox, M., Bouquet, P., Hüllermeier, E., Dignum, V., Dignum, F., and van Harmelen, F., editors, *22nd European Conference on Artificial Intelligence (ECAI 2016)*, pages 787–794. IOS Press. (Cited pages 43 and 45.)
- Aziz, H., Micha, E., and Shah, N. (2023). Group fairness in peer review. In Agmon, N., An, B., Ricci, A., and Yeoh, W., editors, *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, London, United Kingdom, 29 May 2023 - 2 June 2023*, pages 2889–2891. ACM. (Cited page 128.)
- Aziz, H., Moulin, H., and Sandomirskiy, F. (2020). A polynomial-time algorithm for computing a pareto optimal and almost proportional allocation. *Operations Research Letters*, 48(5):573 – 578. (Cited pages 31 and 103.)
- Aziz, H. and Rey, S. (2020). Almost group envy-free allocation of indivisible goods and chores. In Bessiere, C., editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 39–45. ijcai.org. (Cited page 128.)
- Aziz, H., Walsh, T., and Xia, L. (2015b). Possible and necessary allocations via sequential mechanisms. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*, pages 468–474. AAAI Press. (Cited pages 39 and 42.)

- Babaioff, M., Nisan, N., and Talgam-Cohen, I. (2017). Competitive equilibria with indivisible goods and generic budgets. *arXiv preprint arXiv:1703.08150*. (Cited page 50.)
- Bade, S. (2019). Matching with single-peaked preferences. *Journal of Economic Theory*, 180:81–99. (Cited pages 54, 55, and 80.)
- Baert, Q., Caron, A., Morge, M., Routier, J., and Stathis, K. (2021). An adaptive multi-agent system for task reallocation in a mapreduce job. *J. Parallel Distributed Comput.*, 153:75–88. (Cited page 127.)
- Ballester, M. A. and Haeringer, G. (2011). A characterization of the single-peaked domain. *Social Choice and Welfare*, 36(2):305–322. (Cited page 24.)
- Barberà, S., Bossert, W., and Pattanaik, P. K. (2004). *Ranking Sets of Objects*, pages 893–977. Springer US, Boston, MA. (Cited page 21.)
- Barman, S. and Krishnamurthy, S. (2019). On the proximity of markets with integral equilibria. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33:1748–1755. (Cited pages 31 and 103.)
- Bartholdi, J. J., Tovey, C. A., and Trick, M. A. (1989). Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6:157–165. (Cited page 12.)
- Baumeister, D., Bouveret, S., Lang, J., Nguyen, N.-T., Nguyen, T. T., and Rothe, J. (2014). Scoring Rules for the Allocation of Indivisible Goods. In *21st European Conference on Artificial Intelligence (ECAI’14)*, pages 75–80, Prague, Czech Republic. (Cited page 12.)
- Beauprez, E., Caron, A., Morge, M., and Routier, J. (2022). Échange de tâches pour la réduction de la durée moyenne de réalisation. In Camps, V., editor, *SMA et Smart Cities - Trentièmes Journées Francophones sur les Systèmes Multi-Agents, JFSMA 2022, Saint-Etienne, France, June 27-29, 2022*, pages 19–28. Cépaduès. (Cited page 127.)
- Bei, X., Qiao, Y., and Zhang, S. (2017). Networked fairness in cake cutting. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 3632–3638. (Cited page 93.)
- Belahcene, K., Chevaleyre, Y., Maudet, N., Labreuche, C., Mousseau, V., and Ouerdane, W. (2018). Accountable Approval Sorting. In *27th International Joint Conference on Artificial Intelligence and 23rd European Conference on Artificial Intelligence (IJCAI-ECAI 2018)*, Stockholm, Sweden. (Cited page 126.)
- Benabbou, N., Chakraborty, M., Elkind, E., and Zick, Y. (2019). Fairness towards groups of agents in the allocation of indivisible items. In Kraus, S., editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 95–101. ijcai.org. (Cited page 128.)
- Bensana, E., Lemaître, M., and Verfaillie, G. (1999). Earth observation satellite management. *Constraints An Int. J.*, 4(3):293–299. (Cited page 15.)

- Berger, B., Cohen, A., Feldman, M., and Fiat, A. (2022). Almost full efx exists for four agents. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(5):4826–4833. (Cited page 31.)
- Beynier, A., Bouveret, S., Lemaître, M., Maudet, N., Rey, S., and Shams, P. (2019a). Efficiency, sequenceability and deal-optimality in fair division of indivisible goods. In Elkind, E., Veloso, M., Agmon, N., and Taylor, M. E., editors, *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019*, pages 900–908. International Foundation for Autonomous Agents and Multiagent Systems. (Cited pages 41, 62, and 99.)
- Beynier, A., Chevaleyre, Y., Gourvès, L., Harutyunyan, A., Lesca, J., Maudet, N., and Wilczynski, A. (2019b). Local Envy-Freeness in House Allocation Problems. *Autonomous Agents and Multi-Agent Systems*, 33(5):591–627. (Cited pages 15 and 93.)
- Beynier, A., Maudet, N., Rey, S., and Shams, P. (2020). An optimal procedure to check pareto-optimality in house markets with single-peaked preferences. ArXiv 2002.11660. (Cited page 56.)
- Beynier, A., Maudet, N., Rey, S., and Shams, P. (2021). Swap dynamics in single-peaked housing markets. *Autonomous Agents Multi Agent Syst.*, 35(2):20. (Cited page 53.)
- Bianchessi, N., Cordeau, J.-F., Desrosiers, J., Laporte, G., and Raymond, V. (2007). A heuristic for the multi-satellite, multi-orbit and multi-user management of earth observation satellites. *European Journal of Operational Research*, 177:750–762. (Cited page 15.)
- Biswas, A. and Barman, S. (2018). Fair division under cardinality constraints. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 91–97. International Joint Conferences on Artificial Intelligence Organization. (Cited page 99.)
- Black, D. (1948). On the rationale of group decision-making. *Journal of political economy*, 56(1):23–34. (Cited page 23.)
- Black, D., Newing, R. A., McLean, I., McMillan, A., and Monroe, B. L. (1958). *The theory of committees and elections*. Springer. (Cited page 118.)
- Boixel, A. and Endriss, U. (2020). Automated justification of collective decisions via constraint solving. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '20*, page 168–176, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems. (Cited page 126.)
- Boixel, A., Endriss, U., and Haan, R. d. (2022). A calculus for computing structured justifications for election outcomes. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(5):4859–4866. (Cited page 126.)
- Bouveret, S., Endriss, U., and Lang, J. (2010). Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In Coelho, H., Studer, R., and

- Wooldridge, M. J., editors, *ECAI 2010 - 19th European Conference on Artificial Intelligence, Lisbon, Portugal, August 16-20, 2010, Proceedings*, volume 215 of *Frontiers in Artificial Intelligence and Applications*, pages 387–392. IOS Press. (Cited page 21.)
- Bouveret, S. and Lang, J. (2011). A general elicitation-free protocol for allocating indivisible goods. In Walsh, T., editor, *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, pages 73–78, Barcelona, Spain. IJCAI/AAAI. (Cited page 39.)
- Bouveret, S. and Lemaître, M. (2016a). Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2):259–290. (Cited pages 26, 27, 29, 31, 48, 50, 51, 94, 103, 107, 110, 117, and 118.)
- Bouveret, S. and Lemaître, M. (2016b). Efficiency and sequenceability in fair division of indivisible goods with additive preferences. *Proceedings of the Sixth International Workshop on Computational Social Choice (COMSOC'16)*. (Cited pages 43, 44, 45, 48, and 50.)
- Brams, S. J. and Fishburn, P. C. (2002). Fair division of indivisible items between two people with identical preferences: Envy-freeness, pareto-optimality, and equity. *Social Choice and Welfare*, 17:247–267. (Cited page 27.)
- Brams, S. J. and King, D. (2005). Efficient fair division—help the worst off or avoid envy? *Rationality and Society*, 17(4):387–421. (Cited pages 39, 42, 43, and 45.)
- Brams, S. J. and Taylor, A. D. (1996). *Fair Division — From Cake-cutting to Dispute Resolution*. Cambridge University Press. (Cited pages 13, 15, and 20.)
- Brams, S. J. and Taylor, A. D. (2000). *The Win-win Solution. Guaranteeing Fair Shares to Everybody*. W. W. Norton & Company. (Cited page 39.)
- Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A. D., and Moulin, H. (2016). *Handbook of Computational Social Choice*. Cambridge University Press. (Cited page 13.)
- Brânzei, S., Hosseini, H., and Miltersen, P. B. (2015). Characterization and computation of equilibria for indivisible goods. In *International Symposium on Algorithmic Game Theory*, pages 244–255. Springer. (Cited pages 29, 50, and 51.)
- Bredereck, R., Kaczmarczyk, A., and Niedermeier, R. (2018). Envy-free allocations respecting social networks. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '18*, page 283–291, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems. (Cited page 93.)
- Brânzei, S. and Sandomirskiy, F. (2019). Algorithms for competitive division of chores. (Cited pages 31 and 103.)
- Bu, X., Song, J., and Yu, Z. (2023). Efx allocations exist for binary valuations. In Li, M., Sun, X., and Wu, X., editors, *Frontiers of Algorithmics*, pages 252–262, Cham. Springer Nature Switzerland. (Cited page 31.)

- Budish, E. (2011). The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103. (Cited pages 13, 15, 20, 26, 29, 30, 31, and 93.)
- Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., and Wang, J. (2016). The unreasonable fairness of maximum nash welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 305–322, New York, NY, USA. ACM. (Cited pages 13, 14, 22, 30, 34, 86, 88, and 93.)
- Chaudhury, B. R., Garg, J., and Mehlhorn, K. (2020). Efx exists for three agents. In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC '20, page 1–19, New York, NY, USA. Association for Computing Machinery. (Cited pages 30 and 88.)
- Chevaleyre, Y., Dunne, P. E., Endriss, U., Lang, J., Lemaître, M., Maudet, N., Padget, J., Phelps, S., Rodriguez-Aguilar, J. A., and Sousa, P. (2006). Issues in multiagent resource allocation. *Informatica*, 30(1). (Cited page 31.)
- Chevaleyre, Y., Endriss, U., Estivie, S., and Maudet, N. (2007). Reaching envy-free states in distributed negotiation settings. In *2007 IJCAI Conference*, volume 7, pages 1239–1244. (Cited page 37.)
- Chevaleyre, Y., Endriss, U., Estivie, S., and Maudet, N. (2008a). Multiagent resource allocation in k-additive domains: preference representation and complexity. *Annals of Operations Research*, 163(1):49–62. (Cited page 22.)
- Chevaleyre, Y., Endriss, U., Lang, J., and Maudet, N. (2008b). Preference handling in combinatorial domains: From ai to social choice. *AI Magazine*, 29(4):37. (Cited page 21.)
- Chevaleyre, Y., Endriss, U., and Maudet, N. (2010). Simple negotiation schemes for agents with simple preferences: Sufficiency, necessity and maximality. *Autonomous Agents and Multi-Agent Systems*, 20(2):234–259. (Cited page 67.)
- Chevaleyre, Y., Endriss, U., and Maudet, N. (2017). Distributed fair allocation of indivisible goods. *Artificial Intelligence*, 242:1–22. (Cited pages 28, 37, and 93.)
- Ciná, G. and Endriss, U. (2016). Proving classical theorems of social choice theory in modal logic. *Autonomous Agents and Multi-Agent Systems*, 30(5):963–989. (Cited page 12.)
- Condorcet, N. d. (1785). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Cambridge Library Collection - Mathematics. Cambridge University Press. (Cited page 12.)
- Conitzer, V. (2009). Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research*, 35:161–191. (Cited page 72.)
- Conitzer, V., Freeman, R., and Shah, N. (2017). Fair public decision making. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, page

- 629–646, New York, NY, USA. Association for Computing Machinery. (Cited pages 14, 31, 94, and 103.)
- Conitzer, V., Freeman, R., Shah, N., and Vaughan, J. W. (2019). Group fairness for the allocation of indivisible goods. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 1853–1860. AAAI Press. (Cited page 128.)
- Damamme, A., Beynier, A., Chevaleyre, Y., and Maudet, N. (2015). The power of swap deals in distributed resource allocation. In *2015 International Conference on Autonomous Agents And Multiagent Systems*, pages 625–633. (Cited pages 37, 48, 49, 54, 68, and 69.)
- de Keijzer, B., Bouveret, S., Klos, T., and Zhang, Y. (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, pages 98–110, Venice, Italy. Springer Verlag. (Cited page 27.)
- de Keijzer, B., Bouveret, S., Klos, T., and Zhang, Y. (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, pages 98–110, Venice, Italy. Springer Verlag. (Cited page 44.)
- Dean, J. and Ghemawat, S. (2008). Mapreduce: Simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113. (Cited page 127.)
- Deemen, A. V. (2014). On the empirical relevance of condorcet’s paradox. *Public Choice*, 158(3/4):311–330. (Cited page 25.)
- DeMarzo, P. M., Vayanos, D., and Zwiebel, J. (2003). Persuasion bias, social influence, and unidimensional opinions. *The Quarterly journal of economics*, 118(3):909–968. (Cited page 23.)
- Dickerson, J. P., Goldman, J., Karp, J., Procaccia, A. D., and Sandholm, T. (2014). The computational rise and fall of fairness. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI-14)*, pages 1405–1411, Québec City, Québec, Canada. AAAI Press. (Cited pages 22, 90, and 118.)
- Do, V., Corbett-Davies, S., Atif, J., and Usunier, N. (2021). Two-sided fairness in rankings via lorenz dominance. In Ranzato, M., Beygelzimer, A., Dauphin, Y. N., Liang, P., and Vaughan, J. W., editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 8596–8608. (Cited page 126.)

- Durand, F., Mathieu, F., and Noirie, L. (2016). SVVAMP: simulator of various voting algorithms in manipulating populations. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 4353–4354. (Cited pages 25, 117, and 122.)
- Edmonds, J. (1965). Maximum matching and a polyhedron with 0,1-vertices. *Journal of Research of the National Bureau of Standards Section B Mathematics and Mathematical Physics*, page 125. (Cited page 91.)
- Egan, P. J. (2014). “do something” politics and double-peaked policy preferences. *The Journal of Politics*, 76(2):333–349. (Cited page 23.)
- Elkind, E., Lackner, M., and Peters, D. (2017). Structured preferences. In E., U., editor, *Trends in Computational Social Choice*, pages 187–207. Lulu.com. (Cited page 23.)
- Endriss, U., Maudet, N., Sadri, F., and Toni, F. (2006). Negotiating socially optimal allocations of resources. *Journal of artificial intelligence research*. (Cited page 37.)
- Fargier, H., Lang, J., Lemaître, M., and Verfaillie, G. (2004). Partage équitable de ressources communes. (1) un modèle général et son application au partage de ressources satellitaires. *Tech. Sci. Informatiques*, 23(9):1187–1217. (Cited page 15.)
- Feldman, M., Mauras, S., and Ponitka, T. (2023). On optimal tradeoffs between efx and nash welfare. *ArXiv*, abs/2302.09633. (Cited page 31.)
- Fisher, I. (1892). *Mathematical Investigations in the Theory of Value and Prices, and Appreciation and Interest*. Augustus M. Kelley, Publishers. (Cited page 29.)
- Fisher, M. L. (1981). The lagrangian relaxation method for solving integer programming problems. *Management Science*, 27(1):1–18. (Cited page 91.)
- Foley, D. K. (1967). Resource allocation and the public sector. *Yale Economic Essays*, 7(1):45–98. (Cited pages 13 and 27.)
- Freeman, R., Micha, E., and Shah, N. (2021). Two-sided matching meets fair division. In *International Joint Conference on Artificial Intelligence*. (Cited page 126.)
- Freeman, R., Shah, N., and Vaish, R. (2020). Best of both worlds: Ex-ante and ex-post fairness in resource allocation. *Proceedings of the 21st ACM Conference on Economics and Computation*. (Cited page 128.)
- Fürnkranz, J. and Hüllermeier, E. (2010). Preference learning and ranking by pairwise comparison. In *Preference learning*, pages 65–82. Springer. (Cited page 12.)
- Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15. (Cited page 15.)
- Gan, J., Suksompong, W., and Voudouris, A. A. (2019). Envy-freeness in house allocation problems. *Mathematical Social Sciences*, 101(C):104–106. (Cited page 15.)
- Garfinkel, R. S. (1971). An improved algorithm for the bottleneck assignment problem. *Operations Research*, 19(7):1747–1751. (Cited page 73.)

- Garman, M. B. and Kamien, M. I. (1968). The paradox of voting: Probability calculations. *Behavioral Science*, 13(4):306–316. (Cited page 25.)
- Gehrlein, W. V. and Fishburn, P. C. (1976). The probability of the paradox of voting: A computable solution. *Journal of Economic Theory*, 13(1):14 – 25. (Cited page 118.)
- Geist, C. and Peters, D. (2017). Chapter 13 computer-aided methods for social choice theory. In *Trends in Computational Social Choice*. (Cited page 12.)
- Geoffrion, A. (1974). Lagrangean relaxation for integer programming. balinski m, ed. *Approaches to Integer Programming*, pages 82–114. (Cited page 91.)
- Goldman, J. and Procaccia, A. D. (2015). Spliddit: Unleashing fair division algorithms. *SIGecom Exch.*, 13(2):41–46. (Cited pages 88, 89, and 117.)
- Gondran, M. and Minoux, M. (1984). *Graphs and algorithms*. Wiley-Interscience series in discrete mathematics. Wiley. (Cited pages 114 and 116.)
- Gonzales, C. and Perny, P. (2020). Multicriteria decision making. *A Guided Tour of Artificial Intelligence Research*. (Cited page 127.)
- Gourvès, L., Monnot, J., and Tlilane, L. (2014). Near fairness in matroids. In Schaub, T., Friedrich, G., and O’Sullivan, B., editors, *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, pages 393–398. IOS Press. (Cited page 93.)
- Grabisch, M. (1997). k -order additive discrete fuzzy measure and their representation. *Fuzzy Sets and Systems*, 92:167–189. (Cited page 22.)
- Grandi, U. and Endriss, U. (2013). First-order logic formalisation of impossibility theorems in preference aggregation. *Journal of Philosophical Logic*, 42(4):595–618. (Cited page 12.)
- Gross-Humbert, N., Benabbou, N., Beynier, A., and Maudet, N. (2022). Sequential and swap mechanisms for public housing allocation with quotas and neighbourhood-based utilities. *ACM Trans. Economics and Comput.*, 10(4):15:1–15:24. (Cited pages 15 and 128.)
- Gross-Humbert, N., Benabbou, N., Beynier, A., and Maudet, N. (2023). On the notion of envy among groups of agents in house allocation problems. In Gal, K., Nowé, A., Nalepa, G. J., Fairstein, R., and Radulescu, R., editors, *ECAI 2023 - 26th European Conference on Artificial Intelligence, September 30 - October 4, 2023, Kraków, Poland - Including 12th Conference on Prestigious Applications of Intelligent Systems (PAIS 2023)*, volume 372 of *Frontiers in Artificial Intelligence and Applications*, pages 924–931. IOS Press. (Cited page 128.)
- Gross-Humbert, N., Benabbou, N., Beynier, A., and Maudet, N. (2023). Sequential and swap mechanisms for public housing allocation with quotas and neighbourhood-based utilities. *ACM Trans. Econ. Comput.*, 10(4). (Cited page 15.)

- Guibet Lafaye, C. (2006). *Justice sociale et éthique individuelle*. Presses Université Laval. (Cited pages 16 and 95.)
- Hadar, J. and Russell, W. R. (1971). Stochastic dominance and diversification. *Journal of Economic Theory*, 3(3):288–305. (Cited page 127.)
- Halpern, D., Procaccia, A. D., Psomas, A., and Shah, N. (2020). Fair division with binary valuations: One rule to rule them all. In *Workshop on Internet and Network Economics*. (Cited page 25.)
- Hashemi, V. and Endriss, U. (2014). Measuring diversity of preferences in a group. In *Proceedings of the 21st European Conference on Artificial Intelligence*, pages 423–428. (Cited page 72.)
- Heinen, T., Nguyen, N.-T., and Rothe, J. (2015). Fairness and rank-weighted utilitarianism in resource allocation. In Walsh, T., editor, *Algorithmic Decision Theory*, pages 521–536, Cham. Springer International Publishing. (Cited page 83.)
- Hougaard, J. L., Moreno-Ternero, J. D., and Østerdal, L. P. (2014). Assigning agents to a line. *Games and Economic Behavior*, 87:539–553. (Cited page 23.)
- Igarashi, A. and Yokoyama, T. (2023). Kajibuntan: A house chore division app. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(13):16449–16451. (Cited page 15.)
- IPCC (2023). *Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, pages 35–115. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA. (Cited page 11.)
- Irving, R. W. (1985). An efficient algorithm for the “stable roommates” problem. *Journal of Algorithms*, 6(4):577–595. (Cited page 15.)
- Kaczmarczyk, A. (2021). *Algorithmic aspects of resource allocation and multiwinner voting: theory and experiments*. PhD thesis, Technical University of Berlin, Germany. (Cited page 91.)
- Kalinowski, T., Narodytska, N., and Walsh, T. (2013a). A social welfare optimal sequential allocation procedure. In Rossi, F., editor, *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI-13)*, pages 227–233, Beijing, China. IJCAI/AAAI. (Cited page 39.)
- Kalinowski, T., Narodytska, N., Walsh, T., and Xia, L. (2013b). Strategic behavior when allocating indivisible goods sequentially. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI-12)*, pages 452–458, Bellevue, WA. AAAI Press. (Cited page 39.)
- Kasajima, Y. (2013). Probabilistic assignment of indivisible goods with single-peaked preferences. *Social Choice and Welfare*, 41(1):203–215. (Cited page 23.)

- Kohler, D. A. and Chandrasekaran, R. (1971). A class of sequential games. *Operations Research*, 19(2):270–277. (Cited page 38.)
- Kondratev, A. Y. and Nesterov, A. S. (2019). Minimal envy and popular matchings. arXiv preprint arXiv:1902.08003. (Cited page 80.)
- Koutsoupias, E. and Papadimitriou, C. (1999). Worst-case equilibria. In Meinel, C. and Tison, S., editors, *STACS 99*, pages 404–413, Berlin, Heidelberg. Springer Berlin Heidelberg. (Cited page 68.)
- Kyropoulou, M., Suksompong, W., and Voudouris, A. A. (2020). Almost envy-freeness in group resource allocation. *Theoretical Computer Science*, 841:110 – 123. (Cited page 30.)
- Lee, J.-S., Filatova, T., Ligmann-Zielinska, A., Hassani-Mahmooei, B., Stonedahl, F., Lorscheid, I., Voinov, A., Polhill, J. G., Sun, Z., and Parker, D. C. (2015). The complexities of agent-based modeling output analysis. *Journal of Artificial Societies and Social Simulation*, 18(4). (Cited page 79.)
- Lesca, J., Minoux, M., and Perny, P. (2019). The fair OWA one-to-one assignment problem: Np-hardness and polynomial time special cases. *Algorithmica*, 81(1):98–123. (Cited page 83.)
- Leventhal, G. S. (1980). What should be done with equity theory? In *Social exchange*, pages 27–55. Springer. (Cited page 37.)
- Lipton, R., Markakis, E., Mossel, E., and Saberi, A. (2004). On approximately fair allocations of divisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC-04)*, pages 125–131, New York, NY. ACM. (Cited pages 13, 22, 27, 29, 30, 37, 82, 96, 99, 101, and 102.)
- List, C., Luskin, R. C., Fishkin, J. S., and McLean, I. (2013). Deliberation, single-peakedness, and the possibility of meaningful democracy: evidence from deliberative polls. *The journal of politics*, 75(1):80–95. (Cited page 23.)
- Lovász, L. and Plummer, M. D. (2009). *Matching theory*, volume 367. American Mathematical Soc. (Cited page 114.)
- Ma, J. (1994). Strategy-proofness and the strict core in a market with indivisibilities. *International Journal of Game Theory*, 23(1):75–83. (Cited page 54.)
- Mallows, C. L. (1957). Non-Null Ranking Models. I. *Biometrika*, 44(1-2):114–130. (Cited pages 25 and 122.)
- Mattei, N. and Walsh, T. (2013). Preflib: A library for preferences <http://www.preflib.org>. In *2013 Algorithmic Decision Theory Conference*, pages 259–270. Springer. (Cited page 80.)
- Moulin, H. (1988). *Axioms of Cooperative Decision Making*. Econometric Society Monographs. Cambridge University Press. (Cited page 23.)

- Moulin, H. (1990). Uniform externalities: Two axioms for fair allocation. *Journal of Public Economics*, 43(3):305–326. (Cited page 26.)
- Moulin, H. (2003). *Fair Division and Collective Welfare*. MIT Press. (Cited pages 20 and 29.)
- Nguyen, N. T., Nguyen, T. T., Roos, M., and Rothe, J. (2014). Computational complexity and approximability of social welfare optimization in multiagent resource allocation. *Autonomous agents and multi-agent systems*, 28(2):256–289. (Cited page 34.)
- Ogryczak, W. and Śliwiński, T. (2003). On solving linear programs with the ordered weighted averaging objective. *European Journal of Operational Research*, 148:80–91. (Cited page 84.)
- Ogryczak, W. and Tamir, A. (2003). Minimizing the sum of the k largest functions in linear time. *Information Processing Letters*, 85(3):117 – 122. (Cited page 84.)
- Othman, A., Sandholm, T., and Budish, E. (2010). Finding approximate competitive equilibria: efficient and fair course allocation. In van der Hoek, W., Kaminka, G. A., Lespérance, Y., Luck, M., and Sen, S., editors, *Proceedings of the 9th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-10)*, pages 873–880, Toronto, Canada. IFAAMAS. (Cited pages 15, 20, and 29.)
- Parijs, P. V. (1997). *Real Freedom for All: What (if Anything) Can Justify Capitalism?* Number 9780198293576 in OUP Catalogue. Oxford University Press. (Cited pages 16, 93, and 95.)
- Pascual, F., Rzdadca, K., and Skowron, P. (2018). Collective schedules: Scheduling meets computational social choice. In André, E., Koenig, S., Dastani, M., and Sukthankar, G., editors, *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, pages 667–675. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM. (Cited pages 15 and 127.)
- Pazner, E. A. and Schmeidler, D. (1978). Egalitarian equivalent allocations: A new concept of economic equity. *The Quarterly Journal of Economics*, 92(4):671–687. (Cited page 34.)
- Perny, P. and Spanjaard, O. (2003). An axiomatic approach to robustness in search problems with multiple scenarios. In Meek, C. and Kjærulff, U., editors, *UAI '03, Proceedings of the 19th Conference in Uncertainty in Artificial Intelligence, Acapulco, Mexico, August 7-10 2003*, pages 469–476. Morgan Kaufmann. (Cited pages 34 and 84.)
- Perry, J. and Powers, R. C. (2010). Anonymity, monotonicity, and quota pair systems. *Mathematical Social Sciences*, 60(1):57–60. (Cited page 97.)
- Picard, G. (2022). Auction-based and distributed optimization approaches for scheduling observations in satellite constellations with exclusive orbit portions. In Faliszewski, P., Mascardi, V., Pelachaud, C., and Taylor, M. E., editors, *21st International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2022, Auckland, New*

- Zealand, May 9-13, 2022*, pages 1056–1064. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS). (Cited page 15.)
- Plaut, B. and Roughgarden, T. (2018). Almost envy-freeness with general valuations. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18, page 2584–2603, USA. Society for Industrial and Applied Mathematics. (Cited pages 30 and 31.)
- Pozanco, A., Mosca, F., Zehtabi, P., Magazzeni, D., and Kraus, S. (2022). Explaining preference-driven schedules: The expres framework. *Proceedings of the International Conference on Automated Planning and Scheduling*, 32(1):710–718. (Cited page 126.)
- Procaccia, A. D. and Wang, J. (2014). Fair enough: guaranteeing approximate maximin shares. In Babaioff, M., Conitzer, V., and Easley, D. A., editors, *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*, pages 675–692. ACM. (Cited pages 22 and 30.)
- Puppe, C. (2018). The single-peaked domain revisited: A simple global characterization. *Journal of Economic Theory*, 176(C):55–80. (Cited page 23.)
- Rawls, J. (1971). *A Theory of Justice*. Harvard University Press, Cambridge, Mass. Traduction française disponible aux éditions du Seuil. (Cited pages 34 and 82.)
- Rosenschein, J. S. and Zlotkin, G. (1994). *Rules of encounter: designing conventions for automated negotiation among computers*. MIT press. (Cited page 37.)
- Roth, A. and Sotomayor, M. (1992). Two-sided matching: A study in game-theoretic modeling and analysis. *Games and Economic Behavior*, 4(1):161–165. (Cited page 113.)
- Roth, A. E. (1982). Incentive compatibility in a market with indivisible goods. *Economics letters*, 9(2):127–132. (Cited page 54.)
- Roth, A. E., Sönmez, T., and Ünver, M. U. (2005). Pairwise kidney exchange. *Journal of Economic theory*, 125(2):151–188. (Cited page 54.)
- Roussel, S., Picard, G., Pralet, C., and Maqrot, S. (2023). Conflicting bundle allocation with preferences in weighted directed acyclic graphs: Application to orbit slot allocation problems. *Systems*, 11(6). (Cited page 15.)
- Sandholm, T. W. (1998). Contract types for satisficing task allocation: I. theoretical results. In Sen, S., editor, *Proceedings of the AAI Spring Symposium: Satisficing Models*, pages 68–75, Menlo Park, California. AAI Press. (Cited page 37.)
- Scarlett, J., Teh, N., and Zick, Y. (2023). For one and all: Individual and group fairness in the allocation of indivisible goods. In Agmon, N., An, B., Ricci, A., and Yeoh, W., editors, *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, London, United Kingdom, 29 May 2023 - 2 June 2023*, pages 2466–2468. ACM. (Cited page 128.)

- Schmidtlein, M. C. and Endriss, U. (2023). Voting by axioms. In Agmon, N., An, B., Ricci, A., and Yeoh, W., editors, *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, London, United Kingdom, 29 May 2023 - 2 June 2023*, pages 2067–2075. ACM. (Cited pages 33 and 126.)
- Segal-Halevi, E. (2020). Competitive equilibrium for almost all incomes: existence and fairness. *Autonomous Agents and Multi-Agent Systems*, 34. (Cited page 31.)
- Segal-Halevi, E., Nitzan, S., Hassidim, A., and Aumann, Y. (2020). Envy-free division of land. *Math. Oper. Res.*, 45(3):896–922. (Cited page 13.)
- Segal-Halevi, E. and Suksompong, W. (2019). Democratic fair allocation of indivisible goods. *Artificial Intelligence*, 277. (Cited pages 27 and 97.)
- Sen, A. K. (1966). A possibility theorem on majority decisions. *Econometrica*, pages 491–499. (Cited page 24.)
- Shams, P., Beynier, A., Bouveret, S., and Maudet, N. (2020). Fair in the eyes of others. In Giacomo, G. D., Catalá, A., Dilkina, B., Milano, M., Barro, S., Bugarín, A., and Lang, J., editors, *ECAI 2020 - 24th European Conference on Artificial Intelligence, 29 August-8 September 2020, Santiago de Compostela, Spain, August 29 - September 8, 2020 - Including 10th Conference on Prestigious Applications of Artificial Intelligence (PAIS 2020)*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pages 203–210. IOS Press. (Cited page 93.)
- Shams, P., Beynier, A., Bouveret, S., and Maudet, N. (2021). Minimizing and balancing envy among agents using ordered weighted average. In Fotakis, D. and Insua, D. R., editors, *Algorithmic Decision Theory - 7th International Conference, ADT 2021, Toulouse, France, November 3-5, 2021, Proceedings*, volume 13023 of *Lecture Notes in Computer Science*, pages 289–303. Springer. (Cited page 81.)
- Shams, P., Beynier, A., Bouveret, S., and Maudet, N. (2022). Fair in the eyes of others. *J. Artif. Intell. Res.*, 75:913–951. (Cited page 93.)
- Shapley, L. and Scarf, H. (1974). On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37. (Cited pages 34, 37, 54, 57, and 113.)
- Spector, D. (2000). Rational debate and one-dimensional conflict. *The Quarterly Journal of Economics*, 115(1):181–200. (Cited page 23.)
- Sprumont, Y. (1991). The division problem with single-peaked preferences: a characterization of the uniform allocation rule. *Econometrica*, pages 509–519. (Cited page 23.)
- Sprumont, Y. (1996). Axiomatizing ordinal welfare egalitarianism when preferences may vary. *Journal of Economic Theory*, 68(1):77–110. (Cited page 34.)
- Steinhaus, H. (1948). The problem of fair division. *Econometrica*, 16(1):101–104. (Cited pages 13, 14, 20, and 26.)
- Suksompong, W. (2016). Asymptotic existence of proportionally fair allocations. *Mathematical Social Sciences*, 81:62 – 65. (Cited page 119.)

- Suksompong, W. (2021). Constraints in fair division. *SIGecom Exch.*, 19(2):46–61. (Cited page 127.)
- Suryanarayana, S. A., Sarne, D., and Kraus, S. (2022a). Explainability in mechanism design: Recent advances and the road ahead. In Baumeister, D. and Rothe, J., editors, *Multi-Agent Systems*, pages 364–382, Cham. Springer International Publishing. (Cited page 126.)
- Suryanarayana, S. A., Sarne, D., and Kraus, S. (2022b). Explainability in mechanism design: Recent advances and the road ahead. In Baumeister, D. and Rothe, J., editors, *Multi-Agent Systems*, pages 364–382, Cham. Springer International Publishing. (Cited page 126.)
- Thibaut, J. W. and Walker, L. (1975). *Procedural justice: A psychological analysis*. L. Erlbaum Associates. (Cited page 37.)
- Thomson, W. (1983). Problems of fair division and the egalitarian solution. *Journal of Economic Theory*, 31(2):211–226. (Cited page 34.)
- Thomson, W. (2016). Introduction to the theory of fair allocation. In Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D., editors, *Handbook of Computational Social Choice*, page 261–283. Cambridge University Press. (Cited page 31.)
- Tinbergen, J. (1953). *Redelijke Inkomensverdeling*. N. V. DeGulden Pers., Haarlem. (Cited page 13.)
- Van den Bos, K., Lind, E. A., Vermunt, R., and Wilke, H. A. (1997). How do I judge my outcome when I do not know the outcome of others? the psychology of the fair process effect. *Journal of personality and social psychology*, 72(5):1034. (Cited page 37.)
- Varian, H. R. (1974). Equity, Envy and Efficiency. *Journal of Economic Theory*, 9:63–91. (Cited pages 13 and 37.)
- Verfaillie, G. and Lemaître, M. (2001). Selecting and scheduling observations for agile satellites: Some lessons from the constraint reasoning community point of view. In Walsh, T., editor, *Principles and Practice of Constraint Programming - CP 2001, 7th International Conference, CP 2001, Paphos, Cyprus, November 26 - December 1, 2001, Proceedings*, volume 2239 of *Lecture Notes in Computer Science*, pages 670–684. Springer. (Cited page 15.)
- Walras, L. (1874). *Éléments d'économie politique pure ou Théorie de la richesse sociale*. L. Corbaz, 1st edition. (Cited page 29.)
- Walsh, T. (2015). Generating single peaked votes. arXiv:1503.02766. (Cited page 72.)
- Weymark, J. A. (1981). Generalized gini inequality indices. *Mathematical Social Sciences*, 1(4):409 – 430. (Cited page 83.)
- Yager, R. R. (1988). On ordered weighted averaging aggregation operators in multicriteria decision making. *IEEE Transactions on Systems, Man, and Cybernetics*, 18:183–190. (Cited page 83.)

- Yager, R. R. (1998). New modes of owa information fusion. *International Journal of Intelligent Systems*, 13(7):661–681. (Cited page 13.)
- Young, H. (1974). An axiomatization of borda’s rule. *Journal of Economic Theory*, 9(1):43–52. (Cited page 12.)
- Young, H. P. (1994). *Equity in Theory and Practice*. Princeton University Press. (Cited page 20.)
- Zahedi, Z., Sengupta, S., and Kambhampati, S. (2023). ‘why didn’t you allocate this task to them?’ negotiation-aware explicable task allocation and contrastive explanation generation. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS ’23, page 2292–2294, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems. (Cited page 126.)

Appendix A

Determining if a solution is Pareto Optimal

We will here explain how we managed to computationally check if some solution π of an add-MARA instance I (we recall that a solution is an allocation of the objects to the agents) is Pareto-Optimal or not.

The idea is that if π is not Pareto-Optimal then it means that there is some other allocation π' Pareto dominating π which implies that for every agent $a_i \in \mathcal{N}$ we have $u_i(\pi'_i) \geq u_i(\pi_i)$ and there exists one agent $a_j \in \mathcal{N}$ such that $u_j(\pi'_j) > u_j(\pi_j)$. By relying on this property we have simply used Mixed Integer Linear Programming looking for this potential π' (note that there can be several allocations Pareto dominating π but finding a single one is enough to prove π is not PO). If such an allocation exists it will be returned by the following MILP :

$$\begin{aligned} & \max \sum_{i=1}^n \sum_{j=1}^m z_i^j u_i^j \\ & \begin{cases} \sum_{j=1}^m z_i^j u_i^j \geq u_i(\pi_i) & \forall i \in \llbracket 1, n \rrbracket \\ \sum_{i=1}^n z_i^j = 1 & \forall j \in \llbracket 1, m \rrbracket \\ z_i^j \in \{0, 1\} & \forall i \in \llbracket 1, n \rrbracket \quad \forall j \in \llbracket 1, m \rrbracket \end{cases} \end{aligned}$$

In this MILP, z_i^j are boolean variables where $z_i^j = 1$ if object o_j is allocated to agent a_i and 0 otherwise.

Proposition A.1. *If $\sum_{i=1}^n \sum_{j=1}^m z_i^j u_i^j = \sum_{i=1}^n u_i(\pi_i)$ then π is Pareto-Optimal, otherwise π is not and the MILP returns an allocation π' that Pareto dominates it.*

Proof. Note that first constraint corresponds to the fact that the allocation π' returned by our MILP should be such that for every agent $a_i \in \mathcal{N}$ we have $u_i(\pi'_i) \geq u_i(\pi_i)$. So, if we also have that there exists one agent $a_j \in \mathcal{N}$ such that $u_j(\pi'_j) > u_j(\pi_j)$ then the returned allocation indeed Pareto dominates π . For this part we consider two cases:

- $\sum_{i=1}^n \sum_{j=1}^m z_i^j u_i^j = \sum_{i=1}^n u_i(\pi_i)$

In this case, it means the MILP has returned an allocation that is either π itself or another allocation but such that each agent gets exactly the same utility. Basically, it means that π was indeed Pareto-Optimal.

- $$\sum_{i=1}^n \sum_{j=1}^m z_i^j u_i^j > \sum_{i=1}^n u_i(\pi_i)$$

In this case, it means the MILP has returned an allocation π' that is such that there exists one agent $a_j \in \mathcal{N}$ such that $u_j(\pi'_j) > u_j(\pi_j)$. Moreover, we had already shown that for every agent $a_i \in \mathcal{N}$ we have $u_i(\pi'_i) \geq u_i(\pi_i)$. Hence, π is not Pareto-Optimal and is Pareto dominated by π' .

Note that the case where $\sum_{i=1}^n \sum_{j=1}^m z_i^j u_i^j > \sum_{i=1}^n u_i(\pi_i)$ is not possible because it would contradict the first set of constraints.

□

Appendix B

MILP characterizing the EF1

Note that we do not need any objective function as the constraints themselves are enough to characterize the set of EF1 allocations.

$$\left\{ \begin{array}{ll} \sum_{j=1}^m u(i, o_j)(z_i^j - z_k^j + p_{ik}^j) \geq 0 & \forall i, k \in \llbracket 1, n \rrbracket \\ \sum_{i=1}^n p_{ik}^j \leq 1 & \forall i, k \in \llbracket 1, n \rrbracket \\ \sum_{i=1}^n z_i^j = 1 & \forall j \in \llbracket 1, m \rrbracket \\ p_{ik}^j \leq z_k^j \quad \forall i, k \in \llbracket 1, n \rrbracket & \forall j \in \llbracket 1, m \rrbracket \\ z_i^j \in \{0, 1\} \quad \forall i \in \llbracket 1, n \rrbracket & \forall j \in \llbracket 1, m \rrbracket \\ p_{ik}^j \in [0, 1] \quad \forall i, k \in \llbracket 1, n \rrbracket & \forall j \in \llbracket 1, m \rrbracket \end{array} \right.$$

In this MIP, there are nm boolean variables z_i^j from the classic modelization of the MARA problem and n^2m variables p_{ik}^j . We first consider that these are boolean variables in order to explain the modelization. The variables p_{ik}^j equal 1 if agent a_i “forgets” item o_j in order to be envy free towards agent k , 0 otherwise. The variables z_i^j equal 1 if agent a_i possesses item o_j . We explain hereafter each constraint of the MIP.

The first set of constraints starts linking the variables p_{ik}^j with EF1. Indeed, it ensures that the difference, from the point of view of agent a_i (ie its preferences), between her share and agent a_k 's minus the items for which $p_{ik}^j = 1$ is greater or equal than 0. We write this constraint for every couple a_i, a_k of agents. We note that for the allocation to be EF1 we need another constraint on the number of items for which $p_{ik}^j = 1$.

The second set of constraints ensures that the returned allocation is envy free up to one good (EF1) as for each couple a_i, a_k of agents, there is only one item needed for agent a_i not to envy agent a_k . We note that is we replace 1 by zero we will obtain an envy free allocation (if it exists).

The third set of constraints ensures that each item is affected to exactly one agent. Therefore, it ensures both completeness of the allocation and the non sharing of an item.

The fourth set of constraints traduces the fact that agent a_i can forget item o_j from agent a_k 's bundle if and only if the latter possesses it.

This MILP is hence enough to characterize the set of EF1 allocations. However, there is a large number of EF1 allocations and one could want to have control over the returned one.

Proposition B.1 (Continuous relaxation). *We can relax the integrity constraints on the p_{ik}^j variables.*

Proof. We explain here why we can relax the integrity constraint on the p_{ik}^j variables. We first note that the p_{ik}^j variables do not affect the returned allocation as they are just here to force envy freeness up to one good. So, we just have to proof that even with relaxing their integrity we ensure EF1.

Besides, we also note that the p_{ik}^j are not linked between them apart from the dimension j . Indeed, for a certain couple of agents a_i, a_k the second set of constraints keeps the sum of the p_{ik}^j over the objects from exceeding 1.

We separate the three possible cases that could occur by the integrity relaxation. The first case would be that every p_{ik}^j stays an integer. Obviously, this keeps the EF1 property as stated above.

The second case would be that for every couple of agents i, k there is only one $p_{ik}^j > 0$. It is rather natural that the returned allocation will still be EF1. In order to show that we introduce P_{ik}^j which equals 1 if $p_{ik}^j > 0$ and 0 otherwise. We have:

$$\sum_{j=1}^m u(i, o_j)(z_i^j - z_k^j + P_{ik}^j) \geq \sum_{j=1}^m u(i, o_j)(z_i^j - z_k^j + p_{ik}^j) \quad (\text{B.1})$$

And the first set of constraints ensures $\sum_{j=1}^m u(i, o_j)(z_i^j - z_k^j + p_{ik}^j) \geq 0$. Finally by transitivity with equation B.1 we have $\sum_{j=1}^m u(i, o_j)(z_i^j - z_k^j + P_{ik}^j) \geq 0$. The latter inequation is the EF1 property showed above so the second case is not problematic either.

The third and final case would be that for some couple of agents i, k there are several items o_j for which $p_{ik}^j > 0$. We prove EF1 by showing that the third case can be reduced to the second. For a couple of agents a_i, a_k , we note $\mathcal{O}_{i,k}^{>0}$ the set of items o_j for which $p_{ik}^j > 0$. Let us consider among $\mathcal{O}_{i,k}^{>0}$ the item with the highest utility for agent a_i and we note this item o_J . By transferring the values of the p_{ik}^j for which $p_{ik}^j > 0$ to p_{ik}^J we have $p_{ik}^J u(i, o_j) \geq \sum_{j=1}^m p_{ik}^j u(i, o_j)$. So, with the exact same reasoning that for the second case we show EF1. □

Appendix C

MILP characterizing the EFX

From the polyhedron of the EF1 we write the polyhedron of the EFX allocations by simply adding the constraint

$$z_k^h = 1 \implies \sum_j p_{ik}^j u_i^j \leq z_k^h u_i^h \forall i, k \in \llbracket 1, n \rrbracket \forall h \in \llbracket 1, m \rrbracket$$

Indeed, this constraint means that if the total utility an agent a_i has to forget in order not to envy agent a_k is lower or equal to the lowest utility (from the point of view of a_i) of the objects a_k possesses. This constraint can be linearized as follows:

$$\sum_j p_{ik}^j u_i^j - z_k^h u_i^h \leq M(1 - z_k^h)$$

with $M = \max_{i \in \mathcal{N}} \sum_{j=1}^m u_i^j$

$$\left\{ \begin{array}{ll} \sum_{j=1}^m u(i, o_j)(z_i^j - z_k^j + p_{ik}^j) \geq 0 & \forall i, k \in \llbracket 1, n \rrbracket \\ \sum p_{ik}^j \leq 1 & \forall i, k \in \llbracket 1, n \rrbracket \\ \sum_j p_{ik}^j u_i^j - z_k^h u_i^h \leq M(1 - z_k^h) & \forall i, k \in \llbracket 1, n \rrbracket \\ & \forall h \in \llbracket 1, m \rrbracket \\ \sum_{i=1}^n z_i^j = 1 & \forall j \in \llbracket 1, m \rrbracket \\ p_{ik}^j \leq z_k^j & \forall i, k \in \llbracket 1, n \rrbracket & \forall j \in \llbracket 1, m \rrbracket \\ z_i^j \in \{0, 1\} & \forall i \in \llbracket 1, n \rrbracket & \forall j \in \llbracket 1, m \rrbracket \\ p_{ik}^j \in [0, 1] & \forall i, k \in \llbracket 1, n \rrbracket & \forall j \in \llbracket 1, m \rrbracket \end{array} \right.$$

Appendix D

Fair in the eyes of others

Proposition D.1. *Let I be an add-MARA instance. Then, there is an optimal solution with $\mathbf{K} = L$ to $\mathcal{M}_1(I)$ if and only if I is an (L -app envy)-free instance and not an $((L-1)$ -app envy)-free one. Moreover, $\mathcal{M}_1(I)$ does not admit any solution if and only if I is an unanimous envy instance.*

Proof. To prove the proposition, we show that there is an (L -app envy)-free allocation in I if and only if there is a solution to the MIP $\mathcal{M}_1(I)$ such that $\mathbf{K} = L$.

(\Rightarrow) Let I be an instance, and π be an (L -app envy)-free allocation. Then, consider the partial instantiation of the variables such that $\mathbf{z}_i^j = 1$ if and only if $o_j \in \pi_i$. We prove that this partial instantiation extends to a solution of $\mathcal{M}_1(I)$ such that $\mathbf{K} = L$.

First observe that Constraint 5.1 is directly satisfied.

Now, consider any triple of agents (a_k, a_i, a_h) . Suppose that agent a_k thinks a_i should envy a_h . Then in this case, we have $\sum_{j \in \pi_h} u(k, j) > \sum_{j \in \pi_i} u(k, j)$. In other words, $\sum_{j=1}^m u(k, j)(\mathbf{z}_h^j - \mathbf{z}_i^j) > 0$ which is in turn equivalent to $\sum_{j=1}^m u(k, j)(\mathbf{z}_h^j - \mathbf{z}_i^j) \geq 1$ since all utilities are integers. By Constraint 5.2, we thus have that $\mathbf{e}_{kih} = 1$ which implies that Constraint 5.3 is satisfied as well.

Conversely, suppose that agent a_k thinks a_i should not envy a_h . Then, we have $\sum_{j \in \pi_h} u(k, j) \leq \sum_{j \in \pi_i} u(k, j)$. In other words, $\sum_{j=1}^m u(k, j)(\mathbf{z}_h^j - \mathbf{z}_i^j) \leq 0$. By Constraint 5.3, we thus have that $\mathbf{e}_{kih} = 0$ in this case, which in turns implies that Constraint 5.2 is satisfied as well. Hence, we have that $\mathbf{e}_{kih} = 1$ if and only if a_k thinks a_i should envy a_h in π .

Finally, consider any pair of agents (a_i, a_h) . If a_i does not envy a_h then $\mathbf{e}_{i ih} = 0$. As a consequence, \mathbf{x}_{ih} can be null and still satisfy Constraints 5.4 and 5.5 (no matter the value of \mathbf{K} is).

Now suppose that a_i does envy a_h (hence $\mathbf{e}_{i ih} = 1$). Then, we should have $\mathbf{x}_{ih} = 1$ to satisfy Constraint 5.4. Since π is (L -app envy)-free, then at most $L - 1$ agents (including a_i herself) think that a_i should indeed envy a_h , which means that $\sum_{k=1}^n \mathbf{e}_{kih} \leq L - 1$. Instantiating \mathbf{K} to L is hence enough to satisfy Constraint 5.5.

(\Leftarrow) Now suppose that there is a solution to $\mathcal{M}_1(I)$ such that $\mathbf{K} = L$. Then we will prove that the allocation π such that $o_j \in \pi_i$ if and only if $\mathbf{z}_{ij} = 1$ is a valid (L -app envy)-free allocation.

First, according to Constraints 5.1, π is indeed a valid allocation.

Secondly, Constraint 5.2 ensures that if $\mathbf{e}_{kih} = 0$ then $\sum_{j=1}^m u(k, j)(\mathbf{z}_h^j - \mathbf{z}_i^j) \leq 0$, in turn meaning that agent a_k thinks that a_i should not envy a_h . Conversely, Constraint 5.3

ensures that if $\mathbf{e}_{kih} = 1$ then $\sum_{j=1}^m u(k, j)(\mathbf{z}_h^j - \mathbf{z}_i^j) > 0$, in turn meaning that agent a_k thinks that a_i should envy a_h . It also obviously implies that $\mathbf{e}_{iih} = 1$ if and only if a_i envies a_h .

Now consider any pair of agents (a_i, a_h) such that a_i envies a_h . From what precedes, $\mathbf{e}_{iih} = 1$. By Constraint 5.4, $\mathbf{x}_{ih} = 1$. Hence, by Constraint 5.5, $\sum_{k=1}^h \mathbf{e}_{kih} \leq L - 1$. This implies that the total number of agents agreeing with the fact that a_i envies a_h is strictly lower than L . In other words, π is (L -app envy)-free. \square

Proposition D.2. *Let I be an instance. Then, there is an optimal solution with $\mathbf{K} = L$ to $\mathcal{M}_2(I)$ if and only if I is an (L -app non-prop)-free instance and not an $((L - 1)$ -app non-prop)-free one. Moreover, $\mathcal{M}_2(I)$ does not admit any solution if and only if I is an unanimous non-proportional instance.*

Proof. The key here is to show that there is a solution to the MIP $\mathcal{M}_2(I)$ such that $\mathbf{K} = L$ iff the corresponding allocation π such that $\mathbf{z}_i^j = 1$ if and only if $o_j \in \pi_i$ is (L -app non-prop)-free. However this is done in the proof of Proposition 5.14. We also have to show that Constraints 5.6 and 5.7 are indeed a valid translation of the logical equivalence, and that Constraints 5.8 and 5.9 correctly encode the logical OR. The same type of linearization is also done in the proof of Proposition 5.14. \square

Proposition D.3. *The probability for an instance being randomly generated under uniformly distributed preferences to exhibit unanimous envy is upper bounded by $n(n - 1)/2^n$.*

Proof. The probability of the event o_i is strictly preferred to o_j by one agent is $1/2$ if preferences are strict. As preferences are not strict, this probability becomes an upper bound (think for instance if the agent values all the objects the same then the probability to have strict preference between two objects is zero). Hence, the probability of the event o_i is strictly preferred to o_j by all agents is upper bounded by $1/2^{n-1}$ as the preferences between the agents are independent. Assuming, for all pairs of items, these events to be independent (which is not the case, hence an upper bound of the upper bound), we derive our result by summing up over the $n(n - 1)/2$ possible pairs. \square

Note that this value quickly tends towards 0: for instance, for 10 agents, the probability for an instance to exhibit unanimous envy is upper-bounded by 0.088.

Proposition D.4. *The probability for an instance being randomly generated under uniformly distributed preferences to exhibit unanimous non-proportionality is upper bounded by $n/2^n$.*

To prove this property, we will need a small lemma:¹

Lemma D.1. *Let U_1, \dots, U_n be n independent random variables having a uniform distribution over real interval $[a, b]$, and let us denote by \bar{U}_n the empiric mean of U_1, \dots, U_n : $\bar{U}_n = \frac{1}{n} \sum_{i=1}^n U_i$.*

Then we have:

$$P(U_1 < \bar{U}_n) = \frac{1}{2}.$$

¹We warmly thank Olivier François for this result.

Proof. The probability we seek can be reformulated as follows:

$$P(U_1 < \bar{U}_n) = P\left(U_1 < \frac{1}{n} \sum_{i=1}^n U_i\right) = P\left(U_1 < \frac{1}{n-1} \sum_{i=2}^n U_i\right) = P(U_1 < \bar{U}_{n-1}). \quad (\text{D.1})$$

We can notice that the two latter variables, U_1 and \bar{U}_{n-1} are independent. For any two independent variables X and Y , we have that $P(X < Y) = \mathbb{E}[F_X(Y)]$, where F_X is the cumulative distribution function of X . To see this, we can consider the following steps:

$$\begin{aligned} P(X < Y) &= \int P(X < Y | Y = y) f_Y(y) dy = \int P(X < y) f_Y(y) dy \\ &= \int F_X(y) f_Y(y) dy \\ &= \mathbb{E}[F_X(Y)], \end{aligned} \quad (\text{D.2})$$

where the first step is obtained using the law of total probability, and the last step is obtained using the law of unconscious statistician.

Putting Equations (D.1) and (D.2) together, we obtain:

$$P(U_1 < \bar{U}_n) = \mathbb{E}[F_{U_1}(\bar{U}_{n-1})] = \mathbb{E}\left(\frac{\bar{U}_{n-1} - a}{b - a}\right) = \frac{\mathbb{E}(\bar{U}_{n-1}) - a}{b - a} \quad (\text{D.3})$$

Observing that \bar{U}_{n-1} has the same (uniform) law as U_1 , we have that $\mathbb{E}(\bar{U}_{n-1}) = \frac{a+b}{2}$. Injecting this to Equation (D.3) yields:

$$P(U_1 < \bar{U}_n) = \frac{1}{2} \quad (\text{D.4})$$

as expected. \square

We are now ready to prove Proposition D.4.

Proof (Proposition D.4). Let I be a random instance generated under uniformly distributed preferences. According to Proposition 5.17, I is unanimous non-proportional if and only if there exists at least an item o_p such that $u(k, p) < Prop_k$ for all agents a_k . In what follows, we will denote by $U_{k,p}$ the random variable corresponding to $u(k, p)$.

For any item o_p and any agent a_k , $P(U_{k,p} < \frac{1}{n} \sum_{j=1}^n U_{k,j}) = \frac{1}{2}$ by Lemma D.1. All the variables $U_{k,p}$ being independent, we have that:

$$P\left(\bigcap_{k=1}^n \left(U_{k,p} < \frac{1}{n} \sum_{j=1}^n U_{k,j}\right)\right) = \frac{1}{2^n} \text{ for all } o_p.$$

The events $U_{k,p} < \frac{1}{n} \sum_{j=1}^n U_{k,j}$ not being independent, we can only derive an upper bound on the probability for I to have at least one object o_k such that $u(k, p) < Prop_k$ for all agents a_k . Namely:

$$P\left(\exists k \in \llbracket 1, n \rrbracket \mid \bigcap_{k=1}^n \left(U_{k,p} < \frac{1}{n} \sum_{j=1}^n U_{k,j}\right)\right) \leq \frac{n}{2^n},$$

which concludes the proof. \square

Note that once again this value quickly tends towards 0: for instance, for 10 agents, the probability for an instance to exhibit unanimous non-proportionality is upper-bounded by 0.00977.

Résumé

Cette thèse s'inscrit dans le contexte du Choix Social Computationnel. Il s'agit d'un domaine à l'intersection du Choix Social, de l'Informatique et de l'Intelligence Artificielle. Nous nous intéressons plus particulièrement au problème de partage équitable de ressources indivisibles qui consiste à trouver une allocation équitable et efficace d'un ensemble d'objets à un ensemble d'agents. Tandis que la notion d'efficacité est communément traduite par l'exigence minimale de complétude (tous les objets doivent être alloués dans le but de ne gâcher aucun objet) ou la notion plus exigeante de Pareto-Optimalité (une allocation est dite Pareto-Optimale s'il n'existe pas une autre allocation telle que tous les agents sont au moins aussi contents et un agent est strictement plus content), plusieurs notions ont été proposées pour définir l'équité. Une des mesures d'équité la plus importante est l'absence d'envie. Une allocation est dite sans envie ou *envy-free* si aucun agent n'aurait envie de changer ses objets contre ceux d'un autre agent. Cependant, il n'existe pas forcément d'allocation *envy-free* quand on est dans le cadre de ressources indivisibles. Afin de surmonter cette limitation, des relaxations ont été récemment proposées dans la littérature.

Dans cette thèse, nous étudions tout d'abord une famille de procédures décentralisées basée sur des échanges d'objets entre agents. Nous analysons en particulier comment ces procédures se comportent et les propriétés désirables qu'elles montrent. Plus précisément, on étudie les séquences de choix sincères et les cycles d'échanges de ressources. Dans un second temps, nous proposons de nouvelles relaxations de la notion d'absence d'envie (et d'autres mesures d'équité) et les étudions en profondeur. La première relaxation a pour but d'équilibrer l'envie entre les agents (quand elle ne peut être évitée) et se base sur l'Ordered Weighted Average (OWA), un agrégateur habituellement utilisé dans le domaine de l'optimisation multicritère pour traduire l'équité. La deuxième relaxation se concentre sur l'approbation sociale de l'envie et se rapproche plus de la théorie du vote étant donné que les agents votent sur l'envie des autres agents. Nous examinons les aspects computationnels liés à ces nouvelles relaxations, leurs liens avec des notions d'équité et d'efficacité existantes avant de les tester expérimentalement.

Mots clés. Partage équitable, Choix Social Computationnel, Décision multi-agents, Intelligence Artificielle.

Abstract

The work of this thesis is in the scope of Computational Social Choice. It is a field at the intersection of Social Choice, Computer Science and Artificial Intelligence. In particular, we study the problem of Fair Division of Indivisible Goods where the objective is to find a fair and efficient allocation of a set of (valuable) objects among a set of agents. While efficiency is usually brought by the minimal requirement of completeness (all the objects have to be allocated in order not to waste anything), or the more demanding notion of Pareto-Optimality (an allocation is Pareto-Optimal if there is no allocation such that all the agents are not worse off and one agent is strictly better off), several notions have been proposed to define the fairness of an allocation. One of the most prominent fairness measures is called envy-freeness. An allocation is said to be envy-free if no agent would like to exchange her bundle of resources with another agent. However, envy-freeness is not guaranteed to exist when considering indivisible goods so various relaxations have been proposed recently in the literature to overcome this limitation.

In this thesis, we first thoroughly study a family of decentralized allocation procedures related to exchanges of goods. We analyze how these procedures behave and the desirable properties they exhibit. More specifically, we study sequence of sincere choices and cycle exchanges of resources. We then propose new relaxations of the envy-freeness notion (and also of other fairness measures) and thoroughly study them. Our first relaxation aims at balancing the envy among the agents (when it cannot be avoided) and is based on the Order Weighted Average (OWA) aggregator usually used in multi-criteria optimisation to bring fairness. The second relaxation focuses on the social approval of the envy and is more related to voting theory, as it lets agents vote about the envy of the other agents. We investigate computational issues related to these new relaxations, their link with existing fairness and efficiency notions and we experimentally test them.

Keywords. Fair division, Computational Social Choice, Multiagent decision, Artificial Intelligence.