ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ
ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

DOCTORAL THESIS

ALGORITHMIC AND MECHANISM DESIGN ASPECTS OF PROBLEMS WITH LIMITED—OR NO—PAYMENTS

Georgios AMANATIDIS

Advisor: Asst. Prof. Evangelos MARKAKIS

A thesis submitted in fulfillment of the requirements for the degree of Doctor of Philosophy in the

Department of Informatics

September 12, 2017

ALGORITHMIC AND MECHANISM DESIGN ASPECTS OF PROBLEMS WITH LIMITED—OR NO—PAYMENTS

Approved by:

Asst. Prof. Evangelos Markakis, Advisor
Department of Informatics
Athens University of Economics and Business

Asst. Prof. Dimitris Fotakis
School of Electrical and Computer Engineering
National Technical University of Athens

Assoc. Prof. Iordanis Koutsopoulos
Department of Informatics
Athens University of Economics and Business

Assoc. Prof. Aris Pagourtzis
School of Electrical and Computer Engineering
National Technical University of Athens

Asst. Prof. Orestis Telelis
Department of Digital Systems
University of Piraeus

Emer. Prof. Stathis Zachos
School of Electrical and Computer Engineering
National Technical University of Athens

Prof. Vassilis Zissimopoulos
Department of Informatics and Telecommunications
National and Kapodistrian University of Athens

Date Approved: August 28, 2017

"The most effective way to do it, is to do it."

Amelia Earhart

Athens University of Economics and Business
Department of Informatics

Abstract

of the thesis

Algorithmic and Mechanism Design Aspects of Problems with Limited—or no—Payments

by Georgios Amanatidis

The most notable distinction between algorithm design and mechanism design is the notion of truthfulness. Typically, one of the goals of the mechanism designer is to ensure that the agents participating in the mechanism will not have any incentive to misreport their private information. Towards this goal, the payments made to the agents by the mechanism are crucial. However, there is an abundance of scenarios in microeconomics where the payments are restricted (e.g., via budget constraints) or even completely undesirable. This thesis focuses on two such problems that are indicative of the challenges that arise when the payments are limited or absent. In particular, fair division of indivisible items and reverse auctions with hard budget constraints are studied, both from the game-theoretic and the algorithmic point of view.

In the first part of the thesis, we study the problem of computing allocations with maximin share guarantees—a recently introduced fairness notion. Given a set of agents and a set of goods, the objective is to find a partition so that each agent is guaranteed an approximation of his maximin share. Our main algorithmic result is a 2/3-approximation, that runs in polynomial time for any number of agents and items. Furthermore, we undertake a probabilistic analysis and provide a theoretical justification of the experimental evidence reported in the literature indicating that better approximations almost always exist.

From the mechanism design point of view, this is a setting where no monetary transfers are allowed. Even for two agents and a few items, the problem becomes strictly harder than its algorithmic counterpart, and the limitations imposed by truthfulness on the approximability of the problem become apparent. We focus on the case of two players and our main result is a complete characterization of truthful mechanisms that allocate all the items. Applying this result, we derive several consequences

on the design of mechanisms with fairness guarantees, such as maximin share fairness and envy-freeness up to one item.

In the second part of the thesis, we study a family of reverse auctions with a budget constraint. The general algorithmic problem is to purchase a set of resources, which come at a cost, so as not to exceed a given budget and at the same time maximize a given valuation function. This framework captures the budgeted version of several well-known optimization problems, and when the resources are owned by strategic agents the goal is to design truthful and budget-feasible mechanisms.

We obtain mechanisms with significantly improved approximation ratios for several subclasses of submodular valuation functions, like *coverage functions* and *cut functions*. We then provide a general scheme for designing deterministic and randomized mechanisms for a subclass of XOS problems which contains problems whose feasible set forms an *independence system*. Some representative problems include, among others, finding maximum weighted matchings and maximum weighted matroid members. For most of the above, only randomized mechanisms with very high approximation ratios were known prior to our results.

A purely algorithmic byproduct of our work is a polynomial-time $\frac{2e}{e-1}$-approximation algorithm for symmetric submodular maximization subject to a budget constraint. This is the best known factor achieved by a deterministic algorithm assuming only a value oracle for the objective function.

Οικονομικο Πανεπιστημιο Αθηνων

Τμήμα Πληροφορικής

Περίληψη

της διδακτορικής διατριβής

Algorithmic and Mechanism Design Aspects of Problems with Limited—or no—Payments

του Γεωργίου Αμανατίδη

Η πιο αξιοσημείωτη διάκριση μεταξύ της σχεδίασης αλγορίθμων και της σχεδίασης μηχανισμών είναι η έννοια της φιλαλήθειας (truthfulness). Κατά κανόνα, ένας από τους στόχους του σχεδιαστή μηχανισμών είναι να εξασφαλίσει ότι οι παίκτες που συμμετέχουν στον μηχανισμό δεν έχουν κανένα κίνητρο να παραποιήσουν το κομμάτι της εισόδου που αποτελεί προσωπική τους πληροφορία. Για την επίτευξη του στόχου αυτού, οι πληρωμές που πραγματοποιούνται στους παίκτες από τον μηχανισμό είναι κρίσιμες. Ωστόσο, υπάρχουν πολλά σενάρια στη μικροοικονομική θεωρία όπου οι πληρωμές είναι περιορισμένες (π.χ., λόγω ύπαρξης προϋπολογισμού) ή ακόμη και εντελώς ανεπιθύμητες. Η παρούσα διδακτορική διατριβή επικεντρώνεται σε δύο τέτοια προβλήματα που είναι ενδεικτικά των προκλήσεων που προκύπτουν όταν οι πληρωμές είναι περιορισμένες ή απουσιάζουν. Συγκεκριμένα, μελετάται ο δίκαιος διαμοιρασμός μη διαιρετών αντικειμένων και οι αντίστροφες δημοπρασίες με αυστηρούς περιορισμούς προϋπολογισμού, τόσο από την παιγνιοθεωρητική όσο και από την αλγοριθμική σκοπιά.

Στο πρώτο μέρος της διατριβής, μελετάμε το πρόβλημα του υπολογισμού διανομών μη διαιρετών αγαθών, οι οποίες παρέχουν εγγυήσεις ως προς τα μεγιστοελάχιστα μερίδια (maximin shares) των παικτών—μια πρόσφατα ορισμένη έννοια δικαιότητας. Δεδομένου ενός συνόλου παικτών και ενός συνόλου αγαθών, ο στόχος είναι να βρεθεί μια διανομή που να εγγυάται σε κάθε παίκτη μια προσέγγιση του μεγιστοελαχίστου μεριδίου του. Το κύριο αλγοριθμικό αποτέλεσμά μας είναι ένας 2/3-προσεγγιστικός αλγόριθμος, ο οποίος τρέχει σε πολυωνυμικό χρόνο για οποιονδήποτε αριθμό παικτών και αγαθών. Επιπλέον, επιχειρούμε μια πιθανοτική ανάλυση και παρέχουμε μια θεωρητική αιτιολόγηση των πειραματικών δεδομένων που αναφέρονται στη σχετική βιβλιογραφία και που υποδεικνύουν ότι υπάρχουν σχεδόν πάντα καλύτερες προσεγγίσεις.

Από τη σκοπιά της σχεδίασης μηχανισμών, αυτό είναι ένα περιβάλλον στο οποίο δεν επιτρέπονται πληρωμές. Ακόμη και για δύο παίκτες και λίγα αγαθά, το πρόβλημα είναι

αυστηρά δυσκολότερο από το αντίστοιχο αλγοριθμικό και γίνονται προφανείς οι περιορισμοί που επιβάλλει η φιλαλήθεια στην προσεγγισιμότητα του προβλήματος. Εστιάζουμε στην περίπτωση των δύο παικτών και το κύριο αποτέλεσμά μας είναι ένας πλήρης χαρακτηρισμός των φιλαληθών μηχανισμών που κατανέμουν όλα τα αγαθά. Ο χαρακτηρισμός αυτός έχει άμεσες συνέπειες στο σχεδιασμό μηχανισμών με εγγυήσεις δικαιότητας, όπως η δικαιότητα μεγιστοελαχίστου μεριδίου (maximin share fairness) και η απουσία φθόνου εξαιρουμένου το πολύ ενός αντικειμένου (envy-freeness up to one item).

Στο δεύτερο μέρος της διατριβής, μελετάμε μια οικογένεια αντίστροφων δημοπρασιών με περιορισμένο προϋπολογισμό. Το γενικότερο αλγοριθμικό πρόβλημα είναι να αγοράσουμε ένα σύνολο πόρων, καθένας από τους οποίους έχει κάποιο κόστος, έτσι ώστε να μην υπερβούμε έναν δεδομένο προϋπολογισμό και ταυτόχρονα να μεγιστοποιήσουμε μια δεδομένη συνάρτηση αποτίμησης. Αυτό το πλαίσιο περιλαμβάνει τις παραλλαγές πολλών γνωστών προβλημάτων βελτιστοποίησης όπου έχει προστεθεί και ένας περιορισμός προϋπολογισμού. Όταν οι πόροι ανήκουν σε στρατηγικούς παίκτες, ο στόχος είναι να σχεδιαστούν φιλαλήθεις μηχανισμοί που δεν παραβιάζουν τον προϋπολογισμό.

Παίρνουμε μηχανισμούς με σημαντικά βελτιωμένους λόγους προσέγγισης για αρκετές υποκατηγορίες υπομετρικών (submodular) συναρτήσεων, όπως οι *συναρτήσεις κάλυψης* (coverage functions) και οι *συναρτήσεις κοπής* (cut functions). Στη συνέχεια παρέχουμε ένα γενικό σχήμα για τον σχεδιασμό ντετερμινιστικών και τυχαιοποιημένων μηχανισμών για μια υποκλάση των XOS συναρτήσεων που περιέχει προβλήματα των οποίων το σύνολο εφικτών λύσεων σχηματίζει ένα *σύστημα ανεξαρτησίας*. Ορισμένα αντιπροσωπευτικά προβλήματα είναι, μεταξύ άλλων, η εύρεση μέγιστων σταθμισμένων ταιριασμάτων (maximum weighted matchings) και μέγιστων σταθμισμένων μελών μητροειδών (maximum weighted matroid members). Για τα περισσότερα από τα παραπάνω, πριν από τα αποτελέσματά μας, ήταν γνωστοί μόνο τυχαιοποιημένοι μηχανισμοί με πολύ υψηλούς λόγους προσέγγισης.

Ένα καθαρά αλγοριθμικό υποπροϊόν της δουλειάς μας είναι ένας πολυωνυμικός $\frac{2e}{e-1}$-προσεγγιστικός αλγόριθμος για το πρόβλημα της μεγιστοποίησης συμμετρικών υπομετρικών συναρτήσεων υπό περιορισμό προϋπολογισμού. Πρόκειται για τον καλύτερο γνωστό λόγο προσέγγισης που επιτυγχάνεται από ντετερμινιστικό αλγόριθμο, υποθέτοντας μόνο κλήσεις σε ένα μαντείο για την αντικειμενική συνάρτηση.

Acknowledgements

First and foremost, I would like to thank my advisor, Vangelis Markakis. It is hard to avoid clichés here, but Vangelis has truly been an exemplary advisor to me. He guided me with patience, he rekindled my passion for research, and gave me enough space to develop into an independent researcher. Vangelis, for the constant support, the guidance, and the inspiration, I thank you from the bottom of my heart.

I am also thankful to my good friend, and godfather to my son, Vassilis Zikas, for his faith in me. He and my wife compelled me into starting a Ph.D. and I have to admit that it has been way more fun than I had anticipated.

During my Ph.D. I was fortunate to collaborate with many wonderful people aside from Vangelis: Joël Alwen, Nathanaël Barrot, Georgios Birmpas, George Christodoulou, Jérôme Lang, Bernard Ries, Amin Saberi, Krzysztof Sornat, Eftychia Vakaliou, and Vassilis Zikas.

I would like to thank fellow group members Georgios Birmpas and Artemis Tsikiridis for the endless stimulating discussions and the spare cigarettes. Our regular interactions greatly contributed to a pleasant and inspiring work environment and I will miss them.

I would also like to thank my family, my in-laws, and my friends for all their love and encouragement throughout the years. But most of all, I would like to thank my dear wife, Kalliopi, for her unconditional love and support, for being my rock, for making my days brighter. If it wasn't for her, I would have never embarked on this academic journey.

Last, but not least, I would like to thank my son, Dimitris, for all the nights that he stayed up waiting for me to return from the office. Stealing and using his notebooks has been an integral part of working on this thesis!

*This work was partially supported by the European Union (European Social Fund - ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) - Research Funding Program: THALES (*Algorithmic Game Theory *[Mis: 377002] and* Algorithms of Today *[Mis: 379414]).*

To my other half, Kalliopi, and to our son, Dimitris.

Contents

Preface

During the last two decades an exciting research area has emerged, defined by the interplay of algorithms, multi-agent systems, social choice theory, game theory, and complexity. Classic game theory notions have been reexamined through the algorithmic lens, while questions on computational efficiency have been raised and, in many cases, resolved. At the same time, this interaction of economics and computation has breathed new life to countless classic optimization problems, as incentives had to be taken into account and new objectives, like truthfulness or economic efficiency, were added. Although the dividing lines between the subareas that were shaped over the years are blurry, the general area came to be known as *Algorithmic Game Theory* within the theoretical computer science community.

Crucial reasons behind the impressive advancement of this line of interdisciplinary research are the explosive growth of the Internet and the increasing applicability of artificial intelligence. As noted in the seminal position paper of Papadimitriou (2001), "the most novel and defining [characteristic of the Internet] is *its socio-economic complexity*". Indeed, the present-day algorithm designer cannot ignore the socioeconomics of complex networks. Large scale complex networks of autonomous agents create opportunities to do things better, faster, and with a greater profit—financial or social. What makes things particularly interesting from the computer scientist's angle is that most questions that arise are inherently algorithmic.

- How should agents (human or AI) split resources fairly?
- How can the structure of the underlying social network of the potential clients benefit a company for advertising purposes?
- Can a potential employer utilize a large pool of specialized workers without overpaying too much?

This thesis aims to address such questions, mainly from a mechanism design perspective.

A most notable difference of this approach from algorithm design are the *incentives* of the parties involved. Very often, we would like to ensure that the agents participating in a mechanism have no incentive to misreport their private information, i.e., we want to design *truthful* mechanisms. Towards this goal, the payments made to the agents by the mechanism are crucial. However, like computational power, memory, or time, money is a valuable and scarce resource. There is an abundance of scenarios in microeconomics where the payments are restricted via budget constraints, or even completely undesirable.

The focus of my work is on two such problems. Although the general subject is very broad, and our treatment of it is far from being exhaustive, the problems studied here are indicative of the challenges that arise when the payments are limited or absent. In particular, fair division of indivisible items and reverse auctions with hard budget constraints are studied, both from the game-theoretic and the algorithmic point of view.

In Part I we study the problem of computing *fair* allocations when all the available goods are indivisible. In the setting we study, there are n agents with additive valuation functions over a set of m indivisible goods and we want to allocate all the goods in a manner that is considered fair in some sense. This is one of the most standard and fundamental settings in fair division. Given that it is impossible to have any approximation guarantees for classic fairness notions, like proportionality and envy-freeness, in this setting, we mostly deal with *maximin share fairness*, a recently introduced weaker fairness notion. The maximin share of a single agent is the best he can guarantee to himself, if he was to partition the items in any way he prefers into n bundles and then receive his least desirable bundle. The objective then is to find an allocation, so that each player is guaranteed an approximation of his maximin share. In Chapter 2 we study the problem from a purely algorithmic perspective. Our main algorithmic result is a 2/3-approximation algorithm, that runs in polynomial time for any number of agents and items. We also provide a polynomial time 7/8-approximation algorithm for the intriguing three player case, as well as an exact polynomial time algorithm for the case where all item values belong to $\{0, 1, 2\}$. Furthermore, we undertake a probabilistic analysis and provide a theoretical justification of the experimental evidence reported in the literature indicating that better approximations almost always exist. In particular, we prove that in randomly generated instances, maximin share allocations exist with high probability.

From the mechanism design point of view, this is a setting where no monetary transfers are allowed. This is standard in microeconomics, where fair division problems often model scenarios where payments are not desirable, e.g., inheritances or distribution of common goods. In Chapter 3 we study the problem assuming self-interested, rational agents. Even for two agents and a few items, the problem becomes strictly harder than its algorithmic counterpart, and the limitations imposed by truthfulness on the approximability of the problem become apparent. We focus on the case of two players and our main result is a complete characterization of truthful mechanisms that allocate all the items. This fundamental result goes well beyond the quest for fairness and reveals an interesting structure underlying all truthful mechanisms, showing that they can be decomposed into two components: a *selection* part where players pick their best subset among prespecified choices determined by the mechanism, and an *exchange* part where players are offered the chance to exchange certain subsets if it is favorable to do so. Applying this result, we derive several consequences on the design of mechanisms with fairness guarantees, including maximin share fairness and envy-freeness up to one item. Finally, for the case of multiple players, we provide a general class of truthful mechanisms that generalizes the class of truthful mechanisms for two players in a non-trivial way.

In Part II we study a family of procurement (reverse) auctions where the auctioneer has a hard budget constraint. The general algorithmic problem is (for the auctioneer) to purchase a set of resources (each owned by a rational agent), which come at a cost, so as not to exceed a given budget and at the same time maximize a given valuation function. In the setting considered here, the true cost of each resource is private and thus our goal is to design truthful mechanisms that provide a good approximation to the optimal value for the auctioneer, and are *budget-feasible*, i.e., the sum of the payments to the agents does not exceed the budget. This framework captures the budgeted versions of many well-known optimization problems and has been motivated by recent application scenarios including crowdsourcing platforms, where agents can be viewed as workers providing tasks, and influence maximization in social networks, where agents correspond to influential users.

We obtain mechanisms with significantly improved approximation ratios for several subclasses of submodular valuation functions. In Chapter 5 we provide a framework for designing deterministic mechanisms for non-decreasing submodular functions that have well-behaved LP formulations. As a highlight of our approach, we improve the best known factor for *coverage functions* by a factor of 3.

In Chapter 6 the main focus is on symmetric submodular functions, an eminent class of non-monotone objectives that contains *cut functions*. We propose truthful, budget-feasible mechanisms, greatly improving the known approximation ratios for these problems, while for many cases we achieve polynomial running time for the first time. As an example, for the budgeted weighted cut problem we obtain the first deterministic polynomial time mechanism with a 27.25-approximation, and for unweighted cut functions we improve the approximation ratio for randomized mechanisms, from 564 down to 10. Analogous improvements are obtained for arbitrary symmetric submodular functions. In the heart of our approach is a novel combination of (approximate) local search with known mechanisms for non-decreasing submodular functions.

Going beyond submodularity, in Chapter 7, we provide a general scheme for designing deterministic and randomized mechanisms for a subclass of XOS problems which contains problems whose feasible set forms an *independence system*. Some representative problems include, among others, finding maximum weighted matchings and maximum weighted matroid members. For most of such problems, only a randomized 768-approximate mechanism was known prior to our results. Finally, we briefly study the general class of XOS functions, where we improve the current upper bound by a factor of 3.

A purely algorithmic byproduct of our work, in Chapter 6, is a polynomial-time $\frac{2e}{e-1}$-approximation algorithm for symmetric submodular maximization subject to a budget constraint. Assuming only a value oracle for the objective function, this is the best known factor achieved by a deterministic algorithm.

My work on the above topics resulted in the following papers:

1. Approximation Algorithms for Computing Maximin Share Allocations, *42nd International Colloquium on Automata, Languages, and Programming, ICALP 2015, Proceedings, Part I*, pp. 39–51 (with E. Markakis, A.Nikzad, and A.Saberi). [Chapter 2]

2. Truthful Allocation Mechanisms Without Payments: Characterization and Implications on Fairness, *18th ACM conference on Economics and Computation, EC 2017, Proceedings*, pp. 545-562 (with G. Birmpas, G. Christodoulou, and E. Markakis). [Chapter 3]

3. On Truthful Mechanisms for Maximin Share Allocations, *25th International Joint Conference on Artificial Intelligence, IJCAI 2016, Proceedings*, pp. 31-37 (with G. Birmpas and E. Markakis).

4. On Budget-Feasible Mechanism Design for Symmetric Submodular Objectives, *submitted*, (with G. Birmpas and E. Markakis). [Chapters 5 and 6].

5. Coverage, Matching, and Beyond: New Results on Budgeted Mechanism Design, *12th Conference on Web and Internet Economics, WINE 2016, Proceedings*, pp. 414-428 (with G. Birmpas and E. Markakis). [Chapters 5 and 7]

In addition to fair division of indivisible items and reverse auctions with hard budget constraints, I also worked on other topics during my Ph.D. This resulted in the following completed papers:

6. Inequity Aversion Pricing over Social Networks: Approximation Algorithms and Hardness Results, *41st International Symposium on Mathematical Foundations of Computer Science, MFCS 2016, Proceedings*, pp. 9:1-9:13 (with E. Markakis and K. Sornat).

7. Multiple Referenda and Multiwinner Elections Using Hamming Distances: Complexity and Manipulability, *14th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Proceedings*, pp. 715-723 (with N. Barrot, J. Lang, E. Markakis, and B. Ries).

Part I

Fair Allocation of Indivisible Goods

Chapter 1

Introduction

1.1 Fair Division Without Incentives

In Chapter 2 we study a recently proposed fair division problem in the context of allocating indivisible goods. Fair division has attracted the attention of various scientific disciplines, including among others, mathematics, economics, and political science. Ever since the first attempt for a formal treatment by Steinhaus, Banach, and Knaster Steinhaus (1948), many interesting and challenging questions have emerged. Over the past decades, a vast literature has developed (see, e.g., Brams and Taylor, 1996; Robertson and Webb, 1998; Moulin, 2003) and several notions of fairness have been suggested. The area gradually gained popularity in computer science as well, as most of the questions are inherently algorithmic (see, among others, Even and Paz, 1984; Edmonds and Pruhs, 2006; Woeginger and Sgall, 2007) for earlier works and the surveys by Procaccia (2016) and by Bouveret, Chevaleyre, and Maudet (2016) on more recent results.

The objective in fair division problems is to allocate a set of resources to a set of n agents in a way that leaves every agent satisfied. In the continuous case, the available resources are typically represented by the interval [0, 1], whereas in the discrete case, we have a set of distinct, indivisible goods. The preferences of each agent are represented by a valuation function, which is usually an additive function (additive on the set of goods in the discrete case, or a probability distribution on $[0,1]$ in the continuous case). Given such a setup, many solution concepts have been proposed as to what constitutes a fair solution. Some of the standard ones include *proportionality, envy-freeness, equitability* and several variants of them. The most related concept to our work is proportionality, where an allocation is called proportional, if each agent receives a bundle of goods that is worth at least $1/n$ of the total value according to his valuation function.

Interestingly, all the above mentioned solutions and several others can be attained in the continuous case. Apart from mere existence, in some cases we can also have efficient algorithms, see, e.g.,Even and Paz (1984) for proportionality and Aziz and MacKenzie (2016) for some recent progress on envy-freeness. In the presence of indivisible goods however, the picture is quite different. We cannot guarantee existence and it is even NP-hard to decide whether a given instance admits fair allocations. In fact, in most cases it is hard to produce decent approximation guarantees.

Motivated by the question of what can we guarantee in the discrete case, we pay particular attention to a concept recently introduced by Budish (2011), that can be seen as a relaxation of proportionality. The rationale is as follows: suppose that an agent, say agent i, is asked to partition the goods into n bundles and then the rest of the agents choose a bundle before i. In the worst case, agent i will be left with his least valuable bundle. Hence, a risk-averse agent would choose a partition that maximizes the minimum value of a bundle in the partition. This value is called the maximin share of agent i. The objective then is to find an allocation where every agent receives at least his maximin share. Even for this notion, existence is not guaranteed under indivisible goods (Procaccia and Wang, 2014; Kurokawa, Procaccia, and Wang, 2016), despite the encouraging experimental evidence (Bouveret and Lemaître, 2016; Procaccia and Wang, 2014). However, it is possible to have constant factor approximations, as has been recently shown (Procaccia and Wang, 2014) (see also our related work section).

In Chapter 2 we study the above fair division problem mostly from an algorithmic point of view. That is, we try to design algorithms that output allocations such that each agent receives at least a constant fraction of his maximin share. We also try to justify the experimental evidence about the existence of such allocations. In Chapter 3, however, we turn to the game-theoretic version of the problem, where agents may have incentives to misreport their valuations. In fact, to study the fair division problem for two agents, we go way deeper and fully characterize all truthful mechanisms that fully allocate the set of items.

1.2 Item Allocation With Incentives

In Chapter 3 we study a very elementary and fundamental model for allocating indivisible goods from a mechanism design viewpoint. Namely, we consider a set of indivisible items that need to be allocated to a set of players. An outcome of the problem is an allocation of all the items to the players, i.e., a partition into bundles, and each player evaluates an allocation by his own additive valuation function. Our primary motivation originates from the fair division literature, where such models have been considered extensively. However, the same setting also appears in several domains, including job scheduling, load balancing and many other resource allocation problems.

Our focus is on understanding the interplay between truthfulness and fairness in this setting. Hence, we want to identify the effects on fairness guarantees, imposed by eliminating any incentives for the players to misreport their valuation functions. This type of questions has been posed already in previous works and for various notions of fairness, such as envy-freeness, or for the concept of maximin shares (see, among others, Lipton et al., 2004; Caragiannis et al., 2009; Amanatidis, Birmpas, and Markakis, 2016b). However, the results so far have been rather scarce in the sense that a) in most cases, they concern impossibility results which are far from being tight and b) the proof techniques are based on constructing specific families of instances that do not enhance our understanding on the structure of truthful mechanisms, with the exception of Caragiannis et al. (2009) which, however, is only for two players and two items.

In order to comprehend the trade-offs that are inherent between incentives and fairness, we first take a step back and focus solely on truthfulness itself. As is quite common in fair division models, we will not allow any monetary transfers, so that a mechanism simply outputs an allocation of the items. Hence, the question we want to begin with is: *what is the structure of truthful allocation mechanisms?*

There has been already a significant volume of works on characterizing truthful allocation mechanisms for indivisible items, yet there are some important differences from our approach. First, a typical line of work studies this question under the additional assumption of Pareto efficiency or related notions (Pápai, 2000; Klaus and Miyagawa, 2002; Ehlers and Klaus, 2003). The characterization results that have been obtained show that the combination of truthfulness together with Pareto efficiency tends to make the class of available deterministic mechanisms very poor; only some types of dictatorship survive when imposing both criteria. Second, in some cases the analysis is carried out without any restrictions on the class of valuation functions, which again often results in a very limited class of mechanisms (see, e.g., Pápai, 2001). When moving to a specific class, such as the class of additive functions which is usually assumed in fair division, it is conceivable that we can have a much richer class of truthful mechanisms. The results above indicate that the known characterizations of truthful mechanisms are also dependent on further assumptions, which may be well justified in various scenarios, but they are not aligned with the goal of fair division.

1.3 Related Work

For an overview of the classic fairness notions and related algorithmic results, we refer the reader to the books of Brams and Taylor (1996), and Robertson and Webb (1998). Maximin share fairness was introduced by Budish (2011) for ordinal utilities (i.e., agents have rankings over alternatives), building on concepts by Moulin (1990). Later on, Bouveret and Lemaître (2016) defined the notion for cardinal utilities, in the form that we study it here, and provided many important insights as well as experimental evidence. The first constant factor approximation algorithm was given by Procaccia and Wang (2014), achieving a 2/3-approximation but in time exponential in the number of agents.

On the negative side, constructions of instances where no maximin share allocation exists, even for $n = 3$, have been provided both by Procaccia and Wang (2014), and by Kurokawa, Procaccia, and Wang (2016). These elaborate constructions, along with the extensive experimentation of Bouveret and Lemaître (2016), reveal that it has been challenging to produce better lower bounds, i.e., instances where no α-approximation of a maximin share allocation exists, even for α very close to 1. Driven by these observations, a probabilistic analysis, similar in spirit but more general than ours, is carried out by Kurokawa, Procaccia, and Wang (2016). In our analysis in Section 2.4, all values are uniformly drawn from $[0, 1]$; Kurokawa, Procaccia, and Wang (2016) show a similar result with ours but for a a wide range of distributions over $[0, 1]$, establishing that maximin share allocations exist with high probability under all such distributions. However, their analysis, general as it may be, needs very large values of n to guarantee relatively high probability, hence it does not fully justify the experimental results discussed above.

Recently, some variants of the problem have also been considered. Barman and Murthy (2017) gave a constant factor approximation of $1/10$ for the case where the agents have submodular valuation functions. It remains an interesting open problem to determine whether better factors are achievable for submodular, or other non-additive functions. Along a different direction, Caragiannis et al. (2016) introduced the notion of *pairwise maximin share guarantee* and provided approximation algorithms. Although conceptually this is not too far apart from maximin shares, the two notions are incomparable.

A seemingly related problem is that of max-min fairness, also known as the Santa Claus problem (Asadpour and Saberi, 2007; Bansal and Sviridenko, 2006; Bezakova and Dani, 2005). In this problem we want to find an allocation where the value of the least happy person is maximized. With identical agents, this coincides with our problem, but beyond this special case the two problems exhibit very different behavior.

In the game-theoreting setting, the only work we are aware of in which a full characterization is given for truthful mechanisms with indivisible items, additive valuations, and no further assumptions, is by Caragiannis et al. (2009). However, this is only a characterization for two players and two items. Apart from characterizations, there have been several works that try to quantify the effects of truthfulness on several concepts of fairness. For the performance of truthful mechanisms with respect to envy-freeness, see Caragiannis et al. (2009) and Lipton et al. (2004), whereas for max-min fairness see Bezakova and Dani (2005). Coming to more recent results and along the same spirit, Amanatidis, Birmpas, and Markakis (2016b) and Markakis and Psomas (2011) study the notion of maximin share allocations, and a related notion of worst-case guarantees respectively. They obtain separation results, showing that the approximation factors achievable by truthful mechanisms are strictly worse than the known algorithmic (nontruthful) results. Obtaining a better understanding for the structure of truthful mechanisms and how they affect fairness has been an open problem underlying all the above works. For a better and more complete elaboration on fairness and the numerous fairness concepts that have been suggested, we refer the reader to the books of Brams and Taylor (1996), Robertson and Webb (1998), and Moulin (2003) and the recent surveys Bouveret, Chevaleyre, and Maudet (2016) and Procaccia (2016).

There has been a long series of works on characterizing mechanisms with indivisible items beyond the context of fair division. Many works characterize the allocation mechanisms that arise when we combine truthfulness with Pareto efficiency (see, e.g., Pápai, 2000; Klaus and Miyagawa, 2002; Ehlers and Klaus, 2003). Typically, such mechanisms tend to be dictatorial, and it is also well known that economic efficiency is mostly incompatible with fairness (see, e.g., Bouveret, Chevaleyre, and Maudet, 2016). Another assumption that has been used is nonbossiness, which means that one cannot change the outcome without affecting his own bundle. For instance, Svensson (1999) assumes nonbossiness in a setting where each player is interested in acquiring only one item. For general valuations, this also leads to dictatorial algorithms (Pápai, 2001). In most of these works ties are ignored by considering strict preference orders over all subsets of the items, while in some cases it is also allowed for the mechanism not to allocate all the items.

There is a relevant line of work for the setting of divisible goods (see, among others, Chen et al., 2013; Cole, Gkatzelis, and Goel, 2013). We note that for additive valuation functions, a mechanism for divisible items can be interpreted as a randomized mechanism for indivisible items. This connection is already discussed and explored in Guo and Conitzer (2010) and Aziz et al. (2016). In our work, we do not study randomized mechanisms, however it is an interesting question to have characterization results for such settings as well. Along this direction, see Mennle and Seuken (2014) where a relaxed notion of truthfulness is studied.

Related to our work is also the literature on exchange markets. These are models where players are equipped with an initial endowment, e.g., a house or a set of items. For the case where players can have multiple indivisible items as their initial endowment, see Pápai (2003) and Pápai (2007). Exchange markets provide an example where the existing characterizations go well beyond dictatorships and are closely related to the exchange component of our mechanisms.

Finally, for settings with payments, the work of Dobzinski and Sundararajan (2008), and independently of Christodoulou, Koutsoupias, and Vidali (2008), provided a characterization of truthful mechanisms with two players and additive valuations when all items are allocated. However, their characterization does not apply to our setting because they make an additional assumption, namely *decisiveness*. It roughly requires that each player should be able to receive any possible bundle of items, by making an appropriate bid. Their motivation is the characterization of truthful mechanisms with bounded makespan (maximum finishing time) for the scheduling problem, and in their case decisiveness is necessary in order to achieve bounded guarantees. In our case, our motivation is fairness, and decisiveness is a very strong assumption which has the opposite effects of what we need; e.g., assigning the full-bundle to a player is unacceptable in terms of fairness. Finally, Christodoulou and Kovács (2011) give a global characterization of envy-free and truthful mechanisms for settings with payments, when there are multiple players but only two items.

1.4 Preliminaries and Notation

For any $k \in \mathbb{N}$, we denote by $[k]$ the set $\{1,\dots,k\}$. Let $N = [n]$ be a set of n agents and $M = [m]$ be a set of indivisible items. Following the usual setup in the fair division literature, we assume each agent has an additive valuation function $v_i(\cdot)$, so that for every $S \subseteq M$, $v_i(S) = \sum_{j\in S} v_i(\{j\})$. For $j \in M$, we will use v_{ij} instead of $v_i(\{j\})$.

We say that $(T_1, T_2,\dots,T_k)$ is a *partition* of a set S, if $\bigcup_{i\in[k]} T_i = S$, and $T_i \cap T_j = \emptyset$ for any $i, j \in [k]$ with $i \neq j$. Note that, contrary to the usual definition of a partition, we do not require that $T_i \neq \emptyset$ for all $i \in [n]$. Given any subset $S \subseteq M$, an allocation of S to the n agents is a partition $T = (T_1,..., T_n)$. Let $\Pi_n(S)$ be the set of all partitions of a set S into n bundles.

1.4.1 Notation for $n = 2$

By $\mathcal{M}$ we denote the set of all allocations of M to two players.

Specifically for the two player case, the set $\mathcal{V}_m$ of all possible profiles is $\mathbb{R}^m_+ \times \mathbb{R}^m_+$, i.e., we assume that $v_{ij} > 0$ for every $i \in \{1,2\}$ and $j \in M$. For some statements we need the assumption that the players' valuation functions are such that no two sets have the same value. So, let $\mathcal{V}^{\neq}_m$ denote the set of such profiles, i.e.,

$$\mathcal{V}^{\neq}_m = \Big\{(v_1, v_2) \in \mathcal{V}_m \;\Big|\; \forall S,T \subseteq [m] \text{ with } S \neq T, \text{ and } \forall i \in \{1,2\}, \sum_{j\in S} v_{ij} \neq \sum_{j\in T} v_{ij}\Big\}.$$

A *deterministic allocation mechanism with no monetary transfers*, or simply a *mechanism*, for allocating all the items in $M = [m]$, is a mapping $\mathcal{X}$ from $\mathcal{V}_m$ to $\mathcal{M}$. For $n = 2$ this means that for any profile $\mathbf{v}$, the outcome of the mechanism is $\mathcal{X}(\mathbf{v}) = (X_1(\mathbf{v}), X_2(\mathbf{v})) \in \mathcal{M}$, and $X_i(\mathbf{v})$ denotes the set of items player i receives.

Definition 1.4.1. A mechanism $\mathcal{X}$ for two players is *truthful* if for any instance $\mathbf{v} = (v_1, v_2)$, any player $i \in \{1,2\}$, and any v'_i: $v_i(X_i(\mathbf{v})) \geq v_i(X_i(v'_i, \mathbf{v}_{-i}))$.

Since we will repeatedly argue about intersections of $X_i(\mathbf{v})$ with various subsets of M, we use $X^S_i(\mathbf{v})$ as a shortcut for $X_i(\mathbf{v}) \cap S$, where $S \subseteq M$.

1.4.2 Fairness concepts

Several notions have emerged throughout the years as to what can be considered a fair allocation. We define below the concepts that we will examine. We begin with two of the most dominant solution concepts in fair division, namely proportionality and envy-freeness.

Definition 1.4.2. An allocation $T = (T_1,..., T_n)$ is

1. *proportional*, if $v_i(T_i) \geq \frac{1}{n} v_i(M)$, for every $i \in [n]$.

2. *envy-free*, if for every $i, j \in [n]$, $v_i(T_i) \geq v_i(T_j)$.

Proportionality was considered in the very first work on fair division by Steinhaus (1948). Envy-freeness was suggested later by Gamow and Stern (1958), and with a more formal argumentation by Foley (1967) and Varian (1974).

Envy-freeness is a stricter notion than proportionality, but even for the latter existence cannot be guaranteed under indivisible goods. One can also consider approximation versions of these problems as follows: Given an instance I, let $E(I)$ be the minimum possible envy that can be achieved at I, among all possible allocations. We say that a mechanism achieves a ρ-approximation, if for every instance I, it produces an allocation where the envy between any pair of players is at most $\rho E(I)$. Similarly for proportionality, suppose that an instance I admits an allocation where every player receives a value of at least $\frac{c(I)}{n} v_i(M)$ for some $c(I) \leq 1$. Then a ρ-approximation would mean that each player is guaranteed a bundle with value at least $\frac{\rho c(I)}{n} v_i(M)$.

Apart from the approximation versions, the fact that we cannot always have proportional or envy-free allocations gives rise to relaxations of these definitions, with the hope of obtaining more positive results. We describe below three such relaxations, all of which admit either exact or constant-factor approximation algorithms (not necessarily truthful) in polynomial time.

The first such relaxation is the concept of envy-freeness up to one item, where each person may envy another player by an amount which does not exceed the value of a single item in the other player's bundle. Formally:

Definition 1.4.3. An allocation $T = (T_1, ..., T_n)$ is *envy-free up to one item*, if for every pair of agents $i, j \in [n]$, there exists an item $g \in T_j$, such that $v_i(T_i) \geq v_i(T_j \setminus \{g\})$.

It is quite easy to achieve envy-freeness up to one item, e.g., a round-robin algorithm that alternates between the players and gives them in each step their best remaining item suffices. Other algorithms are also known to satisfy this criterion (see Lipton et al., 2004).

A more interesting relaxation from an algorithmic point of view, comes from the notion of maximin share guarantees, recently proposed by Budish (2011). For two players, the maximin share of a player i is the value that he could achieve by being the cutter in a discretized form of the cut and choose protocol. This is a guarantee for player i, if he would partition the items into two bundles so as to maximize the value of the least valued bundle. We define below the approximate version of maximin share fairness for any number of agents.

Definition 1.4.4. Given a set of n agents, and any set $S \subseteq M$, the n-maximin share of an agent i with respect to S, is:

$$\boldsymbol{\mu}_i(n, S) = \max_{T \in \Pi_n(S)} \min_{T_j \in T} v_i(T_j).$$

Note that $\boldsymbol{\mu}_i(n, S)$ depends on the valuation function $v_i(\cdot)$ but is independent of any other function $v_j(\cdot)$ for $j \neq i$. When $S = M$, we refer to $\boldsymbol{\mu}_i(n, M)$ as the maximin share of agent i. The solution concept we study asks for a partition that gives each agent his maximin share.

Definition 1.4.5. Given a set of agents $N = [n]$, and a set of goods M, a partition $T = (T_1, ..., T_n) \in \Pi_n(M)$ is called a ρ-approximate *maximin share allocation* if $v_i(T_i) \geq \rho \cdot \boldsymbol{\mu}_i(n, M)$, for every agent $i \in [n]$. When $\rho = 1$, T is just called a maximin share allocation.

For two players maximin share allocations always exist and even though they are NP-hard to compute, we have a PTAS by reducing this to standard job scheduling problems. Hence each player can receive a value of at least $(1-\epsilon)\boldsymbol{\mu}_i$. For a higher number of players, however, as shown in Procaccia and Wang (2014), maximin share allocations do not always exist. Hence, our focus is on approximation algorithms, i.e., on algorithms that produce ρ-approximate maximin share allocations, for some $\rho \leq 1$.

Before we continue, a few words are in order regarding the appeal of this new concept. First of all, it is very easy to see that having a maximin share guarantee to every agent forms a relaxation of proportionality, see Claim 2.1.1. Given the known impossibility results for proportional allocations under indivisible items, it is worth investigating whether such relaxations are easier to attain. Second, the maximin share guarantee has an intuitive interpretation; for an agent i, it is the value that could be achieved if we run the generalization of the cut-and-choose protocol for multiple agents, with i being the cutter. In other words, it is the value that agent i can guarantee to himself, if he were given the advantage to control the partition of the items into bundles, but not the allocation of the bundles to the agents.

Example 1. Consider an instance with three agents and five items:

	a	b	c	d	e
Agent 1	1/2	1/2	1/3	1/3	1/3
Agent 2	1/2	1/4	1/4	1/4	0
Agent 3	1/2	1/2	1	1/2	1/2

If $M = \{a, b, c, d, e\}$ is the set of items, one can see that $\boldsymbol{\mu}_1(3, M) = 1/2$, $\boldsymbol{\mu}_2(3, M) = 1/4$, $\boldsymbol{\mu}_3(3, M) = 1$. E.g., for agent 1, no matter how he partitions the items into three bundles, the worst bundle will be worth at most 1/2 for him, and he achieves this with the partition $(\{a\}, \{b, c\}, \{d, e\})$. Similarly, agent 3 can guarantee a value of 1 (which is best possible as it is equal to $v_3(M)/n$) by the partition $(\{a, b\}, \{c\}, \{d, e\})$.

Note that this instance admits a maximin share allocation, e.g., $(\{a\}, \{b, c\}, \{d, e\})$, and in fact this is not unique. Note also that if we remove some agent, say agent 2, the maximin values for the other two agents increase. E.g., $\boldsymbol{\mu}_1(2, M) = 1$, achieved by the partition $(\{a, b\}, \{c, d, e\})$. Similarly, $\boldsymbol{\mu}_3(2, M) = 3/2$.

Finally, a related approach was undertaken by Hill (1987). This work examined what is the worst case guarantee that a player can have as a function of the total number of players and the maximum value of an item across all players. Since this fairness notion has a rather complicated definition and is never used for more than two agents in this thesis, we only define it for $n = 2$. In this case, the following function was identified precisely as the guarantee that can be given to each player. Note that the total value of the items is normalized to 1 in this case.

Definition 1.4.6. Let $V_2 : [0,1] \to [0,1/2]$ be the unique nonincreasing function satisfying $V_2(\alpha) = 1/2$ for $\alpha = 0$, whereas for $\alpha > 0$:

$$V_2(\alpha) = \begin{cases} 1 - k\alpha & \text{if } \alpha \in I_k \\ 1 - \frac{(k+1)}{2(k+1)-1} & \text{if } \alpha \in J_k \end{cases}$$

where for any integer $k \geq 1$, $I_k = \left[\frac{k+1}{k(2(k+1)-1)}, \frac{1}{2k-1}\right]$ and $J_k = \left(\frac{1}{2(k+1)-1}, \frac{k+1}{k(2(k+1)-1)}\right)$.

Markakis and Psomas (2011) proved that for two players, there always exists an allocation such that each player i receives at least $V_2(\alpha_i)$, where $\alpha_i = \max_{j \in [m]} v_{ij}$. The approximation version of this notion would be to construct allocations where each player receives a value of at least $\rho V_2(\alpha_i)$. Recently, a stricter variant of this guarantee has been provided by Gourvès, Monnot, and Tlilane (2015) (also see Remark 3.2.9).

Chapter 2

Computing Maximin Share Allocations[1]

In this chapter we deal with the problem of computing maximin share allocations from an algorithmic perspective, i.e., without taking incentives into account.

Our main result, in Section 2.2, is a $(2/3-\varepsilon)$-approximation algorithm, for any constant $\varepsilon > 0$, that runs in polynomial time for any number of agents and any number of goods. That is, the algorithm produces an allocation where every agent receives a bundle worth at least $2/3-\varepsilon$ of his maximin share. Our result improves upon the 2/3-approximation of Procaccia and Wang (2014), which runs in polynomial time only for a constant number of agents. To achieve this, we redesign certain parts of their algorithm, arguing about the existence of appropriate, carefully constructed matchings in a bipartite graph representation of the problem. Before that, in Section 2.1, we provide a much simpler and faster 1/2-approximation algorithm. Despite the worse factor, this algorithm still has its own merit due to its simplicity.

Moreover, we study two special cases, motivated by previous works. The first one is the case of $n = 3$ agents. This is an interesting turning point on the approximability of the problem; for $n = 2$, there always exist maximin share allocations, but adding a third agent makes the problem significantly more complex, and the best known ratio was 3/4 (Procaccia and Wang, 2014). We provide an algorithm with an approximation guarantee of 7/8, by examining more deeply the set of allowed matchings that we can use to satisfy the agents. The second case is the setting where all item values belong to $\{0,1,2\}$. This is an extension of the $\{0,1\}$ setting studied by Bouveret and Lemaître (2016) and we show that there always exists a maximin share allocation, for any number of agents.

Finally, motivated by the apparent difficulty in finding impossibility results on the approximability of the problem, we undertake a probabilistic analysis in Section 2.4. Our analysis shows that in randomly generated instances, maximin share allocations exist with high probability. This may be seen as a justification of the reported experimental evidence (Bouveret and Lemaître, 2016; Procaccia and Wang, 2014), which show that maximin share allocations exist in most cases.

[1] A conference paper containing most results of this chapter appeared in ICALP '15 (Amanatidis et al., 2015).

2.1 Warmup: Useful Properties and a 1/2-Approximation

We find it instructive to provide first a simpler and faster algorithm that achieves a worse approximation of 1/2. In the course of obtaining this algorithm, we also identify some important properties and insights that we will use in the next sections.

We start with an upper bound on our solution for each agent. The maximin share guarantee is a relaxation of proportionality, so we trivially have:

Claim 2.1.1. *For every $i \in N$ and every $S \subseteq M$,*

$$\boldsymbol{\mu}_i(n, S) \leq \frac{v_i(S)}{n} = \frac{\sum_{j \in S} v_{ij}}{n}.$$

Proof. This follows by the definition of maximin share. If there existed a partition where the minimum value for agent i exceeded the above bound, then the total value for agent i would be more than $\sum_{j \in S} v_{ij}$. □

Based on this, we now show how to get an additive approximation. Algorithm 1 below achieves an additive approximation of v_{max}, where $v_{max} = \max_{i,j} v_{ij}$. This simple algorithm, which we will refer to as the *Greedy Round-Robin Algorithm*, has also been discussed by Bouveret and Lemaître (2016), where it was shown that when all item values are in $\{0,1\}$, it produces an exact maximin share allocation. At the same time, we note that the algorithm also achieves envy-freeness up to one item, another solution concept defined by Budish (2011), and further discussed in Caragiannis et al. (2016). Finally, some variations of this algorithm have also been used in other allocation problems, see, e.g., Brams and King (2005), or the protocol in Bouveret and Lang (2011). We discuss further the properties of Greedy Round-Robin in Section 2.4.

In the statement of the algorithm below, the set V_N is the set of valuation functions $V_N = \{v_i : i \in N\}$, which can be encoded as a valuation matrix since the functions are additive.

ALGORITHM 1: Greedy Round-Robin(N, M, V_N)

```
Set S_i = ∅ for each i ∈ N.
Fix an ordering of the agents arbitrarily.
while ∃ unallocated items do
    S_i = S_i ∪ {j}, where i is the next agent to be examined in the current round
      (proceeding in a round-robin fashion) and j is i's most desired item among the
      currently unallocated items.
return (S_1, ..., S_n)
```

Theorem 2.1.2. *If $(S_1, \ldots, S_n)$ is the output of Algorithm 1, then for every $i \in N$,*

$$v_i(S_i) \geq \frac{\sum_{j \in M} v_{ij}}{n} - v_{max} \geq \boldsymbol{\mu}_i(n, M) - v_{max}.$$

Proof. Let $(S_1, \ldots, S_n)$ be the output of Algorithm 1. We first prove the following claim about the envy of each agent towards the rest of the agents:

Claim 2.1.3. *For every $i, j \in N$, $v_i(S_i) \geq v_i(S_j) - v_{max}$.*

Proof of Claim 2.1.3. Fix an agent i, and let $j \neq i$. We will upper bound the difference $v_i(S_j) - v_i(S_i)$. If j comes after i in the order chosen by the algorithm, then the statement of the claim trivially holds, since i always picks an item at least as desirable as the one j picks. Suppose that j precedes i in the ordering. The algorithm proceeds in $\ell = \lceil m/n \rceil$ rounds. In each round k, let r_k and r'_k be the items allocated to j and i respectively. Then

$$v_i(S_j) - v_i(S_i) = (v_{i,r_1} - v_{i,r'_1}) + (v_{i,r_2} - v_{i,r'_2}) + \cdots + (v_{i,r_\ell} - v_{i,r'_\ell}).$$

Note that there may be no item r'_ℓ in the last round if the algorithm runs out of goods but this does not affect the analysis (simply set $v_{i,r'_\ell} = 0$).

Since agent i picks his most desirable item when it is his turn to choose, this means that for two consecutive rounds k and $k+1$ it holds that $v_{i,r'_k} \geq v_{i,r_{k+1}}$. This directly implies that $v_i(S_j) - v_i(S_i) \leq v_{i,r_1} - v_{i,r'_\ell} \leq v_{i,r_1} \leq v_{max}$. ⊡

If we now sum up the statement of Claim 2.1.3 for each j, we get: $n v_i(S_i) \geq \sum_j v_i(S_j) - n v_{max}$, which implies

$$v_i(S_i) \geq \frac{\sum_j v_i(S_j)}{n} - v_{max} = \frac{\sum_{j \in M} v_{ij}}{n} - v_{max} \geq \boldsymbol{\mu}_i(n, M) - v_{max},$$

where the last inequality holds by Claim 2.1.1. □

The next important ingredient is the following monotonicity property, which says that we can allocate a single good to an agent without decreasing the maximin share of other agents. Note that this lemma also follows from Lemma 1 of Bouveret and Lemaître (2016), yet, for completeness, we prove it here as well.

Lemma 2.1.4 (Monotonicity property)**.** *For any agent i and any good j, it holds that*

$$\boldsymbol{\mu}_i(n-1, M \setminus \{j\}) \geq \boldsymbol{\mu}_i(n, M).$$

Proof. Let us look at agent i, and consider a partition of M that attains his maximin share. Let $(S_1, ..., S_n)$ be this partition. Without loss of generality, suppose $j \in S_1$. Consider the remaining partition $(S_2, ..., S_n)$ enhanced in an arbitrary way by the items of $S_1 \setminus \{j\}$. This is a $(n-1)$-partition of $M \setminus \{j\}$ where the value of agent i for any bundle is at least $\boldsymbol{\mu}_i(n, M)$. Thus, we have $\boldsymbol{\mu}_i(n-1, M \setminus \{j\}) \geq \boldsymbol{\mu}_i(n, M)$. □

We are now ready for the 1/2-approximation, obtained by Algorithm 2 below, which is based on using Greedy Round-Robin, but only after we allocate first the most valuable goods. This is done so that the value of v_{max} drops to an extent that Greedy Round-Robin can achieve a multiplicative approximation.

ALGORITHM 2: APX-MMS$_{1/2}(N, M, V_N)$

1 Set $S = M$
2 **for** $i = 1$ *to* $|N|$ **do**
3   Let $\alpha_i = \frac{\sum_{j \in S} v_{ij}}{|N|}$
4 **while** $\exists i, j$ *s.t.* $v_{ij} \geq \alpha_i/2$ **do**
5   Allocate j to i.
6     $S = S \setminus \{j\}$
7     $N = N \setminus \{i\}$
8   Recompute the α_is.
9 Run Greedy Round-Robin on the remaining instance.

Theorem 2.1.5. *Let N be a set of n agents, and let M be a set of goods. Algorithm 2 produces an allocation $(S_1, ..., S_n)$ such that*

$$v_i(S_i) \geq \frac{1}{2}\boldsymbol{\mu}_i(n, M), \ \forall i \in N.$$

Proof. We will distinguish two cases. Consider an agent i who was allocated a single item during the first phase of the algorithm (lines 4 - 8). Suppose that at the time when i was given his item, there were n_1 active agents, $n_1 \leq n$, and that S was the set of currently unallocated items. By the design of the algorithm, this means that the value of what i received is at least

$$\frac{\sum_{j \in S} v_{ij}}{2n_1} \geq \frac{1}{2}\boldsymbol{\mu}_i(n_1, S)$$

where the inequality follows by Claim 2.1.1. But now if we apply the monotonicity property (Lemma 2.1.4) $n - n_1$ times, we get that $\boldsymbol{\mu}_i(n_1, S) \geq \boldsymbol{\mu}_i(n, M)$, and we are done.

Consider now an agent i, who gets a bundle of goods according to Greedy Round-Robin, in the second phase of the algorithm. Let n_2 be the number of active agents at that point, and S be the set of goods that are unallocated before Greedy Round-Robin is executed. We know that v_{max} at that point is less than half the current value of α_i for agent i. Hence by the additive guarantee of Greedy Round-Robin, we have that the bundle received by agent i has value at least

$$\frac{\sum_{j \in S} v_{ij}}{n_2} - v_{max} > \frac{\sum_{j \in S} v_{ij}}{n_2} - \frac{\alpha_i}{2} = \frac{\sum_{j \in S} v_{ij}}{2n_2} \geq \frac{1}{2}\boldsymbol{\mu}_i(n_2, S).$$

Again, after applying the monotonicity property repeatedly, we get that $\boldsymbol{\mu}_i(n_2, S) \geq \boldsymbol{\mu}_i(n, M)$, which completes the proof. □

2.2 A Polynomial Time $\left(\frac{2}{3} - \varepsilon\right)$-Approximation

The main result of this section is Theorem 2.2.1, establishing a polynomial time algorithm for achieving a 2/3-approximation to the maximin share of each agent.

Theorem 2.2.1. *Let N be a set of n agents, and let M be a set of goods. For any constant $\varepsilon > 0$, Algorithm 3 produces in polynomial time an allocation $(S_1, ..., S_n)$, such* *that*

$$v_i(S_i) \geq \left(\frac{2}{3}-\varepsilon\right)\boldsymbol{\mu}_i(n,M),\ \forall i \in N.$$

Our result is based on the algorithm by Procaccia and Wang (2014), which also guarantees to each agent a 2/3-approximation. However, their algorithm runs in polynomial time only for a constant number of agents. Here, we identify the source of exponentiality and take a different approach regarding certain parts of the algorithm. For the sake of completeness, we first present the necessary related results of Procaccia and Wang (2014), before we discuss the steps that are needed to obtain our result.

First of all, we note that even the computation of the maximin share values is already a hard problem. For a single agent i, the problem of deciding whether $\boldsymbol{\mu}_i(n,M) \geq k$ for a given k is NP-complete. However, a PTAS follows by the work of Woeginger (1997). In the original paper, which is in the context of job scheduling, Woeginger gave a PTAS for maximizing the minimum completion time on identical machines. But this scheduling problem is identical to computing a maximin partition with respect to a given agent i. Indeed, from agent i's perspective, it is enough to think of the machines as identical agents (the only input that we need for computing $\boldsymbol{\mu}_i(n,M)$ is the valuation function of i). Hence:

Theorem 2.2.2 (Follows by Woeginger (1997))**.** *Suppose we have a set M of goods to be divided among n agents. Then, for each agent i, there exists a PTAS for approximating $\boldsymbol{\mu}_i(n,M)$.*

A central quantity in the algorithm of Procaccia and Wang (2014) is the *n-density balance parameter*, denoted by ρ_n and defined below. Before stating the definition, we give for clarity the high level idea, which can be seen as an attempt to generalize the monotonicity property of Lemma 2.1.4. Assume that in the course of an algorithm, we have used a subset of the items to "satisfy" some of the agents, and that those items do not have "too much" value for the rest of the agents. If k is the number of remaining agents, and S is the remaining set of goods, then we should expect to be able to "satisfy" these k agents using the items in S. A good approximation in this reduced instance however, would only be an approximation with respect to $\boldsymbol{\mu}_i(k,S)$. Hence, in order to hope for an approximation algorithm for the original instance, we would need to examine how $\boldsymbol{\mu}_i(k,S)$ relates to $\boldsymbol{\mu}_i(n,M)$. Essentially, the parameter ρ_n is the best guarantee one can hope to achieve for the remaining agents, based only on the fact that the complement of the set left to be shared is of relatively small value. Formally:

Definition 2.2.3 (*Procaccia and Wang (2014)*)**.** For any number n of agents, let

$$\rho_n = \max\left\{\lambda \;\middle|\; \begin{array}{l} \forall M, \forall \text{ additive } v_i \in (\mathbb{R}^+)^{2^M}, \forall S \subseteq M, \forall k,\ell \text{ s.t. } k+\ell=n, \\ v_i(M\setminus S) \leq \ell\lambda\boldsymbol{\mu}_i(n,M) \Rightarrow \boldsymbol{\mu}_i(k,S) \geq \lambda\boldsymbol{\mu}_i(n,M) \end{array}\right\}.$$

After a quite technical analysis, Procaccia and Wang calculate the exact value of ρ_n in the following lemma.

Lemma 2.2.4 (Lemma 3.2 of Procaccia and Wang (2014))**.** *For any $n \geq 2$,*

$$\rho_n = \frac{2\lfloor n \rfloor_{odd}}{3\lfloor n \rfloor_{odd} - 1} > \frac{2}{3},$$

where $\lfloor n \rfloor_{odd}$ denotes the largest odd integer less than or equal to n.

We are now ready to state our algorithm, referred to as APX-MMS (Algorithm 3 below). We elaborate on the crucial differences between Algorithm 3 and the result of Procaccia and Wang (2014) after Lemma 2.2.5. At first, the algorithm computes each agent's $(1-\varepsilon')$-approximate maximin value using Woeginger's PTAS, where $\varepsilon' = \frac{3\varepsilon}{4}$. Let $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_n)$ be the vector of these values. Hence, $\forall i,\ \boldsymbol{\mu}_i(n, M) \geq \xi_i \geq (1-\varepsilon')\boldsymbol{\mu}_i(n, M)$. Then, APX-MMS makes a call to the recursive algorithm REC-MMS (Algorithm 4) to compute a $\left(\frac{2}{3}-\varepsilon\right)$-approximate partition. REC-MMS takes the arguments ε', $n = |N|$, $\boldsymbol{\xi}$, S (the set of items that have not been allocated yet), K (the set of agents that have not received a share of items yet), and the valuation functions $V_K = \{v_i | i \in K\}$.

The guarantee provided by REC-MMS is that as long as the already allocated goods are not worth too much for the currently active agents of K, we can satisfy them with the remaining goods. More formally, under the assumption that

$$\forall i \in K, \quad v_i(M \setminus S) \leq (n - |K|)\rho_n \boldsymbol{\mu}_i(n, M), \tag{2.1}$$

which we will show that it holds before each call, REC-MMS$(\varepsilon', n, \boldsymbol{\xi}, S, K, V_K)$ computes a $|K|$-partition of S, so that each agent receives items of value at least $(1-\varepsilon')\rho_n \xi_i$.

The initial call of the recursion is, of course, REC-MMS$(\varepsilon', n, \boldsymbol{\xi}, M, N, V_N)$. Before moving on to the next recursive call, REC-MMS appropriately allocates some of the items to some of the agents, so that they receive value at least $(1-\varepsilon')\rho_n \xi_i$ each. This is achieved by identifying an appropriate matching between some currently unsatisfied agents and certain bundles of items, as described in the algorithm. In particular, the most important step in the algorithm is to first compute the set X^+ (line 6), which is the set of agents that will not be matched in the current call. The remaining active agents, i.e., $K \setminus X^+$, are then guaranteed to get matched in the current round, whereas X^+ will be satisfied in the next recursive calls. In order to ensure this for X^+, REC-MMS guarantees that inequality (2.1) holds for $K = X^+$ and with S being the rest of the items. Note that (2.1) trivially holds for the initial call of REC-MMS, where $K = N$ and $S = M$.

ALGORITHM 3: APX-MMS(ε, N, M, V_N)

$\varepsilon' = \frac{3\varepsilon}{4}$

for $i = 1$ *to* $|N|$ **do**

  Use Woeginger's PTAS to compute a $(1-\varepsilon')$-approximation ξ_i of $\boldsymbol{\mu}_i(|N|, M)$. Let $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_n)$.

return REC-MMS$(\varepsilon', |N|, \boldsymbol{\xi}, M, N, V_N)$

For simplicity, in the description of REC-MMS, we assume that $K = \{1, 2, \ldots, |K|\}$. Also, for the bipartite graph defined in line 5 of the algorithm, by $\Gamma(X^+)$ we denote the set of neighbors of the vertices in X^+.

ALGORITHM 4: REC-MMS$(\varepsilon', n, \boldsymbol{\xi}, S, K, V_K)$

1 **if** $|K| = 1$ **then**
2   Allocate all of S to agent 1.
3 **else**
4   Use Woeginger's PTAS to compute a $(1-\varepsilon')$-approximate $|K|$-maximin partition of S with respect to agent 1 from K, say $(S_1, \ldots, S_{|K|})$.
5   Create a bipartite graph $G = (X \cup Y, E)$, where $X = Y = K$ and $E = \{(i, j) \mid i \in X, j \in Y, v_i(S_j) \geq (1-\varepsilon')\rho_n \xi_i\}$.
6   Find a set $X^+ \subset X$, as described in Lemma 2.2.5.
7   Given a perfect matching A, between $X \setminus X^+$ and a subset of $Y \setminus \Gamma(X^+)$, allocate S_j to agent i iff $(i, j) \in A$ (the matching is a byproduct of line 6).
8   **if** $X^+ = \emptyset$ **then**
9   Output the above allocation.
10   **else**
11   Output the above allocation, together with REC-MMS$(\varepsilon', n, \boldsymbol{\xi}, S^*, X^+, V_{X^+})$, where S^* is the subset of S not allocated in line 7.

To proceed with the analysis, and since the choice of X^+ plays an important role (line 6 of Algorithm 4), we should first clarify what properties of X^+ are needed for the algorithm to work. Lemma 2.2.5 is the most crucial part in the design of our algorithm.

Lemma 2.2.5. *Assume that for n, M, S, K, V_K inequality (2.1) holds and let $G = (X \cup Y, E)$ be the bipartite graph defined in line 5 of* REC-MMS. *Then there exists a subset X^+ of $X \setminus \{1\}$, such that:*

(i) X^+ can be found efficiently.

(ii) There exists a perfect matching between $X \setminus X^+$ and a subset of $Y \setminus \Gamma(X^+)$.

(iii) If we allocate subsets to agents according to such a matching (as described in line 7) and $X^+ \neq \emptyset$, then inequality (2.1) holds for n, M, S^*, X^+, V_{X^+} where $S^* \subseteq S$ is the unallocated set of items, i.e.:

$$\forall i \in X^+, \quad v_i(M \setminus S^*) \leq (n - |X^+|)\rho_n \boldsymbol{\mu}_i(n, M).$$

Before we prove Lemma 2.2.5, we elaborate on the main differences between our setup and the approach of Procaccia and Wang (2014):

Choice of X^+. In Procaccia and Wang (2014), X^+ is defined as $\arg\max_{Z \subseteq K \setminus \{1\}} \{|Z| \mid |Z| \geq |\Gamma(Z)|\}$. Clearly, when n is constant, so is $|K|$, and thus the computation of X^+ is trivial. However, it is not clear how to efficiently find such a set in general, when n is not constant. We propose a definition of X^+, which is efficiently computable and has the desired properties. In short, our X^+ is any appropriately selected counterexample to Hall's Theorem for the graph G constructed in line 5.

Choice of ε. The algorithm works for any $\varepsilon > 0$, but Procaccia and Wang (2014) choose an ε that depends on n, and it is such that $(1-\varepsilon)\rho_n \geq \frac{2}{3}$. This is possible since for any n, $\rho_n \geq \frac{2}{3}\big(1+\frac{1}{3n-1}\big)$. However, in this case, the running time of Woeginger's PTAS (line 4) is not polynomial in n. Here, we consider any fixed ε, independent of n, hence the approximation ratio of $\frac{2}{3}-\varepsilon$.

The formal definition of X^+ is given within the proof of Lemma 2.2.5 that follows.

Proof of Lemma 2.2.5. We will show that either $X^+ = \emptyset$ (in the case where G has a perfect matching), or some set X^+ with $X^+ \in \big\{Z \subseteq X : |Z| > |\Gamma(Z)| \wedge \exists \text{ matching of size } |X \setminus Z| \text{ in } G\setminus\{Z \cup \Gamma(Z)\}\big\}$ has the desired properties. Moreover, we propose a way to find such a set efficiently. We first find a maximum matching B of G. If $|B| = |K|$, then we are done, since for $X^+ = \emptyset$, properties *(i)* and *(ii)* of Lemma 2.2.5 hold, while we need not check *(iii)*. If $|B| < |K|$, then there must be a subset of X violating the condition of Hall's Theorem.[2] Let X_u, X_m be the partition of X in unmatched and matched vertices respectively, according to B, with $X_u \neq \emptyset$, $X_m \neq \emptyset$. Similarly, we define Y_u, Y_m.

We now construct a directed graph $G' = (X \cup Y, E')$, where we direct all edges of G from X to Y, and on top of that, we add one copy of each edge of the matching but with direction from Y to X. In particular, $\forall i \in X, \forall j \in Y$, if $(i,j) \in E$ then $(i,j) \in E'$, and moreover if $(i,j) \in B$ then $(j,i) \in E'$. We claim that the following set satisfies the desired properties

$$X^+ := X_u \cup \{v \in X : v \text{ is reachable from } X_u \text{ in } G'\}.$$

Note that X^+ is easy to compute; after finding the maximum matching in G, and constructing G', we can run a depth-first search in each connected component of G', starting from the vertices of X_u. See also Figure 2.1, after the proof of Theorem 2.2.1 for an illustration.

Given the definition of X^+, we now show property *(ii)*. Back to the original graph G, we first claim that $|X^+| > |\Gamma(X^+)|$. To prove this, note that if $j \in \Gamma(X^+)$ in G, then $j \in Y_m$. If not, then it is not difficult to see that there is an augmenting path from a vertex in X_u to j, which contradicts the maximality of B. Indeed, since $j \in \Gamma(X^+)$, let i be a neighbor of j in X^+. If $i \in X_u$, then the edge (i,j) would enlarge the matching. Otherwise, $i \in X_m$ and since also $i \in X^+$, there is a path in G' from some vertex of X_u to i. But this path by construction of the directed graph G' must consist of an alternation of unmatched and matched edges, hence together with (i,j) we have an augmenting path.

Therefore, $\Gamma(X^+) \subseteq Y_m$, i.e., for any $j \in \Gamma(X^+)$, there is an edge (i,j) in the matching B. But then i has to belong to X^+ by the construction of G' (and since $j \in \Gamma(X^+)$).

[2] The special case of Hall's Theorem (Hall, 1935) used here, states that given a bipartite graph $G = (X \cup Y, E)$, where X, Y are disjoint independent sets with $|X| = |Y|$, there is a perfect matching in G if and only if $|W| \leq |\Gamma(W)|$ for every $W \subseteq X$.

To sum up: for any $j \in \Gamma(X^+)$, there is exactly one distinct vertex i, with $(i,j) \in E$, and $i \in X^+ \cap X_m$, i.e., $|X^+ \cap X_m| \geq |\Gamma(X^+)|$. In fact, we have equality here, because it is also true that for any $i \in X^+ \cap X_m$, there is a distinct vertex $j \in Y_m$ which is trivially reachable from X^+. Hence, $|X^+ \cap X_m| = |\Gamma(X^+)|$. Since $X_u \neq \emptyset$, we have $|X^+| = |X_u| + |X^+ \cap X_m| \geq 1 + |\Gamma(X^+)|$. So, $|X^+| > |\Gamma(X^+)|$.

Also, note that $X^+ \subseteq X \setminus \{1\}$, because for any $Z \subseteq X$ that contains vertex 1 we have $|\Gamma(Z)| = |K| \geq |Z|$. This is due to the fact that for any vertex $j \in Y$, the edge $(1,j)$ is present by the construction, since $v_1(S_j) \geq (1-\varepsilon')\boldsymbol{\mu}_1(k,S) \geq (1-\varepsilon')\rho_n\boldsymbol{\mu}_1(n,M) \geq (1-\varepsilon')\rho_n\xi_1$, for all $1 \leq j \leq |K|$.

We now claim that if we remove X^+ and $\Gamma(X^+)$ from G, then the restriction of B on the remaining graph, still matches all vertices of $X \setminus X^+$, establishing property *(ii)*. Indeed, note first that for any $i \in X \setminus X^+$, it has to hold that $i \in X_m$, since X^+ contains X_u. Also, for any edge $(i,j) \in B$ with $i \in X$ and $j \in \Gamma(X^+)$, we have $i \in X^+$ by the construction of X^+. So, for any $i \in X \setminus X^+$, its pair in B belongs to $Y \setminus \Gamma(X^+)$. Equivalently, B induces a perfect matching between $X \setminus X^+$ and a subset of $Y \setminus \Gamma(X^+)$ (this is the matching A in line 7 of the algorithm).

What is left to prove is that property *(iii)* also holds for X^+. This can be done by the same arguments as in Procaccia and Wang (2014), specifically by the following lemma which can be inferred from their work.

Lemma 2.2.6 (Procaccia and Wang (2014), end of Subsection 3.1)**.** *Assume that inequality* (2.1) *holds for n, M, S, K, V_K, and let G be the graph defined in line* 5*. For any $Z \subseteq X$, if there exists a perfect matching between $X \setminus Z$ and a subset of $Y \setminus \Gamma(Z)$, say Y^*, and there are no edges between Z and Y^* in G, then property (iii) holds as well.*

Clearly, there are no edges between X^+ and $Y \setminus \Gamma(X^+)$. Hence, Lemma 2.2.6 can be applied to X^+, completing the proof. □

Given Lemma 2.2.5, we can now prove the main result of this section, the correctness of APX-MMS.

Proof of Theorem 2.2.1. It is clear that the running time of the algorithm is polynomial. Its correctness is based on the correctness of REC-MMS. The latter can be proven with strong induction on $|K|$, the number of still active agents that REC-MMS receives as input, under the assumption that (2.1) holds before each new call of REC-MMS (which we have established by Lemma 2.2.5). For $|K| = 1$, assuming that inequality (2.1) holds, we have for agent 1 of K:

$$\begin{aligned} v_1(S) = v_1(M) - v_1(M \setminus S) &\geq n\boldsymbol{\mu}_1(n,M) - (n-1)\rho_n\boldsymbol{\mu}_1(n,M) \\ &\geq \boldsymbol{\mu}_1(n,M) \geq \left(\frac{2}{3}-\varepsilon\right)\boldsymbol{\mu}_1(n,M). \end{aligned}$$

For the inductive step, Lemma 2.2.5 and the choice of X^+ are crucial. Consider an execution of REC-MMS during which some agents will receive a subset of items and the rest will form the set X^+ to be handled recursively. For all the agents in X^+ –if any– we are guaranteed $(\frac{2}{3}-\varepsilon)$-approximate shares by property *(iii)* of Lemma 2.2.5 and by the inductive hypothesis. On the other hand, for each agent i that receives a subset S_j of items in line 7, we have

$$v_i(S_j) \geq (1-\varepsilon')\rho_n\xi_i \geq (1-\varepsilon')^2\rho_n\boldsymbol{\mu}_i(n,M) > (1-2\varepsilon')\frac{2}{3}\boldsymbol{\mu}_i(n,M) = \left(\frac{2}{3}-\varepsilon\right)\boldsymbol{\mu}_i(n,M),$$

where the first inequality holds because $(i,j) \in E(G)$. □

In Figure 2.1, we give a simple snapshot to illustrate a recursive call of REC-MMS. In particular, in Subfigure 2.1(a), we see a bipartite graph G that could be the current configuration for REC-MMS, along with a maximum matching. In Subfigure 2.1(b), we see the construction of G', as described in Lemma 2.2.5, and the set X^+. The bold (black) edges in G' signify that both directions are present. The set X^+ consists then of X_u and all other vertices of X reachable from X_u. Finally, Subfigure 2.1(b) also shows the set of agents that are satisfied in the current call along with the corresponding perfect matching, as claimed in Lemma 2.2.5.

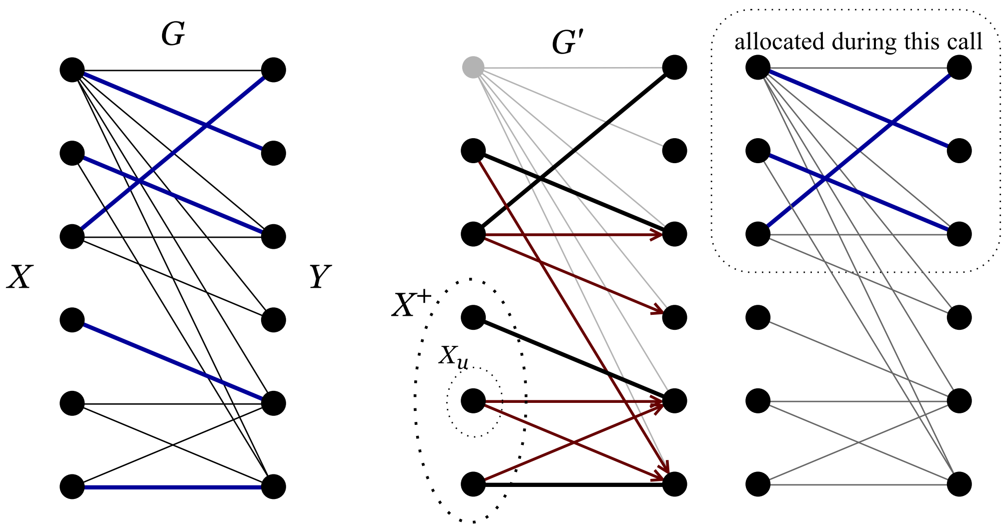

(a) The graph G defined in line 5 of Algorithm 4 shown with a maximum matching (blue edges). Agent 1 is the top vertex of X.

(b) The graph G' defined in the proof of Lemma 2.2.5, where for clarity, agent 1 and his edges are grayed out. The black edges signify that both directions are present, i.e., they correspond to pairs of anti-parallel edges. On the right we show the actual allocation resulting from G.

FIGURE 2.1: Ilustration of G, G' and X^+.

We note that the analysis of the algorithm is tight, given the analysis on ρ_n (see Section 3.3 of Procaccia and Wang (2014)). Improving further on the approximation ratio of 2/3 seems to require drastically new ideas and it is a challenging open problem. We stress that even a PTAS is not currently ruled out by the lower bound constructions (Kurokawa, Procaccia, and Wang, 2016; Procaccia and Wang, 2014). Related to this, in the next section we consider two special cases in which we can obtain better positive results.

2.3 Two Special Cases

In this section, we consider two interesting special cases, where we have improved approximations. The first is the case of $n = 3$ agents, where we obtain a 7/8-approximation, improving on the 3/4-approximation of Procaccia and Wang (2014). The second is the case where all values for the goods belong to $\{0,1,2\}$. This is an extension of the $\{0,1\}$ setting discussed in Bouveret and Lemaître (2016), and we show how to get an exact allocation without any approximation loss.

2.3.1 The Case of Three Agents

For $n = 2$, it is pointed out in Bouveret and Lemaître (2016) that maximin share allocations exist via an analog of the cut and choose protocol. Using the PTAS of Woeginger (1997), we can then have a $(1-\varepsilon)$-approximation in polynomial time. In contrast, as soon as we move to $n = 3$, things become more interesting. It is proven that with 3 agents there exist instances where no maximin share allocation exists (Procaccia and Wang, 2014). The best known approximation guarantee is $\frac{3}{4}$ by observing that the quantity ρ_n, defined in Section 2.2, satisfies $\rho_3 \geq \frac{3}{4}$.

We provide a different algorithm, improving the approximation to $\frac{7}{8} - \varepsilon$. To do this, we combine ideas from both algorithms presented so far in Sections 2.1 and 2.2. The main result of this subsection is as follows:

Theorem 2.3.1. *Let $N = \{1,2,3\}$ be a set of three agents with additive valuations, and let M be a set of goods. For any constant $\varepsilon > 0$, Algorithm 5 produces in polynomial time an allocation (S_1, S_2, S_3), such that*

$$v_i(S_i) \geq \left(\frac{7}{8} - \varepsilon\right) \boldsymbol{\mu}_i(3, M), \ \forall i \in N.$$

The algorithm is shown below. Before we prove Theorem 2.3.1, we provide here a brief outline of how the algorithm works.

Algorithm Outline: First, approximate values for the $\boldsymbol{\mu}_i$s are calculated as before. Then, if there are items with large value to some agent, in analogy to Algorithm 2, we first allocate one of those reducing this way the problem to the simple case of $n = 2$. If there are no items of large value, then the first agent partitions the items as in Algorithm 4. In the case where this partition does not satisfy all three agents, then the second agent repartitions two of the bundles of the first agent. Actually, she tries two different such repartitions, and we show that at least one of them works out. The definition of a bipartite preference graph and a corresponding matching (as in Algorithm 4) is never mentioned explicitly here. However, the main idea (and the difference with Algorithm 4) is that if there are several ways to pick a perfect matching between $X \backslash X^+$ and a subset of $Y \backslash \Gamma(X^+)$, then we try them all and choose the best one. Of course, since $n = 3$, if there is no perfect matching in the preference graph, then $X \setminus X^+$ is going to be just a single vertex, and we only have to examine two possible perfect matchings between $X \setminus X^+$ and a subset of $Y \setminus \Gamma(X^+)$.

ALGORITHM 5: APX-3-MMS$(\varepsilon, M, v_1, v_2, v_3)$

---

1 $\varepsilon' = \frac{8}{7}\varepsilon$

2 Compute a $(1-\varepsilon)$-approximation ξ_i of $\boldsymbol{\mu}_i(3,M)$ for $i \in \{1,2,3\}$.

3 **if** $\exists i \in \{1,2,3\}, j \in M$ *such that* $v_{ij} \geq \frac{7}{8}\xi_i$ **then**

4 Give item j to agent i and divide $M \setminus \{j\}$ among the other two agents in a "cut-and-choose" fashion.

5 **else**

6 Agent 1 computes a $(1-\varepsilon)$-approximate maximin partition of M into three sets, say (A_1, A_2, A_3).

7 **if** $\exists j_2, j_3 \in \{1,2,3\}$ *such that* $j_2 \neq j_3$, $v_2(A_{j_2}) \geq \frac{7}{8}\xi_2$ *and* $v_3(A_{j_3}) \geq \frac{7}{8}\xi_3$ **then**

8 Give set A_{j_2} to agent 2, set A_{j_3} to agent 3, and the last set to agent 1.

9 **else**

10 There are two sets that have value less than $\frac{7}{8}\xi_2$ w.r.t. agent 2, say for simplicity A_2 and A_3.

11 Agent 2 computes $(1-\varepsilon')$-approximate 2-maximin partitions of $A_1 \cup A_2$ and $A_1 \cup A_3$, say (B_1, B_2) and (B_1', B_2') respectively, and discards the partition with the smallest maximin value. Let (D_1, D_2) be the partition he keeps.

12 Agent 3 takes the set he prefers from (D_1, D_2); agent 2 gets the other, and agent 1 gets $M \setminus (D_1 \cup D_2)$.

Proof of Theorem 2.3.1. First, note that for constant ε the algorithm runs in time polynomial in $|M|$. Next, we prove the correctness of the algorithm.

If the output is computed in lines 3-4 then for agent i, as defined in line 3, the value he receives is at least $\frac{7}{8}\xi_i \geq \frac{7}{8}(1-\varepsilon)\boldsymbol{\mu}_i(3,M) > \left(\frac{7}{8}-\varepsilon\right)\boldsymbol{\mu}_i(3,M)$. The remaining two agents i_1, i_2 essentially apply an approximate version of a cut and choose protocol. Agent i_1 computes a $(1-\varepsilon)$-approximate 2-maximin partition of $M \setminus \{j\}$, say (C_1, C_2), then agent i_2 takes the set he prefers among C_1 and C_2, and agent i_1 gets the other. By the monotonicity lemma (Lemma 2.1.4), we know that $\boldsymbol{\mu}_{i_1}(2, M \setminus \{j\}) \geq \boldsymbol{\mu}_{i_1}(3,M)$, and thus no matter which set is left for agent i_1, he is guaranteed a total value of at least $(1-\varepsilon)\boldsymbol{\mu}_{i_1}(3,M) > \left(\frac{7}{8}-\varepsilon\right)\boldsymbol{\mu}_{i_1}(3,M)$. Similarly, we have $\boldsymbol{\mu}_{i_2}(2, M \setminus \{j\}) \geq \boldsymbol{\mu}_{i_2}(3,M)$, and therefore $v_{i_2}(M \setminus \{j\}) \geq 2\boldsymbol{\mu}_{i_2}(3,M)$. Since i_2 chooses before i_1, he is guaranteed a total value that is at least $\boldsymbol{\mu}_{i_2}(3,M) > \left(\frac{7}{8}-\varepsilon\right)\boldsymbol{\mu}_{i_2}(3,M)$.

If the output is computed in lines 6-8 then clearly all agents receive a $(7/8-\varepsilon)$-approximation, since for agent 1 it does not matter which of the A_is he gets.

The most challenging case is when the output is computed in lines 10-12 (starting with the partition from line 6). Then, as before, agent 1 receives a value that is at least a $(7/8-\varepsilon)$-approximation no matter which of the three sets he gets. For agents 2 and 3, however, the analysis is not straightforward. We need the following lemma.

Lemma 2.3.2. *Let N, M, ε be as above, such that for all $j \in M$ we have $v_{2j} < \frac{7}{8}\xi_2$. Consider any partition of M into 3 sets A_1, A_2, A_3 and assume that there are no $j_2, j_3 \in \{1,2,3\}$ such that $j_2 \neq j_3$, $v_2(A_{j_2}) \geq \frac{7}{8}\xi_2$ and $v_3(A_{j_3}) \geq \frac{7}{8}\xi_3$. Then lines 10-12 of Algorithm 5 produce an allocation (S_2, S_3) for agents 2 and 3, such that for $i \in \{2,3\}$: $v_i(S_i) \geq \left(\frac{7}{8}-\varepsilon\right)\boldsymbol{\mu}_i(3,M)$. Moreover, if agent 1 is given set A_k, then $S_2 \cup S_3 = \bigcup_{\ell \in N \setminus k} A_\ell$.*

Clearly, Lemma 2.3.2 completes the proof. □

Before stating the proof of Lemma 2.3.2, we should mention how it is possible to go beyond the previously known $\frac{3}{4}$-approximation. As noted above, ρ_n is by definition the best guarantee we can get, based only on the fact that the complement of the set left to be shared is not too large. As a result, the $\frac{7}{8}$ ratio cannot be guaranteed just by the excess value. Instead, in addition to making sure that the remaining items are valuable enough for the remaining agents, we further argue about how a maximin partition would distribute those items.

There is an alternative interpretation of Algorithm 5 in terms of Algorithm 3. Whenever only a single agent (i.e., agent 1) is going to become satisfied in the first recursive call, we try all possible maximum matchings of the graph G for the calculation of X^+. Then we proceed with the "best" such matching. Here, for $n = 3$, this means we only have to consider two possibilities for the set agent 1 is going to get matched to; it is either A_2 or A_3 (subject to the assumptions in Algorithm 5).

Proof of Lemma 2.3.2. First, recall that $v_2(M) \geq 3\boldsymbol{\mu}_2(3, M) \geq 3\xi_2$. Like in the description of the algorithm we may assume that agent 1 gets set A_3, without loss of generality. Before we move to the analysis we should lay down some facts. Let (B_1, B_2) be agent 2's $(1-\varepsilon')$-approximate maximin partition of $A_1 \cup A_2$ computed in line 11; similarly (B_1', B_2') is agent 2's $(1-\varepsilon')$-approximate maximin partition of $A_1 \cup A_3$. We may assume that $v_2(B_1) \geq v_2(B_2)$. Also, assume that in line 11 of the algorithm we have $(D_1, D_2) = (B_1, B_2)$, i.e., $\min\{v_2(B_1'), v_2(B_2')\} \leq v_2(B_2)$ and $M \setminus (D_1 \cup D_2) = A_3$. The case where $(D_1, D_2) = (B_1', B_2')$ is symmetric. Our goal is to show that $v_2(B_2) \geq \left(\frac{7}{8} - \varepsilon\right)\boldsymbol{\mu}_2(3, M)$. For simplicity, we write $\boldsymbol{\mu}_2$ instead of $\boldsymbol{\mu}_2(3, M)$.

Note, towards a contradiction, that

$$\begin{aligned} v_2(B_2) &< \left(\frac{7}{8} - \varepsilon\right)\boldsymbol{\mu}_2 \Rightarrow \\ (1-\varepsilon')\boldsymbol{\mu}_2(2, A_1 \cup A_2) &< \left(\frac{7}{8} - \varepsilon\right)\boldsymbol{\mu}_2 \Rightarrow \\ (1-\varepsilon')\boldsymbol{\mu}_2(2, A_1 \cup A_2) &< \left(\frac{7}{8} - \frac{7}{8}\varepsilon'\right)\boldsymbol{\mu}_2 \Rightarrow \\ \boldsymbol{\mu}_2(2, A_1 \cup A_2) &< \frac{7}{8}\boldsymbol{\mu}_2 . \end{aligned}$$

Moreover, this means $\min\{v_2(B_1'), v_2(B_2')\} < \left(\frac{7}{8} - \varepsilon\right)\boldsymbol{\mu}_2$ as well, which leads to $\boldsymbol{\mu}_2(2, A_1 \cup A_3) < \frac{7}{8}\boldsymbol{\mu}_2$. So, it suffices to show that either $\boldsymbol{\mu}_2(2, A_1 \cup A_2)$ or $\boldsymbol{\mu}_2(2, A_1 \cup A_3)$ is at least $\frac{7}{8}\boldsymbol{\mu}_2$. This statement is independent of the B_is and in what follows we consider exact maximin partitions with respect to agent 2. Before we proceed, we should make clear that for the case we are analyzing there are indeed exactly two sets in $\{A_1, A_2, A_3\}$ each with value less than $\frac{7}{8}\boldsymbol{\mu}_2$ with respect to agent 2, as claimed in line 10 of the algorithm. Indeed, notice that in any partition of M there is at least one set with value at least $\boldsymbol{\mu}_2$ with respect to agent 2, due to the fact that $v_2(M) \geq 3\boldsymbol{\mu}_2$ and by the definition of a maximin partition. If, however, there were at least 2 sets in $\{A_1, A_2, A_3\}$ with value at least $\frac{7}{8}\xi_2$, then we would be at the case handled in steps 6-8. Hence, there will be exactly two sets each with value less than $\frac{7}{8}\xi_2 \le \frac{7}{8}\boldsymbol{\mu}_2$ for agent 2 and as stated in the algorithm we assume these are the sets A_2, A_3.

Consider a 3-maximin share allocation (A'_1, A'_2, A'_3) of M with respect to agent 2. Let $F_i = A'_i \cap A_3$ for $i = 1,2,3$. Without loss of generality, we may assume that $v_2(F_1) \le v_2(F_2) \le v_2(F_3)$.

If $v_2(F_1) \le \frac{1}{8}\boldsymbol{\mu}_2$, then the partition $(A'_1 \setminus A_3, (A'_2 \cup A'_3) \setminus A_3)$ is a partition of $A_1 \cup A_2$ such that

$$v_2(A'_1 \setminus A_3) = v_2(A'_1) - v_2(F_1) \ge \boldsymbol{\mu}_2 - \frac{1}{8}\boldsymbol{\mu}_2 = \frac{7}{8}\boldsymbol{\mu}_2$$

and

$$v_2((A'_2 \cup A'_3) \setminus A_3) \ge v_2(A'_2) + v_2(A'_3) - v_2(A_3) \ge 2\boldsymbol{\mu}_2 - \frac{7}{8}\boldsymbol{\mu}_2 = \frac{9}{8}\boldsymbol{\mu}_2.$$

So, in this case we conclude that $\boldsymbol{\mu}_2(2, A_1 \cup A_2) \ge \frac{7}{8}\boldsymbol{\mu}_2$.

On the other hand, if $v_2(F_1) > \frac{1}{8}\boldsymbol{\mu}_2$ we are going to show that $\boldsymbol{\mu}_2(2, A_1 \cup A_3) \ge \frac{7}{8}\boldsymbol{\mu}_2$. Towards this we consider a 2-maximin share allocation (C_1, C_2) of A_1 with respect to agent 2 and let us assume that $v_2(C_1) \ge v_2(C_2)$. For a rough depiction of the different sets involved in the following arguments, see Figure 2.2.

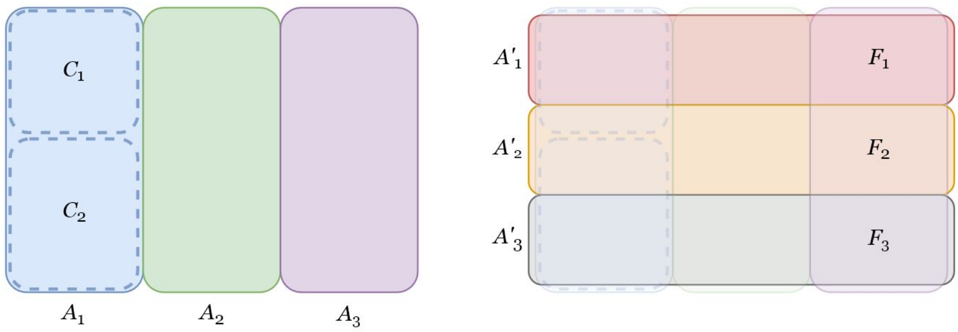

FIGURE 2.2: Assuming that the set of items M is represented by a rectangle, here is a depiction of several sets involved in the proof of Lemma 2.3.2. Recall that (A_1, A_2, A_3) and (A'_1, A'_2, A'_3) are partitions of M, (C_1, C_2) is a partition of A_1, and $F_i = A'_i \cap A_3$ for $i = 1,2,3$.

Claim 2.3.3. *For $C_1, C_2, A_3, F_1, F_2, F_3$ as above, we have*

(i) $v_2(A_3) + v_2(C_2) \ge \frac{7}{8}\boldsymbol{\mu}_2$, *and*

(ii) $v_2(F_1) + v_2(F_2) + v_2(C_1) > \frac{7}{8}\boldsymbol{\mu}_2$.

Proof of Claim 2.3.3. Note that

$$v_2(C_1) + v_2(C_2) + v_2(A_3) = v_2(M) - v_2(A_2) > 3\boldsymbol{\mu}_2 - \frac{7}{8}\boldsymbol{\mu}_2 = \frac{17}{8}\boldsymbol{\mu}_2.$$

If $v_2(A_3) + v_2(C_2) < \frac{7}{8}\boldsymbol{\mu}_2$ then $v_2(C_1) > \frac{10}{8}\boldsymbol{\mu}_2$. Moreover,

$$v_2(A_3) = v_2(F_1) + v_2(F_2) + v_2(F_3) \ge 3v_2(F_1) > \frac{3}{8}\boldsymbol{\mu}_2,$$

so $v_2(A_3) + v_2(C_2) < \frac{7}{8}\boldsymbol{\mu}_2$ implies that $v_2(C_2) < \frac{4}{8}\boldsymbol{\mu}_2$.

Let d denote the difference $v_2(C_1) - v_2(C_2)$; clearly $d > \frac{6}{8}\boldsymbol{\mu}_2$. It is not hard to see that $\min_{j \in C_1} v_{2j} \geq d$. Indeed, suppose there existed some $j \in C_1$ such that $v_{2j} < d$. Then, by moving j from C_1 to C_2 we increase the minimum value of the partition, which contradicts the choice of (C_1, C_2).

Since $v_2(C_1) > \frac{10}{8}\boldsymbol{\mu}_2$ and no item has value more than $\frac{7}{8}\boldsymbol{\mu}_2$ for agent 2, this means that C_1 contains at least two items. Thus, $v_2(C_1) \geq \min_{j \in C_1} v_{2j} > \frac{12}{8}\boldsymbol{\mu}_2$.

Now, for any item $g \in \arg\min_{j \in C_1} v_{2j}$, the partition $(\{g\}, A_1 \setminus \{g\})$ is strictly better than (C_1, C_2), since $v_{2g} > \frac{6}{8}\boldsymbol{\mu}_2 > v_2(C_2)$ and $v_2(A_1 \setminus \{g\}) = v_2(A_1) - v_{2g} \geq v_2(C_1) - v_{2g} > \frac{12}{8}\boldsymbol{\mu}_2 - \frac{6}{8}\boldsymbol{\mu}_2 = \frac{6}{8}\boldsymbol{\mu}_2 > v_2(C_2)$. Again, this contradicts the choice of (C_1, C_2). Hence, it must be that $v_2(A_3) + v_2(C_2) \geq \frac{7}{8}\boldsymbol{\mu}_2$.

The proof of *(ii)* is simpler. Notice that

$$\begin{aligned} v_2(F_1) + v_2(F_2) + v_2(C_1) &\geq v_2(F_1) + v_2(F_1) + \frac{1}{2}v_2(A_1) \\ &> \frac{1}{8}\boldsymbol{\mu}_2 + \frac{1}{8}\boldsymbol{\mu}_2 + \frac{1}{2}(3\boldsymbol{\mu}_2 - \frac{7}{8}\boldsymbol{\mu}_2 - \frac{7}{8}\boldsymbol{\mu}_2) = \frac{7}{8}\boldsymbol{\mu}_2 . \end{aligned}$$ ⊡

Now, if $v_2(C_1) \geq \frac{7}{8}\boldsymbol{\mu}_2$ then *(i)* of Claim 2.3.3 implies that $\min\{v_2(C_1), v_2(A_3 \cup C_2)\} \geq \frac{7}{8}\boldsymbol{\mu}_2$. Similarly, if $v_2(F_3) + v_2(C_2) \geq \frac{7}{8}\boldsymbol{\mu}_2$ then *(ii)* of Claim 2.3.3 implies that $\min\{v_2(F_1 \cup F_2 \cup C_1), v_2(F_3 \cup C_2)\} \geq \frac{7}{8}\boldsymbol{\mu}_2$. In both cases, we have $\boldsymbol{\mu}_2(2, A_1 \cup A_3) \geq \frac{7}{8}\boldsymbol{\mu}_2$. So, it is left to examine the case where both $v_2(C_1)$ and $v_2(F_3) + v_2(C_2)$ are less than $\frac{7}{8}\boldsymbol{\mu}_2$.

Claim 2.3.4. *Let* $C_1, C_2, A_3, F_1, F_2, F_3$ *be as above and* $\max\{v_2(C_1), v_2(F_3 \cup C_2)\} < \frac{7}{8}\boldsymbol{\mu}_2$. *Then* $\min\{v_2(F_1 \cup C_1), v_2(F_2 \cup F_3 \cup C_2)\} \geq \frac{7}{8}\boldsymbol{\mu}_2$.

Proof of Claim 2.3.4. Recall that $v_2(A_1) + v_2(A_3) > \frac{17}{8}\boldsymbol{\mu}_2$. Suppose $v_2(F_1 \cup C_1) < \frac{7}{8}\boldsymbol{\mu}_2$. Then $v_2(F_2 \cup F_3 \cup C_2) > \frac{10}{8}\boldsymbol{\mu}_2$. Since $v_2(F_3 \cup C_2) < \frac{7}{8}\boldsymbol{\mu}_2$ we have $v_2(F_2) > \frac{3}{8}\boldsymbol{\mu}_2$. But then we get the contradiction

$$\frac{7}{8}\boldsymbol{\mu}_2 > v_2(A_3) = v_2(F_1) + v_2(F_2) + v_2(F_3) \geq \frac{1}{8}\boldsymbol{\mu}_2 + \frac{3}{8}\boldsymbol{\mu}_2 + \frac{3}{8}\boldsymbol{\mu}_2 = \frac{7}{8}\boldsymbol{\mu}_2 .$$

Hence, $v_2(F_1 \cup C_1) \geq \frac{7}{8}\boldsymbol{\mu}_2$. Similarly, suppose $v_2(F_2 \cup F_3 \cup C_2) < \frac{7}{8}\boldsymbol{\mu}_2$. Then $v_2(F_1 \cup C_1) > \frac{10}{8}\boldsymbol{\mu}_2$. Since $v_2(C_1) < \frac{7}{8}\boldsymbol{\mu}_2$ we have $v_2(F_1) > \frac{3}{8}\boldsymbol{\mu}_2$. Then we get the contradiction

$$\frac{7}{8}\boldsymbol{\mu}_2 > v_2(A_3) = v_2(F_1) + v_2(F_2) + v_2(F_3) \geq \frac{3}{8}\boldsymbol{\mu}_2 + \frac{3}{8}\boldsymbol{\mu}_2 + \frac{3}{8}\boldsymbol{\mu}_2 = \frac{9}{8}\boldsymbol{\mu}_2 .$$

Hence, $v_2(F_2 \cup F_3 \cup C_2) \geq \frac{7}{8}\boldsymbol{\mu}_2$. ⊡

Claim 2.3.4 implies $\boldsymbol{\mu}_2(2, A_1 \cup A_3) \geq \frac{7}{8}\boldsymbol{\mu}_2$ and this concludes the proof. □

2.3.2 Values in $\{0, 1, 2\}$

Bouveret and Lemaître (2016) consider a binary setting where all valuation functions take values in $\{0, 1\}$, i.e., for each $i \in N$, and $j \in M$, $v_{ij} \in \{0, 1\}$. This can correspond to expressing approval or disapproval for each item. It is then shown that it is always possible to find a maximin share allocation in polynomial time. In fact, they show that the Greedy Round-Robin algorithm, presented in Section 2.1, computes such an allocation in this case.

Here, we extend this result to the setting where each v_{ij} is in $\{0,1,2\}$, allowing the agents to express two types of approval for the items. Enlarging the set of possible values from $\{0,1\}$ to $\{0,1,2\}$ by just one extra possible value makes the problem significantly more complex. Greedy Round-Robin does not work in this case, so a different algorithm is developed.

Theorem 2.3.5. *Let $N=[n]$ be a set of agents and $M=[m]$ be a set of items. If for any $i \in N$, agent i has a valuation function v_i such that $v_{ij} \in \{0,1,2\}$ for any $j \in M$, then we can find, in time $O(nm\log m)$, an allocation $(T_1,\dots,T_n)$ of M so that $v_i(T_i) \geq \boldsymbol{\mu}_i(n,M)$ for every $i \in [n]$.*

To design our algorithm, we make use of an important observation by Bouveret and Lemaître (2016) that allows us to reduce appropriately the space of valuation functions that we are interested in. We say that the agents have *fully correlated* valuation functions if they agree on a common ranking of the items in decreasing order of values. That is, $\forall i \in N$, if $M = \{1,2,...,m\}$, we have $v_{i1} \geq v_{i2} \geq \ldots \geq v_{im}$. Bouveret and Lemaître (2016) show that to find a maximin share allocation for any set of valuation functions, it suffices to do so in an instance where the valuation functions are fully correlated. This family of instances seems to be the difficulty in computing such allocations. Actually, their result preserves approximation ratios as well (with the same proof); hence we state this stronger version. For a valuation function v_i let σ_i be a permutation on the items such that $v_i(\sigma_i(j)) \geq v_i(\sigma_i(j+1))$ for $j \in \{1,\dots,m-1\}$. We denote the function $v_i(\sigma_i(\cdot))$ by $v_i^{\uparrow}$. Note that $v_1^{\uparrow}, v_2^{\uparrow},\dots,v_n^{\uparrow}$ are now fully correlated.

Theorem 2.3.6 (Bouveret and Lemaître (2016))**.** *Let $N=[n]$ be a set of agents with additive valuation functions, $M=[m]$ be a set of goods and $\rho \in (0,1]$. Given an allocation $(T_1,\dots,T_n)$ of M so that $v_i^{\uparrow}(T_i) \geq \rho\boldsymbol{\mu}_i(n,M)$ for every i, one can produce in linear time an allocation $(T_1',\dots,T_n')$ of M so that $v_i(T_i') \geq \rho\boldsymbol{\mu}_i(n,M)$ for every i.*

We are ready to state a high level description of our algorithm. The detailed description, however, is deferred to the end of this subsection. The reason for this is that the terminology needed is gradually introduced through a series of lemmata motivating the idea behind the algorithm and proving its correctness. In fact, the remainder of the subsection is the proof of Theorem 2.3.5. Algorithm 6 in the end summarizes all the steps.

Algorithm Outline: We first construct $v_1^{\uparrow}, v_2^{\uparrow},\dots,v_n^{\uparrow}$ and work with them instead. The Greedy Round-Robin algorithm may not directly work, but we partition the items in a similar fashion, although without giving them to the agents. Then, we show that it is possible to choose some subsets of items and redistribute them in a way that guarantees that everyone can get a bundle of items with enough value. At a higher level, we could say that the algorithm simulates a variant of the Greedy Round-Robin, where for an appropriately selected set of rounds the agents choose in the reverse order. Finally, a maximin share allocation can be obtained for the original v_is, as described in Bouveret and Lemaître (2016).

Proof of Theorem 2.3.5. According to Theorem 2.3.6 it suffices to focus on instances where the valuation functions take values in $\{0,1,2\}$ and are fully correlated. Given such an instance we distribute the m objects into n buckets in decreasing order, i.e., bucket i will get items $i, n+i, 2n+i, \ldots$ Notice that this is compatible with how the Greedy Round-Robin algorithm could distribute the items; however, we do not assign any buckets to any agents yet. We may assume that $m = kn$ for some $k \in \mathbb{N}$; if not, we just add a few extra items with 0 value to everyone. It is convenient to picture the collection of buckets as the matrix

$$B = \begin{pmatrix} (k-1)n+1 & (k-1)n+2 & \cdots & kn \\ \vdots & \vdots & \ddots & \vdots \\ n+1 & n+2 & \cdots & 2n \\ 1 & 2 & \cdots & n \end{pmatrix},$$

since our algorithm will systematically redistribute groups of items corresponding to rows of B.

Before we state the algorithm, we establish some properties regarding these buckets and the way each agent views the values of these bundles. First, we introduce some terminology.

Definition 2.3.7. We say that agent i is

- *satisfied* with respect to the current buckets, if all the buckets have value at least $\boldsymbol{\mu}_i(n, M)$ according to v_i.
- *left-satisfied* with respect to the current buckets, if he is not satisfied, but at least the $n/2$ leftmost buckets have value at least $\boldsymbol{\mu}_i(n, M)$ according to v_i.
- *right-satisfied* if the same as above hold, but for the rightmost $n/2$ buckets.

Now suppose that we see agent i's view of the values in the buckets. A typical view would have the following form (recall the goods are ranked from highest to lowest value):

$$\begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ . & . & . & . & . & . & \cdots & . & . & . \\ 1 & 1 & 1 & 1 & 1 & 0 & \cdots & 0 & 0 & 0 \\ . & . & . & . & . & . & \cdots & . & . & . \\ 1 & 1 & 1 & 1 & 1 & 1 & \cdots & 1 & 1 & 1 \\ 2 & 2 & 2 & 1 & 1 & 1 & \cdots & 1 & 1 & 1 \\ . & . & . & . & . & . & \cdots & . & . & . \\ 2 & 2 & 2 & 2 & 2 & 2 & \cdots & 2 & 2 & 2 \end{pmatrix}$$

A row that has only 2s for i will be called a 2-row for i. A row that has both 2s and 1s will be called a 2/1-row for i, and so forth. An agent can also have a 2/1/0-row. It is not necessary, of course, that an agent will have all possible types of rows in his view. Note, however, that there can be at most one 1/0-row and at most one 2/1-row in her view. We first prove the following lemma for agents that are not initially satisfied.

Lemma 2.3.8. *Any agent not satisfied with respect to the initial buckets must have both a 1/0-row and a 2/1-row in his view of B. Moreover, initially all agents are either satisfied or left-satisfied.*

Proof of Lemma 2.3.8. Let us focus on the multiset of values of an agent that is not satisfied, say i. It is straightforward to see that if i has no 1s, or the number of 2s is a multiple of n (including 0), then agent i gets value $\boldsymbol{\mu}_i(n,M)$ from any bucket. So, i must have a row with both 2s and 1s. If this is a 2/1/0-row, then again it is easy to see that the initial allocation is a maximin share allocation for i. So, i has a 2/1-row. The only case where he does not have a 1/0-row is if the total number of 1s and 2s is a multiple of n.But then the maximum and the minimum value of the initial buckets differ by 1, hence we have a maximin share allocation and i is satisfied.

Next we show that an agent i who is not initially satisfied is left-satisfied. In what follows we only refer to i's view. Buckets B_1 and B_n, indexed by the corresponding columns of B, have maximum and minimum total value respectively. Since i is not satisfied, we have $v_i(B_1) \geq v_i(B_n)+2$, but the way we distributed the items guarantees that the difference between any two buckets is at most the largest value of an item; so $v_i(B_1) = v_i(B_n)+2$. Moreover, since $v_i(M) \geq n\boldsymbol{\mu}_i(n,M)$ and $v_i(B_n) < \boldsymbol{\mu}_i(n,M)$, we must have $v_i(B_1) > \boldsymbol{\mu}_i(n,M)$. This implies that $v_i(B_1) = \boldsymbol{\mu}_i(n,M)+1$ and $v_i(B_n) = \boldsymbol{\mu}_i(n,M)-1$.

More generally, we have buckets of value $\boldsymbol{\mu}_i(n,M)+1$ (leftmost columns), we have buckets of value $\boldsymbol{\mu}_i(n,M)-1$ (rightmost columns), and maybe some other buckets of value $\boldsymbol{\mu}_i(n,M)$ (columns in the middle). We know that the total value of all the items is at least $n\boldsymbol{\mu}_i(n,M)$, so, by summing up the values of the buckets, we conclude that there must be at most $n/2$ buckets of value $\boldsymbol{\mu}_i(n,M)-1$. Therefore i is left-satisfied. ⊡

So far we may have some agents that could take any bucket and some agents that would take any of the $n/2$ (at least) first buckets. Clearly, if the left-satisfied agents are at most $n/2$ then we can easily find a maximin share allocation. However, there is no guarantee that there are not too many left-satisfied agents initially, so we try to fix this by reversing some of the rows of B. To make this precise, we say that *we reverse the ith row of B* when we take items $(i-1)n+1, (i-1)n+2, \ldots, in$ and we put item in in bucket 1, item $in-1$ in bucket 2, etc.

The algorithm then proceeds by picking a subset of rows of B and reversing them. The rows are chosen appropriately so that the resulting buckets (i.e., the columns of B) can be easily paired with the agents to get a maximin share allocation. First, it is crucial to understand the effect that the reversal of a set of rows has to an agent.

Lemma 2.3.9. *Any agent satisfied with respect to the initial buckets remains satisfied independently of the rows of B that we may reverse. On the other hand, any agent not satisfied with respect to the initial buckets, say agent i, is affected if we reverse* *her* 1/0*-row or her* 2/1*-row. If we reverse only one of those, then i becomes satisfied with respect to the new buckets; if we reverse both, then i becomes right-satisfied. The reversal of any other rows is irrelevant to agent i.*

Proof of Lemma 2.3.9. Fix an agent i. First notice that, due to symmetry, reversing any row that for i is a 2-row, a 1-row, or a 0-row does not improve or worsen the initial allocation from i's point of view. Also, clearly, reversing both the 1/0-row and the 2/1-row of a left-satisfied agent makes him right-satisfied. Similarly, if i is satisfied and has a 2/1/0-row, or has a 2/1-row but no 1/0-row, or has a 1/0-row but no 2/1-row, then reversing those keeps i satisfied.

The interesting case is when i has both a 1/0-row and a 2/1-row. If i is satisfied, then even removing her 1/0-row leaves all the buckets with at least as much value as the last bucket; so reversing it keeps i satisfied. A similar argument holds for i's 2/1-row as well. If i is not satisfied, then the difference of the values of the first and the last bucket will be 2. Like in the proof of Lemma 2.3.8, the number of columns that have 1 in i's 1/0-row and 2 in i's 2/1-row (i.e., total value $\boldsymbol{\mu}_i(n,M)+1$) are at least as many as the columns that have 0 in i's 1/0-row and 1 in i's 2/1-row (i.e., total value $\boldsymbol{\mu}_i(n,M)-1$). So, by reversing her 1/0-row, the values of all the "worst" (rightmost) buckets increase by 1, the values of some of the "best" (leftmost) buckets decrease by 1, and the values of the buckets in the middle either remain the same or increase by 1. The difference between the best and the worst buckets now is 1 (at most), so this is a maximin share allocation for i and he becomes satisfied. Due to symmetry, the same holds for reversing i's 2/1-row only. ⊡

Now, what Lemma 2.3.9 guarantees is that when we reverse some of the rows of the initial B, we are left with agents that are either satisfied, left-satisfied, or right-satisfied. If the rows are chosen so that there are at most $n/2$ left-satisfied and at most $n/2$ right-satisfied agents, then there is an obvious maximin share allocation: to any left-satisfied agent we arbitrarily give one of the first $n/2$ buckets, to any right-satisfied agent we arbitrarily give one of the last $n/2$ buckets, and to each of the remaining agents we arbitrarily give one of the remaining buckets. In Lemma 2.3.10 below, we prove that it is easy to find which rows to reverse to achieve that.

We use a graph theoretic formulation of the problem for clarity. With respect to the initial buckets, we define a graph $G=(V,E)$ with $V=[k]$, i.e., G has a vertex for each row of B. Also, for each left-satisfied agent i, G has an edge connecting i's 1/0-row and 2/1-row. We delete, if necessary, any multiple edges to get a simple graph with n edges at most. We want to color the vertices of G with two colors, "red" (for reversed rows) and "blue" (for non reversed), so that the number of edges having both endpoints red is at most $n/2$ and at the same time the number of edges having both endpoints blue is at most $n/2$. Note that if we reverse the rows that correspond to red vertices, then the agents with red endpoints become right-satisfied, the agents with blue endpoints remain left-satisfied and the agents with both colors become satisfied.

Moreover, the initially satisfied agents are not affected, and we can find a maximin share allocation as previously discussed. This is illustrated in Figure 2.3 below.

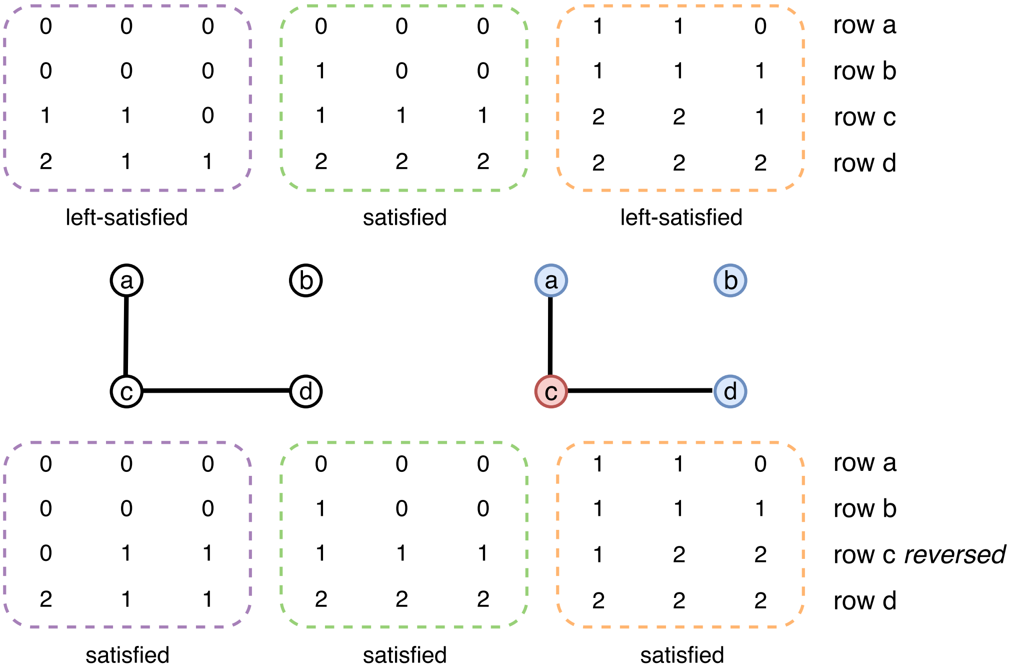

FIGURE 2.3: Assuming an instance with 3 agents and 11 items, the tables on top are the three different views on the initial buckets. This results in the graph shown in the middle—before and after the coloring. By reversing row c that corresponds to a red vertex, every agent becomes satisfied and thus any matching of the columns to the agents defines an MMS allocation.

Lemma 2.3.10. *Given graph G defined above, in time $O(k+n)$ we can color the vertices with two colors, red and blue, so that the number of edges with two red endpoints is less than $n/2$ and the number of edges with two blue endpoints is at most $n/2$.*

Proof of Lemma 2.3.10. lor vertices red, one at a time, until the number of edges with two blue endpoints becomes at most $|E|/2$ for the first time. Assume this happens after recoloring vertex u. Before turning u from blue to red, the number of edges with at most one blue endpoint was strictly less than $|E|/2$. Also, the recoloring of u did not force any of the edges with two blue endpoints to become edges with two red endpoints. So, the number of edges with two red endpoints after the recoloring of u is at most equal to the number of edges with at most one blue endpoint before the recoloring of u, i.e., less than $|E|/2$. To complete the proof, notice that $|E| \leq n$. For the running time, notice that each vertex changes color at most once and when this happens we only need to examine the adjacent vertices in order to update the counters on each type of edges (only red, only blue, or both). ◻

Lemma 2.3.10 completes the proof of correctness for Algorithm 6 that is summarized below. For the running time notice that $v_1^{\uparrow}, \ldots, v_n^{\uparrow}$ can be computed in $O(nm \log m)$, since we get $v_i^{\uparrow}$ by sorting $v_{i1}, \ldots, v_{im}$. Also step 5 can be computed in $O(nm)$; for each agent i we scan the first column of B to find his (possible) 1/0-row and 2/1-row, and then in $O(n)$ we check whether he is left-satisfied by checking that the positions that have 1 in i's 1/0-row and 2 in i's 2/1-row are at least as many as the positions that have 0 in i's 1/0-row and 1 in i's 2/1-row. □

ALGORITHM 6: EXACT-MMS$_{0,1,2}(m, v_1, \ldots, v_n)$

1 Let $k = \lceil \frac{m}{n} \rceil$. Add $kn - m$ dummy items with value 0 for everyone.
2 **if** $v_1, \ldots, v_n$ *are not fully correlated* **then**
3 Compute $v_1^{\uparrow}, \ldots, v_n^{\uparrow}$ and use them instead.
4 Construct a $k \times n$ matrix B so that B_{ij} is the $(i-1)n + j$th item.
5 Find the set of left-satisfied agents and their corresponding 1/0-rows and 2/1-rows.
6 Construct a graph $G = ([k], E)$ with $E = \{\{i, j\} | \exists$ left-satisfied agent that i and j are her 1/0-row and 2/1-row$\}$.
7 Color the vertices of G with two colors, red and blue, so that the number of edges having both endpoints red, and the number of edges having both endpoints blue, each is $\leq n/2$.
8 Reverse the rows of B that correspond to red vertices, and keep track of who is satisfied, left-satisfied, or right-satisfied.
9 Arbitrarily give some of the first $n/2$ buckets (columns of B) to each of the left-satisfied agents and some of the last $n/2$ buckets to each of the right-satisfied agents. Arbitrarily give the rest of the buckets to the satisfied agents.
10 **if** $v_1^{\uparrow}, \ldots, v_n^{\uparrow}$ *were used* **then**
11 Based on the allocation in step 9 compute and return a maximin share allocation for the original v_is as described in Bouveret and Lemaître (2016).
12 **else**
13 Return the allocation in step 9.

2.4 A Probabilistic Analysis

As argued in the previous works (Bouveret and Lemaître, 2016; Procaccia and Wang, 2014), it has been quite challenging to prove impossibility results. Setting efficient computation aside, what is the best ρ for which a ρ-approximate allocation does exist? All we know so far is that $\rho \neq 1$ by the elaborate constructions by Kurokawa, Procaccia, and Wang (2016) and Procaccia and Wang (2014). However, extensive experimentation by Bouveret and Lemaître (2016) (and also by Procaccia and Wang (2014)), showed that in all generated instances, there always existed a maximin share allocation. Motivated by these experimental observations and by the lack of impossibility results, we present a probabilistic analysis, showing that indeed we expect that in most cases there exist allocations where every agent receives his maximin share. In particular, we analyze the Greedy Round-Robin algorithm from Section 2.1 when each v_{ij} is drawn from the uniform distribution over $[0, 1]$.

Recently, Kurokawa, Procaccia, and Wang (2016) show similar results for a large set of distributions over $[0, 1]$, including $U[0, 1]$. Although, asymptotically, their results yield a theorem that is more general than ours, we consider our analysis to be of independent interest, since we have much better bounds on the probabilities for the special case of $U[0,1]$, even for relatively small values of n.

For completeness, before stating and proving our results, we include the version of Hoeffding's inequality we are going to use.

Theorem 2.4.1 (Hoeffding (1963))**.** *Let $X_1, X_2, \ldots, X_n$ be independent random variables with $X_i \in [0,1]$ for $i \in [n]$. Then for the empirical mean $\bar{X} = \frac{1}{n}(X_1 + \ldots + X_n)$ we have* $\mathrm{P}\big(\bar{X} - \mathrm{E}[\bar{X}] \geq t\big) \leq \exp(-2nt^2)$.

We start with Theorem 2.4.2. Its proof is based on tools like Hoeffding's and Chebyshev's inequalities, and on a careful estimation of the probabilities when $m < 3n$. Note that for $m \geq 2n$, the theorem provides an even stronger guarantee than the maximin share (by Claim 2.1.1).

Theorem 2.4.2. *Let $N = [n]$ be a set of agents and $M = [m]$ be a set of goods, and assume that the v_{ij}s are i.i.d. random variables that follow $U[0,1]$. Then, for $m \geq 2n$ and large enough n, the Greedy Round-Robin algorithm allocates to each agent i a set of goods of total value at least $\frac{1}{n}\sum_{j=1}^{m} v_{ij}$ with probability $1 - o(1)$. The $o(1)$ term is $O(1/n)$ when $m > 2n$ and $O(\log n/n)$ when $m = 2n$.*

Proof. In what follows we assume that agent 1 chooses first, agent 2 chooses second, and so forth. We consider several cases for the different ranges of m. We first assume that $2n \leq m < 3n$.

It is illustrative to consider the case of $m = 2n$ and examine the nth agent that chooses last. Like all the agents in this case, he receives exactly two items; let Y_n be the total value of those items. From his perspective, he sees $n+1$ values chosen uniformly from $[0,1]$, picks the maximum of those, then u.a.r. $n-1$ of the rest are removed, and he takes the last one as well. If we isolate this random experiment, it is as if we take $Y_n = \max\{X_1, ..., X_{n+1}\} + X_Y$, where $Y \sim U\big(\{1,2,...,n+1\} \setminus \{\mu\}\big)$, $\mu \in \arg\max\{X_1, ..., X_{n+1}\}$, $X_i \sim U[0,1]$ $\forall i \in [n+1]$, and all the X_is are independent. We estimate now the probability $\mathrm{P}(Y_n \leq a)$ for $1 < a < 2$. We will set a to a particular value in this interval later on. In fact, we bound this probability using the corresponding probability for $Z_n = \max\{X_1, ..., X_{n+1}\} + X_{Y'}$, where $Y' \sim U\{1,2,...,n+1\}$. For Z_n we have

$$\begin{aligned}
\mathrm{P}(Z_n \leq a) &= \sum_{i=1}^{n+1} \int_0^a \mathrm{P}\left(\max_{1 \leq j \leq n+1} X_j \leq t \wedge Y = i \wedge X_i \leq a - t\right) dt \\
&= (n+1) \int_0^a \mathrm{P}\left(\max_{1 \leq j \leq n+1} X_j \leq t \wedge Y = 1 \wedge X_1 \leq a - t\right) dt \\
&= \int_0^a \mathrm{P}\left(\max_{1 \leq j \leq n+1} X_j \leq t \wedge X_1 \leq a - t\right) dt \\
&= \int_0^a \mathrm{P}(X_1 \leq t \wedge X_1 \leq a - t \wedge X_2 \leq t \wedge \ldots \wedge X_{n+1} \leq t) dt \\
&= \int_0^{a/2} \mathrm{P}(X_1 \leq t \wedge X_2 \leq t \wedge \ldots \wedge X_{n+1} \leq t) dt + \\
&\quad + \int_{a/2}^1 \mathrm{P}(X_1 \leq a - t \wedge X_2 \leq t \wedge \ldots \wedge X_{n+1} \leq t) dt +
\end{aligned}$$

$$+ \int_1^a \mathrm{P}(X_1 \le a - t \wedge X_2 \le t \wedge \ldots \wedge X_{n+1} \le t) dt$$
$$= \int_0^{a/2} t^{n+1} dt + \int_{a/2}^1 (a-t) t^n dt + \int_1^a (a-t) dt.$$

Also, by the definition of Y' we have $\mathrm{P}(Y' \notin \arg\max\{X_1, ..., X_{n+1}\}) = n/(n+1)$. Therefore, for Y_n we get

$$\begin{aligned}
\mathrm{P}(Y_n \le a) &= \mathrm{P}(Z_n \le a \mid Y' \notin \arg\max\{X_1, ..., X_{n+1}\}) \\
&= \frac{\mathrm{P}(Z_n \le a \ \wedge \ Y' \notin \arg\max\{X_1, ..., X_{n+1}\})}{\mathrm{P}(Y' \notin \arg\max\{X_1, ..., X_{n+1}\})} \\
&\le \frac{\mathrm{P}(Z_n \le a)}{\mathrm{P}(Y' \notin \arg\max\{X_1, ..., X_{n+1}\})} = \frac{n+1}{n} \mathrm{P}(Z_n \le a) \\
&= \frac{n+1}{n} \left(\int_0^{a/2} t^{n+1} dt + \int_{a/2}^1 (a-t) t^n dt + \int_1^a (a-t) dt \right),
\end{aligned}$$

where for the inequality we used the fact that $\mathrm{P}(A \cap B) \le \mathrm{P}(A)$ for any events A, B.

A similar analysis for the jth agent yields

$$\mathrm{P}(Y_j \le a) \le \frac{2n-j+1}{n} \left(\int_0^{a/2} t^{2n-j+1} dt + \int_{a/2}^1 (a-t)^{n-j+1} t^n dt + \int_1^a (a-t)^{n-j+1} dt \right).$$

In the more general case where $m = 2n + \kappa(n)$, $0 \le \kappa(n) < n$, we have a similar calculation for the agents that receive only two items in the Greedy Round-Robin algorithm, as well as for the first two items of the first $\kappa(n)$ agents (who receive three items each). Let Y_i be the total value agent i receives, and W_i be the value of his first two items. Of course, for the last $2n$ players, $Y_i = W_i$. Also, recall that $\sum_{j=1}^m v_{ij} = v_i(M)$. We now relate the probability that we are interested in estimating, with the probabilities $\mathrm{P}(Y_i \le a)$ that we have already bounded. We will then proceed by setting α appropriately. We have

$$\begin{aligned}
\mathrm{P}\left(\exists i \text{ such that } Y_i < \frac{1}{n} \sum_{j=1}^m v_{ij}\right) &\le \sum_{i=1}^n \mathrm{P}\left(Y_i < \frac{v_i(M)}{n} \right) \\
&= \sum_{i=1}^n \mathrm{P}\left(Y_i < \min\left\{ \frac{v_i(M)}{n}, a \right\} \vee \frac{v_i(M)}{n} > \max\{Y_i, a\} \right) \\
&\le \sum_{i=1}^n \mathrm{P}\left(Y_i < \min\left\{ \frac{v_i(M)}{n}, a \right\} \right) + \sum_{i=1}^n \mathrm{P}\left(\frac{v_i(M)}{n} > \max\{Y_i, a\} \right) \\
&\le \sum_{i=1}^n \mathrm{P}(Y_i < a) + \sum_{i=1}^n \mathrm{P}\left(\frac{v_i(M)}{n} > a \right).
\end{aligned}$$

To upper bound the first sum we use the W_is, i.e., we do not take into account the third item that the first $\kappa(n)$ agents receive. By the definition of Y_i, W_i, for these first $\kappa(n)$ agents we have $\mathrm{P}(Y_i < a) \le \mathrm{P}(W_i < a)$, while for the remaining agents we have $\mathrm{P}(Y_i < a) = \mathrm{P}(W_i < a)$. Note that the bounds for $\mathrm{P}(Y_i \le a)$ calculated above, here hold for $\kappa(n) + 1 \le i \le n$. For $1 \le i \le \kappa(n)$ the same bounds hold for $\mathrm{P}(W_i \le a)$.

$$\sum_{i=1}^n \mathrm{P}(Y_i < a) \le \sum_{i=1}^{\kappa(n)} \mathrm{P}(W_i < a) + \sum_{i=\kappa(n)+1}^{n} \mathrm{P}(Y_i < a)$$

$$
\begin{aligned}
&\le \sum_{i=1}^{n} \tfrac{m-i+1}{n} \Big(\int_0^{a/2} t^{m-i+1} dt + \int_{a/2}^{1} (a-t)^{n+\kappa(n)-i+1} t^n dt + \int_1^a (a-t)^{n+\kappa(n)-i+1} dt \Big) \\
&\le 3 \sum_{j=1}^{n} \Big(\int_0^{a/2} t^{n+\kappa(n)+j} dt + \int_{a/2}^{1} (a-t)^{\kappa(n)+j} t^n dt + \int_1^a (a-t)^{\kappa(n)+j} dt \Big) \\
&= 3 \left(\sum_{j=1}^{n} \frac{(a/2)^{n+\kappa(n)+j+1}}{n+\kappa(n)+j+1} + \sum_{j=1}^{n} \int_{a/2}^{1} (a-t)^{\kappa(n)+j} t^n dt + \sum_{j=1}^{n} \int_0^{a-1} u^{\kappa(n)+j} du \right).
\end{aligned}
$$

We are going to bound each sum separately. We set $a = 1 + \frac{\kappa(n)}{2n} + \sqrt{\frac{3\ln n}{n}} = \frac{m}{2n} + \sqrt{\frac{3\ln n}{n}}$. Note that for $n \ge 46$ we have $a \in (1,2)$. Consider the first sum:

$$
\begin{aligned}
\sum_{j=1}^{n} \frac{(a/2)^{n+\kappa(n)+j+1}}{n+\kappa(n)+j+1} &\le \frac{(a/2)^{n+\kappa(n)+2}}{n+\kappa(n)+2} \sum_{i=0}^{n-1} (a/2)^i \\
&< \frac{1}{n+\kappa(n)+2} \cdot \frac{(a/2)^{n+\kappa(n)+2}}{1-a/2} = O(1/n),
\end{aligned}
$$

where we got $O(1/n)$ because the bound is at most $\frac{3}{n}$ for $n \ge 57$ and for any value of $\kappa(n)$.

Next, we deal with the second sum:

$$
\begin{aligned}
\sum_{j=1}^{n} \int_{a/2}^{1} (a-t)^{\kappa(n)+j} t^n dt &< \int_{a/2}^{1} t^n \left(\sum_{j=0}^{\infty} (a-t)^{\kappa(n)+j} \right) dt \le \int_{a/2}^{1} t^n (a/2)^{\kappa(n)} \frac{1}{1-a+t} dt \\
&\le \frac{(a/2)^{\kappa(n)}}{1-a/2} \int_{a/2}^{1} t^n dt = \frac{(a/2)^{\kappa(n)}}{1-a/2} \left(\frac{1-(a/2)^{n+1}}{n+1} \right) = O(1/n).
\end{aligned}
$$

Here, for $n \ge 58$ the bound is at most $\frac{10}{n}$ for any $\kappa(n)$.

Finally, for the third sum, we rewrite it as

$$
\sum_{j=1}^{n} \int_0^{a-1} u^{\kappa(n)+j} du = \sum_{j=1}^{n} \frac{(a-1)^{\kappa(n)+j+1}}{\kappa(n)+j+1} = \sum_{i=\kappa(n)+2}^{n+\kappa(n)+1} \frac{(a-1)^i}{i}.
$$

We are going to bound each term separately. Consider the case where $\kappa(n) \ge 5\sqrt{n}$. For $n \ge 64$, it can be shown that $\frac{1}{5}\left(\frac{\kappa(n)}{2n} + \sqrt{\frac{3\ln n}{n}} \right)^{5\sqrt{n}} < \frac{10}{n^{3/2}}$. So,

$$
\sum_{i=\kappa(n)+2}^{n+\kappa(n)+1} \frac{(a-1)^i}{i} \le \sum_{i=1}^{n} \frac{(a-1)^{\kappa(n)}}{\kappa(n)} \le n \cdot \frac{\left(\frac{\kappa(n)}{2n} + \sqrt{\frac{3\ln n}{n}} \right)^{5\sqrt{n}}}{5\sqrt{n}} \le n \cdot \frac{10}{n^2} = \frac{10}{n}.
$$

On the other hand, when $\kappa(n) < 5\sqrt{n}$, we have $a - 1 < \frac{2.5+\sqrt{3\ln n}}{\sqrt{n}}$. For $n \ge 59$ and $j \ge 10$ it can be shown that $\frac{1}{j}\left(\frac{2.5+\sqrt{3\ln n}}{\sqrt{n}} \right)^j < \frac{30}{n^2}$. Of course, for $3 \le j \le 9$ it is true that $\frac{1}{j}\left(\frac{2.5+\sqrt{3\ln n}}{\sqrt{n}} \right)^j = o(1/n)$, and particularly for $n \ge 59$ the sum $\sum_{i=3}^{9} \frac{1}{j}\left(\frac{2.5+\sqrt{3\ln n}}{\sqrt{n}} \right)^j$ is bounded by $\frac{25}{n}$. In general, it is to be expected to have relatively large hidden constants when m is very close to $2n$. This changes quickly though; when $\kappa(n) > 21$

the whole sum is less than $1/n$. In any case, if $\kappa(n) > 0$

$$\sum_{i=\kappa(n)+2}^{n+\kappa(n)+1} \frac{(a-1)^i}{i} \le \sum_{i=3}^{n+2} \frac{(a-1)^i}{i} \le \sum_{i=3}^{9} \tfrac{1}{i}\left(\tfrac{2.5+\sqrt{3\ln n}}{\sqrt{n}}\right)^i + \sum_{i=10}^{n+2} \tfrac{1}{i}\left(\tfrac{2.5+\sqrt{3\ln n}}{\sqrt{n}}\right)^i \le$$
$$\le O(1/n) + (n-7)\frac{30}{n^2} = O(1/n)\,.$$

However, if $\kappa(n) = 0$, we have

$$\sum_{i=2}^{n+1} \frac{(a-1)^i}{i} = \left(\tfrac{2.5+\sqrt{3\ln n}}{\sqrt{n}}\right)^2 + \sum_{i=3}^{n+1} \frac{(a-1)^i}{i} = O\left(\tfrac{\log n}{n}\right).$$

So far, we have $\sum_{i=1}^{n} \mathrm{P}(Y_i < a) = O(1/n)$ (or $O(\log n/n)$ when $m = 2n$). In order to complete the proof for this case we use Hoeffding's inequality to bound the probability that the average of the values for any agent is too large.

$$\sum_{i=1}^{n} \mathrm{P}\left(\frac{v_i(M)}{n} > a\right) \le n\cdot\mathrm{P}\left(\frac{v_1(M)}{n} > a\right) = n\cdot\mathrm{P}\left(\frac{v_1(M)}{m} > \frac{n}{m}\left(\frac{m}{2n} + \sqrt{\frac{3\ln n}{n}}\right)\right)$$
$$= n\cdot\mathrm{P}\left(\frac{v_1(M)}{m} - \frac{1}{2} > \frac{n}{m}\sqrt{\frac{3\ln n}{n}}\right) \le n\cdot e^{-2m\left(\frac{n}{m}\sqrt{\frac{3\ln n}{n}}\right)^2}$$
$$= n\cdot e^{-2\frac{n}{m}\cdot 3\ln n} \le n\cdot e^{-2\ln n} = \frac{1}{n}\,.$$

Hence,

$$\mathrm{P}\left(\exists i \text{ such that } Y_i < \frac{v_i(M)}{n}\right) = \begin{cases} O\left(\frac{\log n}{n}\right) & \text{if } m = 2n \\ O\left(\frac{1}{n}\right) & \text{if } 2n < m < 3n \end{cases}.$$

The remaining cases are for $m \ge 3n$. We give the proof for $m \ge 4n$. The cases for $3n \le m < 3.5n$ and $3.5n \le m < 4n$ differ in small details but they essentially follow the same analysis. We briefly discuss these cases at the end of the proof.

Assume that $kn \le m < (k+1)n, k \ge 4$. We focus on the agent that choses last, i.e., agent n, who has the smallest expected value. He gets exactly k items, and like before let Y_n be the total value he receives. In order to bound $\mathrm{P}(Y_n < \beta)$ we introduce the random variables Z_n and W_n. Consider the following random experiment involving the independent random variables $X_1, \ldots, X_{m-n+1}$, $X_i \sim U[0,1]$ $\forall i \in [m-n+1]$. Given a realization of the X_is, i.e., some values $x_1, \ldots, x_{m-n+1}$ in $[0,1]$, Z_n is defined similarly to Y_n:

- Initially, $Z_n = 0$.
- While there are still x_is left, take the maximum of the remaining x_is, add it to Z_n, remove it from the available numbers, and then remove the x_is with the $n-1$ highest indices.
- Return Z_n.

On the other hand, $W_n = \sum_{i=1}^{k-1} X_{(m+1-in, m-i(n-1))}$, where $X_{(j,t)}$ is the jth order statistic of $X_1, \ldots, X_t$. That is, W_n is defined as the sum of the largest of all x_is, the second largest of the first $m-n+1$ x_is, the third largest of the first $m-2n+2$ x_is, and so on.

It is not hard to see that always $W_n \le Z_n$ (in fact, each term of W_n is less than or equal to the corresponding term of Z_n) and that Z_n follows the same distribution as Y_n. So, $\mathrm{P}\left(Y_n < \beta\right) = \mathrm{P}\left(Z_n < \beta\right) \le \mathrm{P}\left(W_n < \beta\right)$. Using the fact that the ith order statistic in a sample of size ℓ drawn independently from $U[0,1]$ has expected value $\frac{i}{\ell+1}$ and variance $\frac{i(\ell-i+1)}{(\ell+1)^2(\ell+2)}$ (Gentle, 2009), we get

$$\begin{aligned}
\mathrm{E}[W_n] &= \frac{m-n+1}{m-n+2} + \frac{m-2n+1}{m-2n+3} + \ldots + \frac{m-(k-1)n+1}{m-(k-1)n+k} \\
&\ge \frac{(k-1)n+1}{(k-1)n+2} + \frac{(k-2)n+1}{(k-2)n+3} + \ldots + \frac{n+1}{n+k} \\
&> k-1-\frac{1}{(k-1)n} - \frac{2}{(k-2)n} - \ldots - \frac{k-1}{n} > k-1-\frac{(k-1)H_{k-1}}{n}.
\end{aligned}$$

Moreover, if $X'_i = X_{(m+1-in,m-i(n-1))}$ we have

$$\begin{aligned}
\sigma^2_{W_n} = \mathrm{Var}(W_n) &= \sum_{i=1}^{k-1}\sum_{j=1}^{k-1} \mathrm{Cov}(X_i, X_j) \le \sum_{i=1}^{k-1}\sum_{j=1}^{k-1} \sqrt{\mathrm{Var}(X_i)\mathrm{Var}(X_j)} \le \left(\sum_{i=1}^{k-1} \sqrt{\mathrm{Var}(X_i)}\right)^2 \\
&< \left(\sum_{i=1}^{k-1} \frac{\sqrt{i}}{m-in+i+1}\right)^2 < \left(\sqrt{k-1}\sum_{i=1}^{k-1} \frac{1}{(k-i)n}\right)^2 = \frac{(k-1)H^2_{k-1}}{n^2},
\end{aligned}$$

where H_{k-1} is the $(k-1)$-th harmonic number. Now we can bound the probability that any agent receives value less than $1/n$ of his total value.

$$\mathrm{P}\left(Y_i < \frac{v_i(M)}{n}\right) \le \mathrm{P}\left(Y_n < \frac{v_n(M)}{n}\right) \le \mathrm{P}\left(Y_n < \frac{13k}{20}\right) + \mathrm{P}\left(\frac{v_n(M)}{n} > \frac{13k}{20}\right).$$

Next, using Chebyshev's inequality we have

$$\begin{aligned}
\mathrm{P}\left(Y_n < \frac{13k}{20}\right) \le \mathrm{P}\left(W_n < \frac{13k}{20}\right) &= \mathrm{P}\left(\mathrm{E}[W_n] - W_n > \mathrm{E}[W_n] - \frac{13k}{20}\right) \\
&\le \mathrm{P}\left(|\mathrm{E}[W_n] - W_n| > k-1-\frac{(k-1)H_{k-1}}{n} - \frac{13k}{20}\right) \\
&\le \mathrm{P}\left(|\mathrm{E}[W_n] - W_n| > \frac{\frac{7k}{20} - 1 - \frac{(k-1)H_{k-1}}{n}}{\frac{\sqrt{k-1}H_{k-1}}{n}} \sigma_{W_n}\right) \\
&\le \frac{(k-1)H^2_{k-1}}{\left(\left(\frac{7k-20}{20}\right)n - (k-1)H_{k-1}\right)^2}.
\end{aligned}$$

On the other hand, using Hoeffding's inequality,

$$\begin{aligned}
\mathrm{P}\left(\frac{v_n(M)}{n} > \frac{13k}{20}\right) &= \mathrm{P}\left(\frac{v_n(M)}{m} - \frac{1}{2} > \frac{n}{m}\frac{13k}{20} - \frac{1}{2}\right) \\
&\le \mathrm{P}\left(\frac{v_n(M)}{m} - \frac{1}{2} > \frac{13k}{20(k+1)} - \frac{1}{2}\right) \\
&\le e^{-2m\left(\frac{3k-10}{20(k+1)}\right)^2} \le e^{-2kn\left(\frac{3k-10}{20(k+1)}\right)^2}.
\end{aligned}$$

Finally, we take a union bound to get

$$\begin{aligned}\mathrm{P}\left(\exists i \text{ s.t. } Y_i < \frac{v_i(M)}{n}\right) &\le \sum_{i=1}^{n} \mathrm{P}\left(Y_i < \frac{v_i(M)}{n}\right) \\ &\le n\left(\frac{(k-1)H_{k-1}^2}{\left(\left(\frac{7k-20}{20}\right)n-(k-1)H_{k-1}\right)^2} + e^{-2kn\left(\frac{3k-10}{20(k+1)}\right)^2}\right) = O(1/n)\,.\end{aligned}$$

The exact same proof works when $3n \le m < 3.5n$, but instead of $\frac{3k-10}{20(k+1)}$ in Hoeffding's inequality, we have $\frac{3\cdot 3-5}{20(3+0.5)}$ and of course we should adjust $\mathrm{E}[W_n]$ and $\mathrm{Var}(W_n)$ accordingly. When $3.5n \le m < 4n$ on the other hand, we need to consider three items in W_n instead of two, since two items are not enough anymore to guarantee separation of Y_i and $\frac{1}{n}\sum_{j=1}^{m} v_{ij}$ with high probability. That said, the proof is the same, but we should adjust $\mathrm{E}[W_n]$ and $\mathrm{Var}(W_n)$, and instead of $\frac{13k}{20} = \frac{39}{20}$ we may choose 2.5. □

We now state a similar result for any m, generalizing Theorem 2.4.2 that only holds when $m \ge 2n$. We use a modification of Greedy Round-Robin. While $m < 2n$, the algorithm picks any agent uniformly at random and gives him only his "best" item (phase 1). When the number of available items becomes two times the number of active agents, the algorithm proceeds as usual (phase 2). We note that while for $m \ge 2n$ Theorem 2.4.2 gives the stronger guarantee of $\frac{v_i(M)}{n}$ for each agent i, here we can only have a guarantee of $\boldsymbol{\mu}_i(n, M)$.

Theorem 2.4.3. *Let $N = [n]$, $M = [m]$, and the v_{ij}s be as in Theorem 2.4.2. Then, for any m and large enough n, the Modified Greedy Round-Robin algorithm allocates to each agent i a set of items of total value at least $\boldsymbol{\mu}_i(n, M)$ with probability $1 - o(1)$. The $o(1)$ term is $O(1/n)$ when $m > 2n$ and $O(\log n/n)$ when $m \le 2n$.*

Proof. If $m \ge 2n$ then this is a corollary of Theorem 2.4.2. When $m < 2n$, then for any agent i we have $\max_j\{v_{ij}\} \ge \boldsymbol{\mu}_i(n, M)$. So the first agent that receives only his most valuable item has total value at least $\boldsymbol{\mu}_i(n, M)$. If N_a, M_a are the sets of remaining agents and items respectively, after several agents were assigned one item in phase 1 of the algorithm, then by Lemma 2.1.4, for any agent $i \in N_a$, we have $\boldsymbol{\mu}_i(|N_a|, M_a) \ge \boldsymbol{\mu}_i(n, M)$. If $|M_a| < 2|N_a|$ it is also true that $\max_{j\in M_a} v_{ij} \ge \boldsymbol{\mu}_i(|N_a|, M_a)$, so correctness of phase 1 follows by induction. If $|M_a| = 2|N_a|$, then by Theorem 2.4.2 phase 2 guarantees that with high probability each agent $i \in N_a$ will receive a set of items with total value at least $\frac{1}{|N_a|} v_i(M_a) \ge \boldsymbol{\mu}_i(|N_a|, M_a) \ge \boldsymbol{\mu}_i(n, M)$. □

Remark 1. The implicit constants in the probability bounds of Theorems 2.4.2 and 2.4.3 depend heavily on n and m, as well as on the point one uses to separate Y_i and $\frac{1}{n}\sum_{j=1}^{m} v_{ij}$ in the proof of Theorem 2.4.2. Our analysis gives good bounds for the case $2n \le m < 3n$ without requiring very large values for n (especially when $\kappa(n)$ in the proof of Theorem 2.4.2 is not small). For example, if $m = 2.4n$ an appropriate adjustment of our bounds gives a $o(1)$ term less than $1.7/n$ for $n \ge 41$. When we switch from the detailed analysis of the $2n \le m < 3n$ case to the sloppier general treatment for $m \ge 3n$, there is definitely some loss, e.g., for $m = 4n$ we get that the $o(1)$ term is less than

$130/n$ for $n > 450$. This is corrected relatively quickly as m grows, e.g., for $m = 13n$ the $o(1)$ term can be made less than $8/n$ for $n \geq 59$. One can significantly improve the constants by breaking the interval $kn \leq m < (k+1)n$ into smaller intervals (not unlike the $3n \leq m < 3.5n$ case).

Theorems 2.4.2 and 2.4.3 may leave the impression that n has to be large. Actually, there is no reason why we cannot consider n fixed and let m grow. Following closely the proof of Theorem 2.4.2 for $m \geq 4n$, we get the next corollary. Notice that now we can use $\mathrm{E}[W_n] \geq 0.7k$ and $\sigma^2_{W_n} < k$.

Corollary 2.4.4. *Let $N = [n]$, $M = [m]$, and the v_{ij}s be as in Theorem 2.4.2. Then, for fixed n and large enough m, the Greedy Round-Robin algorithm allocates to each agent i a set of goods of total value at least $\frac{1}{n}\sum_{j=1}^{m} v_{ij}$ with probability $1 - O(1/m)$.*

2.5 Directions for Future Research

The most interesting open question is undoubtedly whether one can improve on the 2/3-approximation. Going beyond 2/3 seems to require a drastically different approach. One idea that may deserve further exploration is to pick in each step of Algorithm 4, the best out of all possible matchings (and not just an arbitrary matching as is done in line 7 of the algorithm). This is essentially what we exploit for the special case of $n = 3$ agents. However, for a larger number of agents, this seems to result in a heavy case analysis without any visible benefits. In terms of non-combinatorial techniques, we are not currently aware of any promising LP-based approach to the problem.

Even establishing better ratios for special cases could still provide new insights into the problem. It would be interesting, for example, to see if we can have an improved ratio for the special case studied in Bansal and Sviridenko (2006) for the Santa Claus problem. In this case of additive functions, the value of a good j takes only two distinct values, 0 or v_j. On the other hand, obtaining negative results seems to be an even more challenging task, given our probabilistic analysis and the results of related works. The negative results (Kurokawa, Procaccia, and Wang, 2016; Procaccia and Wang, 2014) require very elaborate constructions, which still do not yield an inapproximability factor far away from 1. Apart from improving the approximation quality, exploring practical aspects of our algorithms is another direction (see, e.g., Spliddit, 2015).

Finally, we have not addressed here the issues of truthfulness and mechanism design, a stimulating topic that is the focus of the next chapter.

Chapter 3

Truthful Allocation Mechanisms Without Payments[1]

The main result of this chapter is a characterization of deterministic truthful mechanisms that allocate all the items to two players with additive valuations. In doing so, we identify some important allocation properties that every truthful mechanism should satisfy. One such crucial property is the notion of *controlling* items (Definition 3.1.7); we say that a player controls an item, whenever it is possible to report values that will guarantee him this item, regardless of the other player's valuation function. We show that truthfulness implies that every item is controlled by some player. Exploiting this property further, greatly helps us in understanding how a mechanism operates. Consequently, our analysis and the characterization we eventually obtain reveals an interesting structure underlying all truthful mechanisms; they can all be essentially decomposed into two components: (i) a *selection part* where players pick their best subset among prespecified choices determined by the mechanism, and (ii) an *exchange part* where players are offered the chance to exchange certain subsets if it is favorable to do so. Hence, we call them *picking-exchange mechanisms*.

Next, we apply our main result and derive several consequences on the design of mechanisms with (approximate) fairness guarantees. We consider various notions of fairness in Section 3.2, starting our discussion with the more standard ones such as proportionality and envy-freeness, and explaining why such concepts cannot be attained—even approximately—by truthful mechanisms. We then focus on more recently studied relaxations of either envy-freeness or proportionality where positive algorithmic results have been obtained (e.g., finding allocations that are envy-free up to one item, or achieve approximate maximin share guarantees). For these notions, we provide tight bounds on the approximation guarantees of truthful mechanisms, settling some of the open problems in this area (Caragiannis et al., 2009; Amanatidis, Birmpas, and Markakis, 2016b). Interestingly, our results also reveal that the best truthful approximation algorithms for fair division are achieved by *ordinal* mechanisms, i.e., mechanisms that exploit only the relative ranking of the items and not the cardinal information of the valuation functions.

[1] A conference paper containing the results of this chapter appeared in EC '17 (Amanatidis et al., 2017).

The heart of our approach for obtaining lower bounds on the approximability of fairness criteria, is a necessary condition for fairness in view of our notion of control, which we call *no control of pairs*. It states that no player should control more than one item. We show how this condition summarizes minimum requirements for various fairness concepts previously studied in the literature. Although this condition does not offer an alternative fairness criterion, it is a useful tool for showing lower bounds.

Finally, in Section 3.3 we provide a general class of truthful mechanisms for the case of multiple players. This class generalizes picking-exchange mechanisms in a non-trivial way. As indicated by our mechanisms, there is a much richer structure in the case of multiple players. In particular, the notion of control does not convey enough information anymore. Instead, there seem to exist several different levels of control.

3.1 Characterization of Truthful Mechanisms

We present our main characterization result in this section. We start in subsection 3.1.1 with the main definitions and illustrating examples, and then we state our result in subsection 3.1.2 along with a road map of the proof. To avoid repetition, when referring to a truthful mechanism $\mathscr{X}$, we mean a truthful mechanism for allocating all the items in M to two players with additive valuation functions.

3.1.1 A Non-Dictatorial Class of Mechanisms

The main result of this section is that every truthful mechanism is a picking-exchange mechanism (Theorem 3.1.6). Before we make a precise statement, we formally define the types of mechanisms involved and provide illustrating examples.

Picking Mechanisms. We start with a family of mechanisms where players make a selection out of choices that the mechanism offers to them. Given a subset S of items, we define a *set of offers* $\mathscr{O}$ on S, as a nonempty collection of proper subsets of S that exactly covers S (i.e., $\bigcup_{T\in\mathscr{O}} T = S$), and in which there is no common element that appears in all subsets (i.e., $\bigcap_{T\in\mathscr{O}} T = \emptyset$).

Definition 3.1.1. A mechanism $\mathscr{X}$ is a *picking mechanism*[2] if there exists a partition (N_1, N_2) of M, and sets of offers $\mathscr{O}_1$ and $\mathscr{O}_2$ on N_1 and N_2 respectively, such that for every profile $\mathbf{v}$,

$$X_i(\mathbf{v}) \cap N_i \in \operatorname*{arg\,max}_{S\in\mathscr{O}_i} v_i(S).$$

Technical nuances aside, such a mechanism can be implemented by first letting player 1 choose his best offer from $\mathscr{O}_1$ and giving what remains from N_1 to player 2. Then it lets player 2 choose his best offer from $\mathscr{O}_2$ and gives what remains from N_2 to player 1. The following example illustrates a picking mechanism.

[2] Picking mechanisms are a generalization of *truthful* picking sequences for two players (see Bouveret and Lang, 2014).

Example 2. Consider the following mechanism $\mathscr{X}$ on a set $M = \{1,\dots 6\}$, which first partitions M into $N_1 = \{1,2,3,4\}, N_2 = \{5,6\}$ and then constructs the offer sets $\mathscr{O}_1 = \{\{1,2\},\{2,3\},\{4\}\}, \mathscr{O}_2 = \{\{5\},\{6\}\}$. On input $\mathbf{v}$, $\mathscr{X}$ first gives to player 1 his best set—with respect to v_1—among $\{1,2\}$, $\{2,3\}$ and $\{4\}$, and then gives what remains from N_1 to player 2. Next, $\mathscr{X}$ gives to player 2 his best set—according to v_2—among $\{5\}$ and $\{6\}$, and then gives what remains from N_2 to player 1. $\mathscr{X}$ resolves ties lexicographically, e.g., in case of a tie, $\{1,2\}$ is preferred to $\{4\}$.

It is not hard to see that $\mathscr{X}$ is truthful. For the following input v, the circles denote the allocation.

$$v = \begin{pmatrix} 3 & \textcircled{5} & \textcircled{5} & 10 & 4 & \textcircled{2} \\ \textcircled{2} & 3 & 6 & \textcircled{1} & \textcircled{5} & 3 \end{pmatrix}.$$

Exchange Mechanisms. We now move to a quite different class of mechanisms. Let X, Y be two disjoint subsets of M. We call the ordered pair (X,Y) an exchange deal. Moreover, we say that an exchange deal (X,Y) is *favorable with respect to* $\mathbf{v}$ if $v_1(Y) > v_1(X)$ and $v_2(Y) < v_2(X)$, while it is *unfavorable with respect to* $\mathbf{v}$ if $v_1(Y) < v_1(X)$ or $v_2(Y) > v_2(X)$. Let S and T be two disjoint subsets of items and let $S_1, S_2, \dots, S_k$ and $T_1, \dots, T_k$ be two collections of nonempty and pairwise disjoint subsets of S and T respectively. We say then that the set of exchange deals $D = \{(S_1,T_1),(S_2,T_2),\dots,(S_k,T_k)\}$ on (S,T) is *valid*.

Definition 3.1.2. A mechanism $\mathscr{X}$ is an *exchange mechanism*[3] if there exists a partition (E_1,E_2) of M, and a valid set of exchange deals $D = \{(S_1,T_1),\dots,(S_k,T_k)\}$ on (E_1,E_2), such that for every profile $\mathbf{v}$, there exists a set of indices $I = I(\mathbf{v}) \subseteq [k]$ for which

$$X_1(\mathbf{v}) = \Big(E_1 \setminus \bigcup_{i\in I} S_i\Big) \cup \bigcup_{i \in I} T_i\,, \quad X_2(\mathbf{v}) = M \setminus X_1\,.$$

Moreover, I contains the indices of every favorable exchange deal with respect to $\mathbf{v}$, but no indices of unfavorable exchange deals.

On a high level, an exchange mechanism initially partitions the items into endowments for the players, and then examines a list of possible exchange deals. Every exchange that improves both players is performed, while every exchange that reduces the value of even one player is avoided. The mechanism may also perform other exchanges where one player is indifferent and the other player can be either indifferent or improved. Whether such exchange deals are materialized or not is up to the tie-breaking rule employed by the mechanism. The following example illustrates an exchange mechanism.

Example 3. Let $M = \{1,\dots 5\}$, and consider the following mechanism $\mathscr{Y}$, with $E_1 = \{1,2,3\}$, $E_2 = \{4,5\}$, and a valid set of exchange deals $D = \{(\{2,3\},\{4\})\}$ on (E_1,E_2): One can think of such a mechanism as if $\mathscr{Y}$ initially reserves the set E_1 for player 1 and the set E_2 for player 2. Then it examines whether exchanging $\{2,3\}$ with $\{4\}$ strictly improves

[3] If we think about E_1, E_2 as fixed a priori, then exchange mechanisms are a generalization of fixed deal exchange rules in general exchange markets for two players (see Pápai, 2007).

both players, and performs the exchange only if the answer is yes. Mechanism $\mathscr{Y}$ is an example of an *exchange mechanism* with only one possible *exchange deal*. Again, one can see that no player has an incentive to lie.

For the following input v, the circles denote the allocation produced.

$$v = \begin{pmatrix} \textcircled{6} & 2 & 3 & \textcircled{7} & 1 \\ 1 & \textcircled{6} & \textcircled{1} & 4 & \textcircled{7} \end{pmatrix}.$$

Picking-Exchange Mechanisms Finally, we define the class of picking-exchange mechanisms which is a generalization of both picking and exchange mechanisms.

Definition 3.1.3. A mechanism $\mathscr{X}$ is a *picking-exchange mechanism* if there exists a partition (N_1, N_2, E_1, E_2) of M, sets of offers $\mathscr{O}_1$ and $\mathscr{O}_2$ on N_1 and N_2 respectively, and a valid set of exchange deals $D = \{(S_1, T_1), \dots, (S_k, T_k)\}$ on (E_1, E_2), such that for every profile $\mathbf{v}$, $X_i(\mathbf{v}) \cap N_i \in \arg\max_{S \in \mathscr{O}_i} v_i(S)$ and $X_1(\mathbf{v}) \cap (E_1 \cup E_2) = (E_1 \setminus \bigcup_{i \in I} S_i) \cup \bigcup_{i \in I} T_i$, where $I = I(\mathbf{v}) \subseteq [k]$ contains the indices of all favorable exchange deals, but no indices of unfavorable exchange deals.

It is helpful to think that a picking-exchange mechanism runs independently a picking mechanism on $N_1 \cup N_2$ and an exchange mechanism on $E_1 \cup E_2$, like in Example 4. Although this is true under the assumption that the players' valuation functions are such that no two sets have the same value, it is not true for general additive valuations. The reason is that the tie-breaking for choosing the offers from $\mathscr{O}_1$ and $\mathscr{O}_2$ may not be independent from the decision of whether to perform each exchange that is neither favorable nor unfavorable.

The following example illustrates a picking exchange mechanism.

Example 4. Let $M = \{1, \dots, 11\}$, and consider the mechanism $\mathscr{Z}$ that partitions M into $N_1 = \{1, 2, 3, 4\}, N_2 = \{5, 6\}$, $E_1 = \{7, 8, 9\}$ and $E_2 = \{10, 11\}$, and is the combination of $\mathscr{X}$ and $\mathscr{Y}$ from the previous two examples: On input $\mathbf{v}$, $\mathscr{Z}$ runs $\mathscr{X}$ on $N_1 \cup N_2$ and $\mathscr{Y}$ on $E_1 \cup E_2$. It outputs the union of the outputs of $\mathscr{X}$ and $\mathscr{Y}$.

For the following input v, the circles denote the final allocation.

$$v = \begin{pmatrix} 3 & \textcircled{5} & \textcircled{5} & 10 & 4 & \textcircled{2} & \textcircled{6} & 2 & 3 & \textcircled{7} & 1 \\ \textcircled{2} & 3 & 6 & \textcircled{1} & \textcircled{5} & 3 & 1 & \textcircled{6} & \textcircled{1} & 4 & \textcircled{7} \end{pmatrix}.$$

3.1.2 Truthfulness and Picking-Exchange Mechanisms

Essentially, we show that a mechanism is truthful if and only if it is a picking-exchange mechanism. We begin with the easier part of our characterization, namely that under the assumption that each valuation function induces a strict preference relation over all possible subsets, every picking-exchange mechanism is truthful. Recall that the set of such profiles is denoted by $\mathcal{V}_m^{\neq}$.

Theorem 3.1.4. *When restricted to $\mathcal{V}_m^{\neq}$, every picking-exchange mechanism $\mathscr{X}$ for allocating m items is truthful.*

Remark 3.1.5. For simplicity, Theorem 3.1.4 is stated for a subclass of additive valuation functions. However, it holds for general additive valuations as long as the mechanism uses a sensible tie-breaking rule (e.g., label-based or welfare-based).[4]

We are now ready to state the main result of this work.

Theorem 3.1.6. *Every truthful mechanism $\mathcal{X}$ can be implemented as a picking-exchange mechanism.*

The rest of this subsection is a road map to the proof of Theorem 3.1.6. The proof is long and technical, so for the sake of presentation, it is broken down to several lemmata. In order to illustrate the high-level ideas, the proofs of those lemmata are deferred to Appendix A.

For the rest of this subsection we assume a truthful mechanism $\mathcal{X}$ for allocating all the items in $M = [m]$ to two players with additive valuation functions. Every statement is going to be with respect to this $\mathcal{X}$.

The Crucial Notion of Control

We begin by introducing the notions of *strong desire* and of *control*, which are of key importance for our characterization. We say that player i *strongly desires* a set S if each item in S has more value for him than all the items of $M \setminus S$ combined, i.e., if for every $x \in S$ we have $v_{ix} > \sum_{y \in M \setminus S} v_{iy}$.

Definition 3.1.7. We say that player i *controls* a set S with respect to $\mathcal{X}$, if every time he strongly desires S he gets it whole, i.e., for every $\mathbf{v} = (v_1, v_2)$ in which player i strongly desires S, then we have that $S \subseteq X_i(\mathbf{v})$.

Clearly, given $\mathcal{X}$, any set S can be controlled by at most one player.

The following is a key lemma for understanding how truthful mechanisms operate. The lemma together with Corollary 3.1.9 below show that every item is controlled by some player under any truthful mechanism.

Lemma 3.1.8 (Control Lemma). *Let $S \subseteq M$. If there exists a profile $\mathbf{v} = (v_1, v_2)$ such that both players strongly desire S, and $S \subseteq X_i(\mathbf{v})$ for some $i \in \{1, 2\}$, then player i controls every $T \subseteq S$ with respect to $\mathcal{X}$.*

Proof. Let $\mathbf{v} = (v_1, v_2)$ be a profile such that both players strongly desire S and $S \subseteq X_1(\mathbf{v})$ (the case where $S \subseteq X_2(\mathbf{v})$ is symmetric). We first prove the statement for $T = S$. Let $\mathbf{v}' = (v'_1, v'_2)$ be any profile in which player 1 strongly desires S, i.e., $v'_{1x} > \sum_{y \in M \setminus S} v'_{1y}, \forall x \in S$. Initially, consider the intermediate profile $\mathbf{v}^* = (v_1, v'_2)$. If $S \cap X_2(\mathbf{v}^*) \neq \emptyset$ then player 2 would deviate from profile $\mathbf{v}$ to $\mathbf{v}^*$ in order to strictly improve his total utility. So by truthfulness we derive that $S \subseteq X_1(\mathbf{v}^*)$. Similarly, in the profile $\mathbf{v}'$, if $S \cap X_2(\mathbf{v}') \neq \emptyset$ then

[4] Describing all such tie-breaking rules seems to be an interesting, nontrivial question for future work, but not our main focus here. It is not hard to see, though, that there exist tie-breaking rules that make a picking-exchange mechanism nontruthful, e.g., break ties on offers of player 1 so that the value that player 2 gets from N_1 is minimized.

player 1 would deviate from $\mathbf{v}'$ to $\mathbf{v}^*$ in order to strictly improve. Thus by truthfulness we have $S \subseteq X_1(\mathbf{v}')$. We conclude that player 1 controls S.

Now, suppose that $\mathbf{v}'' = (v_1'', v_2'')$ is any profile in which player 1 strongly desires $T \subsetneq S$. If $T \nsubseteq X_1(\mathbf{v}'')$ then player 1 could strictly improve his utility by playing v_1' from before (i.e., he declares that he strongly desires S) and getting $S \supsetneq T$. Thus, by truthfulness, $T \subseteq X_1(\mathbf{v}'')$, and we conclude that player 1 controls T. □

Notice here that the existence of sets that are controlled by some player is always guaranteed. Specifically, each singleton $\{x\}$ is always controlled (only) by one of the players. Indeed, when both players strongly desire $\{x\}$, it is always the case that $\{x\} \subseteq X_i(\mathbf{v})$ for some $i \in \{1,2\}$. This is summarized in the following corollary.

Corollary 3.1.9. *Let $\mathscr{X}$ be a truthful mechanism for allocating the items in M to two players with additive valuations. For every $x \in M$ there exists $i \in \{1,2\}$ such that only player i controls $\{x\}$ with respect to $\mathscr{X}$.*

Aside from its use in the current proof, the corollary has implications on fairness, that will be explored in Section 3.2.

Identifying the Components of a Mechanism

Our goal now is to determine the "exchange component" and the "picking component" of mechanism $\mathscr{X}$. Every picking-exchange mechanism is completely determined by the seven sets N_1, N_2, $\mathscr{O}_1$, $\mathscr{O}_2$, E_1, E_2, and D mentioned in Definition 3.1.3 (plus a deterministic tie-breaking rule). Below we try to identify these sets. Later we show that the mechanism's behavior is identical to that of a picking-exchange mechanism defined by them.

To proceed, we will need to consider the collection of all maximal sets controlled by each player. For $i \in \{1,2\}$, let

$$\mathscr{A}_i = \{S \subseteq M \mid \text{player } i \text{ controls } S \text{ and for any } T \supsetneq S,\ i \text{ does not control } T\}.$$

Clearly, every set controlled by player i is a subset of an element of $\mathscr{A}_i$. According to Lemma 3.1.8, if we consider the set $C_i = \bigcup_{S \in \mathscr{A}_i} S$, i.e., the union of all the sets in $\mathscr{A}_i$, this is exactly the set of items that are controlled—as singletons—by player i.

Corollary 3.1.10. *The sets C_1 and C_2 define a partition of M.*

Using the $\mathscr{A}_i$s and the C_is, we define the sets of interest that determine the mechanism. We begin with $E_i = \bigcap_{S \in \mathscr{A}_i} S$ for $i \in \{1,2\}$. As we are going to see eventually in Lemma 3.1.18, the "exchange component" of $\mathscr{X}$ is observed on $E_1 \cup E_2$.

Defining the corresponding valid set of exchange deals D is trickier, and we need some terminology. Recall that $X_i^S(\mathbf{v}) = X_i(\mathbf{v}) \cap S$. For $S \subseteq E_1$ and $T \subseteq E_2$, we say that (S,T) is a *feasible exchange*, if there exists a profile $\mathbf{v}$, such that $X_1^{E_1 \cup E_2}(\mathbf{v}) = (E_1 \setminus S) \cup T$. In such a case, each of S and T is called *exchangeable*. An exchangeable set S is called *minimally exchangeable* if any $S' \subsetneq S$ is not exchangeable. Finally, a feasible exchange (S,T) is a *minimal feasible exchange*, if at least one of S and T is minimally exchangeable. Now let

$$D = \{(S,T) \mid (S,T) \text{ is a minimal feasible exchange with respect to } \mathcal{X}\}.$$

Of course, at this point it is not clear whether D is well defined as a valid set of exchange deals, and this is probably the most challenging part of the characterization.

Next, we define $N_i = C_i \setminus E_i$ and $\mathcal{O}_i = \{S \setminus E_i \mid S \in \mathcal{A}_i\}$ for $i \in \{1,2\}$. As shown in Lemmata 3.1.11 and 3.1.12, we identify the "picking component" of $\mathcal{X}$ on $N_1 \cup N_2$, and $\mathcal{O}_i$ will correspond to the set of offers.

Note that by Corollary 3.1.10 and the above definitions, (N_1, N_2, E_1, E_2) is a partition of M. The intuition behind breaking C_i into N_i and E_i is that player i has different levels of control on those two sets. The fact that E_i is contained in every maximal set controlled by player i will turn out to mean that $\mathcal{X}$ gives the ownership of E_i to player i. On the other hand, the control of player i on N_i is much more restricted as shown below.

Cracking the Picking Component

The first step is to show that the $\mathcal{O}_i$s defined above, greatly restrict the possible allocations of the items of $N_1 \cup N_2$. In particular, whatever player i receives from N_i must be contained in some set of $\mathcal{O}_i$.

Lemma 3.1.11. *For every profile* $\mathbf{v}$ *and every* $i \in \{1,2\}$, *there exists* $S \in \mathcal{O}_i$ *such that* $X_i^{N_i}(\mathbf{v}) \subseteq S$.

The idea behind the proof of Lemma 3.1.11 is that by receiving some $X_i^{N_i}(\mathbf{v})$ not contained in any set of $\mathcal{O}_i$, player i is able to extend his control to subsets not contained in C_i, thus leading to contradiction. The proof, as many of the proofs of the remaining lemmata, includes the careful construction of a series of profiles, where in each step one has to argue about how the allocation does or does not change.

Given the restriction implied by Lemma 3.1.11, next we can prove that the subset of N_i that player i receives must be the best possible from his perspective, hence the mechanism behaves as a picking mechanism on each N_i. Intuitively, suppose that player 1 receives a subset S of N_1 which is not an element of $\mathcal{O}_1$. By Lemma 3.1.11, S is contained in an element S' of $\mathcal{O}_1$. Since player 1 controls S', this means that he gave up part of his control to gain something that he was not supposed to. Actually, it can be shown that it is the case where player 2 also gave part of his control (either on N_2 or E_2). This mutual transfer of control, combined with truthfulness, eventually leads to profiles where some of the items must be given to both players at the same time, hence a contradiction.

Lemma 3.1.12. *For every profile* $\mathbf{v}$ *and every* $i \in \{1,2\}$ *we have* $X_i^{N_i}(\mathbf{v}) \in \arg\max_{S \in \mathcal{O}_i} v_i(S)$.

Now we know that $\mathscr{X}$ behaves as the "right" picking-exchange mechanism on $N_1 \cup N_2$. For most of the rest of the proof we would like to somehow ignore this part of $\mathscr{X}$ and focus on $E_1 \cup E_2$.

Separating the Two Components

As mentioned right after Definition 3.1.3, there is some kind of independence between the two components of a picking-exchange mechanism, at least when restricted on $\mathcal{V}_m^{\neq}$. This independence should be present in $\mathscr{X}$ as well; in fact we are going to exploit it to get rid of $N_1 \cup N_2$ until the last part of the proof.

Lemma 3.1.13. *Let $\mathbf{v} = (v_1, v_2), \mathbf{v}' = (v'_1, v'_2) \in \mathcal{V}_m^{\neq}$ such that $v_{ij} = v'_{ij}$ for all $i \in \{1,2\}$ and $j \in E_1 \cup E_2$. Then $X_1^{E_1 \cup E_2}(\mathbf{v}) = X_1^{E_1 \cup E_2}(\mathbf{v}')$.*

The lemma states that assuming strict preferences over all subsets, the allocation of $E_1 \cup E_2$ does not depend on the values of either player for the items in $N_1 \cup N_2$. What allows this separation is the complete lack of ties in the restricted profile space.

Without loss of generality we may assume that $E_1 \cup E_2 = [\ell]$. We can define a mechanism $\mathscr{X}_E$ for allocating the items of $[\ell]$ to two players with valuation profiles in $\mathcal{V}_\ell^{\neq}$ as

$$\mathscr{X}_E(\mathbf{v}) = (X_1^{E_1 \cup E_2}(\mathbf{v}'), X_2^{E_1 \cup E_2}(\mathbf{v}')), \text{ for every } \mathbf{v} \in \mathcal{V}_\ell^{\neq},$$

where $\mathbf{v}'$ is any profile in $\mathcal{V}_m^{\neq}$ with $v_{ij} = v'_{ij}$ for all $i \in \{1,2\}$ and $j \in [\ell]$. This new mechanism is just the projection of $\mathscr{X}$ on $E_1 \cup E_2$ restricted on a domain where it is well-defined. The truthfulness of $\mathscr{X}_E$ on $\mathcal{V}_\ell^{\neq}$ follows directly from the truthfulness of $\mathscr{X}$ on $\mathcal{V}_m^{\neq}$. Moreover, it is easy to see that player i controls E_i with respect to $\mathscr{X}_E$, for $i \in \{1,2\}$.

The plan is to study $\mathscr{X}_E$ instead of $\mathscr{X}$, show that $\mathscr{X}_E$ is an exchange mechanism, and finally sew the two parts of $\mathscr{X}$ back together and show that everything works properly for any profile in $\mathcal{V}_m$. One issue here is that maybe the set of feasible exchanges with respect to $\mathscr{X}_E$ is greatly reduced, in comparison to the set of feasible exchanges with respect to $\mathscr{X}$, because of the restriction on the domain. In such a case, it will not be possible to argue about exchanges in D that are not feasible anymore. It turns out that this is not the case; the set of possible allocations (of $E_1 \cup E_2$) is the same, whether we consider profiles in $\mathcal{V}_m$ or in $\mathcal{V}_m^{\neq}$.

Lemma 3.1.14. *For every profile $\mathbf{v} \in \mathcal{V}_m$ there exists a profile $\mathbf{v}' \in \mathcal{V}_m^{\neq}$ such that $\mathscr{X}(\mathbf{v}) = \mathscr{X}(\mathbf{v}')$.*

In particular, the set of feasible exchanges on $E_1 \cup E_2$ is exactly the same for $\mathscr{X}$ and $\mathscr{X}_E$, and thus we will utilize the following set of exchanges.

$$D = \{(S,T) \mid (S,T) \text{ is a minimal feasible exchange with respect to } \mathscr{X}_E\}.$$

Cracking the Exchange Component

In the attempt to show that $\mathcal{X}_E$ is an exchange mechanism, the first step is to show that D is indeed a valid set of exchange deals.

Lemma 3.1.15. *D is a valid set of exchange deals on (E_1, E_2).*

The above lemma involves three main steps. First we show that each minimally exchangeable set is involved in exactly one exchange deal. Then, we guarantee that minimally exchangeable sets can be exchanged only with minimally exchangeable sets, and finally, we show that minimally exchangeable sets are always disjoint. There is a common underlying idea in the proofs of these steps: whenever there exist two feasible exchanges that overlap in any way, we can construct a profile where both of them are favorable but the two players disagree on which of them is best. On a high level, each player can "block" his least favorable of the conflicting exchanges, and this leads to violation of truthfulness.

Lemma 3.1.15 implies that every exchangeable set $S \subseteq E_1$ can be decomposed as $S = W \cup \bigcup_{i\in I} S_i$, where $W = S \setminus \bigcup_{i\in I} S_i$ does not contain any minimally exchangeable sets. Ideally, we would like two things. First, the set W in the above decomposition to always be empty, i.e., every exchangeable set should be a union of minimally exchangeable sets. Second, we want every union of minimally exchangeable subsets of E_1 to be exchangeable only with the corresponding union of minimally exchangeable subsets of E_2, and vice versa. It takes several lemmata and a rather involved induction to prove those. A key ingredient of the inductive step is a carefully constructed argument about the value that each player must gain from any exchange.

Lemma 3.1.16. *For every exchangeable set $S \subseteq E_1$, there exists some $I \subseteq [k]$ such that $S = \bigcup_{i\in I} S_i$. Moreover, S is exchangeable with $T = \bigcup_{i\in I} T_i$ and only with T.*

Finally, we have all the ingredients to fully describe $\mathcal{X}_E$ as an exchange mechanism on $E_1 \cup E_2$ and set of exchange deals D.

Lemma 3.1.17. *Given any profile $\mathbf{v} \in \mathcal{V}_\ell^{\neq}$, each exchange in D is performed if and only if it is favorable, i.e., $X_1^{E_1\cup E_2}(\mathbf{v}) = (E_1 \setminus \bigcup_{i\in I} S_i) \cup \bigcup_{i\in I} T_i$, where $I \subseteq [k]$ contains exactly the indices of all favorable exchange deals in D.*

Putting the Mechanism Back Together

As a result of Lemma 3.1.17 (combined, of course, with Lemmata 3.1.12 and 3.1.13), the characterization is complete for truthful mechanisms defined on $\mathcal{V}_m^{\neq}$. For general additive valuation functions, however, we need a little more work. This is to counterbalance the fact that in the presence of ties the allocations of $N_1 \cup N_2$ and $E_1 \cup E_2$ may not be independent.

By Lemmata 3.1.14 and 3.1.16, we know that for any $\mathbf{v} \in \mathcal{V}_m$, $X_1^{E_1\cup E_2}(\mathbf{v})$ is the result of some exchanges of D taking place. There are two things that can go wrong: $\mathcal{X}$ performs an unfavorable exchange, or it does not perform a favorable one. In either of these cases it is possible to construct some profile in $\mathcal{V}_m^{\neq}$ that leads to contradiction. Hence we have the following lemma.

Lemma 3.1.18. *Given any profile* $\mathbf{v} \in \mathcal{V}_m$, $X_1^{E_1 \cup E_2}(\mathbf{v}) = (E_1 \setminus \bigcup_{i \in I} S_i) \cup \bigcup_{i \in I} T_i$, *where* $I \subseteq [k]$ *contains the indices of all favorable exchange deals in* D, *but no indices of unfavorable exchange deals.*

Clearly, Lemma 3.1.18, together with Lemma 3.1.12 concludes the proof of Theorem 3.1.6.

3.1.3 Immediate Implications of Theorem 3.1.6

As mentioned in Sections 1.2 and 1.3, there are several works characterizing truthful mechanisms in combination with other notions, such as Pareto efficiency, nonbossiness, and neutrality (these results are usually for unrestricted, not necessarily additive valuations). Pareto efficiency means that there is no other allocation where one player strictly improves and none of the others are worse-off. Non-bossiness means that a player cannot affect the outcome of the mechanism without changing his own bundle of items. Finally, neutrality refers to a mechanism being consistent with a permutation on the items, i.e., permuting the items results in the corresponding permuted allocation.

Although such notions are not our main focus, the purpose of this short discussion is twofold. On one hand, we illustrate how our characterization immediately implies a characterization for mechanisms that satisfy these extra properties under additive valuations, and on the other hand we see how these properties are either incompatible with fairness or irrelevant in our context.

To begin with, nonbossiness comes for free in our case, since we have two players and all the items must be allocated. Neutrality and Pareto efficiency, however, greatly reduce the space of available mechanisms. Note that it makes more sense to study neutral mechanisms when the valuation functions induce a strict preference order over all sets of items.

Corollary 3.1.19. *Every neutral, truthful mechanism* $\mathcal{X}$ *on* $\mathcal{V}_m^{\neq}$ *can be implemented as a picking-exchange mechanism, such that*

1. *there exists* $i \in \{1,2\}$ *such that* $E_i = [m]$, *or*
2. *there exists* $i \in \{1,2\}$ *such that* $N_i = [m]$ *and* $\mathcal{O}_i = \{S \subseteq [m] \mid |S| = \kappa\}$ *for some* $\kappa < m$.

Corollary 3.1.20. *Every Pareto efficient, truthful mechanism* $\mathcal{X}$ *can be implemented as a picking-exchange mechanism, such that*

1. *there exists* $i \in \{1,2\}$ *such that* $E_i = [m]$, *or*
2. *there exists* $j \in [m]$ *such that* $E_{i_1} = \{j\}$, $E_{i_2} = [m] \setminus \{j\}$, *where* $\{i_1, i_2\} = \{1,2\}$, *and* $D = \{(E_1, E_2)\}$, *or*
3. *there exists* $i \in \{1,2\}$ *such that* $N_i = [m]$ *and* $\mathcal{O}_i = \{S \subseteq N_i \mid |S| = m-1\}$.

It is somewhat surprising that the resulting mechanisms are a strict superset of dictatorships, even when we impose both properties together. Pareto efficiency, however, allows only mechanisms that are rather close to being dictatorial, and thus cannot guarantee fairness of any type. On the other hand, most of the mechanisms defined and studied in Section 3.2 are neutral, yet neutrality is not implied by the fairness concepts we consider, nor the other way around.

3.2 A Necessary Fairness Condition and its Implications

In this section, we explore some implications of Theorem 3.1.6 on fairness properties, i.e., on the design of mechanisms where on top of truthfulness, we would like to achieve fairness guarantees.

In Section 3.2.1 we show that the Control Lemma implies that truthfulness prevents any bounded approximation for envy-freeness and proportionality. Then, we move on describing a necessary fairness condition, in terms of our notion of "control", that summarizes a common feature of several relaxations of fairness and provide a restricted version of our characterization that follows this fairness condition. This will allow us, in Section 3.2.2, to examine what this new class of mechanisms can achieve in each of these fairness concepts.

3.2.1 Implications of the Control Lemma.

Control of singletons.

The basic restriction that truthfulness imposes to every mechanism (leading to poor results for some fairness concepts) comes from Corollary 3.1.9, an immediate corollary of the Control Lemma, stating that every single item is controlled by some player.

We begin by studing how the above corollary affects two of the most researched notions in the fair division literature, namely *proportionality* and *envy-freeness*. It is well known that even without the requirement for truthfulness, it is impossible to achieve any of these two objectives, simply because in the presence of indivisible goods, envy-free or proportional allocations may not exist.[5]

This leads to the definition of approximation versions of these two concepts for settings with indivisible goods. For example, one could try to construct algorithms such that for every instance, an approximation to the minimum possible envy admitted by the instance is guaranteed. Similarly, approximate proportionality can be considered, i.e., find allocations that achieve an approximation to the best possible value that an instance can guarantee to all agents. See also the discussion in Subsection 1.4.2 on defining the approximation versions of these problems. Note that if time complexity

[5]Consider, for instance, a profile where both players desire only the first item and have a negligible value for the other items. Then one of the players will necessarily remain unsatisfied and receive a value close to zero, no matter what the allocation is.

is not an issue, we can always identify the allocation with the best possible envy or with the best possible proportionality, achievable by a given instance.

We are now ready to state our first application, showing that truthfulness prohibits us from having any approximation to the minimum envy or to proportionality. This greatly improves the conclusions of Lipton et al. (2004) and Caragiannis et al. (2009) that truthful mechanisms cannot attain the optimal minimum envy allocation.

Application 3.2.1. *For any truthful mechanism that allocates all the items to two players with additive valuations, the approximation achieved for either proportionality or the minimum envy is arbitrarily bad (i.e., not lower bounded by any positive function of m).*

Proof. Consider a setting with m items, and a truthful mechanism $\mathcal{X}$. Suppose now that item 1 is controlled by player 1 with respect to $\mathcal{X}$. This means that in the profile $\mathbf{v} = ([m\ 1\ 1\ \ldots\ 1], [m^d\ 1\ 1\ \ldots\ 1])$ player 1 must obtain item 1, and player 2 ends up with a negligible fraction of his total value for large enough d. The optimal solution would be to assign the first item to the second player and the last m items to the first player, which provides an envy-free and proportional allocation. We conclude that the approximation guarantee that can be obtained by a truthful mechanism is arbitrarily high. □

So far, the conclusion is that even approximate proportionality or envy-freeness are quite stringent and incompatible with truthfulness because of the Control Lemma. The next step would be to relax these notions. There have been already a few approaches on relaxing proportionality and envy-freeness under indivisible goods, leading to solutions such as the maximin share fairness, envy-freeness up to one item (Budish, 2011), as well as the type of worst-case guarantees proposed by Hill (1987) (recall Definitions 1.4.3, 1.4.5 and 1.4.6 in Subsection 1.4.2). The fact that a truthful mechanism $\mathcal{X}$ yields control of singletons does not seem to have such detrimental effects on these notions. However, if even a single pair of items is controlled by a player, the same situation arises.

Control of pairs

We propose the following *necessary* (but not sufficient) condition that captures a common aspect of all these relaxations of fairness. This allows us to treat all the above concepts of fairness in a unified way.

Definition 3.2.2. We say that a mechanism $\mathcal{X}$ *yields control of pairs* if there exists $i \in \{1,2\}$ and $S \subseteq [m]$ with $|S| = 2$, such that player i controls S with respect to $\mathcal{X}$.

The following lemma states that in order to obtain impossibility results for the above concepts, it is enough to focus on mechanisms with control of pairs.

Lemma 3.2.3. *In order to achieve (either exactly or within a bounded approximation) the above mentioned relaxed fairness criteria, a truthful mechanism that allocates all the items to two players with additive valuations cannot yield control of pairs.*

So now we are ready to move to a complete characterization of truthful mechanisms that do not yield control of pairs. Of course such mechanisms are picking-exchange mechanisms, but our fairness condition allows only singleton offers, and the exchange part is completely degenerate.

Definition 3.2.4. A mechanism $\mathcal{X}$ for allocating all the items in $[m]$ to two players is a *singleton picking-exchange mechanism* if it is a picking-exchange mechanism where for each $i \in \{1,2\}$ at most one of N_i and E_i is nonempty, $|E_i| \leq 1$, and

$$\mathcal{O}_i = \begin{cases} \{\{x\} \mid x \in N_i\} & \text{when } N_i \neq \emptyset \\ \{\emptyset\} & \text{otherwise} \end{cases}$$

i.e., the sets of offers contain all possible singletons.

Hence, typically, in a singleton picking-exchange mechanism player i receives from $N_i \cup E_i$ only his best item. Moreover, for $m \geq 3$, no exchanges are allowed.[6]

Lemma 3.2.5. *Every truthful mechanism for allocating all the items to two players with additive valuation functions that does not yield control of pairs can be implemented as a singleton picking-exchange mechanism.*

Proof. From theorem 3.1.6 we know that every truthful mechanism can be implemented as a picking-exchange mechanism. So consider such a mechanism and let us examine the structure of sets N_i, E_i and $\mathcal{O}_i$. Notice that by the definition of picking-exchange mechanisms, each player i controls $N_i \cup E_i$. If both N_i, E_i are nonempty, or $|E_i| > 1$, or $\mathcal{O}_i$ contains a non-singleton set, then the respective player has control over some pair of items. Thus we can conclude that every possible mechanism can be implemented as a singleton picking-exchange mechanism. □

It is interesting to note that, in contrast to Application 3.2.1, proving Lemma 3.2.5 without Theorem 3.1.6 is not straightforward. In fact, it requires a partial characterization which (on a high level) is similar to characterizing the picking component of general mechanisms.

3.2.2 Applications to Relaxed Notions of Fairness

It is now possible to apply Lemma 3.2.5 on each fairness notion separately, and characterize every truthful mechanism that achieves each criterion. Regarding the remaining proofs of this section, it suffices to focus only on singleton picking-exchange mechanisms. Indeed, by Theorem 3.1.6 we know that every truthful mechanism can be implemented as a picking-exchange mechanism, and by Lemmata 3.2.3 and 3.2.5 only the singleton picking-exchange mechanisms among them may achieve some fairness guarantee.

[6] The only exceptions—and the only such mechanisms where both E_1 and E_2 are nonempty—are two mechanisms for the degenerate case of $m = 2$, e.g., $N_1 = N_2 = \emptyset$, $\mathcal{O}_1 = \mathcal{O}_2 = \{\emptyset\}$, $E_1 = \{a\}$, $E_2 = \{b\}$ and $D = \{(\{a\},\{b\})\}$, where $\{a,b\} = \{1,2\}$.

Envy-freeness up to one item. We start with a relaxation of envy-freeness. Below we provide a complete description of the mechanisms that satisfy this criterion.

Application 3.2.6. *For $m \le 3$, every singleton picking-exchange mechanism achieves envy-freeness up to one item. For $m = 4$ every singleton picking-exchange mechanism with $|N_1| = |N_2| = 2$ achieves envy-freeness up to one item. Finally, for $m \ge 5$ there is no truthful mechanism that allocates all the items to two players and achieves envy-freeness up to one item.*

Proof. Initially it is easy to see that when $m = 1$ or $m = 2$, the statement holds in a trivial way for every singleton picking-exchange mechanisms. Indeed, in every instance each player gets at most one item and thus the value a player derives in the worst case is greater or equal to the value of the empty set (bundle of the other player minus an item).

In the case of $m = 3$, in any instance one player gets one item and the other player two items. The singleton picking-exchange mechanism guarantees that the player who gets one item is allocated with at least his second best in terms of value, so the value he derives is always greater or equal to the value of his least desirable item (bundle of the other player minus an item). On the other hand, the player who is allocated with two items always derives value greater or equal to the value of the empty set.

Finally, in the case of $m = 4$ with $|N_1| = |N_2| = 2$, every player gets two items at every instance. The singleton picking-exchange mechanism guarantees that each player will receive at least his second best item, the value of which is greater or equal to the value of his third or fourth best item.

On the other hand, in case of $m \ge 5$ consider profile $v_1 = [1+\epsilon, 1, ..., 1] \cup [1, \delta, ..., \delta]$, $v_2 = [1, \delta, ..., \delta] \cup [1+\epsilon, 1, ..., 1]$ where $1 \gg \epsilon \gg \delta > 0$. The first vector of values is for N_1 (or E_1) and the second is for N_2 (or E_1); notice that it is possible for one of them to be empty. We only examine singleton picking-exchange mechanisms. It is easy to see that in such a case, by the pigeonhole principle, no singleton picking-exchange mechanism can achieve envy-freeness up to one item for both players. □

Maximin share fairness and related notions. For maximin share allocations a truthful mechanism was suggested by Amanatidis, Birmpas, and Markakis (2016b) for any number of items and any number of players. For two players, their mechanism is the singleton picking-exchange mechanism with $N_1 = [m]$ and produces an allocation that guarantees to each player a $\frac{1}{\lfloor m/2 \rfloor}$-approximation of his maximin share. It was left as an open problem whether a better truthful approximation exists. Here we show that this approximation is tight; in fact, almost any other singleton picking-exchange mechanism performs strictly worse. Note that the best previously known lower bound for two players was 1/2.

Application 3.2.7. *For any m there exists a singleton picking-exchange mechanism that guarantees to player i a $\lfloor \max\{2, m\}/2 \rfloor^{-1}$-approximation of $\boldsymbol{\mu}_i$, for $i \in \{1, 2\}$. There is* *no truthful mechanism that allocates all the items to two players and achieves a better guarantee with respect to maximin share fairness.*

Proof. We only need to prove that among all the singleton picking-exchange mechanisms there is no better approximation ratio than $\lfloor m/2 \rfloor^{-1}$ for $m \geq 3$. Consider profile $v_1 = [1+\epsilon, 1, ..., 1] \cup [|N_1 \cup E_1|, \delta, ..., \delta]$, $v_2 = [|N_2 \cup E_2|, \delta, ..., \delta] \cup [1+\epsilon, 1, ..., 1]$, where $1 \gg \epsilon \gg \delta > 0$. The first vector of values is for N_1 (or E_1) and the second is for N_2 (or E_1); notice that it is possible for one of them to be empty.

It is easy to see that when both $N_1 \cup E_1$, $N_2 \cup E_2$ are nonempty, then $\boldsymbol{\mu}_i \geq |N_i \cup E_i|$ while they both receive value that is slightly greater than 1. Therefore, no singleton picking-exchange mechanism can achieve a better approximation ratio than $\lfloor m/2 \rfloor^{-1}$ for both players.

On the other hand, if $N_1 \cup E_1 = \emptyset$ (the other case is symmetric) then this is the mechanism in Amanatidis, Birmpas, and Markakis, 2016b that achieves exactly $\lfloor m/2 \rfloor^{-1}$. □

Regarding now allocations that guarantee an approximation of the function $V_2(\alpha_i)$ defined by Hill (1987) (recall Definition 1.4.5 in Subsection 1.4.2), the singleton picking-exchange mechanism with $N_1 = [m]$ was also suggested by Markakis and Psomas (2011) as a $\frac{1}{\lfloor m/2 \rfloor}$-approximation of $V_2(\alpha_i)$.[7] This comes as no surprise, since there exists a strong connection between maximin shares and the function V_n, especially for two players. This is illustrated in the following corollary, where both the positive and the negative results coincide with the ones for the maximin share fairness.

Application 3.2.8. *For any m there exists a singleton picking-exchange mechanism that guarantees to player i a $\lfloor \max\{2, m\}/2 \rfloor^{-1}$-approximation of $V_2(\alpha_i)$, for $i \in \{1,2\}$, where $\alpha_i = \max_{j \in [m]} v_{ij}$. There is no truthful mechanism that allocates all the items to two players and achieves a better guarantee with respect to the $V_2(\alpha_i)$s.*

Again, the best previously known lower bound for two players was constant, namely 2/3 due to Markakis and Psomas (2011). In Applications 3.2.7 and 3.2.8, it is stated that there exists a $\frac{1}{\lfloor m/2 \rfloor}$-approximate singleton picking-exchange mechanism. It is interesting that *any* singleton picking-exchange mechanism does not perform much worse. Following the corresponding proofs, we have that even the worst singleton picking-exchange mechanism achieves a $\frac{1}{m-1}$-approximation in each case.

Remark 3.2.9. Gourvès, Monnot, and Tlilane (2015) introduced a variant of V_n, called W_n, and showed that there always exists an allocation such that each player i receives $W_n(\alpha_i) \geq V_n(\alpha_i)$ (where the inequality is often strict). Since the definition of W_n is rather involved even for $n = 2$, we omit a formal discussion about it. However, it is not hard to show that for every valuation function v_i we have $V_2(\alpha_i) \leq W_2(\alpha_i) \leq \boldsymbol{\mu}_i$ and thus the analog of Application 3.2.8 holds.

[7]The approximation factor in Markakis and Psomas, 2011 is expressed in terms of $V_2(1/m)$, but it simplifies to $\lfloor m/2 \rfloor^{-1}$.

Remark 3.2.10. Amanatidis, Birmpas, and Markakis (2016b) made the following interesting observation: every single known truthful mechanism achieving a bounded approximation of maximin share fairness is *ordinal*, in the sense that it only needs a ranking of the items for each player rather than his whole valuation function. Finding truthful mechanisms that explicitly take into account the players' valuation functions in order to achieve better guarantees was posed as a major open problem. Note that, weird tie-breaking aside, all singleton picking-exchange mechanisms are ordinal! Therefore, from the mechanism designer's perspective, it is impossible to exploit the extra cardinal information given as input and at the same time maintain truthfulness and some nontrivial fairness guarantee.

3.3 Truthful Mechanisms for Many Players

We introduce a family of non-dictatorial, truthful mechanisms for any number of players. Our mechanisms are defined recursively; in analogy to serial dictatorships, the choices of a player define the sub-mechanism used to allocate the items to the remaining players. Here, however, this serial behavior is observed "in parallel" in several sets of a partition of M.

A *generalized deal* between k players is a collection of (up to $k(k-1)$) exchange deals between pairs of players. A set D of generalized deals is called *valid* if all the sets involved in all these exchange deals are nonempty and pairwise disjoint. Given a profile $\mathbf{v} = (v_1, v_2, \dots, v_n)$ we say that a generalized deal is *favorable* if it strictly improves all the players involved, while it is *unfavorable* if there exists a player involved whose utility strictly decreases.

Definition 3.3.1. A mechanism $\mathcal{X}$ for allocating all the items in $[m]$ to n players is called a *serial picking-exchange mechanism* if

1. when $n = 1$, $\mathcal{X}$ always allocates the whole $[m]$ to player 1.

2. when $n \geq 2$, there exist a partition $(N_1, \dots, N_n, E_1, \dots, E_n)$ of $[m]$, sets of offers $\mathcal{O}_i$ on N_i for $i \in [n]$, a valid set D of generalized deals, and a mapping f from subsets of M to serial picking-exchange mechanisms for $n-1$ players, such that for every profile $\mathbf{v} = (v_1, \dots, v_n)$ we have for all $i \in [n]$:

   - $X_i^{N_i}(\mathbf{v}) \in \arg\max_{S \in \mathcal{O}_i} v_i(S)$,
   - $X_i^{E}(\mathbf{v})$, where $E = \bigcup_{j \in [n]} E_j$, is the result of starting with E_i and performing some of the deals in D, including all the favorable deals but no unfavorable ones,
   - the items of $N_i \setminus X_i^{N_i}(\mathbf{v})$ are allocated to players in $[n] \setminus \{i\}$ using the serial picking-exchange mechanism $f\big(N_i \setminus X_i^{N_i}(\mathbf{v})\big)$.

Clearly, serial picking-exchange mechanisms generalize picking-exchange mechanisms studied in Section 3.1. The following example illustrates how such a mechanism looks like for three players.

Example 5. Suppose that we have three players with additive valuations. For simplicity, assume that each player's valuation induces a strict preference over all possible subsets of items. Let $M = [100]$ be the set of items, and consider the following relevant ingredients of our mechanism:

- $N_1 = \{1,2,\dots,20\}$, $\mathcal{O}_1 = \{\{1,2,3\}, N_1 \setminus \{1\}\}$
- $N_2 = \{21,22,\dots,50\}$, $\mathcal{O}_2 = \{S \subseteq N_2 \mid |S| = 6\}$
- $N_3 = \{51,52,\dots,70\}$, $\mathcal{O}_3 = \{\{51,\dots,60\},\{61,\dots,70\}\}$
- $E_1 = \{71,\dots,80\}$, $E_2 = \{81,\dots,90\}$, $E_3 = \{91,\dots,100\}$
- $D = \big\{\big[(\{75,79\},\{83\})^{1,3}\big], \big[(\{71\},\{88\})^{1,2}, (\{72,80\},\{95\})^{1,3}, (\{85\},\{99,100\})^{2,3}\big]\big\}$
- f is a mapping from subsets of M to picking-exchange mechanisms (for 2 players)

The above sets are the analog of the corresponding sets of a picking-exchange mechanism. The deals, however, are a bit more complex. E.g., by $\big[(\{71\},\{88\})^{1,2}, (\{72,80\},\{95\})^{1,3}, (\{85\},\{99,100\})^{2,3}\big]$ we denote the deal in which:

- player 1 gives item 71 to player 2 and items 72, 80 to player 3
- player 2 gives item 88 to player 1 and item 85 to player 3
- player 3 gives item 95 to player 1 and items 99, 100 to player 2

The mapping f suggests which truthful mechanism should be used every time there are items left to be allocated to only two players.

We are ready to describe our mechanism $\mathcal{X}$:

1. The mechanism gives endowments E_1, E_2, E_3 to the three players and then performs each exchange deal that strictly improves all the players involved.
2. Then, for each $i \in \{1,2,3\}$, the mechanism gives to player i his best set in $\mathcal{O}_i$, say S_i.
3. Finally, for each $i \in \{1,2,3\}$, $\mathcal{X}$ uses mechanism $f(N_i \setminus S_i)$ to allocate the items of $N_i \setminus S_i$ to players in $\{1,2,3\} \setminus i$.

Like picking-exchange mechanisms, serial picking-exchange mechanisms are truthful, given an appropriate tie-breaking rule (e.g., a label-based tie-breaking rule). To bypass a general discussion about tie-breaking, however, we may assume that each player's valuation induces a strict preference over all subsets of M. We denote by $\mathcal{V}^{\neq}_{n,m}$ the set of profiles that only include such valuation functions. Following almost the same proof, however, we have that for general additive valuations every serial picking-exchange mechanism is truthful when using label-based tie-breaking.

Theorem 3.3.2. *When restricted to $\mathcal{V}^{\neq}_{n,m}$, every serial picking-exchange mechanism $\mathcal{X}$ for allocating m items to n players is truthful.*

3.4 Directions for Future Research

A natural question to ask is whether our characterization can be extended for more than two players. Characterizing the truthful mechanisms without money for any number of additive players is, undoubtedly, a fundamental open problem. However, as indicated by Definition 3.3.1, there seems to be a much richer structure when one attempts to describe such mechanisms, even though serial picking-exchange mechanisms are only a subset of nonbossy truthful mechanisms. In particular, the notion of control that was crucial for identifying the structure of truthful mechanisms for two players does not convey enough information anymore. Instead, there seem to exist several different levels of control, and understanding this structure still remains a very interesting and intriguing question. It is also not clear if more positive results can arise when payments are allowed. Similar mechanism design questions also remain open for a related problem studied by Markakis and Psomas, 2011.

Part II

Procurement Auctions with Budget Constraints

Chapter 4

Introduction

4.1 Budget-Feasible Mechanism Design

In this second part, we study a class of mechanism design problems under a budget constraint. Consider a reverse auction setting, where a single buyer wants to select a subset, among a set A of agents, for performing some tasks. Each agent i comes at a cost c_i, in the case that he is chosen. The buyer has a budget B and a valuation function $v(\cdot)$, so that $v(S)$ is the derived value if $S \subseteq A$ is the chosen set. The purely algorithmic version then asks to maximize the generated value subject to the constraint that the total cost of the selected agents should not exceed B (often referred to as a "hard" budget constraint).

The purely algorithmic version of the problem results in natural "budgeted" versions of known optimization problems. Since these problems are typically NP-hard, our focus is on approximation algorithms. Most importantly, in the setting considered here, the true cost of each agent is private information and we would like to design mechanisms that elicit truthful reporting by all agents. Hence, our ideal goal is to have truthful mechanisms that achieve a good approximation to the optimal value for the auctioneer, and are *budget feasible*, i.e., the sum of the payments to the agents does not exceed the prespecified budget. This framework of budget feasible mechanisms is motivated by recent application scenarios including crowdsourcing platforms, where agents can be viewed as workers providing tasks (e.g., Anari, Goel, and Nikzad, 2014; Goel, Nikzad, and Singla, 2014), and influence maximization in networks, where agents correspond to influential users (see e.g., Singer, 2012, where the chosen objective is a coverage function).

Budget feasibility is a tricky property that makes the problem more challenging, with respect to truthfulness, as it already rules out well known mechanisms such as VCG. Although the algorithmic versions of such problems often admit constant factor approximation algorithms, it is not clear how to appropriately convert them into truthful budget feasible mechanisms. We stress that the question is nontrivial even if we allow exponential time algorithms, since computational power does not necessarily make the problem easier (see the discussion in Dobzinski, Papadimitriou, and Singer (2011)). All these issues create an intriguing landscape, where one needs to strike a balance between the incentives of the agents and the budget constraints.

Budgeted mechanism design was first studied by Singer (2010) when $v(\cdot)$ is an additive or a nondecreasing submodular function. Later on, follow-up works provided refinements and further results for richer classes of functions like XOS and subadditive functions (see the related work section). Although these results shed more light on our understanding of the problem, there are still several interesting issues that remain unresolved both for submodular and non-submodular cases. First, the current results on submodular valuations are not known to be tight. Further, and most importantly, when going beyond submodularity, to XOS functions, we are not even aware of general mechanisms with small approximation guarantees, let alone deterministic polynomial time mechanisms.

Going beyond submodular valuations creates severe challenges in general, and thus any results for general classes of functions are relatively poor. Specific problems have been studied however, with quite promising results. The first attempt with a non-submodular objective was due to Chen, Gravin, and Lu (2011) who gave a $(2+\sqrt{2})$-approximation mechanism for a non-submodular variation of Knapsack, while, recently, Goel, Nikzad, and Singla (2014) studied a budgeted maximization problem with matching constraints, which is not submodular, and they achieve an approximation ratio of $3+o(1)$, but under the large market assumption.[1] Despite such scarce results, we are not aware of mechanisms with good guarantees even for very standard variants of Knapsack with matching or matroid constraints. Such problems are studied in Chapter 7.

Moreover, most existing works make the assumption that the valuation function is non-decreasing, i.e., $v(S) \le v(T)$ for $S \subseteq T$, notable exceptions being the works of Dobzinski, Papadimitriou, and Singer (2011) and Bei et al. (2012). Although monotonicity makes sense in several scenarios, one can think of examples where it is violated. For instance, Dobzinski, Papadimitriou, and Singer (2011) studied the unweighted Budgeted Max Cut problem, as an eminent example of a non-monotone submodular objective function. Moreover, when studying models for influence maximization problems in social networks, adding more users to the selected set may some times bring negative influence (Borodin, Filmus, and Oren, 2010) (some combinations of users may also not be compatible or well fitted together). To further motivate the study of non-monotone submodular objectives, consider the following well-studied sensor placement problem (Caselton and Zidek, 1984; Cressie, 1993; Krause, Singh, and Guestrin, 2008): assume that we want to monitor some spatial phenomenon (e.g., the temperature of a specific environment), modeled as a Gaussian process. We may place sensing devices on some of the prespecified locations, but each location has an associated cost. A criterion for finding an optimal such placement, suggested by Caselton and Zidek (1984) for the unit cost case, is to maximize the *mutual information* between chosen and non chosen locations, i.e., we search for the subset of locations that minimizes the uncertainty about the estimates in the remaining space.

[1] A market is said to be large if the number of participants is large enough that no single person can affect significantly the market outcome, i.e., $\max_i c_i/B = o(1)$.

Such mutual information objectives are submodular but not monotone. In addition, it is straightforward to modify this problem to model participatory crowdsensing scenarios where users have incentives to lie about the true cost of installing a sensor.

It becomes apparent that we would like to aim for truthful mechanisms with good performance for subclasses of non-monotone functions. At the moment, the few results known for arbitrary non-monotone submodular functions have very large approximation ratios and often superpolynomial running time. Even worse, in most cases, we do not even know of deterministic mechanisms (see Table 6.1). In trying to impose more structure so as to have better positive results, there is an interesting observation to make: the non-monotone examples mentioned so far, i.e., cut functions and mutual information functions, are *symmetric*[2] *submodular*, a prominent subclass of non-monotone submodular functions, where the value of a set S equals the value of its complement. This subclass has received already considerable attention in operations research (see e.g., Fujishige, 1983; Queyranne, 1998, where more examples are also provided). We therefore find that symmetric submodular functions form a suitable starting point for the study of non-monotone functions. This is the subject of Chapter 6.

4.2 Related Work

The study of budget feasible mechanisms, as considered here, was initiated by Singer (2010), who gave a randomized constant factor approximation mechanism for nondecreasing submodular functions. Later, Chen, Gravin, and Lu (2011) significantly improved these approximation ratios, obtaining a randomized, polynomial time mechanism achieving a 7.91-approximation and a deterministic one with a 8.34-approximation. Their deterministic mechanism does not run in polynomial time in general, but it can be modified to do so for special cases at the expense of its performance (see the beginning of Section 5.2). As an example, Singer (2012) followed a similar approach to obtain a deterministic, polynomial time, 31.03-approximation mechanism for the unweighted version of Budgeted Max Coverage, a class that we also consider in Section 5.2. Along these lines, Horel, Ioannidis, and Muthukrishnan (2014) consider another family of submodular functions and give a deterministic, polynomial time, constant approximation for the so-called Experimental Design Problem, under a mild relaxation on truthfulness. For subadditive functions, Dobzinski, Papadimitriou, and Singer (2011) suggested a randomized $O(\log^2 n)$-approximation mechanism, and they gave the first constant factor mechanisms for non-monotone submodular objectives, specifically for *cut functions*. The factor for subadditive functions was later improved to $O(\frac{\log n}{\log\log n})$ by Bei et al. (2012), who also gave a randomized $O(1)$-approximation mechanism for XOS functions, albeit in exponential time, and further initiated the

[2] In some works on mechanism design, symmetric submodular functions have a different meaning and refer to the case where $v(S)$ depends only on $|S|$. Here we have adopted the terminology of earlier literature on submodular optimization (e.g., Fujishige, 1983).

Bayesian analysis in this setting. Further improved $O(1)$-approximation mechanisms for XOS functions have also been obtained in Leonardi et al. (2016). There is also a line of related work under the *large market* assumption (where no participant can significantly affect the market outcome), which allows for polynomial time mechanisms with improved performance (see, e.g., Singla and Krause, 2013; Anari, Goel, and Nikzad, 2014; Goel, Nikzad, and Singla, 2014; Balkanski and Hartline, 2016; Jalaly and Tardos, 2017).

A somewhat complementary line of work involves the design of frugal mechanisms. These are mechanisms where one cares for minimizing the total amount of payments that are required by the mechanism, for finding a good solution. A series of results has been obtained over the years on designing frugal mechanisms (see, e.g., Archer and Tardos, 2007; Chen et al., 2010; Karlin, Kempe, and Tamir, 2005; Kempe, Salek, and Moore, 2010). Since here we have a hard budget constraint that should never be exceeded, results from this area do not generally transfer to our setting.

There is also a plethora of works on auctions that take budgets into account, from the bidder's point of view, motivated mainly by sponsored search auctions, see among others, Borgs et al., 2005; Dobzinski, Lavi, and Nisan, 2012; Goel, Mirrokni, and Leme, 2012 for some representative problems that have been tackled. Although these are fundamentally different problems than ours, they do highlight the difficulties that arise in the presence of budget constraints.

On maximization of submodular functions subject to knapsack or other type of constraints, there is a vast literature, going back several decades (see, e.g., Nemhauser, Wolsey, and Fisher, 1978; Wolsey, 1982). More recently, Lee et al. (2010) provided the first constant factor randomized algorithm for submodular maximization under k matroid and k knapsack constraints, with factors $k+2+\frac{1}{k}$ and 5 respectively. The problem was also studied by Gupta et al. (2010) who proposed a randomized algorithm, which achieves a $(4+\alpha)$-approximation[3] in case of knapsack constraints, where α is the approximation guarantee of the unconstrained submodular maximization. Later on, Chekuri, Vondrák, and Zenklusen (2014) suggested a randomized 3.07-approximation algorithm improving the previously known results. Finally, Feldman, Naor, and Schwartz (2011) and Kulik, Shachnai, and Tamir (2013) proposed their own randomized algorithms when there are knapsack constraints, achieving an e-approximation.[4]

4.3 Preliminaries and Notation

We use $A = [n] = \{1, 2, \ldots, n\}$ to denote a set of n agents. Each agent i is associated with a private cost c_i, denoting the cost for participating in the solution. We consider a procurement auction setting, where the auctioneer is equipped with a valuation

[3] In the case of symmetric submodular functions the algorithm gives a deterministic 6-approximation.

[4] The algorithm of Kulik, Shachnai, and Tamir (2013) can be derandomized without any performance loss, but only assuming an additional oracle for the extension by expectation, say V, of the objective function v. When only an oracle for v is available, estimation of V by sampling is required in general.

function $v: 2^A \to \mathbb{Q}_{\geq 0}$ and a budget $B > 0$. For $S \subseteq A$, $v(S)$ is the value derived by the auctioneer if the set S is selected (for singletons, we will often write $v(i)$ instead of $v(\{i\})$). Therefore, the algorithmic goal in all the problems we study is to select a set S that maximizes $v(S)$ subject to the constraint $\sum_{i \in S} c_i \leq B$. We assume oracle access to v via value queries, i.e., we assume the existence of a polynomial time value oracle that returns $v(S)$ when given as input a set S.

Throughout our work, we consider valuation functions that are non negative, i.e., $v(S) \geq 0$ for any $S \subseteq A$, and we make the natural assumption that $v(\emptyset) = 0$. We will focus on valuations that come from two natural classes of functions, namely submodular and XOS functions.

Definition 4.3.1. A valuation function, defined on 2^A for some set A, is

(i) non-decreasing, if $v(S) \leq v(T)$ for any $S \subseteq T \subseteq A$.

(ii) submodular, if $v(S \cup \{i\}) - v(S) \geq v(T \cup \{i\}) - v(T)$ for any $S \subset T \subset A$, and $i \not\in T$.

(iii) symmetric submodular, if it is submodular and moreover, $v(S) = v(A \setminus S)$ for any $S \subseteq A$.

(iv) XOS or *fractionally subadditive*, if there exist non-negative additive functions $\alpha_1, ..., \alpha_r$, for some finite r, such that $v(S) = \max\{\alpha_1(S), \alpha_2(S), ..., \alpha_r(S)\}$.

We note that the class XOS is a strict superclass of non-decreasing submodular valuations. Also, it is easy to see that v cannot be both symmetric and non-decreasing unless it is a constant function. In fact, if this is the case and $v(\emptyset) = 0$, then $v(S) = 0$, for all $S \subseteq A$. We also state an alternative definition of a submodular function, which will be useful later on.

Theorem 4.3.2 (Nemhauser, Wolsey, and Fisher (1978))**.** *A set function v is submodular if and only if for all $S, T \subseteq A$ we have $v(T) \leq v(S) + \sum_{i \in T \setminus S} (v(S \cup \{i\}) - v(S)) - \sum_{i \in S \setminus T} (v(S \cup T) - v(S \cup T \setminus \{i\}))$.*

We often need to argue about optimal solutions of sub-instances, from an instance we begin with. Given a cost vector $\mathbf{c}$, and a subset $X \subseteq A$, we denote by $\mathbf{c}_X$ the projection of $\mathbf{c}$ on X, and by $\mathbf{c}_{-X}$ the projection of $\mathbf{c}$ on $A \setminus X$. We also let $\textsc{opt}(X, v, \mathbf{c}_X, B)$ be the value of an optimal solution to the restriction of this instance on X, i.e., $\textsc{opt}(X, v, \mathbf{c}_X, B) = \max_{S: S \subseteq X, \mathbf{c}(S) \leq B} v(S)$. Similarly, $\textsc{opt}(X, v, \mathbf{c}_X, \infty)$ denotes the value of an optimal solution to the unconstrained version of the problem restricted on X. For the sake of readability, we usually drop the valuation function and the cost vector, and write $\textsc{opt}(X, B)$ or $\textsc{opt}(X, \infty)$.

Finally, in Chapter 6 we make one further assumption: we assume that there is at most one item whose cost exceeds the budget. As shown in Lemma B.1.1 and Corollary B.1.2 in Appendix B, this is without loss of generality.

Local Optima and Local Search. Given $v: 2^A \to \mathbb{Q}$, a set $S \subseteq A$ is called a $(1+\epsilon)$-*approximate local optimum* of v, if $(1+\epsilon)v(S) \geq v(S \setminus \{i\})$ and $(1+\epsilon)v(S) \geq v(S \cup \{i\})$ for every $i \in A$. When $\epsilon = 0$, S is called an *exact local optimum* of v. Note that if v is symmetric submodular, then S is a $(1+\epsilon)$-approximate local optimum if and only if $A \setminus S$ is a $(1+\epsilon)$-approximate local optimum.

Approximate local optima produce good approximations in unconstrained maximization of general submodular functions (Feige, Mirrokni, and Vondrák, 2011). However, here they are of interest for a quite different reason that becomes apparent in Lemmata 6.1.1 and 6.2.1. We can efficiently find approximate local optima using the local search algorithm APPROX-LOCAL-SEARCH of Feige, Mirrokni, and Vondrák (2011). Note that this is an algorithm for the unconstrained version of the problem, when there are no budget constraints.

APPROX-LOCAL-SEARCH(A, v, ϵ) (Feige, Mirrokni, and Vondrák, 2011)

$S = \{i^*\}$, where $i^* \in \arg\max_{i \in A} v(i)$
while *there exists some a such that* $\max\{v(S \cup \{a\}), v(S \setminus \{a\})\} > (1+\epsilon/n^2)v(S)$ **do**
  if $v(S \cup \{a\}) > (1+\epsilon/n^2)v(S)$ **then**
    $S = S \cup \{a\}$
  else
    $S = S \setminus \{a\}$
return S

If we care to find an exact local optimum, we can simply set $\epsilon = 0$. In this case, however, we cannot argue about the running time of the algorithm in general.

Lemma 4.3.3 (inferred from Feige, Mirrokni, and Vondrák (2011))**.** *Given a submodular function* $v: 2^{[n]} \to \mathbb{Q}_{\geq 0}$ *and a value oracle for* v, APPROX-LOCAL-SEARCH(A, v, ϵ) *outputs a* $\left(1 + \frac{\epsilon}{n^2}\right)$*-approximate local optimum using* $O\left(\frac{1}{\epsilon} n^3 \log n\right)$ *calls to the oracle.*

Mechanism Design. Each agent here only has his cost as private information, hence we are in the domain of single-parameter problems. A mechanism $\mathcal{M} = (f, p)$ in our context consists of an outcome rule f and a payment rule p. Given a vector of cost declarations, $\mathbf{b} = (b_i)_{i \in A}$, where b_i denotes the cost reported by agent i, the outcome rule of the mechanism selects the set $f(\mathbf{b})$. At the same time, it computes payments $p(\mathbf{b}) = (p_i(\mathbf{b}))_{i \in A}$ where $p_i(\mathbf{b})$ denotes the payment issued to agent i. Hence, the final utility of agent i is $p_i(\mathbf{b}) - c_i$.

The main properties we want to ensure for our mechanisms are the following.

Definition 4.3.4. A mechanism $\mathcal{M} = (f, p)$ is

1. *truthful*, if for any cost vector $\mathbf{c} = (c_i)_{i \in A}$, any player $i \in A$, and any b_i: $p_i(\mathbf{c}) \geq p_i(b_i, \mathbf{c}_{-i})$. That is, reporting c_i is a dominant strategy for every agent i.
2. *individually rational*, if $p_i(\mathbf{b}) \geq 0$ for every $i \in A$, and $p_i(\mathbf{b}) \geq c_i$, for every $i \in f(\mathbf{b})$.
3. *budget feasible*, if $\sum_{i \in A} p_i(\mathbf{b}) \leq B$ for every $\mathbf{b}$.

Compare the definition of truthfulness above with Definition 1.4.1. When referring to randomized mechanisms, the notion of truthfulness we use is *universal truthfulness*, which means that the mechanism is a probability distribution over deterministic truthful mechanisms.

For single-parameter problems we use the characterization by Myerson (1981) for deriving truthful mechanisms. In particular, we say that an outcome rule f is *monotone*, if for every agent $i \in A$, and any vector of cost declarations $\mathbf{b}$, if $i \in f(\mathbf{b})$, then $i \in f(b'_i, \mathbf{b}_{-i})$ for $b'_i \leq b_i$. This simply means that if an agent i is selected in the outcome by declaring cost b_i, then by declaring a lower cost he should still be selected. Myerson's lemma below implies that monotone algorithms admit truthful payment schemes.

Lemma 4.3.5. *Given a monotone algorithm f, there is a unique payment scheme p such that (f, p) is a truthful and individually rational mechanism, given by*

$$p_i(b) = \begin{cases} \sup_{b_i \in [c_i, \infty)} \{b_i : i \in f(b_i, b_{-i})\}, & \text{if } i \in f(b) \\ 0, & \text{otherwise} \end{cases}$$

Lemma 4.3.5 is known as Myerson's lemma, and the payments are often referred to as *threshold payments*, since they indicate the threshold at which an agent stops being selected. Myerson's lemma simplifies the design of truthful mechanisms by focusing only on constructing monotone algorithms and not having to worry about the payment scheme. Nevertheless, in the setting we study here budget feasibility clearly complicates things further. For all the algorithms presented in the next sections, we always assume that the underlying payment scheme is given by Myerson's lemma.

Chapter 5

Mechanisms for Non-Decreasing Submodular Objectives[1]

We begin by optimizing existing truthful, budget-feasible mechanisms for non-decreasing submodular valuation functions following the analysis of Jalaly and Tardos (2017). Then we proceed to provide a framework for designing deterministic mechanisms that run in polynomial time, given that the objective has a "well-behaved" LP formulation. Note that previously known mechanisms give no such guarantee. We apply our approach on *coverage functions*, a notable subclass of non-decreasing submodular functions, in Section 5.2.1. This class has already received attention in previous works (Singer, 2010; Singer, 2012), motivated by problems related to influence maximization in social networks. Our mechanism reduces roughly by a factor of 3 (from 31.03 to 10.03) the known approximation of Singer (2012) and also generalizes it to the weighted version of coverage functions.

In the mechanisms we design for non-monotone submodular functions, in Chapter 6, we repeatedly make use of truthful, budget-feasible mechanisms for non-decreasing submodular functions as subroutines. Therefore, this chapter paves the way for the next one.

5.1 Optimizing Existing Mechanisms

The best known truthful, budget-feasible mechanisms for non-decreasing submodular objectives are due to Chen, Gravin, and Lu (2011). Here, we follow the improved analysis of Jalaly and Tardos (2017) for the approximation ratio of the randomized mechanism RAND-MECH-SM of Chen, Gravin, and Lu (2011), stated below.

RAND-MECH-SM$(A, v, \mathbf{c}, B)$ (Chen, Gravin, and Lu, 2011)

Set $A' = \{i \mid c_i \le B\}$ and $i^* \in \arg\max_{i \in A'} v(i)$
with probability $\frac{2}{5}$ **return** i^*
with probability $\frac{3}{5}$ **return** GREEDY-SM$(A, v, \mathbf{c}, B/2)$

The mechanism GREEDY-SM is a greedy algorithm that picks agents according to their ratio of marginal value over cost, given that this cost is not too large. For the sake of presentation, we assume the agents are sorted in descending order with respect to

[1] A conference paper containing a preliminary version of the results of this chapter appeared in WINE '16 (Amanatidis, Birmpas, and Markakis, 2016a).

this ratio. The marginal value of each agent is calculated with respect to the previous agents in the ordering, i.e., $1 = \arg\max_{j \in A} \frac{v(j)}{c_j}$ and $i = \arg\max_{j \in A \setminus [i-1]} \frac{v([j]) - v([j-1])}{c_j}$ for $i \geq 2$.

GREEDY-SM$(A, v, \mathbf{c}, B/2)$ (Chen, Gravin, and Lu, 2011)

Let $k = 1$ and $S = \emptyset$
while $k \leq |A|$ **and** $v(S \cup \{k\}) > v(S)$ **and** $c_k \leq \frac{B}{2} \cdot \frac{v(S \cup \{k\}) - v(S)}{v(S \cup \{k\})}$ **do**
  $S = S \cup \{k\}$
  $k = k + 1$
return S

Lemma 5.1.1 (inferred from Chen, Gravin, and Lu (2011) and Jalaly and Tardos (2017))**.** GREEDY-SM *is monotone. Assuming a non-decreasing submodular function* v *and the payments of Myerson's lemma,* GREEDY-SM$(A, v, \mathbf{c}, B/2)$ *is budget-feasible and outputs a set* S *such that* $v(S) \geq \frac{1}{3} \cdot \textsc{opt}(A, B) - \frac{2}{3} \cdot v(i^*)$.

A derandomized version of RAND-MECH-SM is also provided by Chen, Gravin, and Lu (2011). It has all the desired properties, while suffering a small loss on the approximation factor. Here, following the improved analysis of Jalaly and Tardos (2017) for the ratio of RAND-MECH-SM, we fine-tune this derandomized mechanism to obtain MECH-SM that has a better approximation guarantee.

MECH-SM$(A, v, \mathbf{c}, B)$

Set $A' = \{i \mid c_i \leq B\}$ and $i^* \in \arg\max_{i \in A'} v(i)$
if $(2 + \sqrt{6}) \cdot v(i^*) \geq \textsc{opt}(A \setminus \{i^*\}, B)$ **then**
  return i^*
else
  return GREEDY-SM$(A, v, \mathbf{c}, B/2)$

The next theorem summarizes the properties of RAND-MECH-SM and MECH-SM.

Theorem 5.1.2 (inferred from Chen, Gravin, and Lu (2011), Jalaly and Tardos (2017), and Lemma 6.3.3)**.**

i. RAND-MECH-SM *runs in polynomial time, it is universally truthful, individually rational, budget-feasible, and has approximation ratio* 5.

ii. MECH-SM *is deterministic, truthful, individually rational, budget-feasible, and has approximation ratio* $3 + \sqrt{6}$.

Proof. Monotonicity (and thus truthfulness and individual rationality) and budget-feasibility of both mechanisms directly follow from Chen, Gravin, and Lu (2011). What is left to show are the approximation guarantees.

The key fact here is the following lemma from Jalaly and Tardos (2017). Note that the lemma also follows from Lemma 6.3.3 for $\epsilon = 0$ and $\beta = 0.5$. For the rest of this proof, S will denote the outcome of GREEDY-SM$(A, v, \mathbf{c}, B/2)$.

Lemma 5.1.3 (Jalaly and Tardos (2017))**.** $\textsc{opt}(A, B) \leq 3 \cdot v(S) + 2 \cdot v(i^*)$.

For RAND-MECH-SM, if X denotes the outcome of the mechanism, then directly by Lemma 5.1.3 we have $\mathrm{E}(v(X)) \geq \frac{3}{5}v(S) + \frac{2}{5}v(i^*) \geq \frac{1}{5}\textsc{opt}(A,B)$, thus proving the approximation ratio.

For MECH-SM we consider two cases:

If i^* is returned by the mechanism, then $(2+\sqrt{6})\cdot v(i^*) \geq \textsc{opt}(A\setminus\{i^*\},B) \geq \textsc{opt}(A,B) - v(i^*)$, and therefore $\textsc{opt}(A,B) \leq (3+\sqrt{6})\cdot v(i^*)$.

On the other hand, if S is returned, then $(2+\sqrt{6})\cdot v(i^*) < \textsc{opt}(A\setminus\{i^*\},B) \leq \textsc{opt}(A,B)$. Combining this with Lemma 5.1.3 we have $\textsc{opt}(A,B) \leq 3\cdot v(S) + \frac{2}{2+\sqrt{6}}\textsc{opt}(A,B)$ and therefore $\textsc{opt}(A,B) \leq \frac{3(2+\sqrt{6})}{\sqrt{6}}\cdot v(S) = (3+\sqrt{6})\cdot v(S)$. □

5.2 Polynomial-Time Deterministic Mechanisms

In both mechanisms stated in Section 5.1, either an agent i^* of maximum value or GREEDY-SM$(A, v, \mathbf{c}, B/2)$ is returned. In RAND-MECH-SM this is done according to a probability distribution, while in MECH-SM the decision depends on the comparison of i^* with an optimal solution at the instance $A\setminus\{i^*\}$ with budget B. As a result, MECH-SM is not guaranteed to run in polynomial time, since we need to compute $\textsc{opt}(A\setminus\{i^*\},B)$, and more often than not, submodular maximization problems turn out to be NP-hard. An obvious question here is whether we can use an approximate solution instead, but it is not hard to see that by doing so we might sacrifice truthfulness. As a way out, Chen, Gravin, and Lu (2011) mention that instead of $\textsc{opt}(A\setminus\{i^*\},B)$, an optimal solution to a fractional relaxation of the problem can be used. Intuitively, this would maintain truthfulness because no losing agent can force the mechanism to run GREEDY-SM without lowering his bid below his current cost. This is due to the fact that $\textsc{opt}_f$ is nonincreasing with respect to the bid of each agent. Although this trick does not always make the mechanism run in polynomial time, it helps in specific cases.

Suppose that for a specific submodular objective, the budgeted maximization problem can be expressed as an ILP, the corresponding LP relaxation of which can be solved in polynomial time. Further, suppose that for any instance I and any budget B, the optimal fractional solution $\textsc{opt}_f(I,B)$ is within a constant factor of the optimal integral solution $\textsc{opt}(I,B)$. Specifically, suppose that the valuation function is such that $\textsc{opt}_f(I,B) \leq \rho\cdot\textsc{opt}(I,B)$, for any I and any B. Then replacing $\textsc{opt}(A\setminus\{i^*\},B)$ by $\textsc{opt}_f(A\setminus\{i^*\},B)$ in MECH-SM still gives a truthful, constant approximation. In fact, we give a variant of MECH-SM below, where the constants have been appropriately tuned, so as to optimize the achieved approximation ratio.

Theorem 5.2.1. *Let $v(\cdot)$ be a non-decreasing submodular function, $A' = \{i \in A \mid c_i \leq B\}$, and consider a relaxation of our problem for which we have an exact algorithm. Moreover, suppose that $\textsc{opt}_f(A', v, \mathbf{c}_{A'}, B) \leq \rho\cdot\textsc{opt}(A', v, \mathbf{c}_{A'}, B) = \rho\cdot\textsc{opt}(A, v, \mathbf{c}, B)$ for any instance, where $\textsc{opt}_f$ and $\textsc{opt}$ denote the value of an optimal solution to the relaxed and the original problem respectively. Then* MECH-SM-FRAC *is deterministic, truthful,* *individually rational, budget-feasible, and has approximation ratio* $\rho+2+\sqrt{\rho^2+4\rho+1}$. *Also, it runs in polynomial time as long as the exact algorithm for the relaxed problem runs in polynomial time.*

Mech-SM-frac$(A, v, \mathbf{c}, B)$

```
1 Set A' = {i | c_i ≤ B} and i* ∈ argmax_{i∈A'} v(i)
2 if (ρ + 1 + √(ρ² + 4ρ + 1)) · v(i*) ≥ OPT_f(A' \ {i*}, B) then
3     return i*
4 else
5     return Greedy-SM(A, v, c, B/2)
```

Proof. For truthfulness and individual rationality, it suffices to show that the allocation rule is monotone, i.e., a winning agent j remains a winner if he decreases his cost to $c'_j < c_j$. If $j = i^*$ then clearly his bid is irrelevant and he remains a winner. If $j \neq i^*$ and he was a winner, then by reducing the cost to c'_j, the mechanism will still execute Greedy-SM, because $\textsc{opt}_f(A\setminus\{i^*\}, B)$ is higher than before. Hence j remains a winner due to the monotonicity of Greedy-SM (Lemma 5.1.1). This argument also highlights why we cannot in general use an arbitrary approximation algorithm instead of $\textsc{opt}_f(A\setminus\{i^*\}, B)$, since we cannot predict how the solution is affected when the cost changes from c_j to c'_j.

Regarding budget feasibility, under the threshold payment scheme of Lemma 4.3.5, we either have to pay agent i^* the whole budget, or pay the winners of Greedy-SM$(A, B/2)$ the maximum bid that guarantees them to win in Mech-SM-frac. We stress that in the latter case, the payments are upper bounded by the payments induced by running Greedy-SM$(A, B/2)$ alone. This holds because $\textsc{opt}_f(A\setminus\{i^*\}, B)$ is decreasing in the cost of each agent, and so line 2 imposes an extra upper bound on the cost that each agent can report and still be a winner. Hence, budget feasibility follows from the budget feasibility of Greedy-SM$(A, B/2)$ (Lemma 5.1.1).

Finally, for the approximation ratio we let $\alpha = \rho+1+\sqrt{\rho^2+4\rho+1}$ and consider two cases.

If i^* is returned by the mechanism, then

$$\alpha \cdot v(i^*) \geq \textsc{opt}_f(A'\setminus\{i^*\}, B) \geq \textsc{opt}(A'\setminus\{i^*\}, B) = \textsc{opt}(A\setminus\{i^*\}, B) \geq \textsc{opt}(A, B) - v(i^*)\,,$$

and therefore $\textsc{opt}(A, B) \leq (\alpha+1)\cdot v(i^*)$.

On the other hand, if Greedy-SM$(A, B/2)$ is executed, and S is the set of agents returned, then

$$\alpha \cdot v(i^*) < \textsc{opt}_f(A'\setminus\{i^*\}, B) \leq \rho \cdot \textsc{opt}(A'\setminus\{i^*\}, B) \leq \rho \cdot \textsc{opt}(A, B)\,. \tag{5.1}$$

Combining (5.1) with the approximation from Lemma 5.1.1 we have

$$\textsc{opt}(A, B) \leq 3\cdot v(S) + 2\cdot v(i^*) < 3\cdot v(S) + \frac{2\rho}{\alpha}\textsc{opt}(A, B)\,,$$

and therefore $\textsc{opt}(A,B) \le \frac{3\alpha}{\alpha-2\rho} \cdot v(S) = (\alpha+1)\cdot v(S)$, where the last equality is just a matter of calculations. □

Note that when the relaxed problem is the same as the original ($\rho = 1$), then Mech-SM-frac becomes Mech-SM and the two approximation ratios coincide.

5.2.1 Budgeted Max Weighted Coverage

We consider the class of *weighted coverage valuations*, a special class of submodular functions. Their unweighted version was studied by Singer (Singer, 2010; Singer, 2012), motivated by the problem of influence maximization over social networks. Imagine a company that tries to promote a new product and as part of its marketing campaign decides to advertise (or even sell at a promotional price) the product to selected influential nodes. Suppose that each node i, is able to influence some set of other nodes, but this comes at a cost c_i (cost of advertizing and convincing i). Then, if there is a budget available for the campaign, the goal would be to select a set of initial nodes respecting the budget, so as to maximize the (weighted) union of people who are eventually influenced. This gives rise to the following problem.

Budgeted Max Weighted Coverage. Given a set of subsets $\{S_i \mid i \in [m]\}$ of a ground set $[n]$, along with costs $c_1, c_2, \ldots, c_m$, on the subsets, weights $w_1, \ldots, w_n$, on the ground elements, and a positive budget B, find $X \subseteq [m]$ so that $v(X) = \sum_{j \in \bigcup_{i\in X} S_i} w_j$ is maximized subject to $\sum_{i\in X} c_i \le B$.

In the definition above, S_i is the set of people that agent i can influence. On a different note, the problem can also be thought of as a crowdsourcing problem, where each (single-minded) worker i is able to execute only the set of tasks S_i.

Singer (2012) takes an approach similar to what led to Mech-SM-frac, but suggests a different polynomial time mechanism for Budgeted Max Coverage that is deterministic, truthful, budget feasible, and achieves approximation ratio 31.03. Here we generalize and improve this result by showing that it is possible to have a deterministic, truthful, budget feasible, polynomial time 10.03-approximate mechanism for the Budgeted *Weighted* Max Coverage problem.

For all $j \in [n]$ define $T_j = \{i \mid j \in S_i\}$. We begin with a LP formulation of this problem, where without loss of generality we assume that $c_i \le B, \forall i \in [n]$ (otherwise we could just discard any subsets with cost greater than B).

$$\text{maximize:} \quad \sum_{j\in[n]} w_j z_j \tag{5.2}$$

$$\text{subject to:} \quad \sum_{i\in T_j} x_i \ge z_j\,, \qquad \forall j \in [n] \tag{5.3}$$

$$\sum_{i\in[m]} c_i x_i \le B \tag{5.4}$$

$$0 \le x_i, z_j \le 1\,, \qquad \forall i \in [m], \forall j \in [n] \tag{5.5}$$

$$x_i \in \{0,1\}\,, \qquad \forall i \in [m] \tag{5.6}$$

It is not hard to see that (5.2)-(5.6) is a natural ILP formulation for Budgeted Max Weighted Coverage and (5.2)-(5.5) is its linear relaxation. For the rest of this subsection, let $\textsc{opt}(I,B)$ and $\textsc{opt}_f(I,B)$ denote the optimal solutions to (5.2)-(5.6) and (5.2)-(5.5) respectively for instance I and budget B.

To show how these two are related we will use the technique of pipage rounding (Ageev and Sviridenko, 1999; Ageev and Sviridenko, 2004). Although we do not provide a description of the general pipage rounding technique, the proof of Theorem 5.2.2 below is self-contained. We should note here that Ageev and Sviridenko (2004) use the above linear programs (as well as the nonlinear program in the proof of Lemma 5.2.2) to obtain an $\frac{e}{e-1}$-approximation LP-based algorithm that uses pipage rounding on a number of different instances of the problem.[2] However, in their algorithm $\textsc{opt}(I,B)$ is never compared directly to $\textsc{opt}_f(I,B)$, and therefore we cannot get the desired bound from there.

Theorem 5.2.2. *Given the fractional relaxation* (5.2)-(5.5) *for Budgeted Max Weighted Coverage, we have that for any instance I and any budget B*

$$\textsc{opt}_f(I,B) \leq \frac{2e}{e-1} \cdot \textsc{opt}(I,B)\,.$$

Proof. Given any feasible solution x, z to (5.2)-(5.5), the value of (5.2) is upper bounded by $L(x) = \sum_{j\in[n]} w_j \cdot \min\{1, \sum_{i\in T_j} x_i\}$, since $z_j \leq \min\{1, \sum_{i\in T_j} x_i\}$ for any $j \in [n]$. In particular, if x^*, z^* is an optimal (fractional) solution to (5.2)-(5.5), then the value of (5.2) is exactly $L(x^*) = \textsc{opt}_f(I,B)$.

Next we consider the nonlinear program

$$\text{maximize:} \quad F(x) = \sum_{j\in[n]} w_j \Big(1 - \prod_{i\in T_j}(1-x_i)\Big) \tag{5.7}$$

$$\text{subject to:} \quad \sum_{i\in[m]} c_i x_i \leq B \tag{5.8}$$

$$0 \leq x_i \leq 1 \;, \forall\, i \in [m] \tag{5.9}$$

and we observe that $F(x) \geq (1-1/e)L(x)$ for any feasible vector x. This follows from the fact that $(1-(1-1/k)^k) \geq (1-1/e)$ for any $k \geq 1$, and the following inequality, derived in Goemans and Williamson (1994) (Lemma 3.1 in their work):

$$1 - \prod_{i\in[k]}(1-y_i) \geq (1-(1-1/k)^k)\min\Big\{1, \sum_{i\in[k]} y_i\Big\}.$$

So, if x^*, z^* is an optimal solution to (5.2)-(5.5) we have

$$F(x^*) \geq (1-1/e)L(x^*) = (1-1/e)\,\textsc{opt}_f(I,B)\,.$$

However, x^* may have several fractional coordinates. Our next step is to transform x^* to a vector x' that has at most one fractional coordinate and at the same time

[2] In fact, Ageev and Sviridenko (2004) study the hitting set version of this problem, but both problems have essentially the same linear program formulation.

$F(x') \geq F(x^*)$. To this end, we show how to reduce the fractional coordinates by (at least) one in any feasible vector with at least two such coordinates.

Consider a feasible vector x, and suppose x_i and x_j are two non integral coordinates. Let $x_\varepsilon^{i,j}$ be the vector we get if we replace x_i by $x_i + \varepsilon$ and x_j by $x_j - \varepsilon c_i / c_j$ and leave every other coordinate of x the same. Note that the function $\bar{F}(\varepsilon) = F(x_\varepsilon^{i,j})$, with respect to ε, is either linear or a polynomial of degree 2 with positive leading coefficient. That is, $\bar{F}(\varepsilon)$ is convex.

Notice now that $x_\varepsilon^{i,j}$ always satisfies the budget constraint (5.8), and also satisfies (5.9) as long as $\varepsilon \in \left[\max\left\{-x_i, (x_j - 1)c_j / c_i\right\}, \min\left\{1 - x_i, x_j c_j / c_i\right\}\right]$. Due to convexity, $\bar{F}(\varepsilon)$ attains a maximum on one of the endpoints of this interval, say at ε^*. Moreover, at either endpoint at least one of $x_i + \varepsilon^*$ and $x_j - \varepsilon^* c_i / c_j$ is integral. That is, $x_{\varepsilon^*}^{i,j}$ has at least one more integral coordinate than x and $F(x_{\varepsilon^*}^{i,j}) \geq F(x)$.

Hence, initially $x \leftarrow x^*$. As long as there exist two non integral coordinates x_i and x_j we set $x \leftarrow x_{\varepsilon^*}^{i,j}$ as described above. This procedure runs for at most $n-1$ iterations, and outputs a feasible vector x' that is integral or almost integral and $F(x') \geq F(x^*)$.

At this point we should note that when all the x_is are integral then the objectives (5.2) and (5.7) have the same value if we set $z_j = \min\{1, \sum_{i \in T_j} x_i\}$ for all $j \in [n]$. Specifically, if x is any feasible integral vector, we have $F(x) \leq \textsc{opt}(I, B)$. So, if x' is integral then

$$\textsc{opt}_f(I, B) = L(x^*) \leq \frac{e}{e-1} \cdot F(x^*) \leq \frac{e}{e-1} \cdot F(x') \leq \frac{e}{e-1} \cdot \textsc{opt}(I, B),$$

and we are done. Thus, suppose that x' has exactly one fractional coordinate, say x_r. Let x^i be the vector we get if we set x_r to $i \in \{0, 1\}$ and leave every other coordinate of x' the same. Note that $x^1 - x^0$ corresponds to the vector that has 1 in the rth coordinate and 0 everywhere else. Clearly, $\textsc{opt}(I, B) \geq \max\left\{F(x^0), F(x^1 - x^0)\right\}$ (where $x^1 - x^0$ is feasible since we have discarded subsets with $c_i > B$). Moreover, F on integral vectors is submodular, and hence subadditive, therefore $F(x^1) - F(x^0) \leq F(x^1 - x^0)$. Finally, it is easy to see that F is increasing with respect to any single coordinate, so $F(x^1) \geq F(x')$. Combining all the above we get

$$\begin{aligned}\textsc{opt}_f(I, B) = L(x^*) &\leq \frac{e}{e-1} \cdot F(x^*) \leq \frac{e}{e-1} \cdot F(x') \leq \frac{e}{e-1} \cdot F(x^1) \\ &\leq \frac{e}{e-1} \cdot \left(F(x^0) + F(x^1 - x^0)\right) \leq \frac{2e}{e-1} \cdot \textsc{opt}(I, B),\end{aligned}$$

thus completing the proof. □

Combining Theorems 5.2.1 and 5.2.2 we get the following result.

Corollary 5.2.3. *There exists a deterministic, truthful, individually rational, budget feasible* 10.03*-approximate mechanism for Budgeted Max Weighted Coverage that runs in polynomial time.*

Chapter 6

Going Beyond Monotonicity: Symmetric Submodular Objectives[1]

The main focus of this chapter is on symmetric submodular functions, a prominent subclass of non-monotone submodular functions. As mentioned in Queyranne (1998), cut functions form a canonical example of this class. Consequently, we use the budgeted Max Cut problem throughout the chapter as an illustrative example of how our more general approach could be refined for concrete objectives that have a well-behaved LP formulation.

In Section 6.1 we obtain a purely algorithmic result, namely a $\frac{2e}{e-1}$-approximation for symmetric submodular functions under a budget constraint. We believe this is of independent interest, as it is the best known factor achieved by a deterministic algorithm (there exists already a randomized e-approximation) assuming only a value oracle for the objective function.

In Sections 6.2 and 6.3 we propose truthful, budget feasible mechanisms for arbitrary symmetric submodular functions, where previously known results regarded only randomized exponential mechanisms. We manage to significantly improve the known approximation ratios of such objectives by providing both randomized and deterministic mechanisms of exponential time. Moreover, we extend the general scheme of Section 5.2 for producing constant factor approximation mechanisms that run in polynomial time when the objective function is well-behaved. These results provide partial answers to some of the open questions discussed in Dobzinski, Papadimitriou, and Singer (2011).

In the same sections we also pay particular attention on the weighted and unweighted versions of Budgeted Max Cut. For the weighted version we obtain the first deterministic polynomial time mechanism with a 27.25-approximation (where only an exponential randomized algorithm was known with a 768-approximation), while for the unweighted version we improve the approximation ratio for polynomial randomized mechanisms, from 564 down to 10, and for polynomial deterministic mechanisms, from 5158 down to 27.25.

Our contributions in mechanism design are summarized in Table 6.1. We also stress that our mechanisms for general symmetric submodular functions use the

[1] The results of this chapter appear in Amanatidis, Birmpas, and Markakis (2017).

value query model for oracle access to v, which is a much weaker requirement than the demand query model assumed in previous works.

Regarding the technical contribution of our work, the core idea of our approach is to exploit a combination of local search with mechanisms for non-decreasing submodular functions. The reason local search is convenient for symmetric submodular functions is that it produces two local optima, and we can then prove that the function $v(\cdot)$ is non-decreasing within each local optimum. This allows us to utilize mechanisms for non-decreasing submodular functions on the two subsets and then prove that one of the two solutions will attain a good approximation. The running time is still an issue under this approach, since finding an exact local optimum is not guaranteed to terminate fast. However, even by finding *approximate* local optima, within each of them the objective remains almost non-decreasing in a certain sense. This way we are still able to appropriately adjust our mechanisms and obtain provable approximation guarantees. To the best of our knowledge, this is the first time that this *"robustness under small deviations from monotonicity"* approach is used to exploit known results for monotone objectives.

	symmetric submod.		unweighted cut		weighted cut	
	rand.	determ.	rand.	determ.	rand.	determ.
Previous work	$768^{*\,\dagger}$	–	564	5158	$768^{*\,\dagger}$	–
This chapter	10^*	10.90^*, $(1+\rho)\left(2+\rho+\sqrt{\rho^2+4\rho+1}\right)$	10	27.25	27.25	

TABLE 6.1: A summary of our results on mechanisms. The asterisk ($*$) indicates that the corresponding mechanism runs in superpolynomial time. The dagger (†) indicates that previously known results do not directly apply here; see Remark 6.0.1. The factor ρ is an upper bound on the ratio of the optimal fractional solution to the integral one, assuming that we can find the former in polynomial time.
The factor 768 is due to Chen, Gravin, and Lu, 2011, while the factors 564 and 5158 are due to Dobzinski, Papadimitriou, and Singer, 2011.

Remark 6.0.1. In Bei et al. (2012) the proposed mechanisms regard XOS and non-decreasing subadditive objectives, but it is stated that their results can be extended for general subadditive functions as well. This is achieved by defining $\hat{v}(S) = \max_{T\subseteq S} v(T)$. It is easily seen that $\hat{v}$ is non-decreasing, subadditive, and any solution that maximizes v is also an optimal solution for $\hat{v}$. Although this is true for subadditive functions, it does not hold for submodular functions. In particular, if v is submodular, then $\hat{v}$ is not necessarily submodular. Therefore the results of Chen, Gravin, and Lu (2011) cannot be extended to our setting, even when time complexity is not an issue. An example of how $\hat{v}$ may fail to be submodular when v is a cut function is given in Appendix B.2, where we also discuss how it can be derived from Gupta, Nagarajan, and Singla (2017) that in such cases $\hat{v}$ is actually XOS. This also implies that the best previously known approximation factor for the class of symmetric submodular functions was indeed 768 (inherited by the use of $\hat{v}$).

Recall that for this chapter we assume that there is at most one item whose cost exceeds the budget. This is without loss of generality, as shown in Appendix B.1.

6.1 The Core Idea: A Simple Algorithm for Symmetric Submodular Objectives

This section deals with the algorithmic version of the problem: given a symmetric submodular function v, the goal is to find $S \subseteq A$ that maximizes $v(S)$ subject to the constraint $\sum_{i\in S} c_i \le B$. The main result is a deterministic $\frac{2e}{e-1}$-approximation algorithm for symmetric submodular functions. For this section only, the costs and the budget are assumed to be integral.

Since our function is not monotone, we cannot directly apply the result of Sviridenko (2004), which gives an optimal simple greedy algorithm for non-decreasing submodular maximization subject to a knapsack constraint. Instead, our main idea is to combine appropriately the result of Sviridenko (2004) with the local search used for unconstrained symmetric submodular maximization (Feige, Mirrokni, and Vondrák, 2011). At a high level, what happens is that local search produces an approximate solution S for the unconstrained problem, and while this does not look related to our goal at first sight, v is "close to being non-decreasing" on both S and $A\setminus S$. This becomes precise in Lemma 6.1.1 below, but the point is that running a modification of the greedy algorithm of Sviridenko (2004), on both S and $A\setminus S$ will now produce at least one good enough solution.

LS-GREEDY$(A, v, \mathbf{c}, B, \epsilon)$

```
S = APPROX-LOCAL-SEARCH(A, v, ε/4)
T_1 = GREEDY-ENUM-SM(S, v, c_S, B)
T_2 = GREEDY-ENUM-SM(A \ S, v, c_{A\S}, B)
Let T be the best solution among T_1 and T_2
return T
```

The first component of our algorithm is the local search algorithm of Feige, Mirrokni, and Vondrák (2011) (see Section 4.3). By Lemma 4.3.3 and the fact that v is symmetric, both S and $A\setminus S$ are $\left(1+\frac{\epsilon}{4n^2}\right)$-approximate local optima. We can now quantify the crucial observation that v is close to being non-decreasing within the approximate local optima S and $A\setminus S$. Actually, we only need this property on the local optimum that contains the best feasible solution.

Lemma 6.1.1. *Let S be a $\left(1+\frac{\epsilon}{4n^2}\right)$-approximate local optimum and consider*

$$X \in \operatorname*{arg\,max}_{Z\in\{S, A\setminus S\}} \textsc{opt}(Z, B)\,.$$

Then, for every $T \subsetneq X$ and every $i \in X\setminus T$, we have $v(T\cup\{i\}) - v(T) > -\frac{\epsilon}{n}\textsc{opt}(X, B)$.

Before proving Lemma 6.1.1, we begin with a simple fact and a useful lemma. We note that Fact 6.1.2 and Lemma 6.1.3 below require only subadditivity. Submodularity is used later, within the proof of Lemma 6.1.1.

Fact 6.1.2. *For any $S \subseteq A$, $\max\{\textsc{opt}(S,B), \textsc{opt}(A\setminus S, B)\} \ge 0.5\,\textsc{opt}(A,B)$ since $\textsc{opt}(S,B) + \textsc{opt}(A\setminus S,B) \ge \textsc{opt}(A,B)$ by subadditivity.*

Lemma 6.1.3. *For any $S \subseteq A$, $\textsc{opt}(S,\infty) \leq 2n \cdot \textsc{opt}(A,B)$.*

Proof. Recall that $|\{i \in A \mid c_i > B\}| \leq 1$. Let $S^* \subseteq S$ be such that $v(S^*) = \textsc{opt}(S,\infty)$. By subadditivity we have $\textsc{opt}(S,\infty) = v(S^*) \leq \sum_{i \in S^*} v(i)$. Consider three cases.

If $\{i \in A \mid c_i > B\} = \emptyset$, then by the fact that every singleton is a feasible solution we have $\sum_{i \in S^*} v(i) \leq n \cdot \max_{i \in A} v(i) \leq n \cdot \textsc{opt}(A,B)$.

If $\{i \in A \mid c_i > B\} = \{x\} \nsubseteq S^*$, then every singleton in $A \setminus \{x\}$ is a feasible solution, and like before we have $\sum_{i \in S^*} v(i) \leq (n-1) \cdot \max_{i \in A \setminus \{x\}} v(i) \leq (n-1) \cdot \textsc{opt}(A,B)$.

If $\{i \in A \mid c_i > B\} = \{x\} \subseteq S^*$, then we need to bound $v(x)$. Since v is symmetric we have $v(x) = v(A \setminus \{x\}) \leq \sum_{i \in A \setminus \{x\}} v(i) \leq (n-1) \cdot \max_{i \in A \setminus \{x\}} v(i)$. Therefore, by using again that every singleton in $A \setminus \{x\}$ is a feasible solution, we have $\sum_{i \in S^*} v(i) \leq v(x) + (n-1) \cdot \max_{i \in A \setminus \{x\}} v(i) \leq (2n-2) \cdot \max_{i \in A \setminus \{x\}} v(i) \leq 2n \cdot \textsc{opt}(A,B)$. □

Proof of Lemma 6.1.1 : By Fact 6.1.2 we have $\textsc{opt}(X,B) \geq 0.5\,\textsc{opt}(A,B)$. Let $T \subseteq X \setminus \{i\}$ for some $i \in X$. By submodularity we have $v(T \cup \{i\}) - v(T) \geq v(X) - v(X \setminus \{i\})$. Since S is a $\left(1 + \frac{\epsilon}{4n^2}\right)$-approximate local optimum and v is symmetric, X is also a $\left(1 + \frac{\epsilon}{4n^2}\right)$-approximate local optimum. As a result, $v(X \setminus \{i\}) \leq \left(1 + \frac{\epsilon}{4n^2}\right) v(X)$ and thus $v(X) - v(X \setminus \{i\}) \geq -\frac{\epsilon}{4n^2} v(X) \geq -\frac{\epsilon}{4n^2} \textsc{opt}(X,\infty) \geq -\frac{\epsilon}{2n} \textsc{opt}(A,B) \geq -\frac{\epsilon}{n} \textsc{opt}(X,B)$, where the second to last inequality follows from Lemma 6.1.3. □

The second component of LS-Greedy is an appropriate modification of the greedy algorithm of Sviridenko (2004) for non-monotone submodular functions. It first enumerates all solutions of size at most 3. Then, starting from each 3-set, it builds a greedy solution, and it outputs the best among these $\Theta(n^3)$ solutions. Here this idea is adjusted for non-monotone functions.

Greedy-Enum-SM($A, v, \mathbf{c}, B$)

Let S_1 be the best feasible solution of cardinality at most 3 (by enumerating them all)
$S_2 = \emptyset$
for *every* $U \subseteq A$ *with* $|U| = 3$ **do**
  $S^0 = U,\ t = 1,\ A^0 = A \setminus U$
  while $A^{t-1} \neq \emptyset$ **do**
    Find $\theta_t = \max\limits_{i \in A^{t-1}} \dfrac{v(S^{t-1} \cup \{i\}) - v(S^{t-1})}{c_i}$, and let i_t be an element of A^{t-1} that attains θ_t
    if $\theta_t \geq 0$ *and* $\sum_{i \in S^{t-1} \cup \{i_t\}} c_i \leq B$ **then**
      $S^t = S^{t-1} \cup \{i_t\}$
    else
      $S^t = S^{t-1}$
    $A^t = A^{t-1} \setminus \{i_t\}$
    $t = t + 1$
  if $v(S^{t-1}) > v(S_2)$ **then**
    $S_2 = S^{t-1}$
return $S \in \arg\max_{X \in \{S_1, S_2\}} v(X)$

By Fact 6.1.2, at least one of S and $A\setminus S$ contains a feasible solution of value at least $0.5\,\textsc{opt}(A,B)$. Lemma 6.1.1 guarantees that in this set, v is very close to a non-decreasing submodular function. This is sufficient for GREEDY-ENUM-SM to perform almost as well as if v was non-decreasing.

Theorem 6.1.4. *For any $\epsilon > 0$, algorithm* LS-GREEDY *achieves a* $\left(\frac{2e}{e-1}+\epsilon\right)$*-approximation.*

Proof. Recall that GREEDY-ENUM-SM runs on both S and $A\setminus S$ and LS-GREEDY returns the best solution of these two. We may assume, without loss of generality that $\textsc{opt}(S,B) = \max\{\textsc{opt}(S,B), \textsc{opt}(A\setminus S,B)\}$ (the case for $A\setminus S$ being symmetric). By Fact 6.1.2 we have $\textsc{opt}(S,B) \geq 0.5\,\textsc{opt}(A,B)$. So, it suffices to show that running GREEDY-ENUM-SM on S outputs a set of value at least $(1-1/e-\epsilon)\,\textsc{opt}(S,B)$.

In what follows we analyze the approximation ratio achieved by GREEDY-ENUM-SM$(S, v, \mathbf{c}, B)$ with respect to $\textsc{opt}(S,B)$. For this, we follow closely the proof of the main result in Sviridenko (2004).

If there is an optimal solution for our problem restricted on S, of cardinality one, two or three, then the set S_1 of GREEDY-ENUM-SM will be such a solution. Hence, assume that the cardinality of any optimal solution is at least four and let S^* be such a solution. If necessary, reorder the elements of $S^* = \{j_1,\ldots,j_{|S^*|}\}$ so that $j_1 = \arg\max_\ell v(\{j_\ell\})$, and $j_{k+1} = \arg\max_{\ell>k}\left[v(\{j_1,\ldots,j_k,j_\ell\}) - v(\{j_1,\ldots,j_k\})\right]$.

Let $Y = \{j_1, j_2, j_3\}$. For notational convenience, we will use the function $g(\cdot) = v(\cdot) - v(Y)$. It is straightforward that $g(\cdot)$ is submodular. Moreover, the following fact follows from Sviridenko (2004).

Fact 6.1.5. $g(X\cup\{i\}) - g(X) \leq \frac{1}{3}v(Y)$, *for any $Y\subseteq X\subseteq S$ and $i\in S^*\setminus X$.*

Consider the execution of the greedy algorithm with initial set $U = Y$. Let t^*+1 be the first time when an element $i_{t^*+1}\in S^*$ is not added to S^{t^*}. In fact, we assume that t^*+1 is the first time when $S^t = S^{t-1}$. (To see that this is without loss of generality, if there is some time $\tau < t^*+1$ such that i_τ is not added to $S^{\tau-1}$, then—by the definition of t^*+1—it must be the case that $i_\tau\notin S^*$. But then, we may consider the instance $(S\setminus\{i_\tau\}, v, \mathbf{c}_{S\setminus\{i_\tau\}}, B)$ instead. We have $v(S^*) = \textsc{opt}(S\setminus\{i_\tau\},B) = \textsc{opt}(S,B)$ and the greedy solution constructed in the iteration where $S_0 = Y$ is exactly the same as before.) We are going to distinguish two cases.

Case 1. For all $t\in[t^*], \theta_t \geq 0$, but $\theta_{t^*+1} < 0$. Using Theorem 4.3.2 for S^* and S^{t^*} we have

$$\begin{aligned}
g(S^*) &\leq g(S^{t^*}) + \sum_{i\in S^*\setminus S^{t^*}}\left(g(S^{t^*}\cup\{i\}) - g(S^{t^*})\right) - \sum_{i\in S^{t^*}\setminus S^*}\left(g(S^{t^*}\cup S^*) - g(S^{t^*}\cup S^*\setminus\{i\})\right)\\
&= g(S^{t^*}) + \sum_{i\in S^*\setminus S^{t^*}}\left(v(S^{t^*}\cup\{i\}) - v(S^{t^*})\right) - \sum_{i\in S^{t^*}\setminus S^*}\left(v(S^{t^*}\cup S^*) - v(S^{t^*}\cup S^*\setminus\{i\})\right)\\
&\leq g(S^{t^*}) + \sum_{i\in S^*\setminus S^{t^*}} c_i\theta_{t^*+1} - |S^{t^*}\setminus S^*|\left(-\frac{\epsilon}{n}\,\textsc{opt}(S,B)\right)\\
&\leq g(S^{t^*}) + \epsilon\,\textsc{opt}(S,B),
\end{aligned}$$

Here, the second to last inequality holds by Lemma 6.1.1 and by the assumptions we have made. That is, for every $i \in S^* \setminus S^{t^*}$, we have that $i \in A^{t^*}$, since we assumed that t^*+1 is the first time when $S^t = S^{t-1}$, hence up until time t^*, A^{t^*} contains all the agents apart from S^{t^*}. This implies that for every $i \in S^* \setminus S^{t^*}$, we have that $v(S^{t^*} \cup \{i\}) - v(S^{t^*}) \le c_i \theta_{t^*+1}$, by the definition of θ_{t^*+1}.

Therefore, we can conclude that

$$v(S^{t^*}) = v(Y) + g(S^{t^*}) \ge v(Y) + g(S^*) - \epsilon\,\mathrm{OPT}(S,B) = (1-\epsilon)\,\mathrm{OPT}(S,B)\,.$$

Case 2. For all $t \in [t^*+1], \theta_t \ge 0$, but $\sum_{i \in S^{t^*} \cup \{i_{t^*+1}\}} c_i > B$ while $\sum_{i \in S^{t^*}} c_i \le B$. Using Theorem 4.3.2 for S^* and each of S^t, $t \in [t^*]$, as well as Lemma 6.1.1, we have

$$\begin{aligned} g(S^*) &\le g(S^t) + \sum_{i \in S^* \setminus S^t} (g(S^t \cup \{i\}) - g(S^t)) - \sum_{i \in S^t \setminus S^*} (g(S^t \cup S^*) - g(S^t \cup S^* \setminus \{i\})) \\ &= g(S^t) + \sum_{i \in S^* \setminus S^t} (v(S^t \cup \{i\}) - v(S^t)) - \sum_{i \in S^t \setminus S^*} (v(S^t \cup S^*) - v(S^t \cup S^* \setminus \{i\})) \\ &\le g(S^t) + \sum_{i \in S^* \setminus S^t} (v(S^t \cup \{i\}) - v(S^t)) - |S^t \setminus S^*| \left(-\frac{\epsilon}{n}\,\mathrm{OPT}(S,B)\right) \\ &\le g(S^t) + \sum_{i \in S^* \setminus S^t} (v(S^t \cup \{i\}) - v(S^t)) + \epsilon\,\mathrm{OPT}(S,B)\,, \end{aligned}$$

and therefore

$$\begin{aligned} g(S^*) - \epsilon\,\mathrm{OPT}(S,B) &\le g(S^t) + \sum_{i \in S^* \setminus S^t} (v(S^t \cup \{i\}) - v(S^t)) \\ &\le g(S^t) + \sum_{i \in S^* \setminus S^t} c_i \theta_{t+1} \\ &\le g(S^t) + \Big(B - \sum_{i \in Y} c_i\Big)\theta_{t+1}\,, \end{aligned}$$

for all $t \in [t^*]$.

For the last part of the proof we need the following inequality of Wolsey (1982).

Theorem 6.1.6 (Wolsey (1982)). *Let k and s be arbitrary positive integers, and $\rho_1, \ldots, \rho_k$ be arbitrary reals with $z_1 = \sum_{i=1}^k \rho_i$ and $z_2 = \min_{t \in [k]} \left(\sum_{i=1}^{t-1} \rho_i + s\rho_t\right) > 0$. Then $z_1/z_2 \ge 1 - (1-1/s)^k \ge 1 - e^{-\frac{k}{s}}$.*

For any τ, define $B_\tau = \sum_{t=1}^{\tau} c_{i_t}$, and let $k = B_{t^*+1}$ and $s = B - \sum_{i \in Y} c_i$. We also define $\rho_1, \ldots, \rho_k$ as follows; for $i \le c_1, \rho_i = \theta_1$ and for $B_\tau < i \le B_{\tau+1}, \rho_i = \theta_{\tau+1}$. It is easy to see that $g(S^{t^*} \cup \{i_{t^*+1}\}) = \sum_{t=1}^{t^*+1} c_{i_t}\theta_{i_t} = \sum_{i=1}^{k} \rho_i$ and similarly $g(S^t) = \sum_{t=1}^{\tau} c_{i_t}\theta_{i_t} = \sum_{i=1}^{B_\tau} \rho_i$.

As noted in Sviridenko (2004), since the ρ_is are nonnegative we have

$$\min_{t \in [k]} \left(\sum_{i=1}^{t-1} \rho_i + s\rho_t\right) = \min_{\tau \in [t^*]} \left(\sum_{i=1}^{B_\tau} \rho_i + s\rho_{B_\tau+1}\right),$$

and therefore

$$g(S^*) - \epsilon\,\mathrm{OPT}(S,B) \le \min_{t \in [k]} \left(\sum_{i=1}^{t-1} \rho_i + s\rho_t\right).$$

So, as a direct application of Theorem 6.1.6 we have

$$\frac{g(S^{t^*} \cup \{i_{t^*+1}\})}{g(S^*) - \epsilon\,\textsc{opt}(S,B)} \geq 1 - e^{-\frac{k}{s}} > 1 - e^{-1}\,.$$

Finally, using the above inequality and Fact 6.1.5, we get

$$\begin{aligned}
v(S^{t^*}) &= v(Y) + g(S^{t^*}) = v(Y) + g(S^{t^*} \cup \{i_{t^*+1}\}) - (g(S^{t^*} \cup \{i_{t^*+1}\}) - g(S^{t^*}))\\
&\geq v(Y) + (1-e^{-1})g(S^*) - (1-e^{-1})\epsilon\,\textsc{opt}(S,B) - (v(S^{t^*} \cup \{i_{t^*+1}\}) - v(S^{t^*}))\\
&\geq v(Y) + (1-e^{-1})g(S^*) - \epsilon\,\textsc{opt}(S,B) - \frac{1}{3}v(Y)\\
&\geq (1-e^{-1}-\epsilon)\,\textsc{opt}(S,B)\,.
\end{aligned}$$

Since in both cases the final output T^* of the algorithm has value at least $v(S^{t^*})$, this implies that

$$v(T^*) \geq (1-e^{-1}-\epsilon)\,\textsc{opt}(S,B) \geq \frac{1-e^{-1}-\epsilon}{2}\,\textsc{opt}(A,B) > \left(\frac{e-1}{2e} - \epsilon\right)\textsc{opt}(A,B)\,,$$

thus concluding the analysis of the performance of the algorithm. □

Theorem 6.1.4 suggests that a straightforward composition of two well known greedy algorithms achieves a good approximation for any symmetric submodular objective. We believe this is of independent interest and could be useful for other problems involving submodular optimization. From a mechanism design perspective, however, algorithm LS-Greedy fails to be monotone and thus it cannot be used directly in the subsequent sections. In the next two sections, we remedy the problem by removing the enumeration part of the algorithm.

6.2 A First Take on Mechanism Design

Utilizing the algorithmic approach of Section 6.1 to get truthful mechanisms is not straightforward. One of the reasons is that LS-Greedy is not monotone. We note that the algorithm Greedy-Enum-SM without the enumeration part *is* monotone even for general objectives, but, to further complicate things, it is not guaranteed to be budget-feasible or have a good performance anymore. Instead of computing approximate local optima like in Section 6.1, in this section we bypass most issues by computing exact local optima. The highlights of this simplified approach are polynomial mechanisms for unweighted cut functions with greatly improved guarantees.

The price we have to pay, however, is that in general, finding exact local optima is not guaranteed to run in polynomial time (Schäffer and Yannakakis, 1991). Still, these mechanisms deepen our understanding of the problem. As mentioned in the Introduction, the problem seems to remain hard even when the running time is not an issue, and many existing mechanisms for various classes of functions are not polynomial. In particular, there are no better known mechanisms—even running in exponential time—for symmetric submodular objectives. We are going to deal further with the issue of running time in Section 6.3.

Below we give a randomized mechanism that reduces the known factor of 768 down to 10, as well as the first deterministic $O(1)$-approximation mechanism for symmetric submodular objectives. In both mechanisms, local search produces a local maximum S for the unbudgeted problem and then the budgeted problem is solved optimally on both S and $A\setminus S$. As shown in Lemma 6.2.1, v is non-decreasing on both S and $A\setminus S$. Thus, running one of the mechanisms RAND-MECH-SM or MECH-SM, as described in the beginning of the chapter, on $T \in \arg\max_{X\in\{S,A\setminus S\}} \textsc{opt}(X,B)$, directly implies a good solution. Since the resulting randomized and deterministic mechanisms are very similar, we state them together for succinctness.

RAND-MECH-SYMSM$(A, v, \mathbf{c}, B)$ (*resp.* DET-MECH-SYMSM$(A, v, \mathbf{c}, B)$)

$S =$ APPROX-LOCAL-SEARCH$(A, v, 0)$ //find an exact local optimum
if $\textsc{opt}(S,B) \geq \textsc{opt}(A\setminus S,B)$ **then**
  return RAND-MECH-SM$(S, v, \mathbf{c}_S, B)$ (*resp.* MECH-SM$(S, v, \mathbf{c}_S, B)$)
else
  return RAND-MECH-SM$(A\setminus S, v, \mathbf{c}_{A\setminus S}, B)$ (*resp.* MECH-SM$(A\setminus S, v, \mathbf{c}_{A\setminus S}, B)$)

The next simple lemma is crucial for the performance of both mechanisms for arbitrary submodular functions, and it shows how local search helps us exploit known results for non-decreasing submodular functions.

Lemma 6.2.1. *Let A be a set and v be an arbitrary submodular function defined on 2^A. If S is a local maximum of v, then v is submodular and non-decreasing when restricted on 2^S.*

Proof. The fact that v is submodular when restricted on 2^S is trivial. Suppose now that the statement is not true and that v is not non-decreasing on 2^S. That is, there exist $T, T' \subseteq S$ such that $T \subsetneq T'$ and $v(T) > v(T')$. Let $T'\setminus T = \{i_1, \dots, i_r\}$.

By Theorem 4.3.2 we have

$$v(T) \leq v(T') - \sum_{j\in[r]} \left(v(T') - v(T'\setminus\{i_j\})\right),$$

and therefore

$$\sum_{j\in[r]} \left(v(T') - v(T'\setminus\{i_j\})\right) \leq v(T') - v(T) < 0.$$

We conclude that there is some $\ell \in [r]$ such that $v(T') - v(T'\setminus\{i_\ell\}) < 0$. Then, by submodularity and the fact that $T'\setminus\{i_\ell\} \subseteq S\setminus\{i_\ell\}$, we get

$$v(S) - v(S\setminus\{i_\ell\}) \leq v(T') - v(T'\setminus\{i_\ell\}) < 0.$$

But then $v(S\setminus\{i_\ell\}) > v(S)$, which contradicts the fact that S is a local maximum of v. So, it must be the case that v is non-decreasing on the subsets of S. □

Since v is symmetric, if S is a local optimum, so is $A \setminus S$. Lemma 6.2.1 suggests that we can use the mechanism RAND-MECH-SM (resp. MECH-SM) on S and $A \setminus S$, to get the following implications.

Theorem 6.2.2. *i. The mechanism* RAND-MECH-SYMSM *is universally truthful, individually rational, budget-feasible, and has approximation ratio* 10.
ii. The mechanism DET-MECH-SYMSM *is deterministic, truthful, individually rational, budget-feasible, and has approximation ratio* $6+2\sqrt{6}$.

Proof. The fact that both mechanism RAND-MECH-SYMSM and DET-MECH-SYMSM are budget-feasible follows from the budget-feasibility of RAND-MECH-SM and MECH-SM respectively, established in Chen, Gravin, and Lu (2011).

For truthfulness and individual rationality, it suffices to show that the allocation rule is monotone. Let us look first at DET-MECH-SYMSM. Consider an agent i with true cost c_i in an instance I where i is included in the winning set. Note first that the local search step is not affected by the costs, hence no player can influence the local optimum. Suppose now that it was the case that $\textsc{opt}(S,B) \geq \textsc{opt}(A \setminus S, B)$, hence $i \in S$. If player i now declares a lower cost, then the optimal solution within S can only get better, hence the mechanism will run MECH-SM on S as before. Since MECH-SM is monotone, player i will again be selected in the solution. We conclude that the outcome rule is monotone and DET-MECH-SYMSM is truthful.

To prove that the randomized mechanism RAND-MECH-SYMSM is universally truthful we use similar arguments. We fix the random bits of the mechanism and we consider a winning agent i like before. Again, no player can influence the outcome of local search. Suppose it is the case that $\textsc{opt}(S,B) \geq \textsc{opt}(A \setminus S, B)$, hence $i \in S$. If i declares a lower cost, then the optimal solution within S improves, and RAND-MECH-SM will still run on S. Since RAND-MECH-SM is universally truthful, it is monotone given the random bits, and player i will again be a winner. We conclude that RAND-MECH-SYMSM is universally truthful.

To argue now about the approximation ratio, suppose that we are in the case that $\textsc{opt}(S,B) \geq \textsc{opt}(A \setminus S, B)$ (the other case being symmetric). We know by Fact 6.1.2 that $\textsc{opt}(S,B) \geq 0.5 \cdot \textsc{opt}(A,B)$. Hence, since we run either MECH-SM or RAND-MECH-SM on S, we will get twice their approximation ratio. This implies a ratio of 10 for RAND-MECH-SYMSM and a ratio of $6+2\sqrt{6}$ for DET-MECH-SYMSM. □

Lower bounds on the approximability have been obtained by Chen, Gravin, and Lu (2011) for additive valuations. Since additive functions are not symmetric, these lower bounds do not directly apply here. However, it is not hard to construct symmetric submodular functions that give the exact same bounds.

Lemma 6.2.3. *Independent of complexity assumptions, there is no deterministic (resp. randomized) truthful, budget feasible mechanism that can achieve an approximation ratio better than* $1+\sqrt{2}$ *(resp.* 2*).*

Proof. The lower bounds of Chen, Gravin, and Lu (2011) are both for additive objectives (Knapsack). It suffices to show that given an instance $(A, v, \mathbf{c}, B)$ of Knapsack, we can construct an equivalent $(A', v', \mathbf{c}', B')$ instance of Budgeted Max Weighted Cut. This is straightforward. Consider a graph G with vertex set $A \cup x$ and edge set $\{(i, x) \mid i \in A\}$. For $i \in A$, vertex i has cost $c'_i = c_i$, while vertex x has cost $c'_x = B + 1$. Edge (i, x) has weight $v(i)$. Finally, for $S \subseteq A \cup x$, $v'(S)$ is equal to the weight of the cut defined by S.

This correspondence between items and vertices creates a natural correspondence between solutions. It is clear that each feasible solution of the Knapsack instance essentially defines a feasible solution to the Budgeted Max Weighted Cut instance of the same value and vice versa. In particular $\textsc{opt}(A, v, \mathbf{c}, B) = \textsc{opt}(A', v', \mathbf{c}', B)$. We conclude that any lower bound for Knapsack gives a lower bound for Budgeted Max Weighted Cut. □

Clearly, both mechanisms presented in this section require superpolynomial time in general (due to their first two lines), unless $P = NP$. In both cases, instead of $\textsc{opt}(\cdot, B)$ we could use the optimal solution of a fractional relaxation of the problem, at the expense of somewhat worse guarantees. This does not completely resolve the problem, although this way local search becomes the sole bottleneck. For certain objectives, however, we can achieve similar guarantees in polynomial time. Unweighted cut functions are the most prominent such example, and it is the focus of the next subsection.

6.2.1 Unweighted Cut Functions

We begin with the definition of the problem when v is a cut function:

Budgeted Max Weighted Cut. Given a complete graph G with vertex set $V(G) = [n]$, non-negative weights w_{ij} on the edges, non-negative costs c_i on the nodes, and a positive budget B, find $X \subseteq [n]$ so that $v(X) = \sum_{i \in X} \sum_{j \in [n] \setminus X} w_{ij}$ is maximized subject to $\sum_{j \in X} c_j \le B$.

For convenience, we assume the problem is defined on a complete graph as we can use zero weights to model any graph. In this subsection, we focus on the unweighted version, where all weights are equal to either 0 or 1. We call this special case *Budgeted Max Cut*. The weighted version is considered in Subsection 6.3.1.

The fact that local search takes polynomial time to find an exact local optimum for the unweighted version (Kleinberg and Tardos, 2006) does not suffice to make RAND-MECH-SYMSM a polynomial time mechanism, since one still needs to compute $\textsc{opt}(S, B)$ and $\textsc{opt}(A \setminus S, B)$. However, a small modification so that RAND-MECH-SM(S, B) and RAND-MECH-SM$(A \setminus S, B)$ are returned with probability $1/2$ each yields a randomized 10-approximate polynomial time mechanism.

RAND-MECH-UCUT$(A, v, \mathbf{c}, B)$

1 $S = $ APPROX-LOCAL-SEARCH$(A, v, 0)$ //find an exact local optimum
2 with probability 1/2 **return** RAND-MECH-SM$(S, v, \mathbf{c}_S, B)$
3 with probability 1/2 **return** RAND-MECH-SM$(A \setminus S, v, \mathbf{c}_{A\setminus S}, B)$

Theorem 6.2.4. *RAND-MECH-UCUT is a randomized, universally truthful, individually rational, budget-feasible mechanism for Budgeted Max Cut that has approximation ratio* 10 *and runs in polynomial time.*

Proof. Clearly, for the unweighted version of Max Cut, the mechanism runs in polynomial time. The fact that the mechanism is truthful, individually rational and budget feasible follows from the same arguments as in the proof of Theorem 6.2.2.

Finally, let

$$X_S = \text{GREEDY-SM}(S, v, \mathbf{c}_S, B/2), \quad X_{A\setminus S} = \text{GREEDY-SM}(A\setminus S, v, \mathbf{c}_{A\setminus S}, B/2),$$
$$i_S \in \operatorname*{arg\,max}_{\{i\in S \mid c_i \le B\}} v(i), \quad \text{and} \quad i_{A\setminus S} \in \operatorname*{arg\,max}_{\{i\in A\setminus S \mid c_i \le B\}} v(i).$$

Since we run mechanism RAND-MECH-SM with probability 1/2 on S and $A\setminus S$, it is not hard to see that each of X_S and $X_{A\setminus S}$ is returned with probability 3/10, while each of i_S and $i_{A\setminus S}$ is returned with probability 2/10. If X denotes the outcome of the mechanism, then using subadditivity and Lemma 5.1.3, we get

$$\begin{aligned} \mathrm{E}(v(X)) &= \frac{3}{10}v(X_S) + \frac{2}{10}v(i_S) + \frac{3}{10}v(X_{A\setminus S}) + \frac{2}{10}v(i_{A\setminus S}) \\ &\ge \frac{1}{10}\textsc{opt}(S,B) + \frac{1}{10}\textsc{opt}(A\setminus S, B) \\ &\ge \frac{1}{10}\textsc{opt}(A,B), \end{aligned}$$

thus establishing an approximation ratio of 10. □

To obtain a deterministic mechanism for Budgeted Max Cut, we will use an LP-based approach, similar to the one used in Subsection 5.2.1, and we will run MECH-SM-FRAC on an appropriate local maximum. Recall that for this we will first need to compare the value of an optimal solution of a fractional relaxation to the value of an optimal solution of the original problem. Ageev and Sviridenko (1999) studied a different Max Cut variant, but we follow a similar approach to obtain the desired bound for our problem as well. We begin with a linear program formulation of the problem. Our analysis is carried out for the weighted version of the problem, as we are going to reuse some results in Subsection 6.3.1, which deals with weighted cut functions. To be more precise, we want to argue about a variant of the problem where we may only be allowed to choose the solution from a specified subset of A. That is, we formulate below the sub-instance $I = (X, v, \mathbf{c}_X, B)$ of $(A, v, \mathbf{c}, B)$, where in fact $X \subseteq A' = \{i \in A \mid c_i \le B\}$.

We associate a binary variable x_i for each vertex i, and the partition of $A = [n]$ according to the value of the x_is defines the cut. There is also a binary variable z_{ij}

for each edge $\{i,j\}$ which is the indicator variable of whether $\{i,j\}$ is in the cut.

$$\begin{aligned}
\text{maximize:}\quad & \sum_{i\in[n]}\sum_{j\in[n]\setminus[i]} w_{ij}z_{ij} && && (6.1)\\
\text{subject to:}\quad & z_{ij}\le x_i+x_j\ , && \forall i\in[n],\forall j\in[n]\setminus[i] && (6.2)\\
& z_{ij}\le 2-x_i-x_j\ , && \forall i\in[n],\forall j\in[n]\setminus[i] && (6.3)\\
& \sum_{i\in[n]} c_i x_i \le B && && (6.4)\\
& x_i = 0\ , && \forall i\in[n]\setminus X && (6.5)\\
& 0\le x_i, z_{ij}\le 1\ , && \forall i\in[n],\ \forall j\in[n]\setminus[i] && (6.6)\\
& x_i\in\{0,1\}\ , && \forall i\in[n] && (6.7)
\end{aligned}$$

It is not hard to see that (6.1)-(6.7) is a natural ILP formulation for Budgeted Max Weighted Cut and (6.1)-(6.6) is its linear relaxation. Let $\textsc{opt}(I)$ and $\textsc{opt}_f(I)$ denote the optimal solutions to (6.1)-(6.7) and (6.1)-(6.6) respectively for instance I. To show how these two are related we will again use the technique of pipage rounding (Ageev and Sviridenko, 1999; Ageev and Sviridenko, 2004). The proof of the next theorem takes the same approach as the proof of Theorem 5.2.2.

Theorem 6.2.5. *Given the fractional relaxation* (6.1)-(6.6) *for Budgeted Max Weighted Cut, we have that* $\textsc{opt}_f(I)\le(2+2\beta_I)\cdot\textsc{opt}(I)$, *for any instance* I, *where* β_I *is such that* $\max_{i\in A'} v(i)\le\beta_I\cdot\textsc{opt}(I)$.

Proof. We begin with a nonlinear program such that if all the x_is are integral then the objectives (6.1) and (6.8) have the same value.

$$\begin{aligned}
\text{maximize:}\quad & F(x)=\sum_{i\in[n]}\sum_{j\in[n]\setminus[i]} w_{ij}(x_i+x_j-2x_ix_j) && (6.8)\\
\text{subject to:}\quad & \sum_{i\in[n]} c_i x_i\le B && (6.9)\\
& x_i=0\ ,\quad \forall i\in[n]\setminus X && (6.10)\\
& 0\le x_i\le 1\ ,\quad \forall i\in[n] && (6.11)
\end{aligned}$$

So, if x is any feasible integral vector to our problem, we have $F(x)\le\textsc{opt}(I)$. Moreover, given any feasible solution x,z to (6.1)-(6.6), the value of (6.1) is upper bounded by $L(x)=\sum_{i\in[n]}\sum_{j\in[n]\setminus[i]} w_{ij}\min\{x_i+x_j,2-x_i-x_j\}$, since $z_{ij}\le\min\{x_i+x_j,2-x_i-x_j\}$ for any $i\in[n], j\in[n]\setminus[i]$.

Next, we show that $F(x)\ge 0.5L(x)$. This follows from the inequality

$$2(a+b-2ab)\ge\min\{a+b,2-a-b\}\ \text{for any}\ a,b\in[0,1],$$

proven by Ageev and Sviridenko (1999). For completeness we prove it here as well. Notice that by replacing a and b by $1-a$ and $1-b$ respectively both sides of the inequality remain exactly the same. Therefore, it suffices to prove $2(a+b-2ab)\ge a+b$ for any $a,b\in[0,1]$ such that $a+b\le 1$ (since otherwise, $1-a+1-b\le 1$). This however is equivalent to $a+b\ge 4ab$ which is true since $a+b\ge(a+b)^2\ge 4ab$.

Hence, if x^*, z^* is an optimal fractional solution to (6.1)-(6.6), then the value of (6.1) is $\textsc{opt}_f(I) = L(x^*)$ and thus $F(x^*) \geq 0.5L(x^*) = 0.5\,\textsc{opt}_f(I)$. However, x^* may have several fractional coordinates. Our next step is to transform x^* to a vector x' that has at most one fractional coordinate and at the same time $F(x') \geq F(x^*)$. To this end, we show how to reduce the fractional coordinates by (at least) one in any feasible vector with at least two such coordinates.

Consider a feasible vector x, and suppose x_i and x_j are two non integral coordinates. Note that $i, j \in X$. Let $x_\epsilon^{i,j}$ be the vector we get if we replace x_i by $x_i + \epsilon$ and x_j by $x_j - \epsilon c_i / c_j$ and leave every other coordinate of x the same. Note that the function $\bar{F}(\epsilon) = F(x_\epsilon^{i,j})$, with respect to ϵ, is either linear or a polynomial of degree 2 with positive leading coefficient. That is, $\bar{F}(\epsilon)$ is convex.

Notice now that $x_\epsilon^{i,j}$ always satisfies the budget constraint (6.9), and also satisfies (6.11) as long as $\epsilon \in \left[\max\{-x_i, (x_j - 1)c_j / c_i\}, \min\{1 - x_i, x_j c_j / c_i\}\right]$. Due to convexity, $\bar{F}(\epsilon)$ attains a maximum on one of the endpoints of this interval, say at ϵ^*. Moreover, at either endpoint at least one of $x_i + \epsilon^*$ and $x_j - \epsilon^* c_i / c_j$ is integral. That is, $x_{\epsilon^*}^{i,j}$ has at least one more integral coordinate than x and $F(x_{\epsilon^*}^{i,j}) \geq F(x)$.

So, initially $x \leftarrow x^*$. As long as there exist two non integral coordinates x_i and x_j we set $x \leftarrow x_{\epsilon^*}^{i,j}$ as described above. This happens at most $n-1$ times, and outputs a feasible vector x' with at most one non-integral coordinate, and with $F(x') \geq F(x^*)$. We have then the following implications:

$$\textsc{opt}_f(I) = L(x^*) \leq 2 \cdot F(x^*) \leq 2 \cdot F(x')\,. \tag{6.12}$$

If x' is integral, then by (6.12) we have $\textsc{opt}_f(I) \leq 2 \cdot F(x') \leq 2 \cdot \textsc{opt}(I)$, and we are done. So, suppose that x'_r is the only fractional coordinate of x'. Let x^0 and x^1 be the vectors we get if we set x'_r to 0 or 1 respectively and leave every other coordinate of x' the same. Notice that for $a, b \in [0,1]$ the inequality $(1-a)(b-1) \leq 0$ implies $a + b - ab \leq 1$ and therefore $a + b - 2ab \leq 1$, so we have

$$F(x') - F(x^0) = \sum_{j \in [n] \setminus r} w_{rj}(x'_r - 2x'_r x'_j) \leq \sum_{j \in [n] \setminus r} w_{rj}(x'_r + x'_j - 2x'_r x'_j) \leq \sum_{j \in [n] \setminus r} w_{rj} = F(x^1 - x^0)\,,$$

and thus

$$F(x') \leq F(x^0) + F(x^1 - x^0)\,. \tag{6.13}$$

Combining (6.12)and (6.13), we have $\textsc{opt}_f(I) \leq 2 \cdot F(x') \leq 2 \cdot \left(F(x^0) + F(x^1 - x^0)\right)$.

Using the fact that F is upper bounded by $\textsc{opt}(I)$ on integral vectors, we have that $F(x^0) \leq \textsc{opt}(I)$. Observe now also that $F(x^1 - x^0) = \sum_{j \in [n] \setminus r} w_{rj} = v(r) \leq \beta_I\, \textsc{opt}(I)$, by the definition of β_I. Hence, overall we get

$$\textsc{opt}_f(I) \leq (2 + 2\beta_I)\,\textsc{opt}(I)\,,$$

thus completing the proof. □

Note that there always exists some $\beta_I \leq 1$ for every instance I, hence, the above theorem implies a worst-case upper bound of 4. Now, we may modify DET-MECH-SYMSM to use OPT_f instead of OPT, and MECH-SM-FRAC instead of MECH-SM. This results in the following deterministic mechanism that runs in polynomial time.

DET-MECH-UCUT$(A, v, \mathbf{c}, B)$

```
1 Set A' = {i | c_i ≤ B} and i* ∈ argmax_{i∈A'} v(i)
2 if 26.25 · v(i*) ≥ OPT_f(A' \ {i*}, B) then
3     return i*
4 else
5     S = APPROX-LOCAL-SEARCH(A, v, 0)
6     if OPT_f(S ∩ A', B) ≥ OPT_f(A' \ S, B) then
7         return MECH-SM-FRAC(S, v, c_S, B)
8     else
9         return MECH-SM-FRAC(A \ S, v, c_{A\S}, B)
```

Theorem 6.2.6. DET-MECH-UCUT *is a deterministic, truthful, individually rational, budget-feasible mechanism for Budgeted Max Cut that has approximation ratio* 27.25 *and runs in polynomial time.*

Proof. Clearly the mechanism runs in polynomial time.

For truthfulness and individual rationality, it suffices to show that the allocation rule is monotone, i.e., a winning agent j remains a winner if he decreases his cost to be $c'_j < c_j$. If $j = i^*$ and he wins in line 2, then his bid is irrelevant and he remains a winner. If $j \neq i^*$, or $j = i^*$ but he wins after line 3, we may assume he wins at line 7 (the case of line 9 is symmetric). When bidding $c'_j < c_j$, the decision of the mechanism in line 2 does not change (if $j \neq i^*$ then OPT_f is improved, if $j = i^*$ nothing changes). Further, since he cannot influence the local search, j is still in S and MECH-SM-FRAC is executed. By the monotonicity of MECH-SM-FRAC we have that j is still a winner. Therefore, the mechanism is monotone.

If the winner is i^* in line 3, then his payment is B. Otherwise, budget-feasibility follows from the budget-feasibility of MECH-SM-FRAC and the observation that the comparison in line 2 only gives additional upper bounds on the payments of winners from MECH-SM-FRAC.

It remains to prove the approximation ratio. We consider two cases. Let $\alpha = 26.25$. If i^* is returned in line 3, then

$$\alpha \cdot v(i^*) \geq \text{OPT}_f(A' \setminus \{i^*\}, B) \geq \text{OPT}(A' \setminus \{i^*\}, B) = \text{OPT}(A \setminus \{i^*\}, B) \geq \text{OPT}(A, B) - v(i^*),$$

and therefore $\text{OPT}(A, B) \leq (\alpha + 1) \cdot v(i^*) = 27.25 \cdot v(i^*)$.

On the other hand, if $X = \text{MECH-SM-FRAC}(S, B)$ is returned in line 7, then by Theorem 6.2.5 with factor 4 we have

$$\alpha \cdot v(i^*) < \text{OPT}_f(A' \setminus \{i^*\}, B) \leq 4 \cdot \text{OPT}(A' \setminus \{i^*\}, B) = 4 \cdot \text{OPT}(A \setminus \{i^*\}, B) \leq 4 \cdot \text{OPT}(A, B).$$

Therefore, $v(i^*) < \frac{4}{\alpha}\textsc{opt}(A,B)$ and for the remaining steps of the mechanism we can use Theorem 6.2.5 with factor $2+8/\alpha$.

At line 6 it must be the case that $\textsc{opt}_f(S\cap A',B) \geq \textsc{opt}_f(A'\setminus S,B)$. Thus,

$$\left(2+\frac{8}{\alpha}\right)\textsc{opt}(S,B) = \left(2+\frac{8}{\alpha}\right)\textsc{opt}(S\cap A',B) \geq \textsc{opt}_f(S\cap A',B) \geq \textsc{opt}_f(A'\setminus S,B)$$
$$\geq \textsc{opt}(A'\setminus S,B) = \textsc{opt}(A\setminus S,B) \geq \textsc{opt}(A,B) - \textsc{opt}(S,B)\,.$$

Therefore $\textsc{opt}(S,B) \geq \frac{\alpha}{3\alpha+8}\textsc{opt}(A,B)$. By Theorem 5.2.1 we have

$$\left(4+8/\alpha+\sqrt{(2+8/\alpha)^2+32/\alpha+9}\right)v(X) \geq \textsc{opt}(S,B) \geq \frac{\alpha}{3\alpha+8}\textsc{opt}(A,B)\,,$$

and by substituting α and doing the calculations we get $\textsc{opt}(A,B) \leq 27.25\cdot v(X)$.

The case where $X = \textsc{Mech-SM-frac}(A\setminus S,B)$ is returned in line 9 is symmetric to the case above and it need not be considered separately. We conclude that $\textsc{opt}(A,B) \leq 27.25\cdot v(X)$. □

6.3 Symmetric Submodular Objectives Revisited

Suppose that for a symmetric submodular function v, an optimal fractional solution can be found efficiently and that $\textsc{opt}_f(A',B) \leq \rho\cdot\textsc{opt}(A,B)$ for any instance, where $\textsc{opt}_f$ and $\textsc{opt}$ denote the value of an optimal solution to the relaxed and the original problem respectively, and $A' = \{i\in A \mid c_i \leq B\}$.

A natural question is whether the approach taken for unweighted cut functions can be fruitful for other symmetric submodular objectives. In the mechanisms of Subsection 6.2.1, however, the complexity of local search can be a bottleneck even for objectives where an optimal fractional solution can be found fast and it is not far from the optimal integral solution. So, we now return to the idea of Section 6.1, where local search runs in polynomial time and produces an approximate local maximum; unfortunately, the nice property of monotonicity in each side of the partition (Lemma 6.2.1) does not hold any longer.

This means that the approximation guarantees of such mechanisms do not follow in any direct way from existing work. Moreover, budget-feasibility turns out to be an even more delicate issue since it crucially depends on the (approximate) monotonicity of the valuation function. Specifically, when a set X only contains a very poor solution to the original problem, every existing proof of budget feasibility for the restriction of v on X completely breaks down. Since we cannot expect that an approximate local maximum S and its complement $A\setminus S$ both contain a "good enough" solution to the original problem, we need to make sure that Greedy-SM never runs on the wrong set.

The mechanism Det-Mech-UCut for the unweighted cut problem seems to take care of this and we are going to build on it, in order to propose mechanisms for arbitrary symmetric submodular functions. To do so we replace the constant 26.25 that appears in Det-Mech-UCut by $\alpha = (1+\rho)\left(2+\rho+\sqrt{\rho^2+4\rho+1}\right)-1$ and we find an approximate local maximum instead of an exact local maximum. Most importantly, in order to achieve budget-feasibility we introduce a modification of MECH-SM-FRAC, called MECH-SM-FRAC-VAR, that runs GREEDY-SM with a slightly reduced budget.

MECH-SM-FRAC-VAR$(A, v, \mathbf{c}, B, \gamma)$

Set $A' = \{i \mid c_i \le B\}$ and $i^* \in \arg\max_{i \in A'} v(i)$
if $\left(\rho + 1 + \sqrt{\rho^2 + 4\rho + 1}\right) \cdot v(i^*) \ge \text{OPT}_f(A' \setminus \{i^*\}, B)$ **then**
    return i^*
else
    return GREEDY-SM$(A, v, \mathbf{c}, \gamma B/2)$

Now we are ready to state mechanism DET-MECH-SYMSM-FRAC. The parameter ϵ' that appears in the description of the mechanism is determined by the analysis of the mechanism and only depends on the constants ρ and ϵ. Clearly the DET-MECH-SYMSM-FRAC runs in polynomial time.

DET-MECH-SYMSM-FRAC$(A, v, \mathbf{c}, B)$

Set $A' = \{i \mid c_i \le B\}$ and $i^* \in \arg\max_{i \in A'} v(i)$
if $\alpha \cdot v(i^*) \ge \text{OPT}_f(A' \setminus \{i^*\}, B)$ **then**
    return i^*
else
    $S =$ APPROX-LOCAL-SEARCH(A, v, ϵ')
    if $\text{OPT}_f(S \cap A', B) \ge \text{OPT}_f(A' \setminus S, B)$ **then**
        return MECH-SM-FRAC-VAR$(S, v, \mathbf{c}_S, B, (1 - (\alpha + 2)\epsilon'))$
    else
        return MECH-SM-FRAC-VAR$(A \setminus S, v, \mathbf{c}_{A \setminus S}, B, (1 - (\alpha + 2)\epsilon'))$

Theorem 6.3.1 below shows that for any objective for which we can establish a constant upper bound ρ on the ratio of the fractional and the integral optimal solutions, we have constant factor approximation mechanisms that run in polynomial-time.

Theorem 6.3.1. *For any $\epsilon > 0$, DET-MECH-SYMSM-FRAC is a deterministic, truthful, individually rational, budget-feasible mechanism for symmetric submodular valuations, that has approximation ratio $\alpha + 1 + \epsilon$ and runs in polynomial time.*

We view Theorem 6.3.1 as the most technically demanding result of this chapter. There are several steps involved in the proof, since we need the good properties of GREEDY-SM to still hold even for objectives that are not exactly non-decreasing. The remaining subsection is dedicated to the proof of the theorem.

First we have to pave the way for the proof and this means we need to prove that certain mechanisms work even for objectives that are not exactly non-decreasing. To make this precise, given a ground set A, a budget B and a constant $\epsilon \ge 0$, we say that a set function v is (B, ϵ)-*quasi-monotone* (or just quasi-monotone) on a set $X \subseteq A$ if for every $T \subsetneq X$ and every $i \in X \setminus T$, we have $v(T \cup \{i\}) - v(T) \ge -\frac{\epsilon}{n} \text{OPT}(X, B)$. Clearly, $(B, 0)$-quasi-monotone on X just means non-decreasing on X.

The main lemmata needed for our proofs are about GREEDY-SM, as it all boils down to the monotonicity, budget-feasibility, and approximation ratio of this simple mechanism. As mentioned in Lemma 5.1.1, GREEDY-SM is monotone, since any item out of the winning set remains out of the winning set if it increases its cost.

Lemma 6.3.2. *Suppose v is a (B,ϵ)-quasi-monotone submodular function on A such that $U\cdot v(S) \geq \textsc{opt}(A,B)$, where S is the set output by* GREEDY-SM$(A, v, \mathbf{c}, (1-U\epsilon)B/2)$ *and U is a constant.*[2] *Assuming the payments of Myerson's lemma, S is budget-feasible.*

Proof. Recall that before the description of GREEDY-SM we assumed—without loss of generality—that agents are sorted in descending order with respect to their ratio of marginal value over cost, i.e., $1 = \arg\max_{j\in A} \frac{v(j)}{c_j}$ and $i = \arg\max_{j\in A\setminus[i-1]} \frac{v([j])-v([j-1])}{c_j}$ for $i \geq 2$. Suppose that $S = [\ell]$, i.e., $1,2,\dots,\ell$ are added in S in that order. Let $S_0 = \emptyset = [0]$ and $S_i = [i]$ for $1 \leq i \leq \ell$. We are going to show that the payment to agent i is upper bounded by $\frac{B\cdot(v([i])-v([i-1])}{v(S)}$, and then budget feasibility directly follows from $\sum_{i\in[\ell]} \frac{B\cdot(v([i])-v([i-1])}{v(S)} = B$.

Suppose this upper bound does not hold for every agent. That is, there exists some $j \in [\ell]$ such that agent j bids $b_j > \frac{B\cdot(v([j])-v([j-1]))}{v(S)}$ and is still included in the output S' of GREEDY-SM$(A, v, (\mathbf{c}_{-j}, b_j), (1-U\epsilon)B/2)$. Let $\mathbf{b} = (\mathbf{c}_{-j}, b_j)$ and notice that up to agent $j-1$, agents are added to S' in the same order as they do in S, but after that the ordering might be affected. Also for $i \in [\ell]$ we have $\frac{v([i])-v([i-1])}{c_i} \geq \frac{v([\ell])-v([\ell-1])}{c_\ell} \geq \frac{2\cdot v([\ell])}{(1-U\epsilon)B}$ and therefore

$$\mathbf{b}(S\setminus\{j\}) = \mathbf{c}(S\setminus\{j\}) \leq \mathbf{c}(S) = \sum_{i\in[\ell]} c_i \leq \frac{(1-U\epsilon)B}{2\cdot v([\ell])} \sum_{i\in[\ell]} (v([i]) - v([i-1]) = \frac{(1-U\epsilon)B}{2}.$$

For GREEDY-SM$(A, v, (\mathbf{c}_{-j}, b_j), (1-U\epsilon)B/2)$ let S'_{j-1} be the chosen set right before j is added and $S'_j = S'_{j-1} \cup \{j\}$. This implies

$$j \in \arg\max_{i\in[n]} \frac{v(S'_{j-1}\cup\{i\}) - v(S'_{j-1})}{b_i} \quad \text{and} \quad \frac{v(S'_j) - v(S'_{j-1})}{b_j} \geq \frac{2\cdot v(S'_j)}{(1-U\epsilon)B}.$$

By Theorem 4.3.2 we have

$$\begin{aligned} v(S) - v(S'_j) &\leq \sum_{i\in S\setminus S'_j} (v(S'_j\cup\{i\}) - v(S'_j)) - \sum_{i\in S'_j\setminus S} (v(S'_j\cup S) - v(S'_j\cup S\setminus\{i\})) \\ &\leq \sum_{i\in S\setminus S'_j} (v(S'_j\cup\{i\}) - v(S'_j)) - |S'_j\setminus S|\Big(-\frac{\epsilon}{n}\textsc{opt}(A,B)\Big) \\ &\leq \sum_{i\in S\setminus S'_j} (v(S'_j\cup\{i\}) - v(S'_j)) + \epsilon U v(S). \end{aligned}$$

[2]Typically, in our case, U is a constant associated with the constant ρ that determines how the optimal solution to the relaxed problem is bounded by the optimal solution to the original problem. In particular, throughout this work, U is upper bounded by $(1+\rho)\big(2+\rho+\sqrt{\rho^2+4\rho+1}\big)+1$.

If $S \setminus S'_j = \emptyset$ then we directly get $v(S'_j) \ge (1-\epsilon U)v(S)$, otherwise we have

$$\begin{aligned}
(1-\epsilon U)v(S) - v(S'_j) &\le \sum_{i \in S \setminus S'_j} b_i \cdot \frac{v(S'_j \cup \{i\}) - v(S'_j)}{b_i} \\
&\le \max_{i \in S \setminus S'_j} \frac{v(S'_j \cup \{i\}) - v(S'_j)}{b_i} \cdot \sum_{i \in S \setminus S'_j} b_i \\
&\le \max_{i \in [n]} \frac{v(S'_j \cup \{i\}) - v(S'_j)}{b_i} \cdot \mathbf{b}(S \setminus S'_j) \\
&\le \max_{i \in [n]} \frac{v(S'_{j-1} \cup \{i\}) - v(S'_{j-1})}{b_i} \cdot \mathbf{b}(S \setminus \{j\}) \\
&\le \frac{v(S'_{j-1} \cup \{j\}) - v(S'_{j-1})}{b_j} \cdot \frac{(1-U\epsilon)B}{2} \\
&\le \frac{v(S_{j-1} \cup \{j\}) - v(S_{j-1})}{b_j} \cdot \frac{(1-U\epsilon)B}{2} \\
&\le \frac{v(S)}{B} \cdot \frac{(1-U\epsilon)B}{2} = \frac{(1-U\epsilon)v(S)}{2},
\end{aligned}$$

where the last inequality follows from the choice of b_j, while the next to last inequality follows from submodularity as $[j-1] = S_{j-1} \subseteq S'_{j-1}$. Therefore, $v(S'_j) \ge (1/2 - \epsilon U + \epsilon U/2)v(S)$.

In any case, we have $v(S'_j) \ge (1/2 - U\epsilon/2)v(S)$ and thus

$$\frac{v(S_j) - v(S_{j-1})}{b_j} \ge \frac{v(S'_j) - v(S'_{j-1})}{b_j} \ge \frac{2v(S'_j)}{(1-U\epsilon)B} \ge \frac{2(1/2 - U\epsilon/2)v(S)}{(1-U\epsilon)B} \ge \frac{v(S)}{B},$$

which contradicts our assumption about b_j. □

Lemma 6.3.2 establishes budget feasibility for quasi-monotone submodular functions. The next step is to make sure that the approximation guarantee does not deteriorate too much.

Lemma 6.3.3. *Suppose v is a (B,ϵ)-quasi-monotone submodular function on A. Let S be the set output by* GREEDY-SM$(A, v, \mathbf{c}, \beta B)$ *and $i^* \in \arg\max_{i \in \{j \in A \mid c_j \le B\}} v(i)$. Then* $\textsc{opt}(A,B) \le \frac{1}{1-\epsilon}\left(\frac{1+\beta}{\beta} v(S) + \frac{1}{\beta} v(i^*)\right)$.

Proof. Like in the proof of Lemma 6.3.2 we assume that agents are sorted in descending order with respect to their ratio of marginal value over cost, and that $S = [\ell]$. Let S^* be an optimal budget-feasible solution, i.e., $v(S^*) = \textsc{opt}(A,B)$. By Theorem 4.3.2 we have

$$\begin{aligned}
v(S^*) - v(S) &\le \sum_{i \in S^* \setminus S} (v(S \cup \{i\}) - v(S)) - \sum_{i \in S \setminus S^*} (v(S \cup S^*) - v(S \cup S^* \setminus \{i\})) \\
&\le \sum_{i \in S^* \setminus S} c_i \cdot \frac{v(S \cup \{i\}) - v(S)}{c_i} - |S \setminus S^*| \left(-\frac{\epsilon}{n} \textsc{opt}(A,B)\right) \\
&\le \max_{i \in S^* \setminus S} \frac{v(S \cup \{i\}) - v(S)}{c_i} \cdot \sum_{i \in S^* \setminus S} c_i + \epsilon v(S^*)
\end{aligned}$$

$$
\begin{aligned}
&= \frac{v(S\cup\{\ell+1\})-v(S)}{c_{\ell+1}}\cdot \mathbf{c}(S^*\setminus S)+\epsilon v(S^*)\\
&\le \frac{v(S\cup\{\ell+1\})}{\beta B}\cdot \mathbf{c}(S^*)+\epsilon v(S^*) \le \frac{v(S\cup\{\ell+1\})}{\beta}+\epsilon v(S^*)\\
&\le \frac{1}{\beta}(v(S)+v(\ell+1))+\epsilon v(S^*) \le \frac{1}{\beta}(v(S)+v(i^*))+\epsilon v(S^*)\,.
\end{aligned}
$$

By rearranging the terms we get $\textsc{opt}(A,B) \le \frac{1}{1-\epsilon}\left(\frac{1+\beta}{\beta}v(S)+\frac{1}{\beta}v(i^*)\right)$. □

We are now ready to state the proof of the theorem.

Proof of Theorem 6.3.1 : The proof follows the proof of Theorem 6.2.6. In fact, the monotonicity—and thus truthfulness and individual rationality—of the mechanism follows from that proof and the observation that MECH-SM-FRAC-VAR is monotone even when v is non-monotone. The latter is due to the monotonicity of GREEDY-SM which is straightforward and is briefly discussed before Lemma 6.3.2.

We proceed with the approximation ratio. If i^* is returned in line 3, then

$$\alpha\cdot v(i^*) \ge \textsc{opt}_f(A'\setminus\{i^*\},B) \ge \textsc{opt}(A'\setminus\{i^*\},B) = \textsc{opt}(A\setminus\{i^*\},B) \ge \textsc{opt}(A,B)-v(i^*)\,,$$

and therefore $\textsc{opt}(A,B) \le (\alpha+1)\cdot v(i^*)$.

On the other hand, if X is returned by MECH-SM-FRAC-VAR in line 7, then we have

$$\alpha\cdot v(i^*) < \textsc{opt}_f(A'\setminus\{i^*\},B) \le \rho\cdot\textsc{opt}(A'\setminus\{i^*\},B) \le \rho\cdot\textsc{opt}(A,B).$$

Therefore, $v(i^*) < \frac{\rho}{\alpha}\textsc{opt}(A,B)$.

At line 6 it must be the case $\textsc{opt}_f(S\cap A',B) \ge \textsc{opt}_f(A'\setminus S,B)$. Thus,

$$
\begin{aligned}
\rho\cdot\textsc{opt}(S,B) &= \rho\cdot\textsc{opt}(S\cap A',B) \ge \textsc{opt}_f(S\cap A',B) \ge \textsc{opt}_f(A'\setminus S,B)\\
&\ge \textsc{opt}(A'\setminus S,B) = \textsc{opt}(A\setminus S,B) \ge \textsc{opt}(A,B)-\textsc{opt}(S,B)\,.
\end{aligned}
$$

Therefore $\textsc{opt}(S,B) \ge \frac{1}{\rho+1}\textsc{opt}(A,B)$. Now we need the following lemma about the performance of MECH-SM-FRAC-VAR.

Lemma 6.3.4. *For $\eta = \rho+1+\sqrt{\rho^2+4\rho+1}$ and any $\epsilon''>0$, there is a sufficiently small ϵ' so that $\textsc{opt}(S,B) \le (\eta+1+\epsilon'')$MECH-SM-FRAC-VAR$(S,v,\mathbf{c}_S,B,(1-(\alpha+2)\epsilon'))$. Moreover ϵ' only depends on the constants ρ and ϵ''.*

Proof. Let $\delta = (1-(\alpha+2)\epsilon')/2$. We consider two cases for MECH-SM-FRAC-VAR.

If i^* is returned, then $\eta\cdot v(i^*) \ge \textsc{opt}_f(A'\setminus\{i^*\},B) \ge \textsc{opt}(A'\setminus\{i^*\},B) = \textsc{opt}(A\setminus\{i^*\},B) \ge \textsc{opt}(A,B)-v(i^*)$, and therefore $\textsc{opt}(A,B) \le (\eta+1)\cdot v(i^*)$.

On the other hand, if the outcome X of GREEDY-SM is returned, then $\eta\cdot v(i^*) < \textsc{opt}_f(A'\setminus\{i^*\},B) \le \rho\cdot\textsc{opt}(A'\setminus\{i^*\},B) \le \rho\cdot\textsc{opt}(A,B)$. Combining this with Lemma 6.3.3 we have

$$\textsc{opt}(A,B) \le \frac{1}{1-\epsilon'}\left(\frac{1+\delta}{\delta}v(X)+\frac{1}{\delta}\frac{\rho}{\eta}\textsc{opt}(A,B)\right),$$

or equivalently

$$\textsc{opt}(A,B) \le \frac{(1+\delta)\eta}{(1-\epsilon')\delta\eta-\rho}\, v(X)\,.$$

Since $\lim_{\epsilon'\to 0}\frac{(1+\delta)\eta}{(1-\epsilon')\delta\eta-\rho} = \frac{3\eta}{\eta-2\rho}$, we have that for sufficiently small ϵ', $\textsc{opt}(A,B) \le \left(\frac{3\eta}{\eta-2\rho}+\epsilon''\right)v(X) = (\eta+1+\epsilon'')\,v(X)$, where the calculations for the last equality are the same as in the proof of the approximation ratio of Mech-SM-frac. □

Now, combining all the above, we have

$$\textsc{opt}(A,B) \le (\rho+1)\,\textsc{opt}(S,B) \le (\rho+1)\left(\rho+2+\epsilon''+\sqrt{\rho^2+4\rho+1}\right)v(X)\,.$$

For $\epsilon'' = \frac{\epsilon}{\rho+1}$ we have $(\rho+1)\left(\rho+2+\epsilon''+\sqrt{\rho^2+4\rho+1}\right) = \alpha+1+\epsilon$, as desired. So, the ϵ' used in the mechanism is the one we get from the proof of Lemma 6.3.4 for $\epsilon'' = \frac{\epsilon}{\rho+1}$.

The case where X is returned in line 9 is symmetric. We conclude that in any case $\textsc{opt}(A,B) \le (\alpha+1+\epsilon)\cdot$Det-Mech-SymSM-frac$(A,B)$.

It remains to show that the mechanism is budget feasible. If the winner is i^* in line 3, then his payment is B. Otherwise, we would like budget-feasibility to follow from the budget-feasibility of Mech-SM-frac-var and the observation that the comparison in line 3 only gives additional upper bounds on the payments of winners from Mech-SM-frac. However, the budget-feasibility of Mech-SM-frac-var depends on the budget-feasibility of Greedy-SM, and according to Lemma 6.3.2 it suffices to have $v(X) \ge \frac{1}{\alpha+2}\cdot\textsc{opt}(A,B)$, where X is the output of Greedy-SM. This, however, follows from the approximation ratio for $\epsilon \le 1$. Thus the mechanism is budget-feasible. □

6.3.1 Weighted Cut Functions

Let us return now to the Max Cut problem, and consider the weighted version. An immediate implication of Theorem 6.3.1 is that we get a deterministic polynomial-time mechanism for Budgeted Max Weighted Cut with approximation ratio 58.72. This is just the result of substituting $\rho = 4$, as suggested by Theorem 6.2.5, in the formula for α.

Corollary 6.3.5. *There is a deterministic, truthful, individually rational, budget-feasible mechanism for Budgeted Max Weighted Cut that has approximation ratio 58.72 and runs in polynomial time.*

However, Theorem 6.2.5 says something stronger: given that $\max_{i\in A'} v(i)$ is small compared to $\textsc{opt}(A,B)$, ρ is strictly smaller than 4. Note that the first step in Det-Mech-SymSM-frac is to compare $\max_{i\in A'} v(i)$ to $\textsc{opt}_f(A'\setminus\{i^*\},B)$. This implies an upper bound on $\max_{i\in A'} v(i)$ in the following steps and we can use it to further fine-tune our mechanism. In particular, by setting $\alpha = 26.245$ instead of $(1+4)\left(2+4+\sqrt{16+16+1}\right)-1 = 57.72$ in Det-Mech-SymSM-frac, we can prove the following improved result that matches the approximation guarantee for unweighted cut functions.

Theorem 6.3.6. *There is a deterministic, truthful, individually rational, budget-feasible mechanism for Budgeted Max Weighted Cut that has approximation ratio 27.25, and runs in polynomial time.*

Proof. The proof is identical with the proof of Theorem 6.3.1 with the exception of the analysis of the approximation ratio which borrows the ideas of the proof of Theorem 6.2.6. So we only focus on the approximation ratio.

If i^* is returned in line 3, then $\textsc{opt}(A,B) \le (\alpha+1)\cdot v(i^*)$ like before.

Otherwise, using Theorem 6.2.5, we have

$$\alpha \cdot v(i^*) < \textsc{opt}_f(A' \setminus \{i^*\}, B) \le 4 \cdot \textsc{opt}(A' \setminus \{i^*\}, B) \le 4 \cdot \textsc{opt}(A,B).$$

Therefore, $v(i^*) < \frac{4}{\alpha}\textsc{opt}(A,B)$ and for the remaining steps of the mechanism we can use Theorem 6.2.5 with factor $\rho' = 2 + 8/\alpha$.

Without loss if generality, assume that X is returned by MECH-SM-FRAC-VAR in line 7. At line 6 it must be the case $\textsc{opt}_f(S \cap A', B) \ge \textsc{opt}_f(A' \setminus S, B)$. Thus, following the previous analysis, $\textsc{opt}(S,B) \ge \frac{1}{\rho'+1}\textsc{opt}(A,B)$.

Combining this inequality with Lemma 6.3.4 for ρ' we get that for $\epsilon'' = \frac{\epsilon}{\rho'+1}$ there is an ϵ' such that

$$\textsc{opt}(A,B) \le (\rho'+1)\left(\rho' + 2 + \epsilon'' + \sqrt{\rho'^2 + 4\rho' + 1}\right) v(X) = (\alpha + 1 + \epsilon) \cdot v(X)\,.$$

This ϵ' is to be used in the mechanism. Showing that $\alpha = 26.245$ works is only a matter of calculations, and for $\epsilon \le 1/200$ we get a 27.25-approximate solution. □

Chapter 7

Going Beyond Submodular Objectives[1]

Our central result in this chapter is a general scheme for obtaining randomized and deterministic polynomial time approximations for a subclass of XOS problems, that contains the budgeted versions of several well known optimization problems.

We first illustrate our ideas in Section 7.1, on the budgeted matching problem, where $v(S)$ is defined as the maximum weight matching that can be derived from the edges of S. For this problem only a randomized 768-approximation was known (Bei et al., 2012). We show that our approach yields a randomized 3-approximation and a deterministic 4-approximation.

Then in Section 7.2, we show how to generalize our results to problems with a similar combinatorial structure, where the set of feasible solutions forms an *independence system*. These structures are more general than matroids (they do not always satisfy the exchange property) and some representative problems that are captured include finding maximum weighted matroid members, maximum weighted k-D-matchings, and maximum weighted independent sets. For such problems we establish that a ρ-approximation to the algorithmic problem can be converted into a deterministic (resp. randomized), truthful, budget feasible mechanism with an approximation ratio of $2\rho+2$ (resp. $2\rho+1$). Note that essentially the approach for matching extends to problems where the unbudgeted versions are not easy as is the case with matching.

Finally, in Section 7.3 we briefly study the class of XOS functions, where we improve the current upper bound by a factor of 3. Note that that the known factor of 768 has been the benchmark against which most results in this, and the previous, chapter are presented.

Going beyond submodular valuations creates severe challenges. Recently, Goel, Nikzad, and Singla (2014) study a budgeted maximization problem with matching constraints, which is not submodular, and they achieve an approximation ratio of $3+o(1)$ with a deterministic mechanism, but under the large market assumption[2] (their mechanism has an unbounded ratio in general). Essentially, they use the same greedy approach with Singer (2010) and Chen, Gravin, and Lu (2011) but seen as a descending price auction. A very similar mechanism was also briefly discussed in

[1] A conference paper containing most results of this chapter appeared in WINE '16 (Amanatidis, Birmpas, and Markakis, 2016a). A preliminary version of the result of Section 7.3 appears in Amanatidis, Birmpas, and Markakis (2017).

[2] A market is said to be large if the number of participants is large enough that no single person can affect significantly the market outcome, i.e., $\max_i c_i/B = o(1)$.

Anari, Goel, and Nikzad (2014) for Knapsack under the large market assumption. We are building on this idea of gradually decreasing a global upper bound on the payment per value ratio to get most results of this chapter.

7.1 Budgeted Max Weighted Matching

We revisit the following budgeted matching problem.

Budgeted Max Weighted Matching. Given a budget B, and a graph $G = (V, E)$, where each edge $e_i \in E$ has a cost c_i and a value v_i, find a matching M of maximum value subject to $\sum_{i \in M} c_i \leq B$.

Here we study the mechanism design version of the problem, where the values are known to the mechanism and the edges are viewed as single-parameter strategic agents whose cost is private information.[3] Note that in order to formulate the problem to fit the general description given in the beginning of Section 4.3, we can define the valuation function as follows (as also mentioned in Bei et al., 2012): for any subset of edges $S \subseteq E$, $v(S)$ is taken to be the value of the maximum weighted matching of G that only uses edges in S. This function turns out to be XOS, but not submodular.

Claim 7.1.1. *The objective $v(\cdot)$ of weighted matching defined above is XOS, but not submodular.*

Proof. Matching is XOS: Let $\{M_1, M_2, \ldots, M_r\}$ be the finite set of all possible matchings of a given graph G. Now set $\alpha_j(S) = \sum_{i \in S \cap M_j} v_i$, for $j \in \{1, \ldots, r\}$, $S \subseteq E(G)$, and note that each α_j is an additive function. Since $v(S)$ is defined to be the value of the maximum weighted matching of G that only uses edges in S, we have that $v(S) = \max\{\alpha_1(S), \alpha_2(S), \ldots, \alpha_r(S)\}$.

Matching is not submodular: Recall that in the case of Budgeted Max Weighted Matching, the value $v(S)$ of a set S of edges is defined as the value of the maximum weight matching contained in S. To prove that $v(\cdot)$ is not submodular, consider the following example:

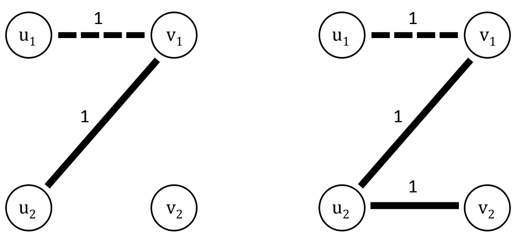

Let $A = \{u_2 v_1\}$, $B = \{u_2 v_1, u_2 v_2\}$ and add the dashed edge $u_1 v_1$ to both sets. Then we have that $v(A \cup \{u_1 v_1\}) - v(A) = 1 - 1 = 0 < 1 = 2 - 1 = v(B \cup \{u_1 v_1\}) - v(B)$. □

Hence, by Bei et al. (2012), there exists a randomized, 768-approximate mechanism for Max Weighted Matching, that is truthful and budget feasible.

[3] The work of Singer (2010) also studies a type of a budgeted matching problem. That objective, however, is OXS (a subclass of submodular objectives), and differs significantly from ours, which is not submodular (Singer, 2016).

We provide both deterministic and randomized polynomial time mechanisms with a much improved approximation ratio, based on selecting an outcome among two candidate solutions. The first solution comes from the greedy mechanism GREEDY-ISK described below. The main idea behind the mechanism is that in each iteration there is an implicit common upper bound on the rate that determines the payment of each winner in the candidate outcome of that iteration. More specifically, if the ith iteration is the final iteration (i.e., the condition in line 5 is true), the common payment per value for each of the winners is upper bounded by $\min\{B/v(M), c_{i-1}/v_{i-1}\}$. This upper bound decreases with each iteration, while the set of active agents is shrinking, until budget feasibility is achieved. At the same time we ensure the mechanism is monotone and returns enough value.

We assume that the mechanism also takes as input a deterministic exact algorithm f for the unbudgeted Max Weighted Matching, e.g., Edmond's algorithm (Edmonds, 1965). Later, in Subsection 7.2 the choice of f will depend on the underlying unbudgeted problem. Finally, note that our mechanisms are named after the generalization we study in Subsection 7.2, namely Independence System Knapsack problems.

GREEDY-ISK(A, v, c, B, f)

1 Set $A = \{i \mid c_i \leq B\}$
2 Possibly rename elements of A so that $\frac{c_1}{v_1} \geq \frac{c_2}{v_2} \geq \ldots \geq \frac{c_m}{v_m}$
3 **for** $i = 1$ *to* m **do**
4     $M = f(A, v)$
5   **if** $v(M) \cdot \frac{c_i}{v_i} \leq B$ **then**
6   **return** M
7   **else**
8         $A = A \setminus \{i\}$

We now exhibit some desirable properties of GREEDY-ISK, starting with truthfulness.

Lemma 7.1.2. *Mechanism* GREEDY-ISK *is monotone, and hence truthful and individually rational.*

Proof. By Lemma 4.3.5, we just need to show that the allocation rule is monotone, i.e., a winning agent remains a winner if he decreases his cost. Initially note that in line 4 the mechanism computes an optimal matching M (without a budget constraint) using only the values of the edges, thus it cannot be manipulated given the set of active edges A.

Fix a vector c_{-j} for the costs of the other agents, and suppose that when agent j declares c_j, he is in the matching M returned in the final iteration, say k, of GREEDY-ISK. Let agent j now report $c'_j < c_j$ to the mechanism. This makes him agent $j' \geq j$ in the new instance, but does not affect the relative ordering of the other agents (although a few of them may move down one position). Therefore, GREEDY-ISK will run exactly as before for each iteration $i < k$ and in the beginning of the kth iteration, it will produce the exact same matching M. Then in line 5, there are 2 cases to examine. If in the initial instance $j > k$, then we have the exact same ratio $\frac{c_k}{v_k}$ to consider, and the algorithm will terminate with M (since it did so in the initial instance). In the second case, $j = k$ in the initial instance. This means that now at the kth iteration, we either have the same agent with the reduced ratio $\frac{c'_k}{v_k}$ (since now $c'_k = c'_j$) or we have the agent who was in position $k+1$ in the initial instance with ratio equal to the original $\frac{c_{k+1}}{v_{k+1}}$. Therefore, the new ratio $\frac{c_k}{v_k}$ that the algorithm considers in this iteration is at most equal to the original ratio $\frac{c_k}{v_k}$. Thus, the condition in line 5 is satisfied, and the mechanism will return M. We conclude that an agent who is in the matching, remains in the matching by decreasing his cost. □

We also make the following remark, which can be derived by the same arguments used in the proof of Lemma 7.1.2. This property is crucial for derandomizing our mechanisms both here and in the next subsection.

Remark 7.1.3. There is no agent i that can manipulate the output set of GREEDY-ISK given that i is guaranteed to be a winner, i.e., fixing c_{-i}, if the winning sets are M and M' when i bids c_i and c'_i respectively, with $i \in M \cap M'$ then $M = M'$.

We move on to prove that the mechanism will never exceed the budget B, by establishing an appropriate upper bound on every winning bid.

Lemma 7.1.4. *Mechanism* GREEDY-ISK *is budget feasible.*

Proof. We will show that the threshold payment of Lemma 4.3.5, for any winning agent i cannot be higher than $\frac{v_i B}{v(M)}$. Fix a vector c_{-i} for all agents other than i and recall that the threshold payment, given c_{-i}, is the maximum cost that i can declare and still be included in the solution. So, towards a contradiction, suppose that agent i declares a cost $c_i > \frac{v_i B}{v(M)}$ and he is a winner. Let j denote the iteration where the mechanism GREEDY-ISK terminates and the matching M is returned. By the construction of the mechanism, and since $i \in M$, we have that $\frac{c_j}{v_j} \geq \frac{c_i}{v_i}$. Since j is the last iteration, we also have by line 5 that $v(M)\frac{c_j}{v_j} \leq B$. Hence $v(M)\frac{c_i}{v_i} \leq v(M)\frac{c_j}{v_j} \leq B$ that leads to the contradiction $c_i \leq \frac{v_i B}{v(M)}$. Therefore, the payment of each winning agent i is bounded by $\frac{v_i B}{v(M)}$, and the total payment of the mechanism is $\sum_{i \in M} p_i \leq \sum_{i \in M} \frac{v_i B}{v(M)} = B$. □

Finally, we analyze the quality of the solution produced by the greedy mechanism.

Lemma 7.1.5. *Mechanism* GREEDY-ISK *produces a matching with value at least* $\frac{1}{2}(v(M^*) - v_{i^*})$, *where* M^* *is an optimal solution to the given instance of Budgeted Max Weighted Matching, and* i^* *has maximum value among the budget feasible edges of* G, *i.e.,* $i^* \in \arg\max_{i \in F} v(i)$ *where* $F = \{i \in E(G) \mid c_i \leq B\}$.

Proof. Let M^* be an optimal budget feasible matching and A be the set of active edges at the final iteration j of GREEDY-ISK when matching M was returned. We have that

$v(M^*) = v(M^* \cap A) + v(M^* \setminus A)$. Since $M^* \cap A$ is a matching with edges from A but M is an optimal such matching, we have that

$$v(M^* \cap A) \leq v(M). \tag{7.1}$$

In addition, notice that if $i \in M^* \setminus A$ then $\frac{c_i}{v_i} \geq \frac{c_{j-1}}{v_{j-1}}$ since $j-1$ is the last edge removed from the set A before the final iteration j. Thus,

$$B \geq \sum_{i \in M^* \setminus A} c_i \geq \sum_{i \in M^* \setminus A} v_i \cdot \frac{c_{j-1}}{v_{j-1}} \geq v(M^* \setminus A) \cdot \frac{c_{j-1}}{v_{j-1}}. \tag{7.2}$$

Now, consider the $(j-1)$th iteration and call M' the matching produced in that iteration. Note that $M' \setminus \{j-1\}$ is a matching containing only edges that are active during iteration j. Therefore, $v(M) \geq v(M' \setminus \{j-1\})$. Moreover, if $j-1 \in M'$ then $v(M') = v(M' \setminus \{j-1\}) + v_{j-1}$, while if $j-1 \notin M'$ then $v(M') = v(M' \setminus \{j-1\}) \leq v(M' \setminus \{j-1\}) + v_{j-1}$. Using also the fact that $j-1$ was not the final iteration we have

$$\frac{v_{j-1}}{c_{j-1}} \cdot B < v(M') \leq v(M' \setminus \{j-1\}) + v_{j-1} \leq v(M) + v_{i^*}. \tag{7.3}$$

By combining (7.2) and (7.3) we get

$$v(M^* \setminus A) \leq v(M) + v_{i^*}. \tag{7.4}$$

Finally, combining (7.4) with (7.1) we get $v(M^*) = v(M^* \cap A) + v(M^* \setminus A) \leq v(M) + v(M) + v_{i^*} = 2v(M) + v_{i^*}$ and therefore $v(M) \geq \frac{1}{2}(v(M^*) - v_{i^*})$. □

We can now state our randomized mechanism for the problem (where the constants below have been optimized to get the best ratio).

Rand-ISK

1 Set $A = \{i \mid c_i \leq B\}$ and $i^* \in \arg\max_{i \in A} v(i)$
2 With probability 1/3 return i^* and with probability 2/3 return Greedy-ISK(A, v, c, B, f)

Theorem 7.1.6. Rand-ISK *is a universally truthful, individually rational, budget feasible, polynomial time randomized mechanism, achieving a 3-approximation in expectation, for the Budgeted Max Weighted Matching problem.*

Proof. Universal truthfulness and individual rationality follow from Lemma 7.1.2 and the fact that the simple mechanism that returns i^* and pays him B is truthful and individually rational. Regarding budget feasibility, just notice that if i^* is returned then the threshold payment is exactly B, otherwise the payments of Greedy-ISK are used, so budget feasibility follows from Lemma 7.1.4. Finally, if M denotes the outcome of Rand-ISK, then directly by Lemma 7.1.5 we have

$$\mathrm{E}(M) \geq \frac{2}{3} \cdot \frac{1}{2}(v(M^*) - v_{i^*}) + \frac{1}{3} v_{i^*} = \frac{1}{3} v(M^*),$$

thus proving the approximation ratio. □

7.1.1 Derandomization

We close this subsection by showing that we can have a deterministic polynomial time mechanism with a slightly worse approximation ratio. It is interesting to note that in contrast to MECH-SM or MECH-SM-FRAC, here i^* is directly compared to its alternative, which is just an approximate solution, without sacrificing truthfulness. This is due to Remark 7.1.3. Note also that although taking the maximum of two truthful algorithms does not always yield a truthful mechanism, we show that this is the case for the mechanism below.

DET-ISK

```
1 Set A = {i | c_i ≤ B} and i* ∈ argmax_{i∈A} v(i)
2 if v_{i*} ≥ GREEDY-ISK(A\{i*}, v, c_{-i*}, B, f) then
3     return i*
4 else
5     return GREEDY-ISK(A\{i*}, v, c_{-i*}, B, f)
```

Theorem 7.1.7. DET-ISK *is a truthful, individually rational, budget feasible, polynomial time deterministic mechanism, achieving a 4-approximation ratio for the Budgeted Max Weighted Matching problem.*

Proof. For truthfulness and individual rationality we show the algorithm is monotone. If i^* wins (lines 2-3) and he decreases his cost, he still wins since his bid is irrelevant to the outcome. On the other hand, suppose that the mechanism reaches line 5 and let $i \in A \setminus \{i^*\}$ be one of the winners. Then, by decreasing his cost, i will remain a winner since the output of GREEDY-ISK will not change (see the proof of Lemma 7.1.2 and Remark 7.1.3 after that), and the same branch of the mechanism DET-ISK will be executed again.

Budget feasibility is straightforward and follows from the same arguments used in the proof of Theorem 7.1.6

For the approximation ratio, we begin with some notation. Let v_G be the value of the matching returned by GREEDY-ISK$(A\setminus\{i^*\}, v, c_{-i^*}, B, f)$, M be the output of DET-ISK(A, v, c, B, f), $\textsc{opt}(S)$ be the value of an optimal solution with respect to the set of edges $S \subseteq E(G)$, and finally let i' be an edge of maximum value in the set $A\setminus\{i^*\}$. Clearly, if $\textsc{opt}$ is the value of an optimal solution to the initial instance, then $\textsc{opt} = \textsc{opt}(A)$. Finally, observe that $\textsc{opt}(A) \le \textsc{opt}(A\setminus\{i^*\}) + v_{i^*} \le 2v_G + v_{i'} + v_{i^*} \le 2v_G + 2v_{i^*}$. Now if $v_{i^*} \ge v_G$ then $\textsc{opt} \le 4v_{i^*} = 4v(M)$, while if $v_{i^*} < v_G$ then $\textsc{opt} \le 4v_G = 4v(M)$. Thus in any case we have that $\textsc{opt} \le 4v(M)$ and this concludes the proof. □

The above analysis of DET-ISK is tight, as shown below, i.e., there exist instances where the value of the optimal solution is arbitrarily close to four times the value of the mechanism's output.

Claim 7.1.8. *The analysis of* DET-ISK *is tight*

Proof. We provide an example where the value achieved by DET-ISK, is almost 1/4 of the value of an optimal solution. Consider the following independence system (U, I):

$U = \{1,2,3,4\}$, $I = 2^U$, where $v_1 = v+2\epsilon$, $v_2 = v$, $v_3 = v$, $v_4 = v+\epsilon$, for $\epsilon > 0$, $c_1 = \delta$, $c_2 = 10$, $c_3 = 10$, $c_4 = \delta$ for $\delta < \frac{5\epsilon}{v} << 10$, and $B = 20+2\delta$. It is easy to check that U is budget feasible and thus U is the optimal solution with value equal to $4v+3\epsilon = \text{OPT}$. Bellow we provide a visualization of such an instance in terms of matching:

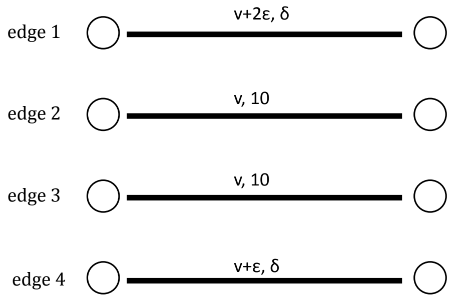

Now let us examine the value of DET-ISK's output: The most valuable item here is 1 with $v_1 = v+2\epsilon$, so $i^* = 1$. On the other hand, GREEDY-ISK orders the remaining items in the following manner: $\frac{c_2}{v_2} \geq \frac{c_3}{v_3} \geq \frac{c_4}{v_4}$. So by running this instance we have that $(3v+\epsilon)\frac{10}{v} = 30 + \frac{10\epsilon}{v} > 20 + \frac{2\cdot 5\epsilon}{v} > 20+2\delta = B$, and thus item 2 is excluded. GREEDY-ISK then moves to the next iteration (items 3 and 4 are active), where $(2v+\epsilon)\frac{10}{v} = 20 + \frac{10\epsilon}{v} = 20 + \frac{2\cdot 5\epsilon}{v} > 20+2\delta = B$, so item 3 is excluded as well. GREEDY-ISK moves to the next iteration (only item 4 is active), where $(v+\epsilon)\frac{\delta}{v+\epsilon} = \delta \leq 20+2\delta = B$, so the output is item 4 with total value $v+\epsilon$.

Now we have that $v_{i^*} = v+2\epsilon > v+\epsilon = \text{GREEDY-ISK}(U\setminus\{i^*\}, B)$ and hence the value of DET-ISK's output is $v+2\epsilon \simeq \frac{1}{4}(4v+3\epsilon) = \frac{1}{4}\text{OPT}$. □

Remark 7.1.9. Chen, Gravin, and Lu (2011) prove lower bounds for Knapsack, namely there is no deterministic (resp. randomized) truthful, budget feasible mechanism for Knapsack that can achieve an approximation ratio better than $1+\sqrt{2}$ (resp. 2). Note that these lower bounds hold here as well, because when the given graph G is a matching to begin with, Budgeted Max Weighted Matching reduces to Knapsack.

7.2 Independence System Knapsack Objectives

Our approach can tackle a number of different problems that have certain structural similarities with Budgeted Max Weighted Matching. Here, we define a class of such problems for which GREEDY-ISK—given an appropriate subroutine f—produces truthful, individually rational, budget feasible mechanisms with good approximation guarantees.

Two crucial properties of the matching problem were used in the previous subsection: (i) every subset of a matching is itself a matching, and (ii) the objective function becomes additive when restricted to matchings. These two properties is all we need, and note that (i) and (ii) are exactly what makes the set of matchings of a graph an *independence system*.

Definition 7.2.1. An *independence system* is a pair (U, I), where U is an arbitrary finite set and $I \subseteq 2^U$ is a family of subsets, whose members are called the independent sets of U and satisfy:

(i) $\emptyset \in I$

(ii) If $B \in I$ and $A \subseteq B$, then $A \in I$.

Below we define a variant of Knapsack where the feasible solutions are constrained to an independence system. This forms a generalization of knapsack problems subject to matroid constraints, which are more common in the literature.

Independence System Knapsack. Given an independence system (U, I) with costs c_i and values v_i on the elements of U, as well as a budget B, find $M \in I$ that maximizes $\sum_{i\in M} v_i$ subject to $\sum_{i\in M} c_i \le B$.

Note that for plain Knapsack $U = [n], I = 2^{[n]}$, while for Budgeted Max Weighted Matching U is the set of edges of a given graph G and I is the set of all matchings of G. There exist several other problems that are special cases of Independence System Knapsack, like

- *Budgeted Max Weighted Forest* where U is the set of edges of a given graph G and I is the set of acyclic subgraphs of G,
- *Budgeted Max Weighted Matroid Member* where (U, I) is a matroid[4] (Budgeted Max Weighted Forest is a special case of this problem),
- *Budgeted Max Independent Set* where U is the set of vertices of a given graph G and I is the set of independence sets of G, and
- *Budgeted Max Weighted k-D-Matching* where U is the set of hyperedges of a k-uniform k-partite hypergraph H and I is the set of all k-dimensional matchings of H.

The following can be easily derived as in the case of Budgeted Max Weighted Matching.

Lemma 7.2.2. *Every problem that can be formulated as an Independence System Knapsack problem belongs to the class XOS.*

Clearly it is not always possible to find an optimal solution to Independence System Knapsack in polynomial time, even if we remove the budget constraint. Putting the running time aside, however, GREEDY-ISK combined with an exact algorithm f for the problem makes RAND-ISK (resp. DET-ISK) a 3-approximate randomized (resp. 4-approximate deterministic) truthful, individually rational, budget feasible mechanism.

Moreover, when the unbudgeted underlying problem is easy—as is the case for Max Weighted Matching, Max Weighted Forest, and Max Weighted Matroid Member—the mechanisms run in polynomial time. Even if the unbudgeted underlying problem

[4]A *matroid* (U, I) is an independence system that also has the *exchange property*:
(iii) If $A, B \in I$ and $|A| < |B|$, then there exists $x \in B \setminus A$ such that $A \cup \{x\} \in I$.

is NP-hard, as long as there is a polynomial time $\rho(n)$-approximation we get $O(\rho(n))$-approximate, truthful, individually rational, budget feasible mechanisms, e.g., for Budgeted Max Weighted k-D-Matching this translates to a $O(k)$-approximation mechanism. Here, n is the size of the input, and we should mention that the independent sets of U may not be explicitly given. Typically we assume an *independence oracle* that decides for any $X \subseteq U$ whether $X \in I$. However, note that in every special case of Independence System Knapsack mentioned above, we are given a combinatorial, succinct representation of I and therefore there is no need to assume access to an oracle.

When using a $\rho(n)$-approximation algorithm we should adjust the probabilities in RAND-ISK, namely we should use $\frac{2\rho(n)}{2\rho(n)+1}$ instead of $2/3$ and $\frac{1}{2\rho(n)+1}$ instead of $1/3$. Moreover, for both mechanisms and without loss of generality, we assume that for every $i \in U$ we have $\{i\} \in I$, or else i can be excluded from the initial set A of active elements that is given as input to the mechanisms.

Theorem 7.2.3. *If a deterministic $\rho(n)$-approximation algorithm f for the unbudgeted version of Independence System Knapsack is given as an auxiliary input to* GREEDY-ISK*, then* RAND-ISK *(resp.* DET-ISK*) becomes a $(2\rho(n)+1)$-approximate randomized (resp. $(2\rho(n)+2)$-approximate deterministic) truthful, individually rational, budget feasible mechanism. Moreover, if f runs in polynomial time so do the mechanisms.*

Proof. The proof of Theorem 7.2.3 follows closely the analysis of Section 7.1. In fact, the proof of truthfulness, individual rationality, and budget feasibility is exactly the same with the proofs of Lemmata 7.1.2 and 7.1.4, if we replace "matching" with "independent set of I" and "edge" with "element of U". The proof of the approximation ratios follows closely the proofs of Lemma 7.1.5 and Theorems 7.1.6 and 7.1.7, so we will only focus on the differences.

Let $M^* \in I$ be an optimal budget feasible independent set and $A \subseteq U$ be the set of active elements at the final iteration j of GREEDY-ISK when the $\rho(n)$-approximate solution M was returned. Also, let M_A be an optimal budget feasible independent set using only elements of A. We have that $v(M^*) = v(M^* \cap A) + v(M^* \setminus A)$. But $M^* \cap A \subseteq M^*$ is an independent set with elements from A so, the analog of (7.1) is now

$$v(M^* \cap A) \leq v(M_A) \leq \rho(n) \cdot v(M).$$

Similarly, the analog of (7.4) is now

$$v(M^* \setminus A) \leq \rho(n) \cdot v(M) + v_{i^*},$$

and thus we get $v(M^*) = v(M^* \cap A) + v(M^* \setminus A) \leq 2\rho(n) \cdot v(M) + v_{i^*}$, or equivalently $v(M) \geq \frac{1}{2\rho(n)}(v(M^*) - v_{i^*})$. Now the approximation ratio of RAND-ISK is straightforward since the expected value of the output of the mechanism is at least

$$\frac{2\rho(n)}{2\rho(n)+1} \cdot \frac{1}{2\rho(n)}(v(M^*) - v_{i^*}) + \frac{1}{2\rho(n)+1} v_{i^*} = \frac{1}{2\rho(n)+1} v(M^*).$$

For the approximation ratio of DET-ISK, using the notation of the proof of Theorem 7.1.7 we have $\text{OPT}(A) \le \text{OPT}(A \setminus \{i^*\}) + v_{i^*} \le 2\rho(n)v_G + v_{i'} + v_{i^*} \le 2\rho(n)v_G + 2v_{i^*}$. If $v_{i^*} \ge v_G$ then $\text{OPT} \le (2\rho(n)+2)v_{i^*}$, while if $v_{i^*} < v_G$ then $\text{OPT} \le (2\rho(n)+2)v_G$. Thus in any case we have $\text{OPT} \le (2\rho(n)+2)v(M)$. □

Combining Theorem 7.2.3 with the polynomial time $(k-1)$-approximation algorithm of Chan and Lau (2012) for Max Weighted k-D-Matching, and the fact that Max Weighted Forest and Max Weighted Matroid Member (given a polynomial time independence oracle) can be solved in polynomial time (see, e.g., Cook et al., 1998), we get the following corollary.

Corollary 7.2.4. *We can obtain*
(i) randomized 3*-approximation mechanisms and deterministic* 4*-approximation mechanisms for Budgeted Max Weighted Forest and Budgeted Max Weighted Matroid Member (as well as Knapsack and Budgeted Max Weighted Matching) that run in polynomial time.*
(ii) randomized 3*-approximation mechanisms and deterministic* 4*-approximation mechanisms for Budgeted Max Weighted Independent Set and Budgeted Max Weighted* k*-D-Matching.*
(iii) for any $k \ge 3$*, a randomized* $(2k-1)$*-approximation mechanism and a deterministic* $2k$*-approximation mechanism for Budgeted Max Weighted* k*-D-Matching that run in polynomial time.*

Remark 7.2.5. Max Weighted Independent Set and Max Weighted k-D-Matching are not submodular, as was the case for Max Weighted Matching. Max Weighted Matroid Member (and thus Max Weighted Forest), on the other hand, is submodular and therefore the results of Chen, Gravin, and Lu (2011) apply. Note, however, that we still get a significant improvement, both on the approximation ratio and on the running time by our approach.

Naturally, Remark 7.1.9 applies here as well. For every problem stated in this section there is no deterministic (resp. randomized) truthful, budget feasible mechanism with better approximation ratio than $1+\sqrt{2}$ (resp. 2). These lower bounds are independent of any complexity assumption.

7.3 An Improved Upper Bound for XOS Objectives

In Bei et al. (2012) a randomized, universally truthful and budget-feasible 768-approximation mechanism was introduced for XOS functions. For several of the results in Part II the best previously known upper bound follows from this result (see also Remark 6.0.1). In this section we slightly modify their mechanism to improve its performance.[5] Although there is not much novelty in this result, it feels appropriate

[5]In a recent unpublished manuscript, Leonardi et al. (2016) also suggest a fine-tuning of this mechanism that yields a factor 436.

to provide this tighter analysis, given that the factor of 768 has been the benchmark against which our results are presented.

Note that in Section 4.3 we defined non-decreasing XOS functions, and everything is stated and proved for those. However, there is a relatively straightforward way to extend any result to general XOS functions, as defined in Gupta, Nagarajan, and Singla (2017); see Remark 6.0.1 and Appendix B.2.

Below we provide the mechanism MAIN-XOS. Our mechanism has the same structure as the one presented in Bei et al. (2012) but we tune its parameters and perform a slightly different analysis in order to improve the approximation factor. The mechanism ADDITIVE-MECHANISM of Chen, Gravin, and Lu (2011) for additive valuation functions is used as a subroutine. ADDITIVE-MECHANISM is a universally truthful, 3-approximate mechanism (see Theorem B.2 in Chen, Gravin, and Lu (2011)). Initially, we revisit the *random sampling* part of the mechanism and modify the threshold bound of line 2:

SAMPLE-XOS

---

1 Pick each item independently at random with probability $\frac{1}{2}$ to form a set T
2 Compute $\textsc{opt}(T, v, \mathbf{c}_T, B)$ and set a threshold $t = \frac{\textsc{opt}(T, v, \mathbf{c}_T, B)}{4.6B}$
3 Find a set $S^* \in \arg\max_{S \subseteq A \setminus T}\{v(S) - t \cdot \mathbf{c}(S)\}$
4 Find an additive function α in the XOS representation of $v(\cdot)$ with $\alpha(S^*) = v(S^*)$
5 **return** ADDITIVE-MECHANISM$(S^*, \alpha, \mathbf{c}_{S^*}, B)$

This part is used as one of the two alternatives of the main mechanism. We modify the probabilities with which the two outcomes occur:

MAIN-XOS

---

1 With probability $p = 0.08$, pick a most valuable item as the only winner and pay him B
2 With probability $1 - p$, run SAMPLE-XOS

By following a similar but more careful analysis, we improve the approximation ratio by a factor of 3 while also retaining its properties. Notice that like the mechanism of Bei et al. (2012), MAIN-XOS is randomized and has superpolynomial running time. In particular, it requires a demand oracle.

Theorem 7.3.1. MAIN-XOS *is universally truthful, individually rational, budget-feasible, and has approximation ratio* 244.

Proof. Universal truthfulness, individual rationality and budget feasibility follow directly from the proof of Theorem 3.1 of Bei et al. (2012).

We proceed in proving the approximation ratio of the mechanism. Let $i^* \in \arg\max_{i \in A} v(i)$ and X_S be the output of SAMPLE-XOS. In addition, let X be an optimal solution, i.e., a subset of A such that $v(X) = \textsc{opt}(A, v, \mathbf{c}, B)$ and $\mathbf{c}(X) \leq B$.

We also need the following two lemmata:

Lemma 7.3.2 (Claim 3.1 in Bei et al. (2012))**.** *For any* $S \subseteq S^*, \alpha(S) - t \cdot \mathbf{c}(S) \geq 0$.

Lemma 7.3.3 (Lemma 2.1 in Bei et al. (2012))**.** *Consider any subadditive function* $v(\cdot)$. *For any given subset* $S \subseteq A$ *and a positive integer* k *assume that* $v(S) \geq k \cdot v(i)$ *for any*

$i \in S$. Further, suppose that S is divided uniformly at random into two groups T_1 and T_2. Then with probability at least $\frac{1}{2}$, we have that $v(T_1) \geq \frac{k-1}{4k} v(S)$ and $v(T_2) \geq \frac{k-1}{4k} v(S)$.

Let $\kappa = 74$. We are going to use κ to set different values for k in Lemma 7.3.3 for different cases. We follow a similar analysis as Bei et al. (2012), but since we use two different values for k, things get a little more complicated. Let X_M to be the output of MAIN-XOS. We have the following three case regarding the value of item i^*:

1. $v(i^*) > \frac{3.8}{\kappa}\textsc{opt}(A, v, \mathbf{c}, B)$. In this case we have that

$$\mathrm{E}(v(X_M)) = p \cdot v(i^*) + (1-p) \cdot \mathrm{E}(v(X_S)) \geq p \cdot v(i^*) \geq p \frac{3.8}{\kappa} \textsc{opt}(A, v, \mathbf{c}, B),$$

and thus, $\textsc{opt}(A, v, \mathbf{c}, B) \leq 243.5 \cdot \mathrm{E}(v(X_M))$.

2. $\frac{1}{\kappa}\textsc{opt}(A, v, \mathbf{c}, B) < v(i^*) \leq \frac{3.8}{\kappa}\textsc{opt}(A, v, \mathbf{c}, B)$. Note that now we can apply Lemma 7.3.3 with $k = \frac{\kappa}{3.8}$. We split this case into two subcases:

   - $\mathbf{c}(S^*) > B$. Since $\mathbf{c}(S^*)$ is more than B we can find a subset $S' \subseteq S^*$, such that $\frac{B}{2} \leq \mathbf{c}(S') \leq B$. By Lemma 7.3.2 we have

$$\alpha(S') \geq t \cdot \mathbf{c}(S') \geq \frac{\textsc{opt}(T, v, \mathbf{c}_T, B)}{4.6 \cdot B} \cdot \frac{B}{2} \geq \frac{\textsc{opt}(T, v, \mathbf{c}_T, B)}{9.2}.$$

   By using Lemma 7.3.3 we have

$$\begin{aligned}\textsc{opt}(T, v, \mathbf{c}_T, B) \geq \textsc{opt}(T \cap X, v, \mathbf{c}_{T\cap X}, B) &\geq \frac{\kappa - 3.8}{4\kappa}\textsc{opt}(X, v, \mathbf{c}_X, B) \\ &= \frac{\kappa - 3.8}{4\kappa}\textsc{opt}(A, v, \mathbf{c}, B).\end{aligned}$$

   Since $\textsc{opt}(S^*, \alpha, \mathbf{c}_{S^*}, B)$ is the value of an optimal solution and S' a particular solution under budget constraint B, we conclude that

$$\textsc{opt}(S^*, \alpha, \mathbf{c}_{S^*}, B) \geq \alpha(S') \geq \frac{\textsc{opt}(T, v, \mathbf{c}_T, B)}{9.2} \geq \frac{(\kappa - 3.8)}{36.8 \cdot \kappa}\textsc{opt}(A, v, \mathbf{c}, B),$$

   with probability at least $\frac{1}{2}$.

   - $\mathbf{c}(S^*) \leq B$. In this case we have that

$$\alpha(S^*) = \textsc{opt}(S^*, \alpha, \mathbf{c}_{S^*}, B) = \textsc{opt}(S^*, v, \mathbf{c}_{S^*}, B) = v(S^*).$$

   Let $S' = X \setminus T$. We have $\mathbf{c}(S') \leq \mathbf{c}(X) \leq B$ and also, by using Lemma 7.3.3, $v(S') \geq \frac{\kappa - 3.8}{4\kappa}\textsc{opt}(A, v, \mathbf{c}, B)$ with probability of at least $\frac{1}{2}$. Then, we have

$$\begin{aligned}\textsc{opt}(S^*, \alpha, \mathbf{c}_{S^*}, B) = v(S^*) &\geq v(S^*) - t \cdot \mathbf{c}(S^*) \geq v(S') - t \cdot \mathbf{c}(S') \\ &\geq \frac{\kappa - 3.8}{4\kappa}\textsc{opt}(A, v, \mathbf{c}, B) - \frac{\textsc{opt}(T, v, \mathbf{c}_T, B)}{4.6 \cdot B} \cdot B \\ &\geq \left(\frac{\kappa - 3.8}{4\kappa} - \frac{1}{4.6}\right)\textsc{opt}(A, v, \mathbf{c}, B),\end{aligned}$$

with probability at least $\frac{1}{2}$.

By substituting κ, it is easy to see that $\frac{\kappa-3.8}{4\kappa} - \frac{1}{4.6} < \frac{(\kappa-3.8)}{36.8\cdot\kappa}$. So, in both cases we have that $\textsc{opt}(S^*, \alpha, \mathbf{c}_{S^*}, B) \geq \left(\frac{\kappa-3.8}{4\kappa} - \frac{1}{4.6}\right)\textsc{opt}(A, v, \mathbf{c}, B)$ with probability at least $\frac{1}{2}$. Now recall that Additive-Mechanism has an approximation factor of at most 3 with respect to $\textsc{opt}(S^*, \alpha, \mathbf{c}_{S^*}, B)$. So we can finally derive that

$$\mathrm{E}(v(X_S)) \geq \frac{1}{3}\textsc{opt}(S^*, \alpha, \mathbf{c}_{S^*}, B) \geq \frac{1}{3}\cdot\frac{1}{2}\cdot\left(\frac{\kappa-3.8}{4\kappa} - \frac{1}{4.6}\right)\cdot\textsc{opt}(A, v, \mathbf{c}, B).$$

Thus the solution that Main-XOS returns has expected value

$$\begin{aligned}\mathrm{E}(v(X_M)) &= p\cdot v(i^*) + (1-p)\cdot\mathrm{E}(v(X_S))\\ &\geq \frac{p}{\kappa}\cdot\textsc{opt}(A, v, \mathbf{c}, B) + \frac{(1-p)}{6}\cdot\left(\frac{\kappa-3.8}{4\kappa} - \frac{1}{4.6}\right)\cdot\textsc{opt}(A, v, \mathbf{c}, B).\end{aligned}$$

By substituting the values for p, κ we get $\textsc{opt}(A, v, \mathbf{c}, B) \leq 243.2\cdot\mathrm{E}(v(X_M))$.

3. $v(i^*) \leq \frac{1}{\kappa}\textsc{opt}(A, v, \mathbf{c}, B)$. The analysis of case 2 holds here as well, so we omit the details. The only difference is that now Lemma 7.3.3 should be applied with $k = \kappa$. The subcase where $\mathbf{c}(S^*) > B$ gives $\textsc{opt}(S^*, \alpha, \mathbf{c}_{S^*}, B) \geq \frac{(\kappa-1)}{36.8\cdot\kappa}\textsc{opt}(A, v, \mathbf{c}, B)$ with probability at least $\frac{1}{2}$, while the subcase where $\mathbf{c}(S^*) \leq B$ gives $\textsc{opt}(S^*, \alpha, \mathbf{c}_{S^*}, B) \geq \left(\frac{\kappa-1}{4\kappa} - \frac{1}{4.6}\right)\textsc{opt}(A, v, \mathbf{c}, B)$ with probability at least $\frac{1}{2}$. By substituting κ, we see that $\frac{(\kappa-1)}{36.8\cdot\kappa} < \frac{\kappa-1}{4\kappa} - \frac{1}{4.6}$. So, in both cases we have that $\textsc{opt}(S^*, \alpha, \mathbf{c}_{S^*}, B) \geq \frac{(\kappa-1)}{36.8\cdot\kappa}\textsc{opt}(A, v, \mathbf{c}, B)$ with probability at least $\frac{1}{2}$. The above analysis gives that the solution returned by Main-XOS has expected value

$$\mathrm{E}(v(X_M)) \geq \frac{(1-p)}{6}\cdot\frac{(\kappa-1)}{36.8\cdot\kappa}\cdot\textsc{opt}(A, v, \mathbf{c}, B).$$

By substituting the values for p, κ we get $\textsc{opt}(A, v, \mathbf{c}, B) \leq 243.3\cdot\mathrm{E}(v(X_M))$.

We conclude that $\textsc{opt}(A, v, \mathbf{c}, B) \leq 244\cdot E(v(X_M))$. □

7.4 Directions for Future Research

There are still many interesting open problems that are worth further exploration in the context of budget-feasible mechanism design. First, for the case of submodular functions, even though we do have a better understanding for designing mechanisms given all the previous works, the current results are still not known to be tight.

In Chapter 6, we have made progress on the design of budget-feasible mechanisms for symmetric submodular objectives. It is a first step towards studying non-monotone submodular objectives, but this is still a largely unexplored territory. Moreover, even for the significantly improved results we obtained, it remains an open problem whether any of these approximation ratios are tight, as is the case for non-decreasing objectives.

For non submodular objectives, and specifically for the XOS class, the picture is way more challenging. We would like to identify more problems that admit better approximation guarantees, even with exponential time mechanisms. A component that seems to be missing at the moment is a characterization of truthful and budget feasible mechanisms. We believe that obtaining characterization results would be crucial in resolving the above questions.

Regarding our techniques, we expect that the idea of utilizing local search in order to identify monotone regions of a general submodular function may have further applications both on mechanism design and on submodular optimization.

Appendix A

Missing Material from Section 3.1

Proof of Theorem 3.1.4. Assume $\mathcal{X}$ is a picking-exchange mechanism with partition (N_1, N_2, E_1, E_2), offer sets $\mathcal{O}_i$ on N_i, for $i \in \{1,2\}$, and set of exchange deals D. Let $\mathbf{v} = (v_1, v_2) \in \mathcal{V}_m^{\neq}$ be a profile, and fix v_2. We are going to show that there is no $\mathbf{v}' = (v_1', v_2) \in \mathcal{V}_m^{\neq}$ such that $v_1(X_1(\mathbf{v}')) > v_1(X_1(\mathbf{v}))$.

For any $\mathbf{v}' = (v_1', v_2) \in \mathcal{V}_m^{\neq}$ there exist the following possibilities:

(a) $X_1(\mathbf{v}') = X_1(\mathbf{v})$. Then clearly $v_1(X_1(\mathbf{v}')) = v_1(X_1(\mathbf{v}))$.

(b) $X_1^{N_1 \cup N_2}(\mathbf{v}') \neq X_1^{N_1 \cup N_2}(\mathbf{v})$, but $X_1^{E_1 \cup E_2}(\mathbf{v}') = X_1^{E_1 \cup E_2}(\mathbf{v})$. Then it must be the case where $X_1^{N_1}(\mathbf{v}') \neq X_1^{N_1}(\mathbf{v})$. Indeed, player 1 has no power over N_2 where the items that he is allocated depend only on the unique best offer to player 2, i.e., $X_1^{N_2}(\mathbf{v}') = X_1^{N_2}(\mathbf{v})$. But this can only mean $v_1(X_1^{N_1}(\mathbf{v}')) < v_1(X_1^{N_1}(\mathbf{v}))$ by the definition of a picking-exchange mechanism and the fact that there are no subsets of equal value. So in total, $v_1(X_1(\mathbf{v}')) < v_1(X_1(\mathbf{v}))$.

(c) $X_1^{N_1 \cup N_2}(\mathbf{v}') = X_1^{N_1 \cup N_2}(\mathbf{v})$, but $X_1^{E_1 \cup E_2}(\mathbf{v}') \neq X_1^{E_1 \cup E_2}(\mathbf{v})$. By the definition of picking-exchange mechanisms, player 1 can never force an exchange that is good for him but not for player 2. That is, by deviating he will lose one or more exchanges that were good for him, and/or force one or more exchanges that were bad for him. We conclude it is the case where $v_1(X_1^{E_1 \cup E_2}(\mathbf{v}')) < v_1(X_1^{E_1 \cup E_2}(\mathbf{v}))$, and therefore, $v_1(X_1(\mathbf{v}')) < v_1(X_1(\mathbf{v}))$.

(d) $X_1^{N_1 \cup N_2}(\mathbf{v}') \neq X_1^{N_1 \cup N_2}(\mathbf{v})$, and $X_1^{E_1 \cup E_2}(\mathbf{v}') \neq X_1^{E_1 \cup E_2}(\mathbf{v})$. By the fact that we are restricted to $\mathcal{V}_m^{\neq}$, we can derive that the "picking part" on $N_1 \cup N_2$ and the "exchange part" on $E_1 \cup E_2$ are independent. So, by cases (b) and (c) above we have $v_1(X_1^{N_1}(\mathbf{v}')) < v_1(X_1^{N_1}(\mathbf{v}))$ and $v_1(X_1^{E_1 \cup E_2}(\mathbf{v}')) < v_1(X_1^{E_1 \cup E_2}(\mathbf{v}))$. Therefore, $v_1(X_1(\mathbf{v}')) < v_1(X_1(\mathbf{v}))$.

We conclude that every picking-exchange mechanism on $\mathcal{V}_m^{\neq}$ is truthful. □

Remark A.0.1. With only slight modifications of the above proof, we have that for general additive valuations every picking-exchange mechanism is truthful when using the following two interesting families of tie-breaking rules:

Tie-breaking with labels. Every set in $\mathcal{O}_1 \cup \mathcal{O}_2$ has a distinct label, and whenever $\arg\max_{S \in \mathcal{O}_i} v_i(S)$ is not a singleton, player i receives the set with the smallest label. Further, every deal in D has a label with five possible values, each indicating one of the following: (i) the exchange takes place every time it is not unfavorable, (ii) it only takes place every time it is not unfavorable and at least one player is strictly improved, (iii) it only takes place every time it is not unfavorable and player 1 is strictly improved, (iv) it only takes place every time it is not unfavorable and player 2 is strictly improved, and (v) it only takes place every time it is favorable.

Welfare maximizing tie-breaking. When $\arg\max_{S\in\mathcal{O}_i} v_i(S)$ is not a singleton, player i receives the set that leaves in N_i as much value as possible for the other player. If there are still ties, labels are used to resolve those. Further, for every deal in D the exchange takes place every time it is not unfavorable and at least one player is strictly improved.

Proof of Lemma 3.1.8. Let $\mathbf{v} = (v_1, v_2)$ be a profile such that both players strongly desire S and $S \subseteq X_1(\mathbf{v})$ (the case where $S \subseteq X_2(\mathbf{v})$ is symmetric). We first prove the statement for $T = S$. Let $\mathbf{v}' = (v'_1, v'_2)$ be any profile in which player 1 strongly desires S, i.e., $v'_{1x} > \sum_{y\in M\setminus S} v'_{1y}, \forall x \in S$. Initially, consider the intermediate profile $\mathbf{v}^* = (v_1, v'_2)$. If $S \cap X_2(\mathbf{v}^*) \neq \emptyset$ then player 2 would deviate from profile $\mathbf{v}$ to $\mathbf{v}^*$ in order to strictly improve his total utility. So by truthfulness we derive that $S \subseteq X_1(\mathbf{v}^*)$. Similarly, in the profile $\mathbf{v}'$, if $S \cap X_2(\mathbf{v}') \neq \emptyset$ then player 1 would deviate from $\mathbf{v}'$ to $\mathbf{v}^*$ in order to strictly improve. Thus by truthfulness we have $S \subseteq X_1(\mathbf{v}')$. We conclude that player 1 controls S.

Now, suppose that $\mathbf{v}'' = (v''_1, v''_2)$ is any profile in which player 1 strongly desires $T \subsetneq S$. If $T \nsubseteq X_1(\mathbf{v}'')$ then player 1 could strictly improve his utility by playing v'_1 from before (i.e., he declares that he strongly desires S) and getting $S \supsetneq T$. Thus, by truthfulness, $T \subseteq X_1(\mathbf{v}'')$, and we conclude that player 1 controls T. □

Proof of Corollary 3.1.10. From the definition of the C_is and Corollary 3.1.9, $C_1 \cup C_2 = M$ follows. On the other hand, if $z \in C_1 \cap C_2$, then there exist a set $A \in \mathcal{A}_1$, such that $z \in A$, and a set $B \in \mathcal{A}_2$ such that $z \in B$. By Lemma 3.1.8, this implies that the singleton $\{z\}$ is controlled by both players, which is a contradiction. Thus, we have $C_1 \cap C_2 = \emptyset$. □

Proof of Lemma 3.1.11. Due to symmetry, it suffices to prove the statement for $i = 1$. If $N_1 = \emptyset$ then the statement is trivially true. So assume $N_1 \neq \emptyset$ and suppose that the statement does not hold. That is, there exists a profile $\mathbf{v} = (v_1, v_2)$ such that for any $S \in \mathcal{O}_1$ we have $X_1^{N_1}(\mathbf{v}) \nsubseteq S$. This means $X_1^{N_1}(\mathbf{v}) \neq \emptyset$. Since the sets in $\mathcal{O}_1$ cover N_1, there exists S' such that $S' \cap X_1^{N_1}(\mathbf{v}) \neq \emptyset$. Let Z be a maximum cardinality such intersection between some $S' \in \mathcal{O}_1$ and $X_1^{N_1}(\mathbf{v})$, and x be any element of $X_1^{N_1}(\mathbf{v})\setminus Z$. Note that x is guaranteed to exist since $X_1^{N_1}(\mathbf{v})$ is not contained in any set of $\mathcal{O}_1$. Also, there is no $S'' \in \mathcal{O}_1$ such that $Z \cup \{x\} \subseteq S''$ due to the maximality of Z.

The generic values that may appear in $\mathbf{v}$ restrict our ability to argue about the allocation, so our first goal is to reach a profile $\mathbf{u}$ that contradicts the lemma's statement, like $\mathbf{v}$, but has appropriately selected values. Then, having $\mathbf{u}$ as a starting point we can create profiles in which the allocations contradict truthfulness.

Now, recall that in profile $\mathbf{v}$, player 1 gets $Z \cup \{x\}$ (notice that he may get more items as well), and consider profiles $\mathbf{v}' = (v_1^{\mathrm{I}}, v_2)$ and $\mathbf{v}'' = (v_1^{\mathrm{II}}, v_2)$, where

	Z	x	$M \setminus (Z \cup \{x\})$
v_1^{I}	$-\ m^2\ -$	m	$-\ 1\ -$

and

	Z	x	$M \setminus (Z \cup \{x\})$
v_1^{II}	$-\ m\ -$	m^2	$-\ 1\ -$

By truthfulness, player 1 continues to get $Z \cup \{x\}$ in both cases, i.e., $Z \cup \{x\} \subseteq X_1^{N_1}(\mathbf{v}')$ and $Z \cup \{x\} \subseteq X_1^{N_1}(\mathbf{v}'')$.

We proceed by changing the values of player 2 this time. Assuming that $M \setminus (Z \cup \{x\}) = \{i_1, i_2, \ldots, i_\ell\}$ let $f_{i_j} = m$ if $i_j \in X_2(\mathbf{v}'')$ and $f_{i_j} = 1$ otherwise. Consider the next profile $\mathbf{u} = (v_1^{\mathrm{II}}, v_2^{\mathrm{I}})$:

	Z	x	$M \setminus (Z \cup \{x\})$
v_1^{II}	$-\ m\ -$	m^2	$-\ 1\ -$
v_2^{I}	$-\ 1\ -$	m^2	$f_{i_1}, \ldots, f_{i_\ell}$

Now notice that player 1 must get item x, since $x \in N_1$ and thus he controls $\{x\}$. On the other hand, since player 2 can not get x he must continue to get at least the items in $X_2(\mathbf{v}'')$ by truthfulness (otherwise he would play v_2 instead). Since this the case, he can not get a strict superset of $X_2(\mathbf{v}'')$ either. Indeed, if this was not the case he would deviate from $\mathbf{v}''$ to $\mathbf{u}$. So we can conclude that $X_2(\mathbf{u}) = X_2(\mathbf{v}'')$.

Now we move to a profile $\mathbf{u}' = (v_1^{\mathrm{I}}, v_2^{\mathrm{II}})$ where eventually player 2 gets item x:

	Z	x	$M \setminus (Z \cup \{x\})$
v_1^{I}	$-\ m^2\ -$	m	$-\ 1\ -$
v_2^{II}	$-\ m^2\ -$	m^2	$f_{i_1}, \ldots, f_{i_\ell}$

In $\mathbf{u}'$, both players strongly desire $Z \cup \{x\}$. But player 1 cannot get both set Z and item x, or by Lemma 3.1.8 he controls $Z \cup \{x\}$ and thus $Z \cup \{x\} \subseteq S$ for some $S \in \mathcal{O}_1$. However, he controls Z, since there exists some $S' \in \mathcal{O}_1$ such that $Z = S' \cap X_1^{N_1}(\mathbf{v}) \subseteq S'$. So, player 1 has to get Z since he strongly desires it, and item x is given to player 2 (probably with other items in $M \setminus (Z \cup \{x\})$.

Finally, consider our final profile $\mathbf{u}'' = (v_1^{\mathrm{I}}, v_2^{\mathrm{I}})$

	Z	x	$M \setminus (Z \cup \{x\})$
v_1^{II}	$-\ m^2\ -$	m	$-\ 1\ -$
v_2^{II}	$-\ 1\ -$	m^2	$f_{i_1}, \ldots, f_{i_\ell}$

By truthfulness, player 2 must get item x, or he would deviate from $\mathbf{u}''$ to $\mathbf{u}'$. However, now player 1 can strictly improve his utility by deviating from profile $\mathbf{u}''$ to $\mathbf{u}$, something that contradicts truthfulness. □

Proof of Lemma 3.1.12. Due to symmetry, it suffices to prove the statement for $i = 1$. If $N_1 = \emptyset$ then the statement is trivially true. So assume $N_1 \neq \emptyset$ and suppose, towards a contradiction, that the statement does not hold. That is, there exists a profile $\mathbf{v} = (v_1, v_2)$ such that $X_1^{N_1}(\mathbf{v}) \notin \arg\max_{S\in\mathcal{O}_1} v_1(S)$. We consider two cases, depending on whether $X_2^{N_2}(\mathbf{v})$ is in $\mathcal{O}_2$ or not. In both cases, we create a series of deviations that eventually contradict truthfulness. Like in the proof of Lemma 3.1.11, our first goal is to reach a profile $\mathbf{u}$ that contradicts the statement, like $\mathbf{v}$, but has appropriately selected values. Using $\mathbf{u}$ as a starting point we create profiles in which the allocations dictated by truthfulness are in conflict.

Case 1. Assume $X_2^{N_2}(\mathbf{v}) \in \mathcal{O}_2$ (note that this includes the case where $\mathcal{O}_2 = \{\emptyset\}$). Intuitively this is the case where the two players trade value between N_1 and E_2.

Consider the profile $\mathbf{v}' = (v_1, v_2^{\mathrm{I}})$, where

	$X_1^{N_1}(\mathbf{v})$	$X_2^{N_1}(\mathbf{v})$	$X_1^{N_2}(\mathbf{v})$	$X_2^{N_2}(\mathbf{v})$	E_1	$X_1^{E_2}(\mathbf{v})$	$X_2^{E_2}(\mathbf{v})$
v_2^{I}	— m —	— m^3 —	— m —	— m^4 —	— 1 —	— m^2 —	— m^4 —

By truthfulness, $X_2(\mathbf{v}') \supseteq X_2^{N_1}(\mathbf{v}) \cup X_2^{N_2}(\mathbf{v}) \cup X_2^{E_2}(\mathbf{v})$. This implies $X_2^{N_2}(\mathbf{v}') = X_2^{N_2}(\mathbf{v})$ due to the maximality of $X_2^{N_2}(\mathbf{v})$ and Lemma 3.1.11, as well as $X_1^{N_1}(\mathbf{v}') \subseteq X_1^{N_1}(\mathbf{v})$. The latter implies that $X_1^{N_1}(\mathbf{v}') \notin \arg\max_{S\in\mathcal{O}_1} v_1(S)$.

Claim A.0.2. $X_1^{E_2}(\mathbf{v}') \neq \emptyset$.

Proof of Claim A.0.2. Suppose $X_1^{E_2}(\mathbf{v}') = \emptyset$ and let $S' \in \arg\max_{S\in\mathcal{O}_1} v_1(S)$. Then player 1, whose total received value in $\mathbf{v}'$ would be strictly less than $v_1(S' \cup (X_1^{N_2}(\mathbf{v})) \cup E_1)$, could force the mechanism to give him at least that by playing

	S'	$N_1 \setminus S'$	N_2	E_1	E_2
v_1^{I}	— m —	— 1 —	— 1 —	— m —	— 1 —

By the definition of N_1, N_2, E_1, and Lemma 3.1.11, player 1 gets S', $N_2 \setminus X_2^{N_2}(\mathbf{v})$, and E_1 (and possibly something from E_2). Since this contradicts truthfulness, it must be the case that $X_1^{E_2}(\mathbf{v}') \neq \emptyset$. (In fact, this settles Case 1 when $E_2 = \emptyset$.) ◁

Next, let $S_1 \in \mathcal{O}_1$ be such that $X_1^{N_1}(\mathbf{v}') \subseteq S_1$ (they could possibly be equal). Consider the profile $\mathbf{u} = (v_1^{\mathrm{II}}, v_2^{\mathrm{I}})$, where

	$X_1^{N_1}(\mathbf{v}')$	$S_1 \setminus X_1^{N_1}(\mathbf{v}')$	$N_1 \setminus S_1$	N_2	$X_1^{E_1}(\mathbf{v}')$	$X_2^{E_1}(\mathbf{v}')$	$X_1^{E_2}(\mathbf{v}')$	$X_2^{E_2}(\mathbf{v}')$
v_1^{II}	— m —	— $1+m^{-1}$ —	— 1 —	— 1 —	— m^3 —	— 1 —	— m^2 —	— 1 —

Notice that S_1 is the unique set in $\arg\max_{S\in\mathcal{O}_1} v_1^{\mathrm{II}}(S)$. By truthfulness, $X_1(\mathbf{u}) \supseteq X_1^{N_1}(\mathbf{v}') \cup X_1^{E_1}(\mathbf{v}') \cup X_1^{E_2}(\mathbf{v}')$.

Claim A.0.3. *$S_1 \nsubseteq X_1(\mathbf{u})$, and therefore $X_1^{N_1}(\mathbf{u}) \notin \arg\max_{S \in \mathcal{O}_1} v_1^{\text{II}}(S)$.*

Proof of Claim A.0.3. Suppose $S_1 \subseteq X_1(\mathbf{u})$. By Lemma 3.1.11 this means $S_1 = X_1^{N_1}(\mathbf{u})$. Then player 2, whose total received value in $\mathbf{u}$ would be strictly less than $v_2^{\text{II}}((N_1 \setminus S_1) \cup X_2^{N_2}(\mathbf{v}') \cup X_2^{E_2}(\mathbf{v}')) + m$, could force the mechanism to give him more than that by playing

	N_1	$X_1^{N_2}(\mathbf{v}')$	$X_2^{N_2}(\mathbf{v}')$	E_1	E_2
v_2^{II}	— 1 —	— 1 —	— m —	— 1 —	— m —

By the definition of N_2, E_2, in $\mathbf{v}'' = (v_1^{\text{II}}, v_2^{\text{II}})$ player 2 gets $X_2^{N_2}(\mathbf{v}')$ and E_2 (and possibly something from N_1 and E_1). Given that, the maximum value that player 1 could achieve in $\mathbf{v}''$ is $v_1^{\text{II}}(S_1 \cup X_1^{N_2}(\mathbf{v}') \cup E_1)$ and there is no subset of $M \setminus (X_2^{N_2}(\mathbf{v}') \cup E_2)$ giving this value other than $S_1 \cup X_1^{N_2}(\mathbf{v}') \cup E_1$. In fact, player 1 can achieve exactly this by increasing his reported value for each item in $S_1 \cup E_1$ to m^3. Thus $X_1(\mathbf{v}'') = S_1 \cup X_1^{N_2}(\mathbf{v}') \cup E_1$ and $v_2^{\text{II}}(X_2(\mathbf{v}'')) = v_2^{\text{II}}((N_1 \setminus S_1) \cup X_2^{N_2}(\mathbf{v}') \cup E_2) \geq v_2^{\text{II}}((N_1 \setminus S_1) \cup X_2^{N_2}(\mathbf{v}') \cup X_2^{E_2}(\mathbf{v}')) + m^2$. Since this contradicts truthfulness, it must be the case that $S_1 \nsubseteq X_1(\mathbf{u})$ (and thus $X_1^{N_1}(\mathbf{u}) \notin \arg\max_{S \in \mathcal{O}_1} v_1^{\text{II}}(S)$). ◁

Claim A.0.3 implies that $S_1 \setminus X_1^{N_1}(\mathbf{u}) \neq \emptyset$. Since the sets in $\mathcal{O}_1$ have empty intersection, there must exist some $T \in \mathcal{O}_1$ such that $S_1 \setminus X_1^{N_1}(\mathbf{u}) \nsubseteq T$. We are going to concentrate most of player 2's value from N_1 on $W = (S_1 \setminus X_1^{N_1}(\mathbf{u})) \setminus T \subseteq X_2^{N_1}(\mathbf{u})$. Notice that $W \neq \emptyset$.

So consider the profile $\mathbf{u}' = (v_1^{\text{II}}, v_2^{\text{III}})$, where

	$N_1 \setminus W$	W	$X_1^{N_2}(\mathbf{v})$	$X_2^{N_2}(\mathbf{v})$	E_1	$X_1^{E_2}(\mathbf{v}')$	$X_2^{E_2}(\mathbf{v}')$
v_2^{III}	— m —	— m^3 —	— m —	— m^4 —	— 1 —	— m^2 —	— m^4 —

By the definition of N_2, E_2 and truthfulness, $X_2(\mathbf{u}') \supseteq W \cup X_2^{N_2}(\mathbf{v}) \cup X_2^{E_2}(\mathbf{v}')$.

Claim A.0.4. $X_1^{E_2}(\mathbf{u}') \neq \emptyset$.

Proof of Claim A.0.4. This proof is very similar to the proof of Claim A.0.2. Suppose $X_1^{E_2}(\mathbf{u}') = \emptyset$. Then player 1, whose total received value in $\mathbf{u}'$ would be strictly less than $v_1^{\text{II}}(S_1 \cup X_1^{N_2}(\mathbf{v}) \cup E_1)$, could force the mechanism to give him at least that by playing

	S_1	$N_1 \setminus S_1$	N_2	E_1	E_2
v_1^{III}	— m —	— 1 —	— 1 —	— m —	— 1 —

Since this contradicts truthfulness, it must be the case where $X_1^{E_2}(\mathbf{u}') \neq \emptyset$. ◁

Before we examine the final profile of the proof, let us consider the following simple profile $\mathbf{u}'' = (v_1^{\text{IV}}, v_2^{\text{IV}})$:

	T	$N_1 \setminus T$	$X_1^{N_2}(\mathbf{v})$	$X_2^{N_2}(\mathbf{v})$	E_1	$X_1^{E_2}(\mathbf{u}')$	$X_2^{E_2}(\mathbf{u}')$
v_1^{IV}	$-\ m\ -$	$-\ 1\ -$	$-\ 1\ -$	$-\ 1\ -$	$-\ 1\ -$	$-\ m^2\ -$	$-\ 1\ -$
v_2^{IV}	$-\ 1\ -$	$-\ 1\ -$	$-\ 1\ -$	$-\ m\ -$	$-\ 1\ -$	$-\ m\ -$	$-\ m\ -$

By the definition of N_2, E_2, in $\mathbf{u}''$ player 2 gets $X_2^{N_2}(\mathbf{v})$ and E_2 (and possibly something from N_1 and E_1). Given that, the maximum value that player 1 could achieve in $\mathbf{u}''$ is $|T| \cdot m + |X_1^{N_2}(\mathbf{v}) \cup E_1|$, and there is no subset of $M \setminus (X_2^{N_2}(\mathbf{v}) \cup E_2)$ giving this value other than $T \cup X_1^{N_2}(\mathbf{v}) \cup E_1$. In fact, player 1 can achieve exactly this by increasing his reported value for each item in $T \cup E_1$ to m^3. Thus $X_1(\mathbf{u}'') = T \cup X_1^{N_2}(\mathbf{v}) \cup E_1$ and $X_2(\mathbf{u}'') = (N_1 \setminus T) \cup X_2^{N_2}(\mathbf{v}) \cup E_2$.

The final profile we need is $\mathbf{u}''' = (v_1^{\mathrm{IV}}, v_2^{\mathrm{III}})$, and the contradiction follows from the allocation of the items in $X_1^{E_2}(\mathbf{u}')$. If $X_1^{E_2}(\mathbf{u}') \not\subseteq X_1(\mathbf{u}''')$ then player 1 has incentive to deviate to profile $\mathbf{u}' = (v_1^{\mathrm{II}}, v_2^{\mathrm{III}})$. So, it must be the case where $X_1^{E_2}(\mathbf{u}') \subseteq X_1(\mathbf{u}''')$, and therefore $v_2^{\mathrm{III}}(X_2(\mathbf{u}''')) \leq v_2^{\mathrm{III}}(M \setminus X_1^{N_2}(\mathbf{u}')) < v_2^{\mathrm{III}}(W \cup X_2^{N_2}(\mathbf{v}) \cup X_2^{E_2}(\mathbf{u}')) + m^2$. On the other hand, notice that $W \subseteq N_1 \setminus T$. Using the allocation for $\mathbf{u}''$ we derived above, by truthfulness we have that $v_2^{\mathrm{III}}(X_2(\mathbf{u}''')) \geq v_2^{\mathrm{III}}(W \cup X_2^{N_2}(\mathbf{v}) \cup E_2) \geq v_2^{\mathrm{III}}(W \cup X_2^{N_2}(\mathbf{v}) \cup X_2^{E_2}(\mathbf{u}')) + m^2$, which is a contradiction.

Case 2. Assume $X_2^{N_2}(\mathbf{v}) \notin \mathcal{O}_2$. Case 1 implies that not only $X_1^{N_1}(\mathbf{v}) \notin \arg\max_{S \in \mathcal{O}_1} v_1(S)$ but $X_1^{N_1}(\mathbf{v}) \notin \mathcal{O}_1$. Intuitively this is the case where the two players trade value between N_1 and N_2. The proof uses a sequence of profiles similar to Case 1.

Consider the profile $\mathbf{v}' = (v_1, v_2^{\mathrm{I}})$, where

	$X_1^{N_1}(\mathbf{v})$	$X_2^{N_1}(\mathbf{v})$	$X_1^{N_2}(\mathbf{v})$	$X_2^{N_2}(\mathbf{v})$	E_1	E_2
v_2^{I}	$-\ 1\ -$	$-\ m^2\ -$	$-\ m\ -$	$-\ m^3\ -$	$-\ 1\ -$	$-\ 1\ -$

By truthfulness, $X_2(\mathbf{v}') \supseteq X_2^{N_1}(\mathbf{v}) \cup X_2^{N_2}(\mathbf{v})$. This implies $X_1^{N_1}(\mathbf{v}') \subseteq X_1^{N_1}(\mathbf{v})$, and thus $X_1^{N_1}(\mathbf{v}') \notin \mathcal{O}_1$. By Case 1, this means that $X_2^{N_2}(\mathbf{v}') \notin \mathcal{O}_2$.

Next, let $S_1 \in \mathcal{O}_1$ be a minimal set of $\mathcal{O}_1$ such that $X_1^{N_1}(\mathbf{v}') \subseteq S_1$. Since $X_1^{N_1}(\mathbf{v}') \notin \mathcal{O}_1$, we have $X_1^{N_1}(\mathbf{v}') \subsetneq S_1$. Consider the profile $\mathbf{u} = (v_1^{\mathrm{I}}, v_2^{\mathrm{I}})$, where

	$X_1^{N_1}(\mathbf{v}')$	$S_1 \setminus X_1^{N_1}(\mathbf{v}')$	$N_1 \setminus S_1$	$X_1^{N_2}(\mathbf{v}')$	$X_2^{N_2}(\mathbf{v}')$	E_1	E_2
v_1^{I}	$-\ m\ -$	$-\ 1 + m^{-1}\ -$	$-\ 1\ -$	$-\ m\ -$	$-\ 1\ -$	$-\ 1\ -$	$-\ 1\ -$

Notice that S_1 is the unique set in $\arg\max_{S \in \mathcal{O}_1} v_1^{\mathrm{I}}(S)$. By truthfulness, $X_1(\mathbf{u}) \supseteq X_1^{N_1}(\mathbf{v}') \cup X_1^{N_2}(\mathbf{v}')$.

Claim A.0.5. *$S_1 \not\subseteq X_1(\mathbf{u})$, and therefore $X_1^{N_1}(\mathbf{u}) \notin \arg\max_{S \in \mathcal{O}_1} v_1^{\mathrm{I}}(S)$.*

Proof of Claim A.0.5. This is similar to the proof of Claim A.0.3.Suppose $S_1 \subseteq X_1(\mathbf{u})$. By Lemma 3.1.11 this means $S_1 = X_1^{N_1}(\mathbf{u})$. Then player 2, whose total received value in $\mathbf{u}$ would be strictly less than $v_2^{\mathrm{I}}(X_2^{N_2}(\mathbf{v}')) + m$, could force the mechanism to give him at least that by playing

	N_1	$N_2 \setminus S_2$	S_2	E_1	E_2
v_2^{II}	$-\ 1\ -$	$-\ 1\ -$	$-\ m\ -$	$-\ 1\ -$	$-\ m\ -$

where $S_2 \in \mathcal{O}_2$ is such that $X_2^{N_2}(\mathbf{v}') \subseteq S_2$. By the definition of N_2, E_2, in $\mathbf{v}'' = (v_1^{\text{I}}, v_2^{\text{II}})$ player 2 gets S_2 and E_2 (and possibly something from N_1 and E_1). Note, however, that $X_2^{N_2}(\mathbf{v}') \notin \mathcal{O}_2$ and thus $X_2^{N_2}(\mathbf{v}') \subsetneq S_2$. Therefore, $v_2^{\text{I}}(X_2(\mathbf{v}'')) \geq v_2^{\text{I}}(S_2) \geq v_2^{\text{I}}(X_2^{N_2}(\mathbf{v}')) + m$. Since this contradicts truthfulness, it must be the case that $S_1 \nsubseteq X_1(\mathbf{u})$ (and thus $X_1^{N_1}(\mathbf{u}) \notin \arg\max_{S \in \mathcal{O}_1} v_1^{\text{I}}(S)$). $\lhd$

This implies that $S_1 \setminus X_1^{N_1}(\mathbf{u}) \neq \emptyset$. Since the sets in $\mathcal{O}_1$ have empty intersection, there must exist some $T \in \mathcal{O}_1$ such that $S_1 \setminus X_1^{N_1}(\mathbf{u}) \nsubseteq T$. We are going to concentrate most of player 2's value from N_1 on $W = (S_1 \setminus X_1^{N_1}(\mathbf{u})) \setminus T \neq \emptyset$. So consider the profile $\mathbf{u}' = (v_1^{\text{I}}, v_2^{\text{III}})$, where

	$N_1 \setminus W$	W	$X_1^{N_2}(\mathbf{u})$	$X_2^{N_2}(\mathbf{u})$	E_1	E_2
v_2^{III}	$-\ 1\ -$	$-\ m^2\ -$	$-\ m\ -$	$-\ m^3\ -$	$-\ 1\ -$	$-\ 1\ -$

By truthfulness, $X_2(\mathbf{u}') \supseteq W \cup X_2^{N_2}(\mathbf{u})$. This implies that $S_1 \nsubseteq X_1(\mathbf{u}')$ and thus $X_1^{N_1}(\mathbf{u}') \notin \arg\max_{S \in \mathcal{O}_1} v_1^{\text{I}}(S)$. By Case 1, this means that $X_2^{N_2}(\mathbf{u}') \notin \mathcal{O}_2$. Therefore, $X_1^{N_2}(\mathbf{u}') \neq \emptyset$.

Now let $S_2' \in \mathcal{O}_2$ is such that $X_2^{N_2}(\mathbf{u}') \subsetneq S_2'$. Before we examine the final profile of the proof, let us consider the following profile $\mathbf{u}'' = (v_1^{\text{II}}, v_2^{\text{IV}})$:

	T	$N_1 \setminus T$	$N_2 \setminus S_2'$	$S_2' \setminus X_2^{N_2}(\mathbf{u}')$	$X_2^{N_2}(\mathbf{u}')$	E_1	E_2
v_1^{II}	$-\ m\ -$	$-\ 1\ -$	$-\ m^2\ -$	$-\ m^2\ -$	$-\ 1\ -$	$-\ 1\ -$	$-\ 1\ -$
v_2^{IV}	$-\ 1\ -$	$-\ 1\ -$	$-\ 1\ -$	$-\ m\ -$	$-\ m\ -$	$-\ 1\ -$	$-\ m\ -$

By the definition of N_2, E_2, in $\mathbf{u}''$ player 2 gets S_2' and E_2 (and possibly something from N_1 and E_1). Given that, the maximum value that player 1 could achieve in $\mathbf{u}''$ is $|T| \cdot m + |N_2 \setminus S_2'| \cdot m^2 + |E_1|$. In fact, player 1 can achieve exactly this by increasing his reported value for each item in $T \cup E_1$ to m^3. Thus $X_1(\mathbf{u}'') = T \cup (N_2 \setminus S_2') \cup E_1$ and $X_2(\mathbf{u}'') = (N_1 \setminus T) \cup S_2' \cup E_2$.

The final profile we need is $\mathbf{u}''' = (v_1^{\text{II}}, v_2^{\text{III}})$, and the contradiction follows from the allocation of the items in $X_1^{N_2}(\mathbf{u}')$. If $X_1^{N_2}(\mathbf{u}') \nsubseteq X_1(\mathbf{u}''')$ then player 1 has incentive to deviate to profile $\mathbf{u}' = (v_1^{\text{I}}, v_2^{\text{III}})$. So, it must be the case where $X_1^{N_2}(\mathbf{u}') \subseteq X_1(\mathbf{u}''')$ and therefore $v_2^{\text{III}}(X_2(\mathbf{u}''')) \leq v_2^{\text{III}}(M \setminus X_1^{N_2}(\mathbf{u}')) < |W| \cdot m^2 + |X_2^{N_2}(\mathbf{u})| \cdot m^3 + m$. On the other hand, notice that $W \subseteq N_1 \setminus T$ and recall that $X_2^{N_2}(\mathbf{u}) \subseteq X_2^{N_2}(\mathbf{u}') \subsetneq S_2'$. Using the allocation for $\mathbf{u}''$ we calculated above, by truthfulness we have that $v_2^{\text{III}}(X_2(\mathbf{u}''')) \geq v_2^{\text{III}}((N_1 \setminus T) \cup S_2') \geq |W| \cdot m^2 + |X_2^{N_2}(\mathbf{u})| \cdot m^3 + m$, which is a contradiction. $\square$

Proof of Lemma 3.1.13. Suppose that this not true. So there are profiles $\mathbf{v} = (v_1, v_2)$, $\mathbf{v}' = (v'_1, v'_2) \in \mathcal{V}_m^{\neq}$ such that $v_{ij} = v'_{ij}$ for all $i \in \{1,2\}$ and $j \in E_1 \cup E_2$, but $X_1^{E_1 \cup E_2}(\mathbf{v}) \neq X_1^{E_1 \cup E_2}(\mathbf{v}')$. In such a case, either $\mathbf{v} = (v_1, v_2), \hat{\mathbf{v}} = (v'_1, v_2)$, or $\hat{\mathbf{v}} = (v'_1, v_2), \mathbf{v}' = (v'_1, v'_2)$ is also a pair of profiles that violates the statement. Without loss of generality we assume that $\mathbf{v}, \hat{\mathbf{v}}$ is such a pair, and that $v_1(X_1^{E_1}(\mathbf{v})) > v_1(X_1^{E_1}(\hat{\mathbf{v}}))$. Now let $S_1, \hat{S}_1 \in \mathcal{O}_1$ be the single best offer in each case. If $S_1 = \hat{S}_1$ then player 1 would deviate from $\hat{\mathbf{v}}$ to $\mathbf{v}$ and strictly improve. So assume that $S_1 \neq \hat{S}_1$ and multiply the values in $E_1 \cup E_2$ for player 1 with a large enough constant K, so that $K\big(\hat{v}_1(X_1^{E_1}(\mathbf{v})) - \hat{v}_1(X_1^{E_1}(\hat{\mathbf{v}}))\big) > \hat{v}_1(N_1 \cup N_2)$.

Call $\mathbf{v}^* = (v_1^*, v_2)$ and $\hat{\mathbf{v}}^* = (v_1'^*, v_2)$ the new profiles and notice that they are still in $\mathcal{V}_m^{\neq}$. Also, it is easy to see that truthfulness implies $X_1(\mathbf{v}) = X_1(\mathbf{v}^*)$ and $X_1(\hat{\mathbf{v}}) = X_1(\hat{\mathbf{v}}^*)$. Indeed, by Lemma 3.1.12, we have $X_1^{N_1 \cup N_2}(\mathbf{v}) = X_1^{N_1 \cup N_2}(\mathbf{v}^*)$, and if it was the case where $X_1^{E_1 \cup E_2}(\mathbf{v}) \neq X_1^{E_1 \cup E_2}(\mathbf{v}^*)$, then player 1 would deviate from profile $\mathbf{v}$ to $\mathbf{v}^*$ or vice versa to strictly improve his utility. The same holds for $\hat{\mathbf{v}}$ to $\hat{\mathbf{v}}^*$.

Now, however, player 1 would deviate from $\hat{\mathbf{v}}^*$ to $\mathbf{v}^*$ in order to improve by at least $\hat{v}_1^*(X_1^{E_1}(\mathbf{v}^*)) - \hat{v}_1^*(X_1^{E_1}(\hat{\mathbf{v}}^*)) - \hat{v}_1^*(N_1 \cup N_2) = K\big(\hat{v}_1(X_1^{E_1}(\mathbf{v})) - \hat{v}_1(X_1^{E_1}(\hat{\mathbf{v}}))\big) - \hat{v}_1(N_1 \cup N_2) > 0$, and this contradicts truthfulness. □

Remark A.0.6. Since we are talking about $\mathcal{X}_E$ in many of the following proofs, it is correct to write $X_i^{E_1 \cup E_2}(\cdot)$, not $X_i(\cdot)$. For the sake of readability, though, we drop the superscript wherever it is not necessary. Similarly, in order to avoid the unnecessary use of extra symbols, we prove the statements for m items, although in Subsection 3.1.2 $\mathcal{X}_E$ is a mechanism on $\ell \leq m$ items.

Remark A.0.7. For most of the following proofs we need to construct profiles in $\mathcal{V}_m^{\neq}$. To facilitate the presentation, however, the valuation functions we construct only use a few powers of m. As a result, the corresponding profiles typically are not in $\mathcal{V}_m^{\neq}$. Still, this is without loss of generality; when defining such valuation functions we can add $2^i/2^\kappa$ to the value of item i, for $i \in [m]$. When $\kappa \in \mathbb{N}$ is large enough (usually $\kappa = m+1$ suffices), our arguments about the allocation are not affected, and a strict preference over all subsets is induced.

Proof of Lemma 3.1.14. Let $\mathbf{v} = (v_1, v_2) \in \mathcal{V}_m$, and consider the intermediate profile $\mathbf{v}^* = (v'_1, v_2)$ where $v'_{1x} = m$, if $x \in X_1(\mathbf{v})$, and $v'_{1x} = 1$ otherwise. By truthfulness, we have that $X_1(\mathbf{v}^*) = X_1(\mathbf{v})$. By defining v'_2 in a similar way (i.e., $v'_{2x} = m$, if $x \in X_2(\mathbf{v})$, and $v'_{2x} = 1$ otherwise), we get the profile $\mathbf{v}' = (v'_1, v'_2)$. Again by truthfulness, we have $\mathcal{X}(\mathbf{v}') = \mathcal{X}(\mathbf{v})$. If $\mathbf{v}^*$ and $\mathbf{v}'$ where defined as described in Remark A.0.7, the same arguments would apply, and moreover, $\mathbf{v}' \in \mathcal{V}_m^{\neq}$. □

Proof of Lemma 3.1.15. To show that D is indeed a valid set of exchange deals, we need to show that for any two distinct deals $(S,T),(S',T') \in D$ we have $S \cap S' = T \cap T' = \emptyset$ and S,T,S',T' are all nonempty. The latter is straightforward due to truthfulness and the fact that all values are positive. The former is done through the next three lemmata, the first of which states that each minimally exchangeable set is involved in exactly one exchange deal.

Lemma A.0.8. *If $S \subseteq E_1$ is a minimally exchangeable set, then there exists a unique $T \subseteq E_2$ such that (S, T) is a feasible exchange.*

The lemma is stated in terms of minimally exchangeable subsets of E_1, but due to symmetry it is true for all minimally exchangeable sets. This is done for the following statements as well, for the sake of readability. The three lemmata are proved right after this proof.

It is implied that every minimally exchangeable set appears in exactly one exchange deal in D. The second lemma, below, guarantees that minimally exchangeable sets can be exchanged only with minimally exchangeable sets.

Lemma A.0.9. *Let $S \subseteq E_1$ be a minimally exchangeable set and (S, T) be the only feasible exchange involving S. Then T is a minimally exchangeable set as well.*

The result of the two lemmata combined is that $D = \{(S_1, T_1), (S_2, T_2), \ldots, (S_k, T_k)\}$, where $S_1, \ldots, S_k, T_1, \ldots, T_k$ are all the minimally exchangeable sets and are all different from each other. What is still needed is that the intersection between any two minimally exchangeable sets is always empty. The third lemma states something stronger (that is indeed needed later in the proof of A.0.12), namely that the intersection between a minimally exchangeable set and any other exchangeable set is always empty, unless the latter contains the former.

Lemma A.0.10. *Let $S \subseteq E_1$ be a minimally exchangeable set and $S' \subseteq E_1$ be an exchangeable set such that $S' \cap S \neq \emptyset$. Then $S \subseteq S'$.*

If the intersection between any two minimally exchangeable sets was nonempty, then by Lemma A.0.10 one is contained in the other, which contradicts minimality. We can conclude that D is a valid set of exchange deals. □

Proof of Lemma A.0.8. Suppose that this does not hold. Without loss of generality, assume that there is some $S_1 \subseteq E_1$ and two profiles $\mathbf{v}^{\mathrm{I}} = (v_1^{\mathrm{I}}, v_2^{\mathrm{I}})$ and $\mathbf{v}^{\mathrm{II}} = (v_1^{\mathrm{II}}, v_2^{\mathrm{II}})$, such that $X_1^{E_1}(\mathbf{v}^{\mathrm{I}}) = E_1 \setminus S_1 = X_1^{E_1}(\mathbf{v}^{\mathrm{II}})$ and $X_1^{E_2}(\mathbf{v}^{\mathrm{I}}) = S_2 \neq S_2' = X_1^{E_2}(\mathbf{v}^{\mathrm{II}})$.

For the sake of readability, let $A = S_2 \setminus S_2'$, $B = S_2 \cap S_2'$, $C = S_2' \setminus S_2$, and $D = M \setminus (S_2 \cup S_2')$. Since $S_2 \neq S_2'$, either $A \neq \emptyset$ or $C \neq \emptyset$. Without loss of generality, suppose that $A \neq \emptyset$. Using this notation, $X_1(\mathbf{v}^{\mathrm{I}}) = (E_1 \setminus S_1) \cup A \cup B$ and $X_2(\mathbf{v}^{\mathrm{I}}) = S_1 \cup C \cup D$, while $X_1(\mathbf{v}^{\mathrm{II}}) = (E_1 \setminus S_1) \cup B \cup C$ and $X_2(\mathbf{v}^{\mathrm{II}}) = S_1 \cup A \cup D$.

We proceed to profile $\mathbf{v}^{\mathrm{III}} = (v_1^{\mathrm{I}}, v_2^{\mathrm{III}})$ by changing the values of player 2:

	$E_1 \setminus S_1$	S_1	A	B	C	D
v_2^{III}	— 1 —	— m^2 —	— 1 —	— 1 —	— m^3 —	— m^3 —

Since the most valuable items of player 2 are those which he was allocated in profile $\mathbf{v}^{\mathrm{I}}$, by truthfulness, he should still get them, but he should not get any other item. Thus $\mathcal{X}_E(\mathbf{v}^{\mathrm{III}}) = \mathcal{X}_E(\mathbf{v}^{\mathrm{I}})$.

We move to profile $\mathbf{v}^{\mathrm{IV}} = (v_1^{\mathrm{III}}, v_2^{\mathrm{III}})$ by changing the values of player 1:

	$E_1 \setminus S_1$	S_1	A	B	C	D
v_1^{III}	$- m^3 -$	$- m -$	$- m^2 -$	$- 1 -$	$- 1 -$	$- 1 -$

By truthfulness we have that player 1 must get $E_1 \setminus S_1$ and A (or else he could deviate to profile $\mathbf{v}^{\mathrm{III}}$ and strictly improve). Since he gets A, an exchange takes place. Due to the minimality of S_1, we can derive that player 2 receives the whole S_1. In addition, player 2 continues to get D, since he strongly desires it and $D \subseteq E_2$. So we can conclude that $(E_1 \setminus S_1) \cup A \subseteq X_1(\mathbf{v}^{\mathrm{IV}})$ and $S_1 \cup D \subseteq X_2(\mathbf{v}^{\mathrm{IV}})$, while we do not care about the allocation of the remaining items.

Now let us return to profile $\mathbf{v}^{\mathrm{II}} = (v_1^{\mathrm{II}}, v_2^{\mathrm{II}})$. Starting from here, we change the values of player 2 and to get profile $\mathbf{v}^{\mathrm{V}} = (v_1^{\mathrm{II}}, v_2^{\mathrm{IV}})$.

	$E_1 \setminus S_1$	S_1	A	B	C	D
v_2^{IV}	$- 1 -$	$- m -$	$- m^2 -$	$- 1 -$	$- 1 -$	$- m^2 -$

By truthfulness, like in profile $\mathbf{v}^{\mathrm{III}}$, we have $\mathcal{X}_E(\mathbf{v}^{\mathrm{V}}) = \mathcal{X}_E(\mathbf{v}^{\mathrm{II}})$.

Next, we proceed to profile $\mathbf{v}^{\mathrm{VI}} = (v_1^{\mathrm{IV}}, v_2^{\mathrm{IV}})$, where

	$E_1 \setminus S_1$	S_1	A	B	C	D
v_1^{IV}	$- m^4 -$	$- m -$	$- m^3 -$	$- m^2 -$	$- m^2 -$	$- 1 -$

Player 2 continues to get A, D since he strongly desires them and $A, D \subseteq E_2$. By the same argument, player 1 gets $E_1 \setminus S_1$. Additionally, we know that an exchange happens (otherwise player 1 would deviate to profile $\mathbf{v}^{\mathrm{V}}$ in order to get the items of $B \cup C$), so player 2 gets set the whole S_1 due to its minimality. Thus we can conclude that $X_1(\mathbf{v}^{\mathrm{VI}}) = (E_1 \setminus S_1) \cup B \cup C$ and $X_2(\mathbf{v}^{\mathrm{VI}}) = S_1 \cup A \cup D$.

Next, we move to profile $\mathbf{v}^{\mathrm{VII}} = (v_1^{\mathrm{IV}}, v_2^{\mathrm{V}})$ by changing player 2 this time:

	$E_1 \setminus S_1$	S_1	A	B	C	D
v_2^{V}	$- 1 -$	$- m^2 -$	$- m -$	$- 1 -$	$- 1 -$	$- m^3 -$

By truthfulness, the allocation does not change, i.e., $X_1(\mathbf{v}^{\mathrm{VII}}) = (E_1 \setminus S_1) \cup B \cup C$ and $X_2(\mathbf{v}^{\mathrm{VII}}) = S_1 \cup A \cup D$.

Finally, we move to profile $\mathbf{v}^{\mathrm{VIII}} = (v_1^{\mathrm{III}}, v_2^{\mathrm{V}})$ by changing the values of player 1 back to the values that he had in profile $\mathbf{v}^{\mathrm{IV}}$. Now recall that $X_2(\mathbf{v}^{\mathrm{IV}}) \supseteq S_1 \cup D$. Since in this profile $S_1 \cup D$ contains player 2's most valuable items, he must continue to get them by truthfulness. This means that there is an exchange. Player 1 however must get some items from A in any exchange; if not he can declare that he strongly desires E_1

and strictly improve. This, however, contradicts the truthfulness of the mechanism, since player 1 can deviate from $\mathbf{v}^{\text{VII}}$ to $\mathbf{v}^{\text{VIII}}$ and become strictly better. □

Proof of Lemma A.0.9. Suppose that this does not hold, i.e., there exists some minimally exchangeable $S_1 \in E_1$, such that (S_1, S_2) is the only feasible exchange involving S_1, but S_2 is not minimally exchangeable. So there exists $S_2' \subseteq S_2$ that is minimally exchangeable. So let S_1' be such that (S_1', S_2') is a feasible exchange (notice that $S_1 \neq S_1'$ by lemma A.0.8).

For the sake of readability, let $A = E_1 \setminus (S_1 \cup S_1')$, $B = S_1' \cap S_1$, $C = S_1 \setminus S_1'$, $D = S_2 \cap S_2'$, $E = S_2'$, and $F = S_2 \setminus S_2'$.

So there is a profile $\mathbf{v}^{\text{I}} = (v_1^{\text{I}}, v_2^{\text{I}})$, where $X_1(\mathbf{v}^{\text{I}}) = (E_1 \setminus S_1) \cup S_2 = A \cup B \cup E \cup F$ and $X_2(\mathbf{v}^{\text{I}}) = (E_2 \setminus S_2) \cup S_1 = C \cup D \cup (E_2 \setminus S_2)$. Also there is another profile $\mathbf{v}^{\text{II}} = (v_1^{\text{II}}, v_2^{\text{II}})$ where $X_1(\mathbf{v}^{\text{II}}) = (E_1 \setminus S_1') \cup S_2' = A \cup D \cup E$ and $X_2(\mathbf{v}^{\text{II}}) = (E_2 \setminus S_2') \cup S_1' = B \cup C \cup F \cup (E_2 \setminus S_2)$.

We start from profile $\mathbf{v}^{\text{I}} = (v_1^{\text{I}}, v_2^{\text{I}})$ and we proceed to profile $\mathbf{v}^{\text{III}} = (v_1^{\text{III}}, v_2^{\text{I}})$ by changing the values of player 1:

	A	B	C	D	E	F	$E_2 \setminus S_2$
v_1^{III}	— m^4 —	— m^4 —	— m —	— m —	— m^3 —	— m^2 —	— 1 —

Since player's 1 most valuable items are those he was allocated in profile $\mathbf{v}^{\text{I}}$, due to the truthfulness of the mechanism, he must continue to get them while not getting any other item. Thus the allocation does not change, i.e., $X_1(\mathbf{v}^{\text{III}}) = A \cup B \cup E \cup F$ and $X_2(\mathbf{v}^{\text{III}}) = C \cup D \cup (E_2 \setminus S_2)$.

Next, move to profile $\mathbf{v}^{\text{IV}} = (v_1^{\text{III}}, v_2^{\text{III}})$ by changing the values of player 2:

	A	B	C	D	E	F	$E_2 \setminus S_2$
v_2^{III}	— m —	— m^4 —	— m —	— m^3 —	— 1—	— m^2 —	— m^5 —

Player 2 must get $E_2 \setminus S_2$ since he strongly desires them and $E_2 \setminus S_2 \subseteq E_2$. Similarly, player 1 gets $A \cup B$. Moreover, we know that an exchange should take place (otherwise player 2 would deviate to $\mathbf{v}^{\text{III}}$ and become strictly better). What can be exchanged from E_1 is a subset of $C \cup D$, and since $C \cup D = S_1$ is minimal, it is exchanged with $S_2 = E \cup F$ (the only set that is exchangeable with S_1, by Lemma A.0.8). Thus we conclude that the allocation here is $X_1(\mathbf{v}^{\text{IV}}) = A \cup B \cup E \cup F$ and $X_2(\mathbf{v}^{\text{IV}}) = C \cup D \cup (E_2 \setminus S_2)$.

Finally we move to profile $\mathbf{v}^{\text{V}} = (v_1^{\text{IV}}, v_2^{\text{III}})$, by changing the values of player 1:

	A	B	C	D	E	F	$E_2 \setminus S_2$
v_1^{IV}	— m^4 —	— m^2 —	— m —	— m —	— m^3 —	— m^2 —	— 1 —

By truthfulness, like above, the allocation does not change, i.e., $X_1(\mathbf{v}^{\text{V}}) = A \cup B \cup E \cup F$ and $X_2(\mathbf{v}^{\text{V}}) = C \cup D \cup (E_2 \setminus S_2)$.

Now let us return to profile $\mathbf{v}^{\text{II}} = (v_1^{\text{II}}, v_2^{\text{II}})$. Starting from this profile we change the values of player 2 to get profile $\mathbf{v}^{\text{VI}} = (v_1^{\text{II}}, v_2^{\text{IV}})$.

	A	B	C	D	E	F	$E_2 \setminus S_2$
v_2^{IV}	— 1 —	— m^3 —	— 1 —	— 1 —	— m—	— m^2 —	— m^4 —

Player 2 must get (at least) $B \cup F \cup (E_2 \setminus S_2)$ since else he could deviate to profile $\mathbf{v}^{\text{II}}$ and become strictly better. Now since player 1 loses B we know that an exchange takes place with some of the available items in E. By the minimality of $E = S'_2$, player 1 gets the whole E and he loses $B \cup C = S'_1$. Thus we can conclude that the allocation here is $X_1(\mathbf{v}^{\text{VI}}) = A \cup D \cup E$, $X_2(\mathbf{v}^{\text{VI}}) = B \cup C \cup F \cup (E_2 \setminus S_2)$.

In order to conclude, we move to profile $\mathbf{v}^{\text{VII}} = (v_1^{\text{IV}}, v_2^{\text{IV}})$ by changing the values of player 1 back to what he played in $\mathbf{v}^{\text{V}}$,

	A	B	C	D	E	F	$E_2 \setminus S_2$
v_1^{IV}	— m^4 —	— m^2 —	— m —	— m —	— m^3 —	— m^2 —	— 1 —
v_2^{IV}	— 1 —	— m^3 —	— 1 —	— 1 —	— m—	— m^2 —	— m^4 —

Player 2 gets $E_2 \setminus S_2$ because he strongly desires it. We also know that an exchange should take place, otherwise player 1 would deviate to $\mathbf{v}^{\text{VI}}$ and strictly improve his total value. As a result, player 2 gets at least one item from set B, or he could increase to m^4 his value for any item in E_2 and improve by getting E_2. However, now player 2 can deviate from profile $\mathbf{v}^{\text{V}}$ to $\mathbf{v}^{\text{VII}}$ and become strictly better, something that contradicts the truthfulness of the mechanism. □

Proof of Lemma A.0.10. Suppose that this does not hold, i.e., there exists a minimally exchangeable set $S_1 \in E_1$ and an exchangeable set $S'_1 \in E_1$, such that $S_1 \cap S'_1 \neq \emptyset$ and $S_1 \not\subseteq S'_1$. Choose S'_1 to be minimal, i.e., if $S''_1 \subsetneq S'_1$ then either $S_1 \cap S''_1 = \emptyset$ or S''_1 is not exchangeable. Let S_2, S'_2 be such that (S_1, S_2), (S'_1, S'_2) are feasible exchanges and S'_2 is minimal in the sense that there is no $S''_2 \subsetneq S'_2$ where (S'_1, S''_2) being a feasible exchange. From Lemmata A.0.8 and A.0.9 we have that $S'_2 \setminus S_2 \neq \emptyset$.

For the sake of readability, let $A = E_1 \setminus (S_1 \cup S'_1)$, $B = S'_1 \setminus S_1$, $C = S'_1 \cap S_1$, $D = S_1 \setminus S'_1$, $E = S_2 \cap S'_2$, $F = S_2 \setminus S'_2$, $G = E_2 \setminus (S_2 \cup S'_2)$, and $H = S'_2 \setminus S_2$.

So there is a profile $\mathbf{v}^{\text{I}} = (E_1 \setminus S_1) \cup S_2 = (v_1^{\text{I}}, v_2^{\text{I}})$, where $X_1(\mathbf{v}^{\text{I}}) = A \cup B \cup E \cup F$, $X_2(\mathbf{v}^{\text{I}}) = (E_2 \setminus S_2) \cup S_1 = C \cup D \cup G \cup H$. There is also a profile $\mathbf{v}^{\text{II}} = (v_1^{\text{II}}, v_2^{\text{II}})$, where $X_1(\mathbf{v}^{\text{II}}) = (E_1 \setminus S'_1) \cup S'_2 = A \cup D \cup E \cup H$, $X_2(\mathbf{v}^{\text{II}}) = (E_2 \setminus S'_2) \cup S'_1 = B \cup C \cup F \cup G$.

We start from profile $\mathbf{v}^{\text{I}} = (v_1^{\text{I}}, v_2^{\text{I}})$ and we proceed to profile $\mathbf{v}^{\text{III}} = (v_1^{\text{I}}, v_2^{\text{III}})$ by changing the values of player 2:

	A	B	C	D	E	F	G	H
v_2^{III}	— 1 —	— 1 —	— m —	— m —	— 1 —	— 1 —	— m^2 —	— m^2 —

By truthfulness, we can conclude that the allocation remains the same, i.e., player 1 gets $A \cup B \cup E \cup F$, while player 2 gets $C \cup D \cup G \cup H$.

Next, we move to profile $\mathbf{v}^{\text{IV}} = (v_1^{\text{III}}, v_2^{\text{III}})$ by changing the values of player 1:

	A	B	C	D	E	F	G	H
v_1^{III}	m^3	m^3	1	1	m^2	m^2	1	m

Again, by truthfulness player 1 gets $A \cup B \cup E \cup F$, and player 2 gets $C \cup D \cup G \cup H$.

We continue by moving to profile $\mathbf{v}^{\text{V}} = (v_1^{\text{III}}, v_2^{\text{IV}})$ by changing the values of player 2:

	A	B	C	D	E	F	G	H
v_2^{IV}	1	m	m^3	1	1	m^2	m^4	1

Player 2 must get G since he strongly desires it and $H \subseteq E_2$. The same goes for player 1 and $A \cup B$. Now we know that an exchange should take place, otherwise player 2 would deviate to $\mathbf{v}^{\text{III}}$ and become strictly better. Since the only available exchangeable set here is $C \cup D = S_1$ (because it is minimal), it is exchanged with set $S_2 = E \cup F$ (the only set exchangeable with S_1 by lemma A.0.8). Thus we conclude that the allocation remains the same, player 1 gets $A \cup B \cup E \cup F$, while player 2 gets $C \cup D \cup G \cup H$.

Next proceed to profile $\mathbf{v}^{\text{VI}} = (v_1^{\text{IV}}, v_2^{\text{IV}})$ by changing the values of player 1:

	A	B	C	D	E	F	G	H
v_1^{IV}	m^4	m	1	1	m^2	m^3	1	m^2

we can derive by truthfulness that player 1 must get (at least) $A \cup F$, or else he would deviate to profile $\mathbf{v}^{\text{V}}$ and improve. Currently, this is all what we need to know for $\mathbf{v}^{\text{VI}}$.

Now let us return to profile $\mathbf{v}^{\text{II}} = (v_1^{\text{II}}, v_2^{\text{II}})$. Starting from here we change the values of player 1 to get profile $\mathbf{v}^{\text{VII}} = (v_1^{\text{V}}, v_2^{\text{II}})$.

	A	B	C	D	E	F	G	H
v_1^{V}	m^2	1	1	m^2	m	1	1	m

By truthfulness, the allocation remains the same, i.e., player 1 gets $A \cup D \cup E \cup H$, while player 2 gets $B \cup C \cup F \cup G$.

We now move to profile $\mathbf{v}^{\text{VIII}} = (v_1^{\text{V}}, v_2^{\text{V}})$ and change the values of player 2.

	A	B	C	D	E	F	G	H
v_2^{V}	1	αm^3	αm^4	1	m^4	m^5	m^5	m^4

The values in $B \cup C$ are set in such a way so that $v_2^{\mathrm{V}}(B \cup C) > v_2^{\mathrm{V}}(E \cup H)$, but $v_2^{\mathrm{V}}(T) < v_2^{\mathrm{V}}(E \cup H)$ for any $T \subsetneq B \cup C$.[1]

Notice that player 2 must get $G \cup F$ since he strongly desires it. The same goes for player 1 and $A \cup D$. We know that an exchange should take place, otherwise player 2 would deviate to $\mathbf{v}^{\mathrm{VII}}$ and improve. In this exchange, values are such that player 2 should get the whole S_1'. Thus we conclude that the allocation remains the same, i.e., player 1 gets $A \cup D \cup E \cup H$, while player 2 gets $B \cup C \cup F \cup G$.

We now move to profile $\mathbf{v}^{\mathrm{IX}} = (v_1^{\mathrm{VI}}, v_2^{\mathrm{V}})$ and change the values of player 1.

	A	B	C	D	E	F	G	H
v_1^{VI}	— m^4 —	— m —	— 1 —	— m^4 —	— m^2 —	— m^3 —	— 1 —	— m^2 —

Again player 2 must get $G \cup F$. Given that, player 1 gets at least $A \cup D \cup E \cup H$, and by truthfulness he cannot receive strictly more items. Therefore, the allocation remains the same, i.e., player 1 gets $A \cup D \cup E \cup H$, while player 2 gets $B \cup C \cup F \cup G$.

We now move to profile $\mathbf{v}^{\mathrm{X}} = (v_1^{\mathrm{IV}}, v_2^{\mathrm{V}})$ by changing the values of player 1 back to what he had at profile $\mathbf{v}^{\mathrm{VI}}$. Recall:

	A	B	C	D	E	F	G	H
v_1^{IV}	— m^4 —	— m —	— 1 —	— 1 —	— m^2 —	— m^3 —	— 1 —	— m^2 —

Like above Player 2 gets $F \cup G$. The same goes for player 1 A. By truthfulness, an exchange must happen and player 1 gets at least the set $E \cup H$ (else he would deviate to $\mathbf{v}^{\mathrm{IX}}$ and improve). Moreover, since player 2 loses $E \cup H$ he must at least get the set $B \cup C$. We conclude that player 1 gets $A \cup E \cup H$, player 2 gets $B \cup C \cup F \cup G$, while we do not care what happens for items in D.

Now notice that player 2 can deviate from profile $\mathbf{v}^{\mathrm{VI}}$ to profile $\mathbf{v}^{\mathrm{X}}$ and become strictly better (recall that at profile $\mathbf{v}^{\mathrm{VI}}$ player 2 loses G, while A, D, E, H all have very small value) and this contradicts truthfulness. □

Proof of Lemma 3.1.16. We begin with a direct implication of the Lemmata A.0.8–A.0.10. Although we are not guaranteed yet that any feasible exchange can be expressed as a union of exchange deals from D as it should, the following corollary is a step towards this direction. Recall that $S_1, \dots, S_k$ and $T_1, \dots, T_k$ are all the minimally exchangeable subsets of E_1 and E_2 respectively, and that (S_i, T_i) is the only feasible exchange involving either one of S_i and T_i, for every $i \in [k]$.

Corollary A.0.11. *For every exchangeable set $S \subseteq E_1$, we have that $S = W \cup \bigcup_{i \in I} S_i$, where $I \subseteq [k]$ with $|I| \geq 1$, while $W = S \setminus \bigcup_{i \in I} S_i$ does not contain any minimally exchangeable sets. Furthermore, this decomposition is unique.*

[1] This is always possible. In particular, if $|B| > 0$ then $\alpha = \frac{|E \cup H| m^4 - m}{(|B|-1)m^3 + |C| m^4}$ works. If $|B| = 0$, then $\alpha = \frac{|E \cup H| m^4 - m}{(|C|-1)m^4}$. In order to apply the idea mentioned in Remark A.0.7, one can multiply the whole profile with the denominator of α.

Ideally, we would like two things. First, the W part in the above decomposition to always be empty, i.e., we want every exchangeable set to be a *u*nion of *m*inimally *e*xchangeable *s*ets (*umes* for short). Second, we want every umes of E_1 to be exchangeable only with the corresponding umes of E_2, and vice versa. To be more precise, we say that an umes $S = \bigcup_{i\in I} S_i$ is *nice* if it is exchangeable with $T = \bigcup_{i\in I} T_i$ and only with T. The definition of a nice umes of E_2 is symmetric. As it turns out, every umes is nice, but it takes a rather involved induction to prove it. Especially the fact that $\left(\bigcup_{i\in I} S_i, \bigcup_{i\in I} T_i\right)$ is exchangeable needs a carefully constructed argument about the value that each player must gain from any exchange (see also Lemma A.0.14).

Lemma A.0.12. *Every umes is nice.*

Given the above lemma, we can now show that the set W in the decomposition of Corollary A.0.11 is always empty. In fact the proof idea is the same as the one for Lemmata A.0.8–A.0.10.

Lemma A.0.13. *Every exchangeable set is an umes.*

The above two lemmata complete the proof. They are proved below, right after Lemma A.0.14. □

For the following lemmata, recall that *umes* is short for union of minimally exchangeable sets!

Lemma A.0.14. *Let (S,T) be a feasible exchange such that S is a nice umes with the property that if $S' \subseteq S$ is exchangeable, then S' is a nice umes. In particular, let $S = \bigcup_{i\in[r]} S_i$, where S_i is minimally exchangeable for all $i \in [r]$. If $\mathbf{v}$ is a profile where (S_i, T_i) is favorable for all $i \in [r]$ then (S,T) gives a lower bound on the value gained from exchanges in profile $\mathbf{v}$ for each player.*

Proof of Lemma A.0.14. Due to symmetry, it suffices to prove the lower bound for player 1. Let $\mathbf{v} = (v_1, v_2)$ be a profile like in the statement, where the values are $v_{i1}, v_{i2}, \ldots, v_{im}$ for $i = 1,2.$. Since (S,T) is a feasible exchange, there exists a profile $\mathbf{v}^{\mathrm{I}} = (v_1^{\mathrm{I}}, v_2^{\mathrm{I}}) \in \mathcal{V}_m^{\neq}$ such that the exchange (S,T) takes place, i.e.,$X_1(\mathbf{v}^{\mathrm{I}}) = (E_1 \setminus S) \cup T$ and $X_2(\mathbf{v}^{\mathrm{I}}) = S \cup (E_2 \setminus T)$. Starting from this profile we will use a series of intermediate profiles in order to reach $\mathbf{v} = (v_1, v_2)$. Initially consider profile $\mathbf{v}^{\mathrm{II}} = (v_1^{\mathrm{II}}, v_2^{\mathrm{I}})$ where we change the values of player 1.

$$v_{1j}^{\mathrm{II}} = \begin{cases} \frac{m \cdot \max_i v_{1i}}{\min_i v_{1i}} \cdot v_{1j} & \text{if } j \in E_1 \setminus S \\ v_{1j} & \text{if } j \in S \cup T \\ \frac{\min_i v_{1i}}{m \cdot \max_i v_{1i}} \cdot v_{1j} & \text{otherwise} \end{cases}$$

In this profile each item in $E_1 \setminus S$ has a value which is higher from the sum of the values in all the other sets. On the other hand, items in $E_2 \setminus T$ have total value less than the value of a single item in the other sets.[2] Since this is the case, player 1 must get $E_1 \setminus S$ since he strongly desires it. In addition, an exchange must take place, or player 1 could deviate to profile $\mathbf{v}^{\mathrm{I}}$ and become strictly better. Thus an exchange takes place and must involve a subset S' of S. Now if S' was a proper subset of S, then it would be a nice umes, i.e., $S' = \bigcup_{j\in I} S_j, I \subsetneq [r]$, and it is exchanged only with $T' = \bigcup_{j\in I} T_j$. However, since exchanges $S_j, T_j, j \in [r]\setminus I$ are also favorable, player 1 would deviate to profile $\mathbf{v}^{\mathrm{I}}$ and become strictly better. Therefore, the exchange involves the whole S, and since S is a nice umes it should be exchanged with T. So the allocation here is $X_1(\mathbf{v}^{\mathrm{II}}) = (E_1 \setminus S) \cup T$, $X_2(\mathbf{v}^{\mathrm{II}}) = S \cup (E_2 \setminus T)$.

By moving to profile $\mathbf{v}^{\mathrm{III}} = (v_1^{\mathrm{II}}, v_2)$ where we change the values of player 2, we have that, once again, player 1 must get the items in $E_1 \setminus S$. Moreover, an exchange must take place, or player 2 could deviate to profile $\mathbf{v}^{\mathrm{II}}$ and become strictly better (recall that he prefers S from T). By following the same arguments as in the previous case, if the exchange involves a proper subset of S, player 2 would deviate to profile $\mathbf{v}^{\mathrm{II}}$ and become strictly better. Hence player 2 gets the whole S, i.e., the allocation here is again $X_1(\mathbf{v}^{\mathrm{III}}) = (E_1 \setminus S) \cup T$ and $X_2(\mathbf{v}^{\mathrm{III}}) = S \cup (E_2 \setminus T)$.

Finally we move to profile $\mathbf{v} = (v_1, v_2)$ by changing the values of player 1. It is easy to see that if there is no exchange that improves player 1 by at least $v_1(T) - v_1(S)$, then he could deviate to profile $\mathbf{v}^{\mathrm{III}} = (v_1^{\mathrm{II}}, v_2)$ and gain exactly that. □

Proof of Lemma A.0.12. We will use induction in the number of minimally exchangeable sets contained in an umes; let us call this number *index* of the umes. Lemmata A.0.8 and A.0.9 imply that every umes of index 1 is nice. That is the basis of our induction.

Assume that every umes of index lower or equal to k is nice and notice that Lemma A.0.10 implies that every exchangeable subset of an umes is also an umes.

Let S be an umes of index $k+1$. In particular, let $S = \bigcup_{i\in[k+1]} S_i$, where for any $i \in [k+1]$ we have that S_i is minimally exchangeable and (S_i, T_i) is a feasible exchange. By the inductive hypothesis we have that both S_1 and $S' = \bigcup_{i=2}^{k+1} S_i$ are nice umes and uniquely exchangeable with S_1 and $T' = \bigcup_{i=2}^{k+1} T_i$ respectively.

We first prove that (S, T) is a feasible exchange. Consider the following profile $\mathbf{v} = (v_1, v_2)$,

	$E_1 \setminus (S' \cup S_1)$	S'	S_1	T'	T_1	$E_2 \setminus (T' \cup T_1)$
v_1	$-\ \Delta\ -$	$-\ \delta\ -$	$-\ \epsilon\ -$	$-\ 1\ -$	$-\ \zeta\ -$	$-\ \delta\ -$
v_2	$-\ \delta\ -$	$-\ n_j\ -$	$-\ 1\ -$	$-\ \theta_j\ -$	$-\ \delta -$	$-\ \Delta\ -$

where $\Delta >> 1 >> \zeta, n_j, \theta_j, \epsilon >> \delta >> \lambda_i$.[3] Regarding the rest values, $|T_1| \cdot \zeta = |S_1| \cdot \epsilon + \lambda_1$ and for all $j \in [k+1]\setminus\{1\}$ we have that $|S_j| \cdot n_j = |T_j| \cdot \theta_j + \lambda_j$. Now notice that S' is a

[2] Notice that the values are chosen in a way such that if $\mathbf{v} \in \mathcal{V}_m^{\neq}$, then $\mathbf{v}^{\mathrm{I}} \in \mathcal{V}_m^{\neq}$ as well.

[3] In order to be able to apply the idea mentioned in Remark A.0.7, one can use m^7 instead of 1, and $\Delta = m^8$, $\delta = m^3$, $\lambda_i = |T_i| \cdot |S_i|$, $n_i = |T_j| \cdot m^4$, $\theta_i = |S_j| \cdot (m^4 - 1)$, $\zeta = |S_1| \cdot m^4$, and $\epsilon = |T_1| \cdot (m^4 - 1)$.

nice umes such that every exchangeable $S'' \subseteq S'$ is a nice umes and for all j, (S_j, T_j) is a favorable exchange with respect to $\mathbf{v}$. Lemma A.0.14 guarantees that in $\mathbf{v}$, player 1 gains at least $v_1(T') - v_1(S') = |T'| - \delta|S'|$ from the exchanges. So player 1 gets a superset of T', i.e., $T' \subseteq X_1^{E_2}(\mathbf{v})$. By lemma A.0.10, this means that $X_1^{E_2}(\mathbf{v})$ is either T' or T.

On the other hand, if we apply lemma A.0.14 for (S_1, T_1) we have that in profile $\mathbf{v}$, player 2 should gain at least $v_2(S_1) - v_2(T_1) = |S_1| - \delta|T_1|$ from the exchanges. So, $X_2^{E_1}(\mathbf{v}) \supseteq S_1$. Since T' is nice, however, we have that $X_1^{E_2}(\mathbf{v}) = T'$ implies $X_2^{E_1}(\mathbf{v}) = S' \not\supseteq S_1$. Therefore, it must be the case where $X_1^{E_2}(\mathbf{v}) \supsetneq T'$ or else player 2 does not get enough value.

We conclude that $X_2^{E_1}(\mathbf{v}) = T$. Now we claim that $X_2^{E_1}(\mathbf{v}) = S$ and therefore (S, T) is a feasible exchange. Indeed, every $S'' \subsetneq S$ that is exchangeable is an umes of index lower or equal to k and therefore is nice. So S'', T cannot be a feasible exchange, due to the fact that S'' has a unique pair $T'' \subsetneq T$.

Next we show that there is no $\hat{T} \neq T$ such that $(S, \hat{T})$ is a feasible exchange. By the proof so far we have that if such a $\hat{T}$ existed, then it is not a subset of T. So suppose that there is a $\hat{T} \neq T$ such that $(S, \hat{T})$is a feasible exchange and let T^* be a minimal such set (that is, if $R \subsetneq T^*$ then (S, R) is not a feasible exchange or $R \subseteq T$).

Thus there are two profiles $\mathbf{v}^{\mathrm{I}} = (v_1^{\mathrm{I}}, v_2^{\mathrm{I}})$ and $\mathbf{v}^{\mathrm{II}} = (v_1^{\mathrm{II}}, v_2^{\mathrm{II}})$ where we have that $X^{E_1}(\mathbf{v}^{\mathrm{I}}) = (E_1 \setminus S) = X_1^{E_1}(\mathbf{v}^{\mathrm{II}})$ and $X_1^{E_2}(\mathbf{v}^{\mathrm{I}}) = T^* \neq T = X_1^{E_2}(\mathbf{v}^{\mathrm{II}})$.

For the sake of readability, let $A = T^* \setminus T$, $B = T^* \cap T$, $C = T \setminus T^*$, $D = E_2 \setminus (T^* \cup T)$.

We start from profile $\mathbf{v}^{\mathrm{I}}$ where the allocation is $X_1(\mathbf{v}^{\mathrm{I}}) = (E_1 \setminus S) \cup A \cup B$, $X_2(\mathbf{v}^{\mathrm{I}}) = S \cup C \cup D$ and we proceed to profile $\mathbf{v}^{\mathrm{III}} = (v_1^{\mathrm{I}}, v_2^{\mathrm{III}})$ by changing the values of player 2:

	$E_1 \setminus S$	S	A	B	C	D
v_2^{III}	— 1 —	— m —	— 1 —	— 1 —	— m^2 —	— m^2 —

By truthfulness, the allocation remains the same, i.e., $X_1(\mathbf{v}^{\mathrm{III}}) = (E_1 \setminus S) \cup A \cup B$, $X_2(\mathbf{v}^{\mathrm{III}}) = S \cup C \cup D$.

Next we move to profile $\mathbf{v}^{\mathrm{IV}} = (v_1^{\mathrm{III}}, v_2^{\mathrm{III}})$ by changing the values of player 1:

	$E_1 \setminus S$	S	A	B	C	D
v_1^{III}	— m^3 —	— m —	— m^2 —	— 1 —	— 1 —	— 1 —

Notice that player 1 must receive $E_1 \setminus S$ since he strongly desires it. The same goes for player 2 and $C \cup D$. Now we know that an exchange should take place and that in this exchange player 1 must get at least set $A = T^* \setminus T$ (otherwise he would deviate to $\mathbf{v}^{\mathrm{III}}$ and become strictly better).

We claim that player 1 gets the whole T^*. If this was not the case then he would get some set $R \supseteq A \neq \emptyset$. Since $R \subsetneq T^*$ and $R \not\subseteq T$ we have that the exchange (S, R) is not feasible due to the minimality of T^*. Thus R is exchanged with some $\hat{S} \subsetneq S$. However $\hat{S}$ is an umes (by Lemma A.0.8) and by inductive hypothesis it is exchangeable only with strict subsets of T which is a contradiction. Similarly, player 2 must get set the whole S, or otherwise he would get some $\hat{S} \subsetneq S$ which is exchangeable only with strict subsets of T, something that can not happen. Thus the allocation here is $X_1(\mathbf{v}^{\mathrm{IV}}) = (E_1 \setminus S) \cup A \cup B$, $X_2(\mathbf{v}^{\mathrm{IV}}) = S \cup C \cup D$.

Next we move to profile $\mathbf{v}^{\mathrm{V}} = (v_1^{\mathrm{III}}, v_2^{\mathrm{IV}})$ by changing the values of player 2.

	$E_1 \setminus S$	S	A	B	C	D
v_2^{IV}	— 1 —	— m^2 —	— 1 —	— 1 —	— m —	— m^3 —

By truthfulness, the allocation remains the same, i.e., $X_1(\mathbf{v}^{\mathrm{V}}) = (E_1 \setminus S) \cup A \cup B$, $X_2(\mathbf{v}^{\mathrm{V}}) = S \cup C \cup D$.

Now let us return to profile $\mathbf{v}^{\mathrm{II}} = (\mathbf{v}_1^{\mathrm{II}}, \mathbf{v}_2^{\mathrm{II}})$. Starting from this profile we change the values of player 2 and get profile $\mathbf{v}^{\mathrm{VI}} = (v_1^{\mathrm{II}}, v_2^{\mathrm{V}})$.

	$E_1 \setminus S$	S	A	B	C	D
v_2^{V}	— 1 —	— m —	— m^2 —	— 1 —	— 1 —	— m^2 —

Since player's 2 most valuable items are those which he was allocated in profile $\mathbf{v}^{\mathrm{II}}$, by truthfulness, the allocation remains the same, i.e., $X_1(\mathbf{v}^{\mathrm{VI}}) = (E_1 \setminus S) \cup B \cup C$, $X_2(\mathbf{v}^{\mathrm{VI}}) = S \cup A \cup D$.

Next we move to profile $\mathbf{v}^{\mathrm{VII}} = (v_1^{\mathrm{IV}}, v_2^{\mathrm{V}})$ by changing the values of player 1.

	$E_1 \setminus S$	S	A	B	C	D
v_1^{IV}	— m^4 —	— m —	— m^3 —	— m^2 —	— m^2 —	— 1 —

Notice that player 1 must get $E_1 \setminus S$ since he strongly desires it. The same goes for player 2 and $A \cup D$. Given that, an exchange takes place and in this exchange player 1 must get the whole $B \cup C = T$ (otherwise he would deviate to $\mathbf{v}^{\mathrm{V}}$ and strictly improve). On the other hand, player 2 must get the whole S, or player 1 would deviate from $\mathbf{v}^{\mathrm{V}}$ to $\mathbf{v}^{\mathrm{VI}}$ and strictly improve. Thus the allocation here remains the same: $X_1(\mathbf{v}^{\mathrm{VII}}) = (E_1 \setminus S) \cup B \cup C$, $X_2(\mathbf{v}^{\mathrm{VII}}) = S \cup A \cup D$.

Next we move to profile $\mathbf{v}^{\mathrm{VIII}} = (v_1^{\mathrm{IV}}, v_2^{\mathrm{VI}})$ by changing the values of player 2.

	$E_1 \setminus S$	S	A	B	C	D
v_2^{VI}	— 1 —	— m^2 —	— m —	— 1 —	— 1 —	— m^3 —

By truthfulness, the allocation remains the same, i.e., $X_1(\mathbf{v}^{\mathrm{VIII}}) = (E_1 \setminus S) \cup B \cup C$, $X_2(\mathbf{v}^{\mathrm{VIII}}) = S \cup A \cup D$.

Finally we move to profile $\mathbf{v}^{\mathrm{IX}} = (v_1^{\mathrm{III}}, v_2^{\mathrm{VI}})$ by changing the values of player 1 back to what he had in profile $\mathbf{v}^{\mathrm{V}}$. Recall:

	$E_1 \setminus S$	S	A	B	C	D
v_1^{III}	— m^3 —	— m —	— m^2 —	— 1 —	— 1 —	— 1 —

Notice that player 1 must get $E_1 \setminus S$ since he strongly desires it. The same goes for player 2 and D. Now if player 1 gets nothing from set A then there is no exchange at all. However, in this case player 2 would deviate to profile $\mathbf{v}^{\mathrm{V}}$ and become strictly better. Thus player 1 should get at least one item from A. As a result, however, player 1 would deviate from profile $\mathbf{v}^{\mathrm{VIII}}$ to $\mathbf{v}^{\mathrm{IX}}$ and become strictly better, something that leads to contradiction.

This completes the inductive step. □

Proof of Lemma A.0.13. Let S be an exchangeable subset of E_1. Then according to corollary A.0.11 $S = \bigcup_{i\in I} S_i \cup W$ for some $I \subseteq [k]$, with $|I| \geq 1$. We are going to show that $W = \emptyset$. So suppose, towards a contradiction, that $W \neq \emptyset$. In fact, choose S so that it is a minimal exchangeable non-umes subset of E_1, i.e., for all $S' \subsetneq S$, S' is either umes or non-exchangeable. In addition, notice that W does not contain any exchangeable sets.

Let T be such that (S,T) is a feasible exchange. In fact let T be a minimal such set, i.e., for all $T' \subsetneq T$, either (S,T') is not a feasible exchange or T' is not exchangeable at all. Finally, let $S^* = \bigcup_{i\in I} S_i$, $T^* = \bigcup_{i\in I} T_i$ and notice that $T \setminus T^* \neq \emptyset$ since otherwise T would be an umes (as an exchangeable subset of an umes, by Lemma A.0.10).

For the sake of readability, let $A = E_1 \setminus S$, $B = T \setminus T^*$, $C = T^* \cap T$, $D = T^* \setminus T$, and $E = E_2 \setminus (T \cup T^*)$.

So there are two profiles, $\mathbf{v}^{\mathrm{I}} = (v_1^{\mathrm{I}}, v_2^{\mathrm{I}})$ where $X_1(\mathbf{v}^{\mathrm{I}}) = A \cup B \cup C$, $X_2(\mathbf{v}^{\mathrm{I}}) = S \cup D \cup E$ and $\mathbf{v}^{\mathrm{II}} = (v_1^{\mathrm{II}}, v_2^{\mathrm{II}})$ where $X_1(\mathbf{v}^{\mathrm{II}}) = A \cup W \cup C \cup D$ and $X_2(\mathbf{v}^{\mathrm{II}}) = S^* \cup B \cup E$.

We start from profile $\mathbf{v}^{\mathrm{I}} = (v_1^{\mathrm{I}}, v_2^{\mathrm{I}})$ and we proceed to profile $\mathbf{v}^{\mathrm{III}} = (v_1^{\mathrm{I}}, v_2^{\mathrm{III}})$ by changing the values of player 2:

	A	S^*	W	B	C	D	E
v_2^{III}	— 1 —	— m^2 —	— m^2 —	— 1 —	— 1 —	— m^3 —	— m^3 —

Since player's 2 most valuable items are those which he was allocated in profile $\mathbf{v}^{\mathrm{I}}$, by truthfulness, the allocation remains the same, i.e., $X_1(\mathbf{v}^{\mathrm{III}}) = A \cup B \cup C$, $X_2(\mathbf{v}^{\mathrm{III}}) = S \cup D \cup E$.

Next we move to profile $\mathbf{v}^{\mathrm{IV}} = (v_1^{\mathrm{III}}, v_2^{\mathrm{III}})$ by changing the values of player 1:

	A	S^*	W	B	C	D	E
v_1^{III}	— m^3 —	— m —	— m —	— m^2 —	— 1 —	— 1 —	— 1 —

Now notice that player 1 must get A since he strongly desires it. The same goes for player 2 and $D \cup E$. Also we know that an exchange should take place and that in this exchange player 1 must get at least $B = T \setminus T^*$ (otherwise he would deviate to $\mathbf{v}^{\mathrm{III}}$).

We claim that player 1 gets the whole T. If this was not the case then he would get some set $R \supseteq B \neq \emptyset$. Since $R \subsetneq T$ and $R \nsubseteq T^*$ we have that the exchange (S,R) is not feasible due to the fact that T is minimal. Thus R should be exchanged with some $\hat{S} \subsetneq S$. However, by the minimality of S, $\hat{S}$ is an umes and it is exchangeable only with strict subsets of T^*, which is a contradiction. On the other hand, player 2 must get the whole S, or otherwise he would get some $\hat{S} \subsetneq S$ which is exchangeable only with strict subsets of T^*, something that can not happen. Thus the allocation is $X_1(\mathbf{v}^{\mathrm{IV}}) = A \cup B \cup C$ and $X_2(\mathbf{v}^{\mathrm{IV}}) = S \cup D \cup E$.

Next we move to profile $\mathbf{v}^{\mathrm{V}} = (v_1^{\mathrm{III}}, v_2^{\mathrm{IV}})$ by changing the values of player 2:

	A	S^*	W	B	C	D	E
v_2^{III}	— 1 —	— m^2 —	— m —	— 1 —	— 1 —	— m^2 —	— m^3 —

Since player's 2 most valuable items are those which he was allocated in profile $\mathbf{v}^{\mathrm{IV}}$, by the truthfulness of the mechanism, he must continue to get them but he can not get any other item. Thus the allocation remains the same, i.e., $X_1(\mathbf{v}^{\mathrm{V}}) = A \cup B \cup C$, $X_2(\mathbf{v}^{\mathrm{V}}) = S \cup D \cup E$.

Now let us return to profile $\mathbf{v}^{\mathrm{II}} = (v_1^{\mathrm{II}}, v_2^{\mathrm{II}})$. Starting from this profile we change the values of player 2 and get profile $\mathbf{v}^{\mathrm{VI}} = (v_1^{\mathrm{II}}, v_2^{\mathrm{IV}})$:

	A	S^*	W	B	C	D	E
v_2^{IV}	— 1 —	— m^2 —	— m —	— m^4 —	— 1 —	— 1 —	— m^4 —

Again, player's 2 most valuable items are those which he was allocated in profile $\mathbf{v}^{\mathrm{II}}$. So, by truthfulness, the allocation remains the same, i.e., $X_1(\mathbf{v}^{\mathrm{VI}}) = A \cup W \cup C \cup D$ and $X_2(\mathbf{v}^{\mathrm{VI}}) = S^* \cup B \cup E$.

Next we move to profile $\mathbf{v}^{\mathrm{VII}} = (v_1^{\mathrm{IV}}, v_2^{\mathrm{IV}})$ by changing the values of player 1:

	A	S^*	W	B	C	D	E
v_1^{IV}	— m^5 —	— m —	— m^2 —	— m^4 —	— m^3 —	— m^3 —	— 1 —

Notice that player 1 must get A and player 2 must get $B \cup E$. Given that, player 1 must get $W \cup C \cup D =$ (otherwise he could deviate to $\mathbf{v}^{\mathrm{V}}$ and strictly improve). Thus the allocation remains the same, i.e., $X_1(\mathbf{v}^{\mathrm{VII}}) = A \cup W \cup C \cup D$ and $X_2(\mathbf{v}^{\mathrm{VII}}) = S^* \cup B \cup E$.

Next we move to profile $\mathbf{v}^{\mathrm{VIII}} = (v_1^{\mathrm{IV}}, v_2^{\mathrm{V}})$ by changing the values of player 2:

	A	S^*	W	B	C	D	E
v_2^{V}	— 1 —	— m^3 —	— m —	— m^2 —	— 1 —	— 1 —	— m^4 —

Again, by truthfulness, the allocation remains the same, i.e., $X_1(\mathbf{v}^{\mathrm{VIII}}) = A \cup W \cup C \cup D$ and $X_2(\mathbf{v}^{\mathrm{VIII}}) = S^* \cup B \cup E$.

Finally we move to profile $\mathbf{v}^{\mathrm{IX}} = (v_1^{\mathrm{III}}, v_2^{\mathrm{V}})$ by changing the values of player 1 back to what he had in profile $\mathbf{v}^{\mathrm{V}}$.

	A	S^*	W	B	C	D	E
v_1^{III}	— m^3 —	— m —	— m —	— m^2 —	— 1 —	— 1 —	— 1 —

Player 1 must get A and player 2 must get E. Now if player 1 gets nothing from B then there will be no exchange. However, in this case player 2 would deviate to profile $\mathbf{v}^{\text{V}}$ and become strictly better. Thus player 1 should get at least one item from B. As a result, player 1 would deviate from profile $\mathbf{v}^{\text{VIII}}$ to $\mathbf{v}^{\text{IX}}$ and strictly improve, something that leads to contradiction. □

Proof of Lemma 3.1.17. Without loss of generality, assume that $(S_1, T_1), \ldots, (S_r, T_r)$ is the set of all favorable exchanges. Then (S, T) where $S = \bigcup_{i \in [r]} S_i$ and $T = \bigcup_{i \in [r]} T_i$ will give a lower bound on the value of each player. Indeed, S is an umes an using Lemmata A.0.14 and A.0.12, we have that player 1 should gain at least $v_1(T) - v_1(S)$, while player 2 should gain at least $v_2(S) - v_2(T)$ from the exchanges.

Since $\mathbf{v} \in \mathcal{V}_m^{\neq}$, it suffices to show that $v_1(X_1(\mathbf{v})) = v_1((E_1 \cup T) \backslash S) = v_1(E_1) + v_1(T) - v_1(S)$. So suppose that $v_1(X_1(\mathbf{v})) > v_1(E_1) + v_1(T) - v_1(S)$ and notice that this also implies that $v_2(X_2(\mathbf{v})) > v_2(E_2) + v_2(S) - v_2(T)$, since otherwise it would be $v_2(X_2(\mathbf{v})) = v_2(E_2) + v_2(S) - v_2(T)$ and we have the desired allocation.

As a result, there exists some $S^* \subseteq X_2^{E_1}(\mathbf{v})$, such that S^* is an umes but (S^*, T^*)—where T^* is the "pair" of S^*—is unfavorable. Without loss of generality, we may assume that $v_1(T^*) < v_1(S^*)$. Now let S' to be the union of all minimally exchangeable sets $S_j \subseteq X_2^{E_1}(\mathbf{v})$ such that $v_1(T_j) < v_1(S_j)$, and notice that $S' \subsetneq X_2^{E_1}(\mathbf{v})$ and $v_1(T') < v_1(S')$.

Let $S^* = X_2^{E_1}(\mathbf{v})$ and $T^* = X_1^{E_2}(\mathbf{v})$. We begin with profile $\mathbf{v} = (v_1, v_2)$ where the allocation is $X_1(\mathbf{v}) = (E_1 \setminus S^*) \cup T^*$ and $X_2(\mathbf{v}) = (E_2 \setminus T^*) \cup S^*$ and we move to profile $\mathbf{v}' = (v_1, v_2')$.

	$E_1 \setminus S^*$	S^*	T^*	$E_2 \setminus T^*$
v_2'	— 1 —	— m —	— 1 —	— m—

By truthfulness, the allocation remains the same, i.e., $X_1(\mathbf{v}') = (E_1 \backslash S^*) \cup T^*$ and $X_2(\mathbf{v}') = (E_2 \setminus T^*) \cup S^*$.

However, now notice that $(S^* \setminus S', T^* \setminus T')$ is a favorable exchange with respect to $\mathbf{v}'$. Moreover, for every minimally exchangeable set $S_i \subseteq S^* \setminus S'$ it holds that (S_i, T_i) is favorable. By using lemma A.0.14 we have that the gain from the exchange in $\mathbf{v}'$ for player 1 must be at least $v_1(T^* \setminus T') - v_1(S^* \setminus S') > v_1(T^*) - v_1(S^*)$ so we arrive at a contradiction. □

Proof of Lemma 3.1.18. Let $\mathbf{v} = (v_1, v_2)$ be a profile in $\mathcal{V}_m$. By Lemmata 3.1.14 and A.0.13, we know that $X_1^{E_1 \cup E_2}(\mathbf{v})$ is the result of some exchanges of D taking place, i.e., $X_1^{E_1 \cup E_2}(\mathbf{v}) = (E_1 \setminus \bigcup_{i \in I} S_i) \cup \bigcup_{i \in I} T_i$, where $I \subseteq [k]$. There are two things that can go wrong: either there exists some $x \in I$ such that (S_x, T_x) is unfavorable, or there exists some $z \in [k] \setminus I$ such that (S_z, T_z) is favorable. We first examine the former case.

Without loss of generality, we may assume that $v_1(T_x) < v_1(S_x)$. Consider the profile $\mathbf{v}' = (v_1, v_2^{\text{I}})$ where

$$v^{\text{I}}_{2j} = \begin{cases} m + 2^{j-m-1} & \text{if } j \in X_2(\mathbf{v}) \\ 1 + 2^{j-m-1} & \text{otherwise} \end{cases}$$

By truthfulness, $X_2(\mathbf{v}') = X_2(\mathbf{v})$. Note also that v_2^{I} induces for player 2 a strict preference over all subsets (see also Remark A.0.7). Moreover, with respect to v_2^{I} the set of "good" minimal exchanges is exactly $\{(S_i, T_i) \mid i \in I\}$.

We now claim that player 1 can deviate and strictly improve his utility, thus contradicting truthfulness. In particular, consider the profile $\mathbf{v}'' = (v_1^{\text{I}}, v_2^{\text{I}})$ where

$$v^{\text{I}}_{1j} = \begin{cases} m + 2^{j-m-1} & \text{if } j \in (X_1(\mathbf{v}') \cup S_x) \setminus T_x \\ 1 + 2^{j-m-1} & \text{otherwise} \end{cases}$$

Again, v_1^{I} induces for player 1 a strict preference over all subsets, and thus $\mathbf{v}'' \in \mathcal{V}_m^{\neq}$. As a result, $\arg\max_{S \in \mathcal{O}_i} v_i^{\text{I}}(S)$ only contains $X_i^{N_i}(\mathbf{v})$, for $i \in \{1,2\}$, and by Lemma 3.1.12 we have $X_1^{N_1 \cup N_2}(\mathbf{v}'') = X_1^{N_1 \cup N_2}(\mathbf{v}')$. Additionally, notice that with respect to $\mathbf{v}''$ the set of favorable minimal exchanges is $\{(S_i, T_i) \mid i \in I \setminus \{x\}\}$. So, by Lemma 3.1.17 we have $X_1^{E_1 \cup E_2}(\mathbf{v}'') = \left(E_1 \setminus \bigcup_{i \in I \setminus \{x\}} S_i\right) \cup \bigcup_{i \in I \setminus \{x\}} T_i = (X_1^{E_1 \cup E_2}(\mathbf{v}') \cup S_x) \setminus T_x$.

So, by deviating from $\mathbf{v}'$ to $\mathbf{v}''$, player 1 improves his utility by $v_1(S_x) - v_1(T_x) > 0$, which contradicts truthfulness. We conclude that there is no $x \in I$ such that (S_x, T_x) is unfavorable with respect to $\mathbf{v}$.

Next, we move on to the second case, i.e., there exists some $z \in [k] \setminus I$ such that (S_z, T_z) is favorable with respect to $\mathbf{v}$. Like in the first case, the valuation functions that we define induce strict preferences over all subsets. Consider the profile $Q = (v_1^{\text{II}}, v_2)$ where

$$v^{\text{II}}_{1j} = \begin{cases} m^2 + 2^{j-m-1} & \text{if } j \in X_1(\mathbf{v}) \setminus S_z \\ m + 2^{j-m-1} & \text{if } j \in T_z \\ 1 + 2^{j-m-1} & \text{otherwise} \end{cases}$$

We know, by Lemmata 3.1.14 and A.0.13, that $X_1^{E_1 \cup E_2}(Q) = \left(E_1 \setminus \bigcup_{i \in J} S_i\right) \cup \bigcup_{i \in J} T_i$ for some $J \subseteq [k]$. By truthfulness, $X_1^{N_1 \cup N_2}(Q) \supseteq X_1^{N_1 \cup N_2}(\mathbf{v})$. In fact, by Lemma 3.1.12, it must be the case where $X_1^{N_1 \cup N_2}(Q) = X_1^{N_1 \cup N_2}(\mathbf{v})$. Again by truthfulness, $X_1^{E_1}(Q) \supseteq X_1^{E_1}(\mathbf{v}) \setminus S_z = E_1 \setminus \bigcup_{i \in I \cup \{z\}} S_i$ and $X_1^{E_2}(Q) \supseteq X_1^{E_2}(\mathbf{v}) = \bigcup_{i \in I \cup \{z\}} T_i$. This implies that $I \subseteq J \subseteq I \cup \{z\}$. If $J = I \cup \{z\}$, then player 1, by deviating from $\mathbf{v}$ to Q, improves his utility by $v_1(T_z) - v_1(S_z) > 0$, which contradicts truthfulness. So, it must be the case where $J = I$.

Now, consider the profile $Q' = (v_1^{\text{II}}, v_2^{\text{II}}) \in \mathcal{V}_m^{\neq}$ where

$$v^{\text{II}}_{2j} = \begin{cases} m + 2^{j-m-1} & \text{if } j \in X_2^{N_1 \cup N_2}(Q) \cup \bigcup_{i \in I \cup \{z\}} S_i \cup \bigcup_{i \notin I \cup \{z\}} T_i \\ 1 + 2^{j-m-1} & \text{otherwise} \end{cases}$$

Since, for $i \in \{1,2\}$, $\arg\max_{S \in \mathcal{O}_i} v_i^{\mathrm{II}}(S)$ only contains $X_i^{N_i}(\mathbf{v})$, by Lemma 3.1.12 we have $X_2^{N_1 \cup N_2}(Q') = X_2^{N_1 \cup N_2}(Q)$. Further, the set of favorable minimal exchanges with respect to Q' is $\{(S_i, T_i) \mid i \in I \cup \{z\}\}$. So, by Lemma 3.1.17 we have $X_2^{E_1 \cup E_2}(Q') = \left(E_2 \setminus \bigcup_{i \in I \cup \{z\}} T_i\right) \cup \bigcup_{i \in I \cup \{z\}} S_i$.

So, by deviating from Q to Q', player 2 improves his utility by $v_2(S_z) - v_2(T_z) > 0$, which contradicts truthfulness. Therefore, there is no $z \in [k] \setminus I$ such that (S_z, T_z) is favorable with respect to $\mathbf{v}$, and this concludes the proof. □

Appendix B

Missing Material from Chapter 6

B.1 Instances with Costs Exceeding the Budget

Consider an instance with a symmetric submodular function, where there exist agents with $c_i > B$. The presence of such agents can create infeasible solutions of very high value and make an analog of Lemma 6.1.1 impossible to prove. It may seem at first sight that we could just discard such agents, since too expensive agents are not included in any feasible solution anyway. Simply discarding them, however, could destroy the symmetry of the function (e.g., if we had a cut function defined on a graph, we could not just remove a node).

Let $\mathscr{I}$ denote the set of all instances of the problem with symmetric submodular functions, and let $\mathscr{J}$ denote the set of all such instances where at most one agent has cost more than B. Given $X \subseteq A$, we let $\mathbf{c}(X) = \sum_{i \in X} c_i$. The next lemma, together with its corollary, show that, when dealing with symmetric submodular functions, we may only consider instances in $\mathscr{J}$ without any loss of generality.

Lemma B.1.1. *Given an instance $I = (A, v, \mathbf{c}, B) \in \mathscr{I}$, we can efficiently construct an instance $J = (A', v', \mathbf{c}', B) \in \mathscr{J}$ such that*

- *Every feasible solution of I is a feasible solution of J and vice versa.*
- *If X is a feasible solution of I, then $v(X) = v'(X)$ and $\mathbf{c}(X) = \mathbf{c}'(X)$. In particular, $\textsc{opt}(J) = \textsc{opt}(I)$.*

Proof. Let $E = \{i \in A \mid c_i > B\}$ be the set of expensive agents and define $A' = (A \setminus E) \cup \{i_E\}$, where i_E is a new agent replacing the whole set E. For $i \in A \setminus E$ we define $c'_i = c_i$, while $c'_{i_E} = B + 1$. Finally, v' is defined as follows

$$v'(T) = \begin{cases} v(T), & \text{if } T \subseteq A \setminus E \\ v((T \setminus \{i_E\}) \cup E), & \text{otherwise} \end{cases}$$

Now suppose X is a budget-feasible solution of I. Then $\mathbf{c}(X) \le B$ and thus $X \subseteq A \setminus E$. But then, by the definition of $\mathbf{c}'$, $\mathbf{c}'(X) = \mathbf{c}(X) \le B$ as well, and therefore X is also a budget-feasible solution of J. Moreover, $v'(X) = v(X)$ by the definition of v'. We conclude that $\textsc{opt}(I) \le \textsc{opt}(J)$.

The proof that every feasible solution of J is a feasible solution of I is almost identical. This implies $\textsc{opt}(J) \le \textsc{opt}(I)$, and therefore $\textsc{opt}(J) = \textsc{opt}(I)$. □

Now, it is not hard to see that we can turn any algorithmic result on $\mathcal{J}$ to the same algorithmic result on $\mathcal{I}$. However, we need a somewhat stronger statement to take care of issues like truthfulness and budget-feasibility. This is summarized in the following corollary.

Corollary B.1.2. *Given a (polynomial time) algorithm* $\textsc{alg}'$ *that achieves a* ρ*-approximation on instances in* $\mathcal{J}$*, we can efficiently construct a (polynomial time)* ρ*-approximation algorithm* $\textsc{alg}$ *that works for all instances in* $\mathcal{I}$*. Moreover, if* $\textsc{alg}'$ *is monotone and budget-feasible on instances in* $\mathcal{J}$*, assuming Myerson's threshold payments, then* $\textsc{alg}$ *is monotone and budget-feasible on instances in* $\mathcal{I}$*.*

Proof. The description of $\textsc{alg}$ is quite straightforward. Given an instance $I = (A, v, \mathbf{c}, B) \in \mathcal{I}$, $\textsc{alg}$ first constructs instance $J = (A', v', \mathbf{c}', B) \in \mathcal{J}$, as described in the proof of Lemma B.1.1. Then $\textsc{alg}$ runs $\textsc{alg}'$ with input J and returns its output. Clearly, if $\textsc{alg}'$ runs in polynomial time, so does $\textsc{alg}$.

If $X = \textsc{alg}'(J) = \textsc{alg}(I)$, then X is feasible with respect to J and $\textsc{opt}(J) \le \rho \cdot v'(X)$. By Lemma B.1.1 we get that X is feasible with respect to I and $\textsc{opt}(I) = \textsc{opt}(J) \le \rho \cdot v'(X) = \rho \cdot v(X)$. This establishes the approximation ratio of $\textsc{alg}$.

Next, assume that $\textsc{alg}'$ is monotone and budget-feasible on instances in $\mathcal{J}$, when using Myerson's threshold payments. Suppose that agent $j \in \textsc{alg}(I)$ reduces his cost from c_j to $b_j < c_j$. This results in a new instance $I_* = (A, v, (b_j, \mathbf{c}_{-j}), B) \in \mathcal{I}$. Since it must be the case where $c_j \le B$, the corresponding instance of $\mathcal{J}$ is $J_* = (A', v', (b_j, \mathbf{c}'_{-j}), B)$. Due to the monotonicity of $\textsc{alg}'$ we have

$$j \in \textsc{alg}(I) = \textsc{alg}'(J) \Rightarrow j \in \textsc{alg}'(J_*) = \textsc{alg}(I_*)\,,$$

and therefore $\textsc{alg}$ is monotone as well.

The budget-feasibility of $\textsc{alg}$ follows from the budget-feasibility of $\textsc{alg}'$ by observing that $i \in \textsc{alg}(A, v, (b_j, \mathbf{c}_{-j}), B)$ if and only if $i \in \textsc{alg}'(A', v', (b_j, \mathbf{c}'_{-j}), B)$ for all $i \in A'$. □

B.2 Regarding Remark 6.0.1

We use the cut function on a very simple graph and show that although v is submodular, $\hat{v}$ is not. Consider the following graph where each edge has unit weight:

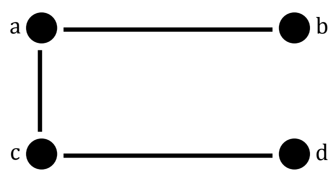

We compute the value of the following sets:

$$\hat{v}(\{a\}) = v(\{a\}) = 2$$

$$\hat{v}(\{a, b\}) = v(\{a\}) = 2$$

$\hat{v}(\{a,c\}) = v(\{a\}) = v(\{c\}) = v(\{a,c\}) = 2$

$\hat{v}(\{a,b,c\}) = v(\{b,c\}) = 3$

Now it is easy to see that although $\{a\} \subseteq \{a,c\}$ we have

$$\hat{v}(\{a\} \cup \{b\}) - \hat{v}(\{a\}) = 0 < 1 = \hat{v}(\{a,c\} \cup \{b\}) - \hat{v}(\{a,c\}).$$

So an interesting question is if $\hat{v}$ can be classified when v is submodular. In Gupta, Nagarajan, and Singla, 2017 the general XOS class was introduced, where it is allowed for a function to be non-monotone (recall here that the XOS class contains only non-decreasing functions by definition). They proved that when v is general XOS then $\hat{v}$ is XOS, while in addition they observed that the class of non-negative submodular functions is a strict subset of the general XOS class. Thus since any non-negative submodular function v is also general XOS we conclude $\hat{v}$ is XOS.

Bibliography

Ageev, A. A. and M. Sviridenko (1999). “Approximation Algorithms for Maximum Coverage and Max Cut with Given Sizes of Parts”. In: *Integer Programming and Combinatorial Optimization, 7th International IPCO Conference, 1999, Proceedings*, pp. 17–30.

— (2004). “Pipage Rounding: A New Method of Constructing Algorithms with Proven Performance Guarantee”. In: *Journal of Combinatorial Optimization* 8.3, pp. 307–328.

Amanatidis, G., G. Birmpas, and E. Markakis (2016a). “Coverage, Matching, and Beyond: New Results on Budgeted Mechanism Design”. In: *Proceedings of the 12th International Conference on Web and Internet Economics, WINE 2016*, pp. 414–428.

— (2016b). “On Truthful Mechanisms for Maximin Share Allocations”. In: *Proceedings of the 25th International Joint Conference on Artificial Intelligence, IJCAI 2016*, pp. 31–37.

— (2017). “On Budget-Feasible Mechanism Design for Symmetric Submodular Objectives”. In: *CoRR* abs/1704.06901. URL: http://arxiv.org/abs/1704.06901.

Amanatidis, G., E. Markakis, A. Nikzad, and A. Saberi (2015). “Approximation Algorithms for Computing Maximin Share Allocations”. In: *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Proceedings, Part I*, pp. 39–51.

Amanatidis, G., G. Birmpas, G. Christodoulou, and E. Markakis (2017). “Truthful Allocation Mechanisms Without Payments: Characterization and Implications on Fairness”. In: *ACM Conference on Economics and Computation, EC ’17*, pp. 545–562.

Anari, N., G. Goel, and A. Nikzad (2014). “Mechanism Design for Crowdsourcing: An Optimal 1-1/e Competitive Budget-Feasible Mechanism for Large Markets”. In: *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014*, pp. 266–275.

Archer, A. and É. Tardos (2007). “Frugal path mechanisms”. In: *ACM Trans. Algorithms* 3.1.

Asadpour, A. and A. Saberi (2007). “An approximation algorithm for max-min fair allocation of indivisible goods”. In: *ACM Symposium on Theory of Computing (STOC)*, pp. 114–121.

Aziz, H. and S. MacKenzie (2016). “A Discrete and Bounded Envy-free Cake Cutting Protocol for Four Agents”. In: *48th ACM Symposium on the Theory of Computing, STOC 2016*, pp. 454–464.

Aziz, H., A. Filos-Ratsikas, J. Chen, S. Mackenzie, and N. Mattei (2016). “Egalitarianism of Random Assignment Mechanisms: (Extended Abstract)”. In: *International Conference on Autonomous Agents & Multiagent Systems, AAMAS 2016*, pp. 1267–1268.

Balkanski, E. and J. D. Hartline (2016). “Bayesian Budget Feasibility with Posted Pricing”. In: *Proceedings of the 25th International Conference on World Wide Web, WWW 2016*. ACM, pp. 189–203.

Bansal, N. and M. Sviridenko (2006). “The Santa Claus Problem”. In: *ACM Symposium on Theory of Computing (STOC)*, pp. 31–40.

Barman, S. and S. K. K. Murthy (2017). “Approximation Algorithms for Maximin Fair Division”. In: *ACM Conference on Economics and Computation, EC ’17*, pp. 647–664.

Bei, X., N. Chen, N. Gravin, and P. Lu (2012). "Budget feasible mechanism design: from prior-free to bayesian". In: *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012*, pp. 449–458.

Bezakova, I. and V. Dani (2005). "Allocating Indivisible Goods". In: *ACM SIGecom Exchanges* 5, pp. 11–18.

Borgs, C., J. T. Chayes, N. Immorlica, M. Mahdian, and A. Saberi (2005). "Multi-unit auctions with budget-constrained bidders". In: *Proceedings 6th ACM Conference on Electronic Commerce EC 2005*, pp. 44–51.

Borodin, A., Y. Filmus, and J. Oren (2010). "Threshold Models for Competitive Influence in Social Networks". In: *Proceedings of the 6th International Workshop on Internet and Network Economics, WINE 2010*, pp. 539–550.

Bouveret, S., Y. Chevaleyre, and N. Maudet (2016). "Fair Allocation of Indivisible Goods". In: *Handbook of Computational Social Choice*. Ed. by F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia. Cambridge University Press. Chap. 12.

Bouveret, S. and J. Lang (2011). "A General Elicitation-free Protocol for Allocating Indivisible Goods". In: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI 2011*, pp. 73–78.

— (2014). "Manipulating picking sequences". In: *ECAI 2014 - 21st European Conference on Artificial Intelligence*, pp. 141–146.

Bouveret, S. and M. Lemaître (2016). "Characterizing conflicts in fair division of indivisible goods using a scale of criteria". In: *Autonomous Agents and Multi-Agent Systems* 30.2. A preliminary version of this work has appeared in AAMAS '14, pp. 259–290.

Brams, S. J. and D. King (2005). "Efficient Fair Division - Help the worst off or avoid envy". In: *Rationality and Society* 17.4, pp. 387–421.

Brams, S. J. and A. D. Taylor (1996). *Fair Division: from Cake Cutting to Dispute Resolution.* Cambridge University press.

Budish, E. (2011). "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes". In: *Journal of Political Economy* 119.6, pp. 1061–1103.

Caragiannis, I., C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou (2009). "On Low-Envy Truthful Allocations". In: *First International Conference on Algorithmic Decision Theory, ADT 2009*, pp. 111–119.

Caragiannis, I., D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang (2016). "The Unreasonable Fairness of Maximum Nash Welfare". In: *ACM Conference on Economics and Computation, EC '16*, pp. 305–322.

Caselton, W. F. and J. V. Zidek (1984). "Optimal monitoring network designs". In: *Statistics & Probability Letters* 2.4, pp. 223–227.

Chan, Y. H. and L. C. Lau (2012). "On linear and semidefinite programming relaxations for hypergraph matching". In: *Mathematical Programming* 135.1-2, pp. 123–148.

Chekuri, C., J. Vondrák, and R. Zenklusen (2014). "Submodular Function Maximization via the Multilinear Relaxation and Contention Resolution Schemes". In: *SIAM J. Comput.* 43.6, pp. 1831–1879.

Chen, N., N. Gravin, and P. Lu (2011). "On the Approximability of Budget Feasible Mechanisms". In: *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011*, pp. 685–699.

Chen, N., E. Elkind, N. Gravin, and F. Petrov (2010). “Frugal Mechanism Design via Spectral Techniques”. In: *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*, pp. 755–764.

Chen, Y., J. K. Lai, D. C. Parkes, and A. D. Procaccia (2013). “Truth, justice, and cake cutting”. In: *Games and Economic Behavior* 77.1, pp. 284–297.

Christodoulou, G., E. Koutsoupias, and A. Vidali (2008). “A characterization of 2-Player Mechanisms for Scheduling”. In: *Algorithms - ESA, 16th Annual European Symposium*, pp. 297–307.

Christodoulou, G. and A. Kovács (2011). “A Global Characterization of Envy-Free Truthful Scheduling of Two Tasks”. In: *Internet and Network Economics - 7th International Workshop, WINE 2011, Proceedings*, pp. 84–96.

Cole, R., V. Gkatzelis, and G. Goel (2013). “Mechanism design for fair division: allocating divisible items without payments”. In: *ACM Conference on Electronic Commerce, EC ’13*. ACM, pp. 251–268.

Cook, W. J., W. H. Cunningham, W. R. Pulleyblank, and A. Schrijver (1998). *Combinatorial Optimization*. John Wiley & Sons, Inc.

Cressie, N. (1993). “Statistics for Spatial Data”. In:

Dobzinski, S., R. Lavi, and N. Nisan (2012). “Multi-unit auctions with budget limits”. In: *Games and Economic Behavior* 74.2, pp. 486–503.

Dobzinski, S., C. H. Papadimitriou, and Y. Singer (2011). “Mechanisms for complement-free procurement”. In: *Proceedings 12th ACM Conference on Electronic Commerce EC 2011*, pp. 273–282.

Dobzinski, S. and M. Sundararajan (2008). “On Characterizations of Truthful Mechanisms for Combinatorial Auctions and Scheduling”. In: *ACM Conference on Electronic Commerce EC 2008)*.

Edmonds, J. (1965). “Maximum Matching and a Polyhedron with 0,1 Vertices”. In: *Journal of Research of the National Bureau of Standards* 69 B, pp. 125–130.

Edmonds, J. and K. Pruhs (2006). “Balanced Allocations of Cake”. In: *Symposium on Foundations of Computer Science (FOCS)*, pp. 623–634.

Ehlers, L. and B. Klaus (2003). “Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems”. In: *Social Choice and Welfare* 21.2, pp. 265–280.

Even, S. and A. Paz (1984). “A note on cake cutting”. In: *Discrete Applied Mathematics* 7, pp. 285–296.

Feige, U., V. S. Mirrokni, and J. Vondrák (2011). “Maximizing Non-monotone Submodular Functions”. In: *SIAM J. Comput.* 40.4, pp. 1133–1153.

Feldman, M., J. Naor, and R. Schwartz (2011). “A Unified Continuous Greedy Algorithm for Submodular Maximization”. In: *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011*. IEEE Computer Society, pp. 570–579.

Foley, D. (1967). “Resource Allocation and the Public Sector”. In: *Yale Economics Essays* 7, pp. 45–98.

Fujishige, S. (1983). “Canonical Decompositions of Symmetric Submodular Systems”. In: *Discrete Applied Mathematics* 5, pp. 175–190.

Gamow, G. and M. Stern (1958). *Puzzle-Math*. Viking press.

Gentle, J.E. (2009). *Computational statistics*. Vol. 308. Springer.

Goel, G., V. S. Mirrokni, and R. P. Leme (2012). "Polyhedral clinching auctions and the adwords polytope". In: *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012*, pp. 107–122.

Goel, G., A. Nikzad, and A. Singla (2014). "Mechanism Design for Crowdsourcing Markets with Heterogeneous Tasks". In: *Proceedings of the Seconf AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2014*.

Goemans, M. X. and D. P. Williamson (1994). "New 3/4-Approximation Algorithms for the Maximum Satisfiability Problem". In: *SIAM Journal of Discrete Mathematics* 7.4, pp. 656–666.

Gourvès, L., J. Monnot, and L. Tlilane (2015). "Worst case compromises in matroids with applications to the allocation of indivisible goods". In: *Theoretical Computer Science* 589, pp. 121–140.

Guo, M. and V. Conitzer (2010). "Strategy-proof allocation of multiple items between two agents without payments or priors". In: *9th International Conference on Autonomous Agents & Multiagent Systems AAMAS 2010*, pp. 881–888.

Gupta, A., V. Nagarajan, and S. Singla (2017). "Adaptivity Gaps for Stochastic Probing: Submodular and XOS Functions". In: *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017*. SIAM, pp. 1688–1702.

Gupta, A., A. Roth, G. Schoenebeck, and K. Talwar (2010). "Constrained Non-monotone Submodular Maximization: Offline and Secretary Algorithms". In: *Internet and Network Economics - 6th International Workshop, WINE 2010, Proceedings*. Springer, pp. 246–257.

Hall, P. (1935). "On Representatives of Subsets". In: *Journal of the London Mathematical Society* s1-10.1, pp. 26–30.

Hill, T. (1987). "Partitioning general probability measures". In: *The Annals of Probability* 15.2, pp. 804–813.

Hoeffding, W. (1963). "Probability Inequalities for Sums of Bounded Random Variables". In: *Journal of the American Statistical Association* 58.301, pp. 13–30.

Horel, T., S. Ioannidis, and S. Muthukrishnan (2014). "Budget Feasible Mechanisms for Experimental Design". In: *LATIN 2014: Theoretical Informatics - 11th Latin American Symposium, Proceedings*, pp. 719–730.

Jalaly, P. and É. Tardos (2017). Simple and Efficient Budget Feasible Mechanisms for Monotone Submodular Valuations, arxiv:1703:10681.

Karlin, A. R., D. Kempe, and T. Tamir (2005). "Beyond VCG: Frugality of Truthful Mechanisms". In: *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), Proceedings*, pp. 615–626.

Kempe, D., M. Salek, and C. Moore (2010). "Frugal and Truthful Auctions for Vertex Covers, Flows and Cuts". In: *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*, pp. 745–754.

Klaus, B. and E. Miyagawa (2002). "Strategy-proofness, Solidarity and Consistency for Multiple Assignment Problems". In: *International Journal of Game Theory* 30.3, pp. 421–435.

Kleinberg, J. and É. Tardos (2006). *Algorithm Design*. Addison Wesley.

Krause, A., A. P. Singh, and C. Guestrin (2008). “Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies”. In: *Journal of Machine Learning Research* 9, pp. 235–284.

Kulik, A., H. Shachnai, and T. Tamir (2013). “Approximations for Monotone and Nonmonotone Submodular Maximization with Knapsack Constraints”. In: *Math. Oper. Res.* 38.4, pp. 729–739.

Kurokawa, D., A. D. Procaccia, and J. Wang (2016). “When Can the Maximin Share Guarantee Be Guaranteed?” In: *30th AAAI Conference on Artificial Intelligence, AAAI 2016*.

Lee, J., V. S. Mirrokni, V. Nagarajan, and M. Sviridenko (2010). “Maximizing Nonmonotone Submodular Functions under Matroid or Knapsack Constraints”. In: *SIAM J. Discrete Math.* 23.4, pp. 2053–2078.

Leonardi, S., G. Monaco, P. Sankowski, and Q. Zhang (2016). “Budget Feasible Mechanisms on Matroids”. In: *CoRR* abs/1612.03150. URL: http://arxiv.org/abs/1612.03150.

Lipton, R. J., E. Markakis, E. Mossel, and A. Saberi (2004). “On approximately fair allocations of indivisible goods”. In: *ACM Conference on Electronic Commerce EC 2004*, pp. 125–131.

Markakis, E. and C.-A. Psomas (2011). “On Worst-Case Allocations in the Presence of Indivisible Goods”. In: *7th Workshop on Internet and Network Economics (WINE 2011)*, pp. 278–289.

Mennle, T. and S. Seuken (2014). “An axiomatic approach to characterizing and relaxing strategyproofness of one-sided matching mechanisms”. In: *ACM Conference on Economics and Computation, EC ’14*, pp. 37–38.

Moulin, H. (1990). “Uniform Externalities: Two Axioms for Fair Allocation”. In: *Journal of Public Economics* 43.3, pp. 305–326.

— (2003). *Fair division and collective welfare*. MIT Press. ISBN: 978-0-262-63311-6.

Myerson, R. (1981). “Optimal Auction Design”. In: *Mathematics of Operations Research* 6.1.

Nemhauser, G. L., L. A. Wolsey, and M. L. Fisher (1978). “An analysis of approximations for maximizing submodular set functions - I”. In: *Math. Program.* 14.1, pp. 265–294.

Papadimitriou, C. H. (2001). “Algorithms, Games, and the Internet”. In: *Automata, Languages and Programming, 28th International Colloquium, ICALP 2001, Proceedings*. Vol. 2076. Lecture Notes in Computer Science. Springer, pp. 1–3.

Pápai, S. (2000). “Strategyproof Multiple Assignment Using Quotas”. In: *Review of Economic Design* 5.1, pp. 91–105.

— (2001). “Strategyproof and Nonbossy Multiple Assignments”. In: *Journal of Public Economic Theory* 3.3, pp. 257–271.

— (2003). “Strategyproof exchange of indivisible goods”. In: *Journal of Mathematical Economics* 39.8, pp. 931–959.

— (2007). “Exchange in a General Market with Indivisible Goods”. In: *Journal of Economic Theory* 132, pp. 208–235.

Procaccia, A. D. (2016). “Cake Cutting Algorithms”. In: *Handbook of Computational Social Choice*. Ed. by F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia. Cambridge University Press. Chap. 13.

Procaccia, A. D. and J. Wang (2014). “Fair enough: guaranteeing approximate maximin shares”. In: *ACM Conference on Economics and Computation, EC ’14*, pp. 675–692.

Queyranne, M. (1998). "Minimizing symmetric submodular functions". In: *Mathematical Programming* 82.1-2, pp. 3–12.

Robertson, J. M. and W. A. Webb (1998). *Cake Cutting Algorithms: be fair if you can*. AK Peters.

Schäffer, A. A. and M. Yannakakis (1991). "Simple Local Search Problems That are Hard to Solve". In: *SIAM J. Comput.* 20.1, pp. 56–87.

Singer, Y. (2010). "Budget Feasible Mechanisms". In: *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*, pp. 765–774.

— (2012). "How to win friends and influence people, truthfully: influence maximization mechanisms for social networks". In: *Proceedings of the Fifth International Conference on Web Search and Web Data Mining, WSDM 2012*, pp. 733–742.

— (2016). Personal communication.

Singla, A. and A. Krause (2013). "Incentives for Privacy Tradeoff in Community Sensing". In: *Proceedings of the First AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2013*. AAAI.

Spliddit (2015). *Provably Fair Solutions*. http://www.spliddit.org/.

Steinhaus, H. (1948). "The Problem of Fair Division". In: *Econometrica* 16, pp. 101–104.

Svensson, L. (1999). "Strategy-proof Allocation of Indivisible Goods". In: *Social Choice and Welfare* 16, pp. 557–567.

Sviridenko, Maxim (2004). "A note on maximizing a submodular set function subject to a knapsack constraint". In: *Oper. Res. Lett.* 32.1, pp. 41–43.

Varian, H. (1974). "Equity, Envy and Efficiency". In: *Journal of Economic Theory* 9, pp. 63–91.

Woeginger, G. (1997). "A polynomial time approximation scheme for maximizing the minimum machine completion time". In: *Operations Research Letters* 20, pp. 149–154.

Woeginger, G. and J. Sgall (2007). "On the Complexity of cake cutting". In: *Discrete Optimization* 4.2, pp. 213–220.

Wolsey, L. A. (1982). "Maximising Real-Valued Submodular Functions: Primal and Dual Heuristics for Location Problems". In: *Math. Oper. Res.* 7.3, pp. 410–425.