Design and Analysis of Auction Mechanisms

Algorithms and Incentives

Artem Tsikiridis

Advisor: Assoc. Prof. Evangelos Markakis

Department of Informatics
Athens University of Economics and Business

This dissertation is submitted for the degree of
Doctor of Philosophy

May 2023

Abstract

In this dissertation, we propose novel algorithms for combinatorial auction environments using an interdisciplinary approach. At the same time, we also analyze the performance of existing auction protocols and highlight design principles that allow for provable performance guarantees.

In the first part of the thesis we study two forward auction paradigms with practical significance: core-selecting mechanisms and multi-unit auctions. We begin with the notion of core-selecting mechanisms, as introduced by Ausubel and Milgrom. Such mechanisms have overall good revenue guarantees, but are known to provide incentives to bidders for misreporting their preferences. Current research has focused on identifying core-selecting mechanisms with minimal incentives to deviate from truth-telling, such as Minimum-Revenue Core-Selecting (MRCS) rules, or proposing truthful mechanisms whose revenue is competitive against core outcomes. Our results contribute to both of these directions. We study the core polytope in more depth and provide new properties and insights that are of independent interest. Utilizing these properties, we then propose a truthful mechanism that is competitive against the MRCS benchmark, the first deterministic core-competitive mechanism for binary single-parameter domains. We also answer an open question from the literature, of whether there exist MRCS non-decreasing mechanisms, in the affirmative. Next, we shift our attention to multi-unit auctions, a class of auctions first studied by Vickrey. We analyze discriminatory price auctions, the natural multi-unit extension of the (non-truthful) first price auction. We consider bidders with capped-additive valuations and establish properties that capture the sources of inefficiency. We derive new lower and upper bounds on the Price of Anarchy of mixed equilibria, showing a complete characterization of inefficient equilibria and a tight upper bound for the case of two bidders. We also show that the Price of Anarchy is strictly worse for multiple bidders and we exhibit a separation result for Bayes Nash equilibria.

In the second part of this dissertation, we study procurement auctions. Firstly, we study a covering problem motivated by spatial models in crowdsourcing markets, where tasks are ordered according to some geographic or temporal criterion. We propose

a truthful mechanism that achieves a bounded approximation guarantee w.r.t. the optimal cost, improving upon the state of the art. For the same objective, we propose a truthful fully polynomial-time approximation scheme (FPTAS) for the case of inputs with a constant number of tasks, a generalization of the minimum knapsack problem. We then focus on the class of budget-feasible procurement auctions, in which agents can provide their service to the auctioneer fractionally or in many levels. We propose two mechanisms, one for each setting. The mechanism for divisible agents improves upon the known state of the art, whereas the mechanism for the multiple levels of service is the first truthful and budget-feasible mechanism that achieves a constant approximation for this setting.

We conclude the dissertation with an extended discussion along with open problems and directions for future research in algorithmic mechanism design for auction environments.

Περίληψη

Σε αυτή τη διατριβή, σχεδιάζουμε νέους αλγορίθμους για περιβάλλοντα συνδυαστικών δημοπρασιών ακολουθώντας μια διεπιστημονική προσέγγιση. Ταυτόχρονα, αναλύουμε την απόδοση υπαρχόντων πρωτοκόλλων δημοπρασιών και αναδεικνύουμε τις σχεδιαστικές αρχές εκείνες που επιτρέπουν εγγυήσεις απόδοσης.

Στο πρώτο κομμάτι της διατριβής μελετάμε δύο υποδείγματα δημοπρασιών σημαντικών ως προς τις πρακτικές εφαρμογές τους: δημοπρασίες πυρήνα (core-selecting auctions) και δημοπρασίες πολλών αντιγράφων ενός αντικειμένου (multi-unit auctions). Αρχικά μελετούμε την έννοια του πυρήνα, όπως ορίστηκε από τους Ausubel και Milgrom. Οι μηχανισμοί αυτοί, παρά το γεγονός πως προσφέρουν συνολικά ικανοποιητικές εγγυήσεις ως προς τα έσοδα του δημοπράτη, παρέχουν κίνητρα στους πλειοδότες να μη δηλώνουν τις πραγματικές τους προτιμήσεις. Επομένως, ένα από τα κύρια ζητούμενα στην βιβλιογραφία είναι ο προσδιορισμός εκείνων των μηχανισμών πυρήνα που παρέχουν στους πλειοδότες τα ελάχιστα δυνατά τέτοια κίνητρα, όπως οι Μηχανισμοί Πυρήνα Ελαχίστων Εσόδων (MRCS), ή η εύρεση μηχανισμών που είναι φιλαλήθεις με έσοδα ανταγωνιστικά ως προς τα αποτελέσματα των μηχανισμών πυρήνα. Τα αποτελέσματά μας συνεισφέρουν ως προς αμφότερες κατευθύνσεις. Αρχικά, μελετούμε το πολύτοπο που σχηματίζει ο πυρήνας σε μεγαλύτερο βάθος και αναδεικνύουμε μερικές νέες ιδιότητες. Χρησιμοποιώντας τις ιδιότητες αυτές, προτείνουμε έναν φιλαλήθη μηχανισμό που είναι ανταγωνιστικός ως προς τα MRCS έσοδα. Ο μηχανισμός αυτός είναι ο πρώτος ντετερμινιστικός, ανταγωνιστικός προς τον πυρήνα μηχανισμός για δυαδικά περιβάλλοντα δημοπρασιών μίας παραμέτρου στη βιβλιογραφία. Ακόμη, δίνουμε μια καταφατική απάντηση στην ερώτηση που είχε τεθεί στην βιβλιογραφία σχετικά με το αν υπάρχουν μη φθίνοντες (non-decreasing) MRCS μηχανισμοί. Στη συνέχεια, επικεντρωνόμαστε στις δημοπρασίες πολλών αντιγράφων ενός αντικειμένου (multi-unit auctions), μια κλάση δημοπρασιών που μελετήθηκε πρώτη φορά από τον Vickrey. Αναλύουμε δημοπρασίες διακριτής τιμής (discriminatory price), οι οποίες αποτελούν φυσική γενίκευση των δημοπρασιών πρώτης τιμής, γνωστών για την μη φιλαλήθειά τους. Στην μελέτη μας, λαμβάνουμε υπόψη με συναρτήσεις αξίας τύπου capped-additive και αποδεικνύουμε ιδιότητες που συλλαμβάνουν τις πηγές μη αποδοτικότητας. Στη συνέχεια εξάγουμε νέα κάτω και άνω φράγματα ως προς το Τίμημα της Αναρχίας των μικτών σημείων ισορροπίας, δείχνοντας έναν πλήρη χαρακτηρισμό των μη αποδοτικών σημείων ισορροπίας καθώς και ένα ακριβές άνω φράγμα για την περίπτωση των δύο παικτών. Τέλος, δείχνουμε πως το Τίμημα της Αναρχίας είναι αυστηρά χειρότερο για την περίπτωση πολλών παικτών και παρουσιάζουμε έναν διαχωρισμό της κλάσης αυτής με την κλάση των Μπεϋζιανών σημείων ισορροπίας κατά Nash.

Στο δεύτερο κομμάτι της διατριβής, μελετάμε δημοπρασίες προμηθειών (procurement auctions). Αρχικά, μελετάμε ένα πρόβλημα κάλυψης που προκύπτει σε γεωγραφικά μοντέλα αγορών πληθοπορισμού, στα οποία οι εργασίες είναι ταξινομημένες με βάση γεωγραφικά ή χρονικά κριτήρια. Σχεδιάζουμε έναν φιλαλήθη μηχανισμό που πετυχαίνει έναν φραγμένο λόγο προσέγγισης σε σχέση με το βέλτιστο κόστος του δημοπράτη, βελτιώνοντας το καλύτερο γνωστό αποτέλεσμα της βιβλιογραφίας. Για την ίδια αντικειμενική συνάρτηση, σχεδιάζουμε έναν φιλαλήθες Πλήρως Πολυωνυμικού Χρόνου Σχήμα Προσέγγισης (FPTAS) για την περίπτωση εισόδων με σταθερό αριθμό εργασιών, μια γενίκευση του προβλήματος minimum knapsack. Στη συνέχεια μελετάμε μια οικογένεια αντίστροφων δημοπρασιών στην οποία ο δημοπράτης έχει περιορισμένο προϋπολογισμό και οι πλειοδότες μπορούν να ανατεθούν να εκτελέσουν το καθήκον τους τμηματικά ή σε πολλά επίπεδα υπηρεσίας. Προτείνουμε δύο μηχανισμούς, έναν για κάθε περιβάλλον. Ο μηχανισμός για το περιβάλλον στο οποίο οι παίκτες μπορούν να προσληφθούν τμηματικά βελτιώνει την υπάρχουσα βιβλιογραφία, ενώ ο μηχανισμός για πολλαπλά επίπεδα υπηρεσίας είναι ο πρώτος γνωστός φιλαλήθης μηχανισμός με σταθερό λόγο προσέγγισης για αυτό το περιβάλλον.

Ολοκληρώνουμε την διατριβή με μια εκτενή συζήτηση, καθώς και με ανοικτά προβλήματα και κατευθύνσεις για μελλοντική έρευνα στον πεδίο του σχεδιασμού μηχανισμών για περιβάλλοντα δημοπρασιών.

Acknowledgements

It is difficult to convey in words my gratitude towards Vangelis Markakis. As my advisor, his support has been unwavering and I will never be able to repay his confidence in me. Thanks to his kindness and his reassuring presence, he has been a great influence on my growth as a researcher and as a person. It has been a pleasure working with him and learning from him. Truly, I cannot thank him enough for everything he has done for me, so I'll just stop here.

I am incredibly grateful to Georgios Amanatidis, who has been a mentor of mine since my undergraduate days, when he was a teaching assistant in my algorithms course. Over the years, Georgios has remained a steadfast source of inspiration and a trusted guide whom I can always rely on. Remarkably, he is always available to answers my questions about anything, always with kindness and generosity. I would also like to thank Georgios for hosting me at the University of Essex for a research visit[1] and for our collaboration ever since.

I owe a debt of gratitude to George Stamoulis for his invaluable support, particularly during my first years as a PhD student. I would also like to express my appreciation to Ioannis Caragiannis for his genuine interest in my development and career. I thank them both for serving on my dissertation committee alongside Antonis Dimakis, Dimitris Fotakis, Aris Pagourtzis, and Alkmini Sgouritsa, whom I would also like to thank for teaching me a lot during our collaboration.

During my PhD studies, aside from Vangelis, Georgios A. and Alkmini, I had the pleasure of collaborating with Georgios Papasotiropoulos, Sophie Klumper and Guido Schäfer. I would also like to thank Guido for hosting me on an internship at CWI in Amsterdam. I am looking forward to returning soon! Finally, I would like to express my gratitude to Timos Sellis, the director of the Archimedes Research Unit of the Athena Research Center, along with its chief scientists, Konstantinos Daskalakis and Christos Papadimitriou, for selecting me in the institute's first round of summer interns.

[1] The visit was made possible by the European Network of Game Theory (GAMENET / COST Action CA16228).

Working in the Theory, Economics and Systems Laboratory of AUEB has been a pleasure, thanks to the wonderful people I have had the privilege of working with, including Aggeliki Anastopoulou, George Darzanos, Georgios Papasotiropoulos, Panagiotis Tsamopoulos, Francis Durand and Christodoulos Santorinaios. I am grateful to them for creating such a pleasant work environment and for making our breaks very enjoyable.

Of course, I must recognize that none of this would have been possible without the support of my friends and family. I am deeply thankful to Epameinondas Koutsoumpas, who has had my back for the past twenty years, as well as Maria Iliadi, Ilias Sarantopoulos, Georgios Papasotiropoulos and Konstantina Rapti, who have provided me constant encouragement throughout our friendship. I would also like to express my gratitude to Dimitris Giannaras, Efthimis Giannaras, Stamatis Paraskevas and Agapi Paravalou and, through them, thank the beautiful community of people at our Kallitechnio, a place I will always call home. Finally, I am deeply grateful to my family for their unconditional love and support in everything I do.

H.F.R.I.
Hellenic Foundation for Research & Innovation

This research work was supported by the Hellenic Foundation for Research and Innovation (HFRI) under the HFRI PhD Fellowship grant (Fellowship Number: 289).

To Liebe,
the seventh daughter of Heim,
who I carry within.

Contents

Chapter 1

Introduction

An auction is a transactional mechanism employed by an entity known as the auctioneer to allocate a set of resources to interested parties known as bidders. Auctions have a long and fascinating history, with instances of their use dating back to ancient civilizations such as Greece and Rome. In medieval Europe, auctions were often used to sell off land or other assets seized from debtors. In the mid-18th century, reputable auction houses began to emerge, such as Sotheby's and Christie's, some of which are still active today.

In recent times, the study of auctions has become a prevalent theme in economic theory. Economists are now examining how auctions can be designed to maximize efficiency, revenue, and fairness. This analysis can be traced back to the seminal work of Vickrey (1961), who was awarded the Nobel Prize in Economics in 1996 for his contributions to the field. Today, auction mechanisms are being used in a multitude of settings by both state-run and private institutions alike. Ad-auctions run by search engines like Google constitute their main source of revenue, while online auction platforms such as eBay and eBid are used by millions of users daily. Other modern applications of auctions include spectrum auctions for TV and telecommunication licenses, allocation of airspace system resources, auctions for financial products, transportation services and procurement auctions (see Chapters 16-21 in Cramton et al. (2006)). Procurement auctions, in particular, are used by governments and other organizations to purchase goods and services from suppliers. These are also referred to as *reverse* auctions, since the auctioneer assumes the role of the buyer and the bidders the roles of the sellers, and have become increasingly important in recent years due to their potential for cost savings and transparency in the procurement process. Theoretical insights from auction theory have been particularly significant in real-world implementations since the early 1990s. For instance, a detailed chronicle by Milgrom (2004), another awardee of the Nobel Prize in Economics, presents the first spectrum auctions run by the US Federal Communications Commission (FCC), which were considered highly successful and incorporated learnings from the auction theory literature. More recently, the FCC Incentive Auction, which was completed in 2017, was also a significant source of income for the US Government and utilized parts of the revenue for deficit reduction.

Contemporary applications of auctions can present challenges and difficulties that are not always captured by classical auction theory. One such challenge arises when multiple items or services are offered simultaneously, and bidders can make offers on combinations of these items. These mechanisms are known as Combinatorial Auctions, and one of their benefits is that they can potentially result in higher revenue than running independent auctions for each item separately (Cramton et al., 2006). However, even for auctions with a small number of items, the decision-making process for the auctioneer can be computationally hard, particularly when there are bids for overlapping subsets of items. In addition to determining the allocation of items, the auctioneer must also compute the payments to be charged to bidders, which can be subject to further constraints. These computational difficulties arise even in relatively simple auction settings and must be carefully considered by auction designers.

Motivated by such questions, an intense interaction has taken place over the last two decades among researchers from computer science, operations research, and economics, which gradually led to the development of a new scientific domain, usually termed Algorithmic Game Theory (AGT). The seminal paper by Nisan and Ronen (1999), which marked the beginning of AGT, studies the computational aspects of game theory and seeks to provide algorithmic solutions to game-theoretic problems, including auctions. Thus, combinatorial auctions form one of the main pillars of these interdisciplinary explorations and are often seen as the paradigmatic problem at the forefront of this domain. To design combinatorial auctions with favorable performance, the crucial issues that need to be addressed, both theoretically and eventually in practice, can be grouped into two categories:

1. *Computational and communication complexity.* When an auction is run, the auctioneer first specifies a set of rules, regarding the bids that can be submitted, which is referred to as the bidding language of the auction. Evidently, there is a trade-off between the expressiveness of the bidding language and the complexity it introduces. A language with many restrictions (on the number or type of bids) does not allow the bidders to fully express their preferences. In a more expressive language, the bidder can communicate her preferences more accurately, but at the expense that the auctioneer receives a much larger input to process

(i.e., high communication complexity). Choosing a suitable bidding language for each application is already an interesting research question. Even more importantly, once the auctioneer has received all the input from the bidders, the next step is to decide how to allocate the items. This is an optimization problem, called the Winner Determination Problem, where the desirable criteria are either to maximize the social welfare (the total value derived for the bidders) or the revenue derived for the auctioneer. In general, this allocation problem is computationally intractable (NP-hard, see Lehmann et al. (2006, 2002) for some of the first approaches, as well as the surveys in Blumrosen and Nisan (2007); Cramton et al. (2006); Nguyen et al. (2013)). To make things even worse, it remains intractable even for some simple versions of the problem. Limiting the bidding language according to certain combinatorial structures (e.g., matroid constraints, submodularity of bidding functions, etc.) leads sometimes to more efficient algorithms. However, it remains a great challenge to come up either with heuristic algorithms that perform well in practice, especially as the number of items grows, or with fast approximation algorithms that provide provable guarantees for approximating the optimal allocation.

2. *Incentive compatibility and strategic behavior.* Along with the allocation of the items, the auctioneer must also assign a payment to each bidder. Ideally, the property that we would like to enforce is that the bidders do not have incentives to strategize and manipulate the auction (by misreporting their preferences). Mechanisms that satisfy this property are called truthful or incentive compatible. Although incentive compatibility is a well-sought theoretical guarantee, it is often the case that truthful mechanisms lack simplicity and prescribe payments that are computationally difficult to compute. In practice, when incentive compatibility is not feasible, researchers have explored mechanisms that induce bidders to report their preferences as close to the truth as possible. However, even with this relaxed version of truthfulness, devising appropriate payment rules can be a challenging task. Consequently, the efficient computation of such payment rules is an active research topic in the field. Another direction in the study of non-truthful auction protocols is the analysis of stable outcomes. Game theory offers various equilibrium concepts, such as Nash equilibria and Bayes-Nash equilibria. The Nash equilibrium is a well-known and stable outcome in the literature. From a Bayesian perspective, bidders in an auction possess incomplete information regarding the types of other bidders. This uncertainty is represented by assigning values to bidders drawn from probability distributions. For a comprehensive

overview, please refer to the survey conducted by Hartline et al. (2013) in 2013. The concept of a Bayes-Nash equilibrium, coined by Harsanyi (1967), extends the idea of a Nash equilibrium to accommodate this type of information asymmetry. Therefore, to assess the performance of a non-truthful auction, it is crucial to compare the social welfare or revenue produced at an equilibrium with the optimal welfare/revenue. This quantification is known as the Price of Anarchy (PoA) and was introduced by Koutsoupias and Papadimitriou (1999). Considerable progress has been made in the study of PoA for combinatorial auctions, and one promising implication is the existence of simple non-truthful mechanisms that perform well concerning the generated welfare. A summary of related results can be found in Roughgarden et al. (2017).

While a considerable amount of research has been conducted on the issues surrounding auction mechanisms, there remain significant challenges in this field. The primary objective of this dissertation is to design and analyze algorithmic mechanisms, for both forward and reverse (procurement) auctions, based on strong theoretical foundations. These mechanisms address both the computational and incentive concerns, while exploring unexplored areas in various auction domains, some of which combinatorial. Although the primary focus is on basic research in algorithmic game theory, it is worth noting that this scientific field has had a direct impact on real-world applications. Therefore, while immediate applicability is not the main goal of this thesis, it is possible that refined versions of our mechanisms could potentially impact such implementations in the future.

1.1 Structure of Dissertation and Main Themes

The dissertation is divided into two parts. The first part focuses on forward auctions, while the second part centers on procurement auctions. In this section, we provide an overview of the main themes explored in each chapter along with our main results.

In Chapter 2, we focus on the notion of the core, as adjusted in the context of auctions by Ausubel and Milgrom (2002). In this context, the core of an auction is a set of payments which ensure that no coalition of bidders together with the auctioneer can switch to a better outcome of higher revenue for the auctioneer. Ausubel and Milgrom formalized this notion after observing several shortcomings of the truthful Vickrey-Clarke-Groves mechanism (Clarke (1971); Groves (1973); Vickrey (1961)), arguing that, despite its good theoretical properties, it also possesses several inherent drawbacks which hinder its implementation in practice. On the other hand, mechanisms that select payments from the core and simultaneously achieve allocations that maximize welfare, known as *core-selecting mechanisms*, have been known to possess good revenue and fairness guarantees and some of their variants have been used in practice especially for spectrum and other public sector auctions. However, these auctions are generally non-truthful. Thus, current research has focused either on identifying core-selecting mechanisms with minimal incentives to deviate from truth-telling, such as the family of Minimum-Revenue Core-Selecting rules, defined by Day and Raghavan (2007), or on proposing truthful mechanisms whose revenue is competitive against core outcomes, as proposed by Goel et al. (2015). In this chapter, we contribute to both of these directions. We start with a comparative statics analysis on the core polytope, which we utilize in two ways. Firstly, we propose a truthful mechanism that is competitive against the minimum revenue in the core for downward-closed single-parameter domains. Secondly, we study the existence of non-decreasing payment rules, meaning that the payment of each bidder is a non-decreasing function of her bid. This property has been advocated by the related literature but it was an open question if there exist non-decreasing mechanisms. We answer the question in the affirmative, by describing a class of rules with this property.

In Chapter 3, we examine multi-unit auctions, with a particular focus on the discriminatory price auction format, which is commonly used in practice. In this auction format, a set of bidders compete to purchase identical units of a good. While the units are allocated in a manner that maximizes social welfare (similar to the VCG auction), each bidder pays the sum of their bids, making it a direct extension of the non-truthful first price auction protocol. We aim to partially characterize the inefficient mixed Nash equilibria that can arise in this auction format and present improved Price of Anarchy (PoA) bounds beyond the current state of the art in the literature. To achieve this, we undertake an equilibrium analysis for bidders with capped-additive valuations (a subclass of submodular valuations). We first establish a series of properties to help us understand the sources of inefficiency, and then use these results to derive new lower and upper bounds on the PoA of mixed equilibria. For the case of two bidders, we achieve a complete characterization of inefficient equilibria and a tight PoA bound. For multiple bidders, we show that the PoA is strictly worse than the two-bidder case, thereby improving the best known lower bound for submodular valuations. Additionally, we present an improved PoA upper bound for the special case where there exists a "high demand" bidder, which surpasses the state of the art in the literature, due to De Keijzer et al. (2013). Furthermore, we also study Bayes-Nash equilibria and reveal a separation result that had been elusive so far. Specifically, we find that even with only two bidders, the PoA for Bayes-Nash equilibria is strictly worse than that for mixed equilibria. Such separation results are not always true (as shown by Christodoulou et al. (2016b) for simultaneous second-price auctions), and they highlight the further inefficiency introduced by the Bayesian model in this context.

In the remaining chapters of this dissertation, we shift our focus to mechanisms for procurement auctions. In Chapter 4, we examine a covering problem that arises in spatial models of crowdsourcing markets, where tasks are ordered according to geographic or temporal criteria. In this problem, each bidder can provide a certain level of contribution for a subset of consecutive tasks, and each task has a demand requirement. The objective is to identify a set of bidders at minimum cost who can meet all the demand constraints. We propose two truthful mechanisms with approximation guarantees against the optimal cost. The first mechanism achieves a bounded approximation guarantee and improves upon the state of the art, which features a mechanism with an arbitrarily large factor in the worst-case. The second mechanism concerns a class of instances that generalizes the minimum knapsack problem. Specifically, we consider inputs with a constant number of tasks and provide a truthful, fully polynomial-time approximation scheme (FPTAS). Additionally, we discuss the connections between our problem and other well-studied optimization problems.

In Chapter 5, we delve into procurement auction scenarios where bidders have a strict budget constraint. The task for the auctioneer is to purchase resources owned by rational bidders, while keeping the cost within a fixed, predetermined budget and maximizing a given valuation function. In this context, the cost of each resource is private, hence our objective is to design truthful mechanisms that provide a good approximation to the optimal value for the auctioneer and are budget-feasible, meaning that the sum of payments to bidders does not exceed the budget. This framework applies to many well-known optimization problems and is relevant in various application scenarios such as crowdsourcing platforms, where bidders are workers offering tasks, and influence maximization in social networks, where agents are influential users. We specifically concentrate on two types of environments. In the first one, agents are completely divisible or can be viewed as offering a fully divisible service to the auctioneer. In the second one, the auctioneer may select each bidder for a certain number of levels of service out of the maximum available. In this chapter, we obtain truthful and budget-feasible mechanisms that improve upon the best-known approximation guarantees for the auctioneer's total value for both settings.

Finally, in Chapter 6 we conclude the dissertation with an extended discussion of the results along with directions for future research.

1.2 Publications

My work on the topics of this dissertation resulted in the following publications:

1. On Core-Selecting and Core-Competitive Mechanisms for Binary Single-Parameter Auctions, *15th Conference on Web and Internet Economics, WINE 2019, Proceedings,* pp. 271-285 (with E. Markakis). [Chapter 2]

2. Towards a Characterization of Worst Case Equilibria in the Discriminatory Price Auction, *17th Conference on Web and Internet Economics, WINE 2021, Proceedings,* pp. 186-204 (with E. Markakis and A. Sgouritsa). [Chapter 3]

3. On Improved Interval Cover Mechanisms for Crowdsourcing Markets, *15th International Symposium on Algorithmic Game Theory, SAGT 2022, Proceedings,* pp. 94-112 (with E. Markakis and G. Papasotiropoulos). [Chapter 4]

In addition to the above, the following work is currently under submission to a leading conference in the field:

4. Budget-feasible Mechanisms for Fractional Agents and Multiple Levels of Service, *Under Submission* (with G. Amanatidis, S. Klumper, V. Markakis and G. Schäfer) [Chapter 5]

I also worked on other topics during my PhD. These resulted in the following publications:

5. Tight Welfare Guarantees for Pure Nash Equilibria of the Uniform Price Auction, *Theory of Computing Systems, Journal* pp. 1451-1469 (with G. Birmpas, E. Markakis and O. Telelis).

6. Integrating Clinical Data from Hospital Databases, *1st International Workshop on Semantic Web Meets Health Data Management, Workshop Proceedings* (with K. Karozos, I. Spartalis, D. Trivela and V. Vassalos).

Chapter 2

Core-selecting Auctions

2.1 Introduction

In this chapter[1], the focus of our work is the study of auction mechanisms, with competitive revenue performance. Undoubtedly, one of the early landmarks within the field of auction theory, is the VCG mechanism. At the same time however, VCG is rarely preferred in more complex real-life auction scenarios, such as the allocation of spectrum or other governmental licences. The shortcomings that have led to this situation have been well summarized by Ausubel and Milgrom (2006), and one of the most prominent drawbacks is the unacceptably low revenue that VCG generates on instances that do not lack competition. The VCG payment corresponds to the externality a bidder imposes on her competitors, and as a result, one can have even zero payments in worst case, giving rise to *free-riders*, see Ausubel and Baranov (2020).

To counterbalance this issue, Ausubel and Milgrom (2002, 2006) adapted the notion of the *core* from the theory of cooperative games and introduced the class of *core-selecting* mechanisms. These mechanisms first select an optimal (welfare-maximizing) allocation as in VCG, but then the payments are set in a way that no coalition of bidders together with the auctioneer can switch to a better outcome, of higher revenue for the auctioneer. It was argued in Ausubel and Milgrom (2006), that a mechanism can be of suboptimal performance in terms of revenue precisely when the payments it assigns are *not* in the core, which is quite common for VCG when the goods exhibit complementarities. Moreover, an outcome that is not in the core can be perceived as unfair by coalitions of bidders, who could be collectively willing to pay more but still were not taken into consideration. Over the last years, core-selecting mechanisms

[1] A conference version with the results of this chapter appeared in WINE '19 (Markakis and Tsikiridis, 2019).

gained even higher support especially among practitioners, due to the fact that they have been successfully implemented for a number of high-profile spectrum auctions, as well as other public sector auctions in several countries, see Day and Cramton (2012) for an exposition of applications.

Given the good performance of core-selecting auctions in terms of revenue and fairness, the next natural question is whether we can have strategyproof payment rules in the core. Interestingly, for complement-free settings, VCG can lie in the core. However, Goeree and Lien (2016) have shown that when complementarities are present, core payments do not generally yield truthful mechanisms. With this negative aspect in mind, research on this topic has focused mainly on two directions. The first direction concerns a game-theoretic analysis of core-selecting mechanisms so as to identify which payments from the core polytope have more desirable incentive properties. As an example of this approach, it has been shown by Day and Raghavan (2007) that selecting a minimum revenue core outcome also minimizes in a certain sense, the total gain from unilateral deviations. When the minimum revenue within the core does not prescribe a unique outcome, a further refinement needs to take place, which is guided again by the incentives to deviate. This has led to the family of quadratic payment rules (see Section 2.5). In parallel to these results, another way to evaluate such mechanisms is by analyzing the performance of their Bayes Nash equilibria, e.g., Ausubel and Baranov (2020). At the moment, the outcomes of these works have not yet led to definite conclusions and there is still an active debate on what are the best core-selecting mechanisms, given also the recent experimental evaluation of Bünz et al. (2022).

The second direction was initiated by Goel et al. (2015) and concerns the design of truthful (hence, not generally core-selecting) mechanisms whose revenue is competitive against a core outcome. The core benchmark was naturally taken to be the minimum revenue core outcome, given the properties highlighted in the previous paragraph. Hence, a mechanism is then called *α-core-competitive*, when it achieves a $1/\alpha$ fraction of the minimum revenue core outcome, for $\alpha \geq 1$. The main results of Goel et al. (2015) involved the design of core-competitive mechanisms for a particular single-parameter domain motivated by online ad auctions. For more general combinatorial auctions, one can also obtain core-competitive mechanisms using the results of Micali and Valiant (2007), where a stronger benchmark[2] has been considered. This approach is still worth

[2] For single-parameter domains, the benchmark Micali and Valiant (2007) suggested is the maximum social welfare achieved after disregarding the highest-valued bidder.

further investigation, as finding the best ratio against the core benchmark has remained open for various domains of interest.

Our Contribution. We focus on binary single-parameter domains, where each bidder is either accepted or rejected in every outcome. We start in Section 2.3, with providing some insights and properties on the geometry of the core polytope. Our aim is to understand how the polytope is affected by a unilateral deviation of a bidder from a given profile. To do this, we need to perform a comparative statics analysis for the constraints of the core. In the remaining of the article, we then make use of the main results of Section 2.3 in two ways. First, in Section 2.4, we derive a deterministic $O(\log n)$-core-competitive strategyproof mechanism, where n is the number of bidders. So far, only a randomized mechanism with the same ratio was known, implied by Micali and Valiant (2007). Our result is the first deterministic core-competitive mechanism for arbitrary single-parameter domains. It also provides a separation between core-competitiveness, and the stronger benchmark of Micali and Valiant (2007), who provided an impossibility result of $\Omega(n)$ for single-parameter environments. Second, in Section 2.5, we focus on the question of identifying more preferred mechanisms among the possible continuum of minimum revenue core-selecting (MRCS) payment rules. This family has been recognized as having better incentive properties among core-selecting mechanisms, and to refine it even further, we study the existence of *non-decreasing* payment rules, meaning that the payment of each bidder is a non-decreasing function of her bid, see Erdil and Klemperer (2010) and Bosshard et al. (2018). This property has been advocated, among others, for minimizing the marginal incentive to deviate, but it has remained an open question if there exist MRCS rules that satisfy it. We provide a positive answer to this question, by describing a subclass of rules possessing the property, which can be seen as a further refinement towards selecting MRCS mechanisms. Overall, we believe our results shed more light on understanding core-selecting and core-competitive mechanisms, and expect that the properties established here can have even broader appeal and applicability.

2.1.1 Related Work

The core in the context of auctions was introduced by Ausubel and Milgrom (2002, 2006), as a suitable formalism to understand settings where the VCG mechanism underperforms in terms of revenue. Ausubel and Milgrom (2002) also proposed core-selection as a standalone auction design goal by introducing an ascending auction format called the ascending-proxy auction, whose equilibrium outcomes are in the core. The topic soon gained popularity both in theory and in practice, and several follow up works emerged afterwards. A series of important works has focused on exploring different core-selecting Pareto-efficient rules that have minimal incentives to deviate or mechanisms that are core-selecting at equilibrium, see e.g., Day and Cramton (2012); Day and Milgrom (2008); Day and Raghavan (2007); Erdil and Klemperer (2010); Ott and Beck (2013); Parkes (2002) and Ausubel and Baranov (2020). The incentives to deviate have been quantified under different metrics and, to our understanding, no final consensus on the most acceptable metric has been reached. Recently, an experimental comparison of Quadratic payment rules (Day and Cramton (2012)), was conducted by Bünz et al. (2022) in an attempt to offer more insights on that front.

Regarding strategyproofness and core-selection, the work of Goeree and Lien (2016) showed that when VCG payments lie in the core, then this is the only truthful mechanism in the core, whereas when VCG is not in the core, there exists no other truthful mechanism that is core-selecting. This reveals a severe incompatibility between truth-telling and core-selection, especially for auction domains that exhibit complementarities. Such domains can arise naturally in spectrum auctions or in auctions related to online advertising. Nevertheless, efforts have been made to characterize the auction environments where the VCG outcome lies in the core, see e.g. Ausubel and Milgrom (2002); Bikhchandani and Ostroy (2002); Parkes (2002); Sano (2011). In Ausubel and Milgrom (2002), it is shown that in domains where the set of feasible allocations form a matroid (e.g. multi-unit auctions Vickrey (1961)), VCG payments always lie in the core, and, therefore, it does not suffer from the shortcomings we have discussed.

Goel et al. (2015), suggested the use of the minimum revenue core-selecting (MRCS) outcome, as a competitive benchmark for the design of truthful mechanisms. In their work, they focus on the so called *Text and Image Ad-Auction*, a special case of knapsack auctions, where k ad slots are being auctioned and each bidder is known to require 1 or k ad slots. They proposed a truthful deterministic mechanism that is $O(\sqrt{\log k})$-core-competitive and a truthful randomized one which is $O(\log\log k)$-core-competitive and these factors are shown to be tight. To our knowledge, this is the only work where a core benchmark has been explicitly used for truthful revenue maximization.

Clearly, the problem of designing truthful mechanisms for maximizing revenue is a fundamental research direction that has attracted considerable attention, especially since the initial works of Fiat et al. (2002); Goldberg and Hartline (2003); Goldberg et al. (2001), see also Hartline (2013). Later on, Guruswami et al. (2005) introduced envy-free pricing, which formed another important approach with several follow up papers. However, these lines of inquiry have mostly focused on environments where goods are substitutes (for which VCG payments are in the core), whereas the core-benchmark is meaningful for environments with complementarities. For such environments, two notable benchmarks have been proposed by Aggarwal and Hartline (2006) for knapsack auctions and by Micali and Valiant (2007) for general combinatorial auctions. We refer the reader to Goel et al. (2015) for a detailed comparison of all these benchmarks with the minimum-core-revenue benchmark. The two main takeaways of these comparisons are that, the mechanism of Aggarwal and Hartline (2006) performs arbitrarily bad against the MRCS benchmark, whereas the benchmark of Micali and Valiant (2007) is stronger than MRCS. Hence being α-competitive in the sense of Micali and Valiant (2007), implies being α-core-competitive, for $\alpha \geq 1$. In the same work, the authors propose a truthful randomized mechanism for general combinatorial auctions that is $O(\log n)$-competitive against their benchmark and show that this result is tight. Moreover, they complement this finding by showing that no deterministic mechanism can be better that $\Omega(n)$-competitive against their benchmark. Their result implies a randomized $O(\log n)$-core-competitive mechanism for the binary single-parameter setting that we study.

Finally, we stress that by definition the core polytope consists of an exponential number of constraints, which makes its use in mathematical programs challenging. Fortunately, a separation oracle was introduced by Day and Raghavan (2007), but still, each call to the separation oracle requires the solution of a welfare optimization problem. Given these considerations, it is often assumed in the core auction literature that a mechanism has oracle access to a welfare optimization algorithm. In these cases, the complexity measure is the number of oracle calls to the welfare optimization problem. Due to Day and Raghavan (2007), one can deduce then a polynomial upper bound for the number of oracle calls required for the computation of a core point. Obviously, when the underlying welfare optimization can be solved in polynomial time, the mathematical program can also be solved in polynomial time. Recently, Niazadeh et al. (2022) presented a faster algorithm for computing approximate, Pareto-efficient core payments using only a quasi-linear number of oracle calls. Other algorithms that perform well in practice but admit no runtime guarantees are proposed by Day and Raghavan (2007) and Bünz et al. (2015).

2.2 Definitions and Preliminaries

2.2.1 Single-Parameter Domains and Mechanisms

Our work focuses on mechanisms for *binary, single-parameter* domains. We consider a set of bidders $N = \{1, 2, \ldots, n\}$, who can express a request for some type of service (e.g., request for obtaining a set of goods, or access to a facility, etc). Each bidder $i \in N$ has a single private parameter $v_i \geq 0$, which denotes the value derived by bidder i if she is granted the service. The environment is binary in the sense that every bidder will be either accepted or rejected. For every subset $S \subseteq N$, we let $\mathcal{F}(S) \subseteq 2^S$ be the set of feasible allocations for the bidders of S, i.e., the collection of subsets of bidders that can be granted service simultaneously. We assume that $\mathcal{F}(N)$ is *downward-closed*, i.e., for every $X \in \mathcal{F}(N)$ and every $Y \subseteq X$ it holds that $Y \in \mathcal{F}(N)$. We also assume that for every $S \subseteq T$, $\mathcal{F}(S) \subseteq \mathcal{F}(T)$.

An auction mechanism $\mathcal{M} = (X, \mathbf{p})$, in this setting, when run on the set N of agents, consists of an *allocation algorithm* $X : \mathbb{R}^n_+ \mapsto 2^N$ and a *payment rule* $\mathbf{p} : \mathbb{R}^n_+ \mapsto \mathbb{R}^n$. Initially, the auctioneer collects the vector of bids $\mathbf{b} = (b_i)_{i \in N}$, where b_i denotes the bid declared by bidder $i \in N$ (which may differ from v_i). We assume that $b_i \in [0, \infty)$ and that there are no further restrictions on the set of allowed bids. Then, given a bidding profile $\mathbf{b}$, the auctioneer runs the allocation algorithm to determine a feasible allocation $X(\mathbf{b}) \in \mathcal{F}(N)$, and the payment rule to determine the payment vector $\mathbf{p}(\mathbf{b}) = (p_1(\mathbf{b}), \ldots, p_n(\mathbf{b}))$, where $p_i(\mathbf{b})$ is the payment requested by bidder i.

We will often need to refer to sub-instances defined by a coalition of bidders. Given a bidding vector $\mathbf{b}$, and a subset of bidders $S \subseteq N$, we denote by $\mathbf{b}_S$ the projection of $\mathbf{b}$ on S, i.e., the vector containing the bids of the members of S. We also denote by $\mathbf{b}_{-i}$ the vector of all bids except for some bidder i. Given a profile $\mathbf{b}$, if we run a mechanism $\mathcal{M} = (X, \mathbf{p})$ on a sub-instance defined by $S \subseteq N$, then $X(\mathbf{b}_S) \in \mathcal{F}(S)$ will denote the resulting allocation and $\mathbf{p}(\mathbf{b}_S)$ will be the corresponding payment vector for the members of S.

We assume that bidders have quasi-linear utilities and hence, given a mechanism $\mathcal{M} = (X, \mathbf{p})$, the final utility of bidder $i \in N$ for a profile $\mathbf{b}$ is $u_i^{\mathcal{M}}(\mathbf{b}) = v_i - p_i(\mathbf{b})$, when $i \in X(\mathbf{b})$, and 0 otherwise (we enforce that losing bidders do not pay anything). We say that $\mathcal{M}$ satisfies individual rationality if for every profile $\mathbf{b}$ and for every bidder $i \in N$, it holds that $u_i^{\mathcal{M}}(\mathbf{b}) \geq 0$. Additionally, a mechanism is truthful, or strategyproof, if for every bidder $i \in N$, every $b_i \geq 0$ and every profile $\mathbf{b}_{-i}$ it holds that $u_i^{\mathcal{M}}(v_i, \mathbf{b}_{-i}) \geq u_i^{\mathcal{M}}(b_i, \mathbf{b}_{-i})$.

Since we are in a single-parameter environment, in order to design truthful mechanisms, we use the characterization of Myerson (1981). In particular, we say that an allocation algorithm X is *monotone* if for every agent $i \in N$ and every profile $\mathbf{b}$, if $i \in X(\mathbf{b})$, then $i \in X(b_i', \mathbf{b}_{-i})$ for $b_i' \geq b_i$. This means that if an agent is selected in an allocation by declaring a bid b_i, then she should also be selected when declaring a higher bid.

Lemma 2.2.1. *Given a monotone allocation algorithm X, there is a unique payment rule $\mathbf{p}$ such that $\mathcal{M} = (X, \mathbf{p})$ is an incentive compatible and individually rational mechanism. For every profile $\mathbf{b}$ and every bidder $i \in N$ this payment is given by $p_i(\mathbf{b}) = \inf_{b_i' \in [0,b_i]} \{b_i' : i \in X(b_i', \mathbf{b}_{-i})\}$ when $i \in X(\mathbf{b})$, and $p_i(\mathbf{b}) = 0$ otherwise.*

Lemma 2.2.1 is known as Myerson's lemma, and the payments are often referred to as *threshold payments*, since they indicate the threshold below which a bidder stops being selected.

2.2.2 Welfare Maximization and VCG Payments

For a mechanism $M = (X, \mathbf{p})$, the social welfare produced when run on a profile $\mathbf{b}$ (from the viewpoint of the mechanism since each b_i may differ from v_i) is equal to $\sum_{i \in X(\mathbf{b})} b_i$. Among the most desirable outcomes in mechanism design is to select allocations that achieve maximum welfare. In particular, for a profile $\mathbf{b} \in \mathbb{R}^n_+$, and for any coalition $S \subseteq N$ the *optimal* allocation with respect to $\mathbf{b}_S$ is defined as

$$X^*(\mathbf{b}_S) := \arg\max_{T \in \mathcal{F}(S)} \sum_{i \in T} b_i \tag{2.1}$$

We will denote by $W(\mathbf{b}_S)$ the maximum social welfare achieved by an optimal allocation. This is also referred to as the *coalitional value* of S: $W(\mathbf{b}_S) := \max_{T \in \mathcal{F}(S)} \sum_{i \in T} b_i = \sum_{i \in X^*(\mathbf{b}_S)} b_i$. When $S = N$, we refer to an optimal allocation by $X^*(\mathbf{b})$ instead of $X^*(\mathbf{b}_N)$, and to the optimal welfare by $W(\mathbf{b})$.

Regarding tie-breaking issues, throughout this work, we assume that a consistent, deterministic tie-breaking rule is used to select an allocation, whenever there are multiple optimal allocations at a given profile. For example a fixed ordering on subsets of bidders would suffice to resolve ties.

Fact 2.2.1. *Given a bidding vector $\mathbf{b}$, the coalitional value is monotone w.r.t. the set of bidders, i.e. for all $S \subset T \subseteq N$, it holds that $W(\mathbf{b}_S) \leq W(\mathbf{b}_T)$.*

A mechanism is called efficient or welfare-maximizing if for every input profile, it outputs an optimal allocation. The VCG mechanism is the most popular example of an efficient mechanism, where for a bidding profile $\mathbf{b}$, the payment of bidder $i \in X^*(\mathbf{b})$ is the externality she imposes to the other bidders (i.e., the loss to their welfare), defined as

$$p_i^{VCG}(\mathbf{b}) = W(\mathbf{b}_{-i}) - \sum_{j \in X^*(\mathbf{b}) \setminus \{i\}} b_j \tag{2.2}$$

For every other bidder $i \notin X^*(\mathbf{b})$, we have $p_i^{VCG}(\mathbf{b}) = 0$. For the settings we study, one can easily check that the VCG mechanism is individually rational and strategyproof.

2.2.3 Core-selecting Payment Rules

The notion of the core as a solution concept originates from cooperative game theory where it captures the fact that coalitions of agents do not have incentives to appeal to a payoff division. To adjust these ideas to the context of auctions, we first define the following quantity, for every coalition $S \subseteq N$ and bidding profile $\mathbf{b}$.

$$\beta(S, \mathbf{b}) := W(\mathbf{b}_S) - \sum_{j \in X^*(\mathbf{b}) \cap S} b_j.$$

This quantity is a generalization of the VCG payment formula, and can be interpreted as the *collective externality* that bidders in $N \setminus S$ impose to the bidders in S. Indeed, with this notation we can restate VCG payments in Equation (2.2) as $p_i^{VCG}(\mathbf{b}) = \beta(N \setminus \{i\}, \mathbf{b})$, for every bidder $i \in X^*(\mathbf{b})$.

Core-selecting payment rules were initially defined in the space of utility vectors by Ausubel and Milgrom (2002). In our work we follow the equivalent formulation of Day and Cramton (2012) that recasts them to the space of payment vectors. For a profile $\mathbf{b}$, the core polyhedron is defined w.r.t. an optimal allocation $X^*(\mathbf{b})$ as follows

$$CORE(\mathbf{b}) = \{\mathbf{p} \in \mathbb{R}^n : \sum_{j \in X^*(\mathbf{b}) \setminus S} p_j \geq \beta(S, \mathbf{b}) \ \forall S \subseteq N, \ p_j = 0 \ \forall j \notin X^*(\mathbf{b})\}. \tag{2.3}$$

Definition 2.2.1. *A payment rule is called core-selecting, if it is individually rational w.r.t. the reported bids, and* $\mathbf{p}(\mathbf{b}) \in CORE(\mathbf{b})$ *for every profile* $\mathbf{b}$*. Furthermore, a mechanism* $\mathcal{M} = (X, \mathbf{p})$ *is a core-selecting mechanism if (i)* $X(\mathbf{b})$ *is a welfare-maximizing allocation for every profile* $\mathbf{b}$*, and (ii)* $\mathbf{p}$ *is a core-selecting payment rule.*

The constraints of the core polytope in (2.3) require that every coalition of bidders pays at least their collective externality or, in other words, the damage their presence inflicts on the remaining bidders. To provide more intuition, another way to view this is that under a core payment vector, and if bidders are truthful, then every coalition S, together with the auctioneer creates a collective utility at least as high as $W(\mathbf{b}_S)$, which is the best they could achieve if they ran an auction among themselves. In more detail, if u_0 is the auctioneer's utility, which equals $\sum_{j \in N} p_j$, the core constraint for S in (2.3) is equivalent to $u_0 + \sum_{j \in S} u_j(\mathbf{b}) \geq W(\mathbf{b}_S)$, under the assumption that bidders are truthful. Using this formulation, and individual rationality, if the outcome of a mechanism is not in the core, this implies that $u_0 < W(\mathbf{b}_S)$. Hence, there was a coalition that could offer the auctioneer a higher revenue and yet this did not happen.

It is easily verifiable that the pay-your-bid auction, where every winning bidder pays her bid, coupled with the optimal allocation, is a core-selecting mechanism. This rule is sometimes mentioned in the literature as the *seller-optimal* core-selecting payment rule since it maximizes the revenue of the auctioneer with respect to the declared bids. Given that core-selecting mechanisms are not truthful in general, see also Goeree and Lien (2016), a natural quest has been to identify payments in the core where the incentives to misreport are minimized. Formalizing this idea, Day and Milgrom (2008) proposed the use of *Pareto-efficient* core payments, which, in the core-literature are also referred to as *bidder-optimal* payment rules.

Definition 2.2.2 (Pareto-efficient core payments Day and Milgrom (2008))**.** *Let* $\mathbf{b}$ *be a bidding profile and* $\mathbf{p} \in CORE(\mathbf{b})$*. We say that* $\mathbf{p}$ *is a Pareto-efficient core payment if for every payment* $\mathbf{p}'$ *such that* $p'_i \leq p_i$ *for every bidder* $i \in X^*(\mathbf{b})$*, and with strict inequality for at least one bidder, we have that* $\mathbf{p}' \notin CORE(\mathbf{b})$*.*

A prominent class of Pareto-efficient payment rules in the literature are the minimum revenue core-selecting (MRCS) rules, i.e., the minimum revenue points in the core, first introduced by Day and Raghavan (2007). An *MRCS* rule assigns payments given a profile $\mathbf{b}$, that are optimal solutions of the linear program:

$$\min_{\mathbf{p} \in \mathbb{R}^n} \left\{ \sum_{j \in N} p_j : \quad \mathbf{p} \in CORE(\mathbf{b}), \quad \mathbf{p} \leq \mathbf{b} \right\}. \tag{2.4}$$

It is trivial to check that this is indeed a Pareto-efficient core payment rule. We denote by $\mathsf{MREV}(\mathbf{b})$ the optimal value of the objective function in (2.4). As shown by Day and Milgrom (2008), the minimum core revenue still gives a better revenue guarantee than VCG, i.e., for a profile $\mathbf{b}$, $\mathsf{MREV}(\mathbf{b}) \geq \sum_{i \in N} p_i^{VCG}(\mathbf{b})$. A further advantage of MRCS rules, established by Day and Raghavan (2007), is that they minimize the *total gains from unilateral deviations*. Finally, an interesting note by Day and Milgrom (2008) is that whenever the VCG payment belongs to the core, it is the unique MRCS rule [3], because it is the unique Pareto-efficient point. Otherwise, the linear program in (2.4) has a continuum of solutions and a secondary refinement is required in practice to select a particular MRCS payment rule in a disciplined way. We continue this discussion in Section 2.5, by studying Quadratic Payment Rules, a class of core payment rules which are often used as such a refinement.

2.2.4 Core-competitive Mechanisms

A different approach has been initiated by Goel et al. (2015) concerning revenue guarantees in relation to the core outcomes. Since core-selecting mechanisms are not always truthful (despite their good incentive properties), Goel et al. (2015) proposed the design of truthful mechanisms whose revenue is competitive against a core outcome. Given the discussion in Section 2.2.3, it is quite natural to use as a core benchmark the revenue attained by the MRCS rules. One can evaluate then truthful mechanisms as follows:

Definition 2.2.3 (Goel et al. (2015))**.** *Let $\mathcal{M} = (X, \mathbf{p})$ be a truthful mechanism. We say that $\mathcal{M}$ is α-core-competitive, with $\alpha \geq 1$, if for any bidding profile $\mathbf{b}$, it assigns a payment vector $\mathbf{p}(\mathbf{b})$, such that $\sum_{i=1}^{n} p_i(\mathbf{b}) \geq \frac{1}{\alpha}\mathsf{MREV}(\mathbf{b})$.*

We will follow this approach in Section 2.4 for single-parameter domains.

2.3 Insights on the Geometry of the Core

The goal of this section is to establish insights and properties for the core polytope, and in particular with regard to how the polytope changes when a single bidder declares a higher bid, i.e., we study the relation between $CORE(\mathbf{b})$ and $CORE(b_i', \mathbf{b}_{-i})$, with $b_i' > b_i$ for some $i \in X^*(\mathbf{b})$. The results we present here will be the key ingredients to prove the two main results of our work in Section 2.4 and Section 2.5.

Throughout this section, we assume that for all payment vectors that we consider, we have set $p_j = 0$ for every $j \notin X^*(\mathbf{b})$, for a profile $\mathbf{b}$.

[3]In this case the total gains from unilateral deviations are actually 0, as VCG is an incentive compatible mechanism.

2.3.1 Warm up: Pareto-efficiency and Individual Rationality within the Core

According to Definition 2.2.1, a core-selecting mechanism must be individually rational with respect to the reported bids. In this section, we show that for Pareto-efficient core-selecting payment rules, we have individual rationality for free, and there is, in fact, no need for the auctioneer to explicitly enforce the IR constraints. We start with Lemma 2.3.1, which is a straightforward characterization of Pareto-efficient payment rules. It simply says that for every winning bidder i, at least one core constraint that contains the payment of i must be satisfied with equality.

Lemma 2.3.1. *Let* $\mathbf{b}$ *be a bidding profile, and* $\mathbf{p} \in CORE(\mathbf{b})$*. The vector* $\mathbf{p}$ *is a Pareto-efficient core payment if and only if for every bidder* $i \in X^*(\mathbf{b})$ *there exists a coalition* $S \subset N$ *with* $i \notin S$ *such that*

$$\sum_{j \in X^*(\mathbf{b}) \setminus S} p_j = \beta(S, \mathbf{b}). \tag{2.5}$$

Proof. ($\Rightarrow$) For every bidder $i \in X^*(\mathbf{b})$, let $S_i \subset N$ be a coalition that satisfies (2.5) and does not include bidder i. Suppose for contradiction that $\mathbf{p}$ is not Pareto-efficient. Then, there exists a bidder $k \in X^*(\mathbf{b})$ and $p'_k < p_k$, such that $(p'_k, \mathbf{p}_{-k}) \in CORE(\mathbf{b})$. However, this vector of payments cannot be feasible since

$$\sum_{j \in X^*(\mathbf{b}) \setminus (S_k \cup \{k\})} p_j \; + \; p'_k < \sum_{j \in X^*(\mathbf{b}) \setminus S_k} p_j = \beta(S_k, \mathbf{b}),$$

which is a violation of the core constraint in (2.3) for the coalition S_k. This implies that $(p'_k, \mathbf{p}_{-k}) \notin CORE(\mathbf{b})$, a contradiction.

($\Leftarrow$) Suppose for contradiction that there exists a bidder $k \in X^*(\mathbf{b})$, such that for every coalition $S \subset N$, that does not include k, Equation (2.5) does not hold, i.e., $\sum_{j \in X^*(\mathbf{b}) \setminus S} p_j > \beta(S, \mathbf{b})$. Then, there exists $p'_k < p_k$ such that the payment $(p'_k, \mathbf{p}_{-k})$ satisfies the core constraint for every coalition $S \subset N$ that does not include bidder k. This implies that $(p'_k, \mathbf{p}_{-k}) \in CORE(\mathbf{b})$ (since the remaining core constraints for coalitions that contain k are satisfied by the fact that $\mathbf{p}$ belongs to the core). But this means that $\mathbf{p}$ is not Pareto-efficient, a contradiction.

□

We now show that Pareto-efficiency within $CORE(\mathbf{b})$ implies individual rationality with respect to $\mathbf{b}$.

Lemma 2.3.2. *A payment rule that for any given profile* $\mathbf{b}$ *prescribes a Pareto-efficient vector of payments* $\mathbf{p} \in CORE(\mathbf{b})$*, satisfies* $p_i \leq b_i$ *for every bidder* $i \in X^*(\mathbf{b})$*.*

Proof. Given a bidding profile $\mathbf{b}$, let $\mathbf{p} \in CORE(\mathbf{b})$ be a Pareto-efficient payment. Fix a bidder $i \in X^*(\mathbf{b})$. Since $\mathbf{p}$ is Pareto-efficient, by Lemma 2.3.1 there exists a coalition $S \subseteq N$, that does not include i, for which $\sum_{j \in X^*(\mathbf{b}) \setminus S} p_j = \beta(S, \mathbf{b})$. We distinguish the following cases:

1. $S = N \setminus \{i\}$. In this case, bidder i is asked to pay precisely her VCG payment since
$$\sum_{j \in X^*(\mathbf{b}) \setminus (N \setminus \{i\})} p_j = p_i = \beta(N \setminus \{i\}, \mathbf{b}) = p_i^{VCG}(\mathbf{b}) \leq b_i$$
and the last inequality holds since VCG is an individually rational mechanism, due to Fact 2.2.1.

2. $S \subset N \setminus \{i\}$. Consider the coalition $S \cup \{i\} \subset N$. Since $\mathbf{p} \in CORE(\mathbf{b})$, by (2.3) we have
$$\begin{aligned}\sum_{j \in X^*(\mathbf{b}) \setminus (S \cup \{i\})} p_j \geq \beta(S \cup \{i\}, \mathbf{b}) &= W(\mathbf{b}_{S \cup \{i\}}) - \sum_{j \in X^*(\mathbf{b}) \cap (S \cup \{i\})} b_j \geq W(\mathbf{b}_S) - \sum_{j \in X^*(\mathbf{b}) \cap S} b_j - b_i \\ &= \beta(S, \mathbf{b}) - b_i = \sum_{j \in X^*(\mathbf{b}) \setminus S} p_j - b_i.\end{aligned}$$
The second inequality follows from Fact 2.2.1 and the last equality from the fact that S satisfies Equation (2.5) by assumption. By rearranging terms we obtain that $p_i \leq b_i$.

□

Lemma 2.3.2 allows us to omit individual rationality constraints and focus only on the core constraints, when reasoning about Pareto-efficient payment rules. Moreover, using the fact that MRCS payments are Pareto-efficient, we can now simplify the linear program of Equation (2.4).

Corollary 2.3.1. *A payment rule is* MRCS *if, given a profile* $\mathbf{b}$*, it assigns payments that are optimal solutions of the linear program*

$$\min_{\mathbf{p} \in \mathbb{R}^n} \left\{ \sum_{j \in N} p_j : \quad \mathbf{p} \in CORE(\mathbf{b}) \right\}. \tag{2.6}$$

2.3.2 The Effects of Unilateral Deviations on the Core

We now aim to understand how the core polytope that forms after a unilateral deviation of a winning bidder is related to the initial core polytope. Initially, we focus on how each of the constraints in the polytope is modified and perform a sensitivity analysis for the term $\beta(S, \mathbf{b})$, the collective externality that appears in the core constraints in (2.3), for every $S \subseteq N$. Hence, for a given profile $\mathbf{b}$, a bidder $i \in X^*(\mathbf{b})$ and a bid $b_i' > b_i$, we are interested in the relationship between $\beta(S, \mathbf{b})$ and $\beta(S, (b_i', \mathbf{b}_{-i}))$.

To proceed, our analysis will be dependent on the following quantity, defined for an input profile $\mathbf{b}$, a bidder $i \in X^*(\mathbf{b})$, and a coalition $S \subseteq N$ with $i \in S$.

$$t_i(\mathbf{b}_{S\setminus\{i\}}) = \min\{z : \exists T \subseteq S, s.t.\ i \in T \text{ and } \sum_{j \in T\setminus\{i\}} b_j + z = W(z, \mathbf{b}_{S\setminus\{i\}})\} \tag{2.7}$$

The term $t_i(\mathbf{b}_{S\setminus\{i\}})$ is the minimum bid i should declare to be included in some optimal allocation in an auction where only the bidders from S are present. This is precisely the Myerson threshold payment, for mechanisms where the allocation algorithm produces an optimal allocation when run on input profile $\mathbf{b}_S$. Namely[4], if $i \in X^*(b_S)$, then $t_i(\mathbf{b}_{S\setminus\{i\}}) = p_i^{VCG}(\mathbf{b}_S)$. The following simple lemma can be easily established for the optimal welfare of coalition S.

Lemma 2.3.3. *Given a bidding vector $\mathbf{b}$, a bidder $i \in X^*(\mathbf{b})$ and a bid b_i' such that $0 \leq b_i' \leq t_i(\mathbf{b}_{S\setminus\{i\}})$, it holds that*

$$W(b_i', \mathbf{b}_{S\setminus\{i\}}) = W(\mathbf{b}_{S\setminus\{i\}}). \tag{2.8}$$

Proof. When $b_i' < t_i(\mathbf{b}_{S\setminus\{i\}})$, by the definition of $t_i(\mathbf{b}_{S\setminus\{i\}})$ in Equation (2.7), bidder i is not included in any optimal allocation when only bidders in S are present. Therefore, since i does not generate any value to the coalition S, her existence might as well be ignored and Equation (2.8) holds.

When $b_i' = t_i(\mathbf{b}_{S\setminus\{i\}})$, using again the definition of $t_i(\mathbf{b}_{S\setminus\{i\}})$ in Equation (2.7), bidder i *is* included in an optimal allocation for the auction among the bidders in S. However, we claim that at the same time there exists another optimal allocation when only bidders in S are present, that does not include i.

Suppose for contradiction that this is not the case. This means that bidder i belongs to *all* optimal allocations among bidders in S, when bidding $t_i(\mathbf{b}_{S\setminus\{i\}})$ against the bids

[4] It can also happen that due to tie-breaking, $X^*(\mathbf{b}_S)$ does not coincide with T from (2.7), and thus $i \notin X^*(\mathbf{b}_S)$, in which case $t_i(\mathbf{b}_{S\setminus\{i\}}) \neq p_i^{VCG}(\mathbf{b}_S) = 0$.

$\mathbf{b}_{S\setminus\{i\}}$. Then, bidder i can bid $t_i(\mathbf{b}_{S\setminus\{i\}}) - \epsilon$, for a sufficiently small $\epsilon > 0$ and remain a member of all optimal allocations among bidders in S. This however, is a contradiction, since $t_i(\mathbf{b}_{S\setminus\{i\}})$ is defined as the minimum bid i can issue and be part of one optimal allocation for set S.

Therefore, since there exists an optimal allocation among bidders in S without i for the profile $(t_i(\mathbf{b}_{S\setminus\{i\}}), \mathbf{b}_{S\setminus\{i\}})$, the coalitional value $W(t_i(\mathbf{b}_{S\setminus\{i\}}), \mathbf{b}_{S\setminus\{i\}})$ can be achieved by the bidders of this particular optimal allocation that does not include i. Hence bidder i can be ignored and Equation (2.8) still holds. □

The following key lemma encapsulates the effects on the collective externality of S by a unilateral deviation of a bidder $i \in S$.

Lemma 2.3.4 (Sensitivity analysis for $\beta(S, \mathbf{b})$)**.** *Let* $\mathbf{b}$ *be a bidding profile. Fix a bidder* $i \in X^*(\mathbf{b})$*, and a coalition* $S \subseteq N$*. Suppose that bidder* i *unilaterally deviates to* $b'_i > b_i$*. Then:*

1. *If* $i \notin S$ *or if* $i \in S$ *and* $b_i \geq t_i(\mathbf{b}_{S\setminus\{i\}})$ *then*
$$\beta(S, (b'_i, \mathbf{b}_{-i})) = \beta(S, \mathbf{b}). \tag{2.9}$$

2. *If* $i \in S$ *and* $b_i < t_i(\mathbf{b}_{S\setminus\{i\}})$ *then*
$$\beta(S, (b'_i, \mathbf{b}_{-i})) = \beta(S, \mathbf{b}) - (\min\{b'_i, t_i(\mathbf{b}_{S\setminus\{i\}})\} - b_i) \tag{2.10}$$

Proof. Since the optimal allocation algorithm is monotone and $i \in X^*(\mathbf{b})$, it holds that $X^*(b'_i, \mathbf{b}_{-i}) = X^*(\mathbf{b})$, for $b'_i > b_i$. We distinguish the following cases concerning bidder i and the coalition S:

1. *$i \notin S$:* Then bidder i has no influence on $\beta(S, \mathbf{b})$. By the monotonicity of the allocation algorithm, we have
$$\beta(S, (b'_i, \mathbf{b}_{-i})) = W(\mathbf{b}_S) - \sum_{j \in X^*(b'_i, \mathbf{b}_{-i}) \cap S} b_j = W(\mathbf{b}_S) - \sum_{j \in X^*(\mathbf{b}) \cap S} b_j = \beta(S, \mathbf{b}).$$

2. *$i \in S$ and $b_i \geq t_i(\mathbf{b}_{S\setminus\{i\}})$:* By the definition of $t_i(\mathbf{b}_{S\setminus\{i\}})$, we know there exists an optimal allocation $T \in \mathcal{F}(S)$ with respect to $\mathbf{b}_S$, and with $i \in T$. Hence $\sum_{j \in T} b_j = W(\mathbf{b}_S) = \sum_{j \in X^*(\mathbf{b})} b_j$. By the monotonicity of the optimal allocation algorithm, it is true that T is also optimal with respect to $(b'_i, \mathbf{b}_{S\setminus\{i\}})$, for all

$b'_i > b_i$. For brevity in the algebraic manipulations below, we denote by X^* the optimal allocation $X^*(\mathbf{b})$. Hence,

$$\beta(S,(b'_i,\mathbf{b}_{-i})) = W(b'_i,\mathbf{b}_{S\setminus\{i\}}) - \sum_{j\in X^*\cap(S\setminus\{i\})} b_j - b'_i = \sum_{j\in T\setminus\{i\}} b_j + b'_i - \sum_{j\in X^*\cap(S\setminus\{i\})} b_j - b'_i$$
$$= \sum_{j\in T\setminus\{i\}} b_j + b_i - \sum_{j\in X^*\cap(S\setminus\{i\})} b_j - b_i = W(\mathbf{b}_S) - \sum_{j\in X^*\cap S} b_j = \beta(S,\mathbf{b}).$$

The second equality holds because $X^*(b'_i,\mathbf{b}_{-i}) = X^*(\mathbf{b})$, and the third equality follows since we argued that T is also optimal for $(b'_i,\mathbf{b}_{S\setminus\{i\}})$.

3. *$i \in S$ and $b_i < t_i(\mathbf{b}_{S\setminus\{i\}})$:* In this case, bidder $i \in X^*(\mathbf{b})$ is not included in any optimal allocation with respect to $\mathbf{b}_S$. We need to consider two subcases. When $b'_i \le t_i(\mathbf{b}_{S\setminus\{i\}})$ we have:

$$\beta(S,(b'_i,\mathbf{b}_{-i})) = W(b'_i,\mathbf{b}_{S\setminus\{i\}}) - \sum_{j\in X^*\cap(S\setminus\{i\})} b_j - b'_i = W(\mathbf{b}_{S\setminus\{i\}}) - \sum_{j\in X^*\cap(S\setminus\{i\})} b_j - b'_i$$
$$= W(\mathbf{b}_S) - \sum_{j\in X^*\cap S} b_j - (b'_i - b_i) = \beta(S,\mathbf{b}) - (b'_i - b_i). \tag{2.11}$$

The second and the third equalities follow from Lemma 2.3.3, since both $b'_i \le t_i(\mathbf{b}_{S\setminus\{i\}})$ and $b_i < t_i(\mathbf{b}_{S\setminus\{i\}})$.

In the second subcase, when $b'_i > t_i(\mathbf{b}_{S\setminus\{i\}})$, the unilateral deviation of i enables her to be included in an optimal allocation among bidders in S. Then, we can see that

$$\beta(S,(b'_i,\mathbf{b}_{-i})) = \beta(S,(t_i(\mathbf{b}_{S\setminus\{i\}}),\mathbf{b}_{-i})) = \beta(S,\mathbf{b}) - (t_i(\mathbf{b}_{S\setminus\{i\}}) - b_i).$$

The first equality above follows by applying Equation (2.9) for the profile $(t_i(\mathbf{b}_{S\setminus\{i\}}),\mathbf{b}_{-i})$, whereas the second equality follows from Equation (2.11), using $b'_i = t_i(\mathbf{b}_{S\setminus\{i\}})$. Summarizing the two subcases, we obtain $\beta(S,(b'_i,\mathbf{b}_{-i})) = \beta(S,\mathbf{b}) - (\min\{b'_i, t_i(\mathbf{b}_{S\setminus\{i\}})\} - b_i)$, which completes the proof.

□

Lemma 2.3.4 enables us to prove the two theorems that follow. The first theorem says that for binary single-parameter domains, when a winning bidder declares a higher bid, the space of core payments can only get larger.

Theorem 2.3.1. *Let $\mathbf{b}$ be a bidding profile and $i \in X^*(\mathbf{b})$. Then, for every $b'_i > b_i$, $CORE(\mathbf{b}) \subseteq CORE(b'_i,\mathbf{b}_{-i})$.*

Proof. Note first that for $b_i' > b_i$, since the optimal allocation algorithm is monotone and $i \in X^*(\mathbf{b})$, it holds that $X^*(b_i', \mathbf{b}_{-i}) = X^*(\mathbf{b})$. Consider now a vector $\mathbf{p}$ in $CORE(\mathbf{b})$. We will show that $\mathbf{p}$ is also a member of $CORE(b_i', \mathbf{b}_{-i})$. This is equivalent to showing that for every $S \subseteq N$, $\mathbf{p}$ satisfies

$$\sum_{j \in X^*(\mathbf{b}) \setminus S} p_j \geq \beta(S, (b_i', \mathbf{b}_{-i})).$$

When $S \subseteq N$ is a coalition such that either $i \notin S$ or $i \in S$ and $b_i \geq t_i(\mathbf{b}_{S \setminus \{i\}})$, then by Lemma 2.3.4, we immediately have

$$\sum_{j \in X^*(\mathbf{b}) \setminus S} p_j \geq \beta(S, \mathbf{b}) = \beta(S, (b_i', \mathbf{b}_{-i})).$$

On the other hand, when $i \in S$ and $b_i < t_i(\mathbf{b}_{S \setminus \{i\}})$, then again by Lemma 2.3.4 (Equation (2.10)), and since $\mathbf{p} \in CORE(\mathbf{b})$, we obtain

$$\sum_{j \in X^*(\mathbf{b}) \setminus S} p_j \geq \beta(S, \mathbf{b}) = \beta(S, (b_i', \mathbf{b}_{-i})) + \min\{b_i', t_i(\mathbf{b}_{S \setminus \{i\}})\} - b_i > \beta(S, (b_i', \mathbf{b}_{-i})),$$

where the last inequality follows from the facts that $b_i' > b_i$ and $t_i(\mathbf{b}_{S \setminus \{i\}}) > b_i$. □

We note that the set inclusion claimed in Theorem 2.3.1 can be strict, i.e., there exists a bidding profile $\mathbf{b}$, a bidder $i \in X^*(\mathbf{b})$ and a $b_i' > b_i$ such that $CORE(\mathbf{b}) \subset CORE(b_i', \mathbf{b}_{-i})$. We refer to Section 2.3.3 for more on this.

The next theorem says that in order for a payment to belong to the enlarged polyhedron $CORE(b_i', \mathbf{b}_{-i})$, albeit not to $CORE(\mathbf{b})$, bidder i should be charged a payment that exceeds her previous bid.

Theorem 2.3.2. *Let* $\mathbf{b}$ *be a bidding profile and fix a bidder* $i \in X^*(\mathbf{b})$. *For* $b_i' > b_i$, *let* $\mathbf{p} \in CORE(b_i', \mathbf{b}_{-i})$ *be a payment vector with* $p_i \leq b_i$. *Then,* $\mathbf{p} \in CORE(\mathbf{b})$.

Proof. Suppose for contradiction that this is not true, i.e., for a deviating bidder $i \in X^*(\mathbf{b})$, there exists a payment profile $\mathbf{p} \in CORE(b_i', \mathbf{b}_{-i})$ with $p_i \leq b_i$, such that $\mathbf{p} \notin CORE(\mathbf{b})$. This implies that there exists a coalition $S \subseteq N$ for which the constraint in (2.3) is violated, i.e.

$$\sum_{j \in X^*(\mathbf{b}) \setminus S} p_j < \beta(S, \mathbf{b}). \tag{2.12}$$

If S is a coalition with $i \notin S$ or $i \in S$ but with $b_i \geq t_i(\mathbf{b}_{S\setminus\{i\}})$, since $\mathbf{p} \in CORE(b'_i, \mathbf{b}_{-i})$, by Equation (2.3) we obtain

$$\sum_{j \in X^*(\mathbf{b})\setminus S} p_j \geq \beta(S, (b'_i, \mathbf{b}_{-i})) = \beta(S, \mathbf{b}),$$

where the last equality is due to (2.9) of the Lemma 2.3.4. However, this contradicts (2.12).

Consider now the case when coalition S is such that $i \in S$ and $b_i < t_i(\mathbf{b}_{S\setminus\{i\}})$. Note that this implies that S cannot be the singleton coalition $\{i\}$, as the minimum bid i must bid to be in an optimal allocation on her own is 0, and we would have $b_i < 0$. Consider the constraint $S \setminus \{i\}$. We have

$$\begin{aligned}\sum_{j \in X^*(\mathbf{b})\setminus S} p_j + p_i &\geq \beta(S \setminus \{i\}, (b'_i, \mathbf{b}_{-i})) \\ &= W(\mathbf{b}_{S\setminus\{i\}}) - \sum_{j \in X^*(\mathbf{b})\cap(S\setminus\{i\})} b_j \\ &= W(\mathbf{b}_S) - \sum_{j \in X^*(\mathbf{b})\cap S} b_j + b_i = \beta(S, \mathbf{b}) + b_i. \end{aligned} \tag{2.13}$$

The inequality follows from applying (2.3) for the constraint corresponding to $S \setminus \{i\}$, and the second equality from Lemma 2.3.3, since $b_i < t_i(\mathbf{b}_{S\setminus\{i\}})$. By combining (2.12) and (2.13) we obtain

$$\beta(S, \mathbf{b}) + b_i < p_i + \beta(S, \mathbf{b}),$$

which is a contradiction. □

Theorem 2.3.2 will be particularly useful in Section 2.5.

2.3.3 A Comment on Revenue Monotonicity of MRCS

Theorem 2.3.1 has the following corollary for MRCS core payments, defined in (2.6).

Corollary 2.3.2. *Let* $\mathbf{b}$ *be a bidding profile. Suppose bidder* $i \in X^*(\mathbf{b})$, *and let* $b'_i > b_i$. *Then*

$$\mathsf{MREV}(b'_i, \mathbf{b}_{-i}) \leq \mathsf{MREV}(\mathbf{b}). \tag{2.14}$$

Proof. Let $\mathbf{p}^* \in CORE(\mathbf{b})$ be an MRCS payment (an optimal solution to the linear program in Equation (2.6)) for the profile $\mathbf{b}$. Moreover, let $\mathbf{p}' \in CORE(b'_i, \mathbf{b}_{-i})$ be an MRCS solution for the profile $(b'_i, \mathbf{b}_{-i})$. By Theorem 2.3.1, it is true that every feasible payment vector $\mathbf{p} \in CORE(\mathbf{b})$ is also in $CORE(b'_i, \mathbf{b}_{-i})$. Therefore, since $\mathbf{p}^* \in CORE(\mathbf{b})$, we have that $\mathbf{p}^* \in CORE(b'_i, \mathbf{b}_{-i})$. Hence,

$$\mathsf{MREV}(\mathbf{b}) = \sum_{j \in N} p^*_j \geq \sum_{j \in N} p'_j = \mathsf{MREV}(b'_i, \mathbf{b}_{-i}).$$

The inequality follows since $\mathbf{p}'$ is an optimal solution for MRCS, a linear program with a minimization objective, for the profile $(b'_i, \mathbf{b}_{-i})$.

□

Corollary 2.3.2 states that a higher willingness to pay by a winning bidder will never lead to an increase of the auctioneer's revenue under MRCS, for all binary single-parameter auctions. This result may look counter-intuitive on a first reading, especially for instances where (2.14) is satisfied with strict inequality. In the literature, this phenomenon is commonly mentioned as a violation of *revenue-monotonicity*. There are several facets in studying revenue monotonicity, as it concerns the effects on the revenue when adding new bidders, or increasing the offers of the current bidders, or more generally when changing some parameter of the auction. The version we consider here is referred to as *bidder revenue monotonicity*, Beck and Ott (2010).

Pareto-efficient rules that assign payments in the core have been known to be susceptible to violating this property. Namely, it has been shown by Beck and Ott (2010); Lamy (2010), that in a multi-parameter domain with at least three items, revenue-monotonicity is violated. Here, we strengthen these results by showing that revenue-monotonicity can be violated in single-parameter auctions as well: we construct an instance with single-minded bidders, where a unilateral bid increase by a winning bidder strictly decreases the MRCS revenue. The proof of the following proposition can be found in Section A.1 of Appendix A.

Proposition 2.3.1. *In binary single-parameter auction environments, there exist examples where MRCS rules violate revenue-monotonicity, i.e., Equation* (2.14) *is satisfied with strict inequality.*

Aside from this discussion, and quite surprisingly, Corollary 2.3.2 also plays a crucial role in the analysis of a core-competitive mechanism that we present in Section 2.4.

2.4 An $O(\log n)$-core-competitive Mechanism

In this section, we present a first application of the properties we derived in Section 2.3. We move away from core-selecting mechanisms with the goal of designing truthful mechanisms that achieve a good revenue approximation with respect to core outcomes. Our main result is a deterministic, truthful mechanism that is also $O(logn)$-core-competitive with respect to the MRCS benchmark. Although we are not analyzing core-selecting mechanisms in this section, the properties of the core, identified in Section 2.3 (namely Corollary 2.3.2 of Theorem 2.3.1), will still come in handy for the analysis of our mechanism.

Mechanism $\hat{\mathcal{M}} = (\hat{X}, \hat{\mathbf{p}})$.

Input: A value profile $\mathbf{v} \in \mathbb{R}^n_+$.
Output: An allocation $\hat{X}(\mathbf{v}) \in \mathcal{F}(N)$ and a vector of payments $\hat{\mathbf{p}}(\mathbf{v})$.

1. Find an optimal allocation $X^*(\mathbf{v})$. Let $m = |X^*(\mathbf{v})|$.
2. Let $s_1, \ldots, s_m$ be an ordering of the bidders in $X^*(\mathbf{v})$, such that $v_{s_1} \geq \cdots \geq v_{s_m}$. Let $k \leq m$ be the largest index such that
$$k \cdot v_{s_k} \geq \frac{\mathsf{MREV}(\mathbf{v})}{H_n}. \tag{2.15}$$
3. Set $\hat{X}(\mathbf{v}) = \{s_1, \ldots, s_k\}$ and $\hat{\mathbf{p}}(\mathbf{v})$ according to Myerson's Lemma (Lemma 2.2.1).

Figure 2.1 An $O(\log n)$-core-competitive and strategyproof mechanism.

The mechanism is described in Figure 2.1, where we have used the real valuation profile for the bidders, $\mathbf{b} = \mathbf{v}$ (since we will establish that the mechanism is truthful). We also denote the n-th harmonic number by $H_n = \sum_{i=1}^{n} 1/i = \Theta(\log n)$. In the first step, we find a welfare-maximizing allocation. However, instead of allocating to all bidders in the optimal solution, in the second step the mechanism disqualifies some bidders with values that do not meet a certain cutoff. In case of ties in step 2, it suffices to have a consistent deterministic tie-breaking rule, e.g., given by an ordering on the set of bidders. The mechanism tries, in some sense, to be as inclusive as possible, as long as the value of the last member of $\hat{X}(\mathbf{v})$ is not too small for the coalition to collectively miss the cutoff.

The main result of this section is the following:

Theorem 2.4.1. *The mechanism $\hat{\mathcal{M}}$ is individually rational, truthful, and $O(\log n)$-core-competitive.*

Sections 2.4.2 and 2.4.3 are devoted to the proof of Theorem 2.4.1. Before proceeding to the proof, we discuss some aspects of the mechanism, and comparisons with other results.

2.4.1 Remarks on Tightness, Complexity and Other Implications

Our mechanism is applicable to all binary single-parameter auction domains with a downward-closed set of feasible allocations $\mathcal{F}(N)$. In particular, for environments where the VCG payments are not in the core, such as environments that exhibit complementarities or where the set $\mathcal{F}(N)$ is not a matroid, our mechanism is the only known deterministic strategyproof mechanism that is competitive with regard to the MRCS benchmark for arbitrary binary single-parameter domains.

Prior to our work, a randomized, exponential, strategyproof mechanism was known that is also $O(\log n)$-core-competitive by Micali and Valiant (2007). Hence, a consequence of our result is that it resolves the question of whether deterministic mechanisms can attain the same performance as randomized ones w.r.t. core-competitiveness. Moreover, the result of Micali and Valiant (2007) is based on establishing competitiveness against a stronger benchmark, which is the maximum welfare that can be achieved when the highest bidder is ignored. We point the reader to Goel et al. (2015) for more detailed comparisons between this benchmark and MRCS. With this in mind, what we find most valuable for our result is that it provides a strict separation on the performance of the two benchmarks, since Micali and Valiant (2007) show that deterministic mechanisms cannot perform better than $\Omega(n)$ for their benchmark, even for single-parameter domains. Hence, our mechanism illustrates that the benchmark of Micali and Valiant (2007) is much more stringent, whereas the MRCS core benchmark is more amenable to multiplicative approximations and might be more suitable for revenue maximization.

Regarding complexity, our mechanism clearly has a worst-case exponential running time, because it requires the computation of an optimal allocation and of $\mathsf{MREV}(\mathbf{b})$. As already discussed in Section 2.1.1, the bottleneck of having to solve the welfare maximization problem for various subsets of bidders is not uncommon in the core auction literature, and it is often assumed that the mechanism has oracle access to a welfare maximization algorithm. Given the results derived by Day and Raghavan (2007) for computing $\mathsf{MREV}(\mathbf{b})$, we can conclude that our mechanism can be implemented with a polynomial number of oracle calls to welfare maximization. Faster algorithms have also been proposed for $\mathsf{MREV}(\mathbf{b})$, e.g., Niazadeh et al. (2022), but these compute ϵ-bidder-optimal core points and hence are not suitable for our mechanism. Finally, the work of Day and Raghavan (2007) implies that for settings where there exist efficient algorithms for welfare optimization, our mechanism is also implementable in polynomial time. As examples, we mention that this is the case (i) for the Text-and-Image setting of Goel et al. (2015), (ii) for Maximum Weight Matching auctions, where bidders represent the edges of a graph and optimal welfare corresponds to a maximum matching[5]. Both of these settings are of interest to us since they possess complementarities, hence there are no truthful mechanisms in the core.

As for tightness, recall that for a given value profile $\mathbf{v}$, the mechanism selects as the set of winning bidders, a subset of the optimal allocation $X^*(\mathbf{v})$. Related to this, for the special case studied by Goel et al. (2015), the authors show that for mechanisms that output as an allocation a subset of an optimal allocation, $O(\log n)$-core-competitiveness is the best one can hope for. This implies that our result is tight, and among such mechanisms, it achieves the best possible core-competitiveness.

2.4.2 Feasibility and Monotonicity of $\hat{X}$

To show that the mechanism always outputs a feasible allocation, we use the fact that for a given $\mathbf{v}$, $\hat{X}(\mathbf{v}) \subseteq X^*(\mathbf{v})$. Since the optimal allocation $X^*(\mathbf{v}) \in \mathcal{F}(N)$ and since we have assumed that $\mathcal{F}(N)$ is downward-closed, then $\hat{X}(\mathbf{v})$ is feasible.

Moreover, we claim that the allocation algorithm $\hat{X}$ always outputs a non-empty allocation, i.e., the cutoff set in (2.15) is always achievable by at least one index $k \in \{1, \dots, |X^*(\mathbf{v})|\}$. To prove this claim, we define first for a vector of values $v_1 \geq v_2 \geq \cdots \geq v_\ell$, the *maximum uniform price revenue* as $\max_{j\in\{1,\dots,\ell\}} j \cdot v_j$. The following is a well known lower bound on the uniform price revenue, proposed by Goldberg et al. (2006).

Lemma 2.4.1 (Due to Goldberg et al. (2006))**.** *Given $v_1 \geq \cdots \geq v_\ell$, it holds that $\max_{j\in\{1,\dots,\ell\}} j \cdot v_j \geq \frac{1}{H_\ell}\sum_{i=1}^{\ell} v_i$.*

Using this, we can now prove a lower bound in terms of the MRCS revenue.

[5]These auctions can be motivated by facility location and franchising considerations. The auctioneer can be seen as a company aiming to place stores that should not be on the same neighborhood or on the same street.

Lemma 2.4.2. *Let $\mathbf{v}$ be a value profile, and $m = |X^*(\mathbf{v})|$. Let $s_1, s_2, \dots s_m$ be an ordering of the bidders in $X^*(\mathbf{v})$ by their value in a non-increasing order. Then* $\max_{j \in \{1,\dots,m\}} j \cdot v_{s_j} \geq \frac{\mathsf{MREV}(\mathbf{v})}{H_n}$.

Proof. Let $\mathbf{p} \in CORE(\mathbf{v})$ be a core payment of minimum revenue, i.e. $\sum_{j \in X^*(\mathbf{v})} p_j = \mathsf{MREV}(\mathbf{v})$. By invoking Lemma 2.4.1 on the values $v_{s_1} \geq v_{s_2} \geq \cdots \geq v_{s_w}$ we have

$$\max_{j \in \{1,\dots,m\}} j \cdot v_{s_j} \geq \frac{\sum_{j \in X^*(\mathbf{v})} v_j}{H_m} \geq \frac{\sum_{j \in X^*(\mathbf{v})} v_j}{H_n} \geq \frac{\sum_{j \in X^*(\mathbf{v})} p_j}{H_n} = \frac{\mathsf{MREV}(\mathbf{v})}{H_n}.$$

The last inequality follows from the fact that the family of MRCS payment rules are individually rational. □

Lemma 2.4.2 directly implies that our proposed mechanism always outputs a non-empty solution, i.e., the cutoff value set in (2.15) will be satisfied by at least one index.
We now show that the allocation algorithm $\hat{X}$ is monotone. Lemma 2.4.3 below will be the key to establish this argument, which is in turn based on Corollary 2.3.2 from Section 2.3. Lemma 2.4.3 states that when a winning bidder increases her bid, the allocation algorithm $\hat{X}$ may only increase the number of bidders it serves.

Lemma 2.4.3. *For every value profile $\mathbf{v}$, bidder $i \in \hat{X}(\mathbf{v})$ and every $v_i' > v_i$, it holds that $|\hat{X}(\mathbf{v})| \leq |\hat{X}(v_i', \mathbf{v}_{-i})|$.*

Proof. Suppose for contradiction that this is not true, i.e., there exists a profile $\mathbf{v}$ with a bidder $i \in \hat{X}(\mathbf{v})$ and a bid $v_i' > v_i$ for which $|\hat{X}(\mathbf{v})| > |\hat{X}(v_i', \mathbf{v}_{-i})|$. Since $i \in X^*(\mathbf{v})$ and due to the fact that the welfare-maximizing algorithm is monotone, it holds that $i \in X^*(v_i', \mathbf{v}_{-i})$ as well. Let $\mathbf{s}$ be the ordering of the players in $X^*(\mathbf{v})$, produced by the mechanism at step 2, on input $\mathbf{v}$, and let $\mathbf{s}'$ be the corresponding ordering of bidders in $X^*(v_i', \mathbf{v}_{-i})$ on input $(v_i', \mathbf{v}_{-i})$. Let $k = |\hat{X}(\mathbf{v})|$ and $k' = |\hat{X}(v_i', \mathbf{v}_{-i})|$. By our assumption, $k' < k$. Bidder i can only be at a lower index in the ranking $\mathbf{s}'$ compared to her position at $\mathbf{s}$, since she has unilaterally deviated to $v_i' > v_i$. This implies that $v_{s_k'} \geq v_{s_k}$. Indeed, either the bidder at position k in $\mathbf{s}'$ has remained the same but with equal or higher value (in case bidder i is at position k) or bidder i has moved up in the ranking and it has displaced some bidder with a higher value, i.e., with an initial index $s_j < s_k$ to position k. However, this yields

$$k \cdot v_{s_k'} \geq k \cdot v_{s_k} \geq \frac{\mathsf{MREV}(\mathbf{v})}{H_n} \geq \frac{\mathsf{MREV}(v_i', \mathbf{v}_{-i})}{H_n}.$$

The second inequality follows from what we have assumed for the execution of the mechanism on input $\mathbf{v}$, whereas the third inequality follows from Equation (2.14) of Corollary 2.3.2. This means that k bidders can still be served on input $(v_i', \mathbf{v}_{-i})$, and hence k' is not the largest index of bidders who can meet the cutoff of (2.15) under $(v_i', \mathbf{v}_{-i})$. This is a contradiction. $\square$

We now prove that the allocation algorithm $\hat{X}$ is monotone.

Lemma 2.4.4. *The allocation algorithm $\hat{X}$ is monotone, i.e., given a profile $\mathbf{v}$, for every bidder $i \in \hat{X}(\mathbf{v})$ and every $v_i' > v_i$ it is true that $i \in \hat{X}(v_i', \mathbf{v}_{-i})$.*

Proof. Given a profile $\mathbf{v}$, fix a bidder $i \in \hat{X}(\mathbf{v})$. Suppose this bidder unilaterally declares a bid $v_i' > v_i$. We will show that bidder i remains in the set of final winners $\hat{X}(v_i', \mathbf{v}_{-i})$, for $v_i' > v_i$. Recall that the mechanism we propose, given a profile $\mathbf{v}$ finds an initial provisional allocation $X^*(\mathbf{v})$ and then selects a $\hat{X}(\mathbf{v}) \subseteq X^*(\mathbf{v})$. Hence, for all $v_i' > v_i$, we need to argue both that $i \in X^*(v_i', \mathbf{v}_{-i})$ and $i \in \hat{X}(v_i', \mathbf{v}_{-i})$.

For the first step, the allocation algorithm $\hat{X}$ calls X^*, and since the welfare-maximizing algorithm X^* is monotone it holds that $i \in X^*(v_i', \mathbf{v}_{-i})$, for all $v_i' > v_i$. For the second step, since bidder i has unilaterally declared a bid $v_i' > v_i$, we can be certain her index in the ranking among bidders in $X^*(\mathbf{v}) = X^*(v_i', \mathbf{v}_{-i})$, can only be lower. By Lemma 2.4.3, we know that $|\hat{X}(\mathbf{v})| \leq |\hat{X}(v_i', \mathbf{v}_{-i})|$, which implies that the allocation algorithm has picked a superset of $\hat{X}(\mathbf{v})$. Hence, bidder i will be a part of the new optimal allocation $\hat{X}(v_i', \mathbf{v}_{-i})$.

$\square$

2.4.3 Payments and Revenue Guarantee

By Lemma 2.4.4 the allocation rule $\hat{X}$ is monotone and hence, by Myerson's Lemma, each bidder must pay her threshold price, to obtain a mechanism that is incentive compatible and individually rational in a single-parameter setting. Hence, with regard to the proof of Theorem 2.4.1, the only statement we are left to prove is that $\hat{\mathcal{M}} = (\hat{X}, \hat{\mathbf{p}})$ is $O(\log n)$-core-competitive. Lemma 2.4.5 provides a relationship that is satisfied by the threshold payment of each winning bidder and that will be crucial to obtain this revenue guarantee.

Lemma 2.4.5. *Given a value profile $\mathbf{v}$, the threshold payment $\hat{p}_i(\mathbf{v})$ of every bidder $i \in \hat{X}(\mathbf{v})$ for the mechanism $\hat{\mathcal{M}} = (\hat{X}, \hat{\mathbf{p}})$ satisfies $\hat{p}_i(\mathbf{v}) \geq p_i^{VCG}(\mathbf{v})$ and, additionally,*

$$\hat{p}_i(\mathbf{v}) \geq \frac{\mathsf{MREV}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})}{|\hat{X}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})| \cdot H_n}. \tag{2.16}$$

Proof. Fix a bidder $i \in \hat{X}(\mathbf{v})$. By definition, her threshold payment $\hat{p}_i(\mathbf{v})$ is the minimum bid $v_i' \leq v_i$ she can unilaterally deviate to, so that $i \in \hat{X}(v_i', \mathbf{v}_{-i})$. Recall that, for every profile $\mathbf{v}$, the first step of the allocation algorithm $\hat{X}$ is to find the optimal allocation $X^*(\mathbf{v})$. Hence, since by Myerson's Lemma, the threshold payments for the algorithm X^* are the VCG payments, we can establish that $\hat{p}_i(\mathbf{v}) \geq p_i^{VCG}(\mathbf{v})$.

To prove (2.16), consider a bid v_i' which also survives step 2, so that $i \in \hat{X}(v_i', \mathbf{v}_{-i})$. In order for i to be included in an optimal allocation, v_i' must be large enough so that Equation (2.15) is satisfied for the profile $(v_i', \mathbf{v}_{-i})$. Suppose $v_i' = \hat{p}_i(\mathbf{v})$. Let $\mathbf{s}'$ be the ordering of the bidders produced by step 2 of the mechanism, and let $k = |\hat{X}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})|$. By Equation (2.15) it holds that

$$k \cdot v_{s_k'} \geq \frac{\mathsf{MREV}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})}{H_n} \Leftrightarrow v_{s_k'} \geq \frac{\mathsf{MREV}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})}{k \cdot H_n}. \tag{2.17}$$

Additionally, since $i \in \hat{X}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})$, her bid cannot be smaller than the bid of the last bidder included in $\hat{X}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})$, as otherwise she would not win. Therefore, $\hat{p}_i(\mathbf{v}) \geq v_{s_k'}$ and the proof follows by combining this fact with Equation (2.17). □

We can now prove the revenue guarantee and conclude the proof of Theorem 2.4.1.

Lemma 2.4.6. *The mechanism is $O(logn)$-core-competitive.*

Proof. Given a vector $\mathbf{v}$, the total revenue of the auctioneer is the sum of the threshold payments of bidders in $\hat{X}(\mathbf{v})$. Recall that for every bidder $i \notin \hat{X}(\mathbf{v})$, $\hat{p}_i(\mathbf{v}) = 0$ by Myerson's Lemma. Hence, we can lower bound the total revenue of the auctioneer as follows:

$$\begin{aligned}\sum_{j \in \hat{X}(\mathbf{v})} \hat{p}_j(\mathbf{v}) &\geq \sum_{i \in \hat{X}(\mathbf{v})} \frac{\mathsf{MREV}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})}{|\hat{X}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})| \cdot H_n} \geq \sum_{i \in \hat{X}(\mathbf{v})} \frac{\mathsf{MREV}(\mathbf{v})}{|\hat{X}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})| \cdot H_n} \\ &\geq \sum_{i \in \hat{X}(\mathbf{v})} \frac{\mathsf{MREV}(\mathbf{v})}{|\hat{X}(\mathbf{v})| \cdot H_n} = \frac{\mathsf{MREV}(\mathbf{v})}{H_n}.\end{aligned}$$

The first inequality follows from Lemma 2.4.5 (Equation (2.16)). To obtain the second inequality, for every bidder $i \in \hat{X}(\mathbf{v})$, we apply Corollary 2.3.2 for the profile $(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})$. Note that the vector $(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})$ satisfies the conditions of Corollary 2.3.2 since, by Lemma 2.4.5 it is also true that $\hat{p}_i(\mathbf{v}) \geq p_i^{VCG}(\mathbf{v})$. Finally, we obtain the third inequality by applying Lemma 2.4.3 regarding the same profile and deviation. □

2.5 A Class of Non-decreasing Quadratic Payment Rules

In this section, we illustrate a second application of our results of Section 2.3, focusing on an important family of quadratic core-selecting payment rules.

2.5.1 Quadratic Payment Rules

As we have mentioned in Section 2.2, when VCG payments are not in the core, the solution space of MRCS payments is always a continuum, in which case the linear program of Equation (2.4) has infinitely many solutions. Even though, as discussed in Section 2.2, all these solutions in this face of the core polytope, have been shown by Day and Raghavan (2007) to minimize the gain of deviating, the question remained whether one of these points should be preferred over others and whether there is a disciplined way to single out a solution. This motivated Erdil and Klemperer (2010) and Day and Cramton (2012) to propose a class of core-selecting mechanisms, based on the idea of picking the point on the minimum revenue face of the core that is the closest in Euclidean distance to a given reference point. This payment rule can be expressed using two mathematical programs: the linear program of Equation (2.6) to compute first $\mathsf{MREV}(\mathbf{b})$, and then a quadratic program, as defined below.

Definition 2.5.1. *Let* $\mathbf{r} \in \mathbb{R}^n_+$. *We call a payment rule* $\mathbf{r}$*-nearest when, for every vector* $\mathbf{b}$*, it assigns the payment*

$$\mathbf{p}^{\mathbf{r}}(\mathbf{b}) = \arg\min_{\mathbf{p}\in\mathbb{R}^n} \left\{ \sum_{j\in X^*(\mathbf{b})} (p_j - r_j)^2 : \quad \mathbf{p} \in CORE(\mathbf{b}), \sum_{j\in X^*(\mathbf{b})} p_j = \mathsf{MREV}(\mathbf{b}) \right\}. \tag{2.18}$$

In words, the quadratic program of (2.18) assigns for a bidding profile $\mathbf{b}$, the MRCS payment in the core that is the closest to a given vector $\mathbf{r}$. Alternatively, this quadratic program can be also defined without the MRCS constraint. In this case, it has been shown by Parkes et al. (2001), that the minimum revenue may not be achieved for certain reference points even for minimization objectives that result to Pareto-efficient payments. In this section, we stick to the version that contains the MRCS constraint. Moreover, since the quadratic program in (2.18) expresses a minimization of Euclidean distance from a convex set to a fixed point, the following well-known fact is true.

Fact 2.5.1. *Given vectors* $\mathbf{r}$ *and* $\mathbf{b}$*, the payment vector* $\mathbf{p}^{\mathbf{r}}(\mathbf{b})$ *is unique.*

A number of vectors have been proposed as the reference point $\mathbf{r}$, for this class of payments. Initially, Day and Cramton (2012) used the VCG payments for a reference point, $\mathbf{r} = \mathbf{p}^{VCG}(\mathbf{b})$, as a refinement of MRCS. The motivation of this choice was the findings of Parkes et al. (2001), who observed that given a profile $\mathbf{b}$ and a payment $\mathbf{p} \in CORE(\mathbf{b})$, the quantity $p_i - p_i^{VCG}(\mathbf{b})$ represents the bidder's "residual incentive to misreport". Hence, minimizing this quantity (or rather, its square) seemed a reasonable choice w.r.t. incentives. In parallel to this, Erdil and Klemperer (2010), developed a different perspective of what $\mathbf{r}$ should be. They leaned more towards constant payment rules with reference points that do not depend on the bidding profile, as their goal was to minimize *marginal* incentives to deviate. One well-studied and intuitive example is the $\mathbf{0}$-nearest mechanism: pick the point in MRCS that is closest to $\mathbf{0}$. Yet another perspective was given by Ausubel and Baranov (2020), who have proposed the $\mathbf{b}$-nearest payment rule, i.e., the MRCS payments closest to the actual bid. Overall, quadratic rules form a family of core-selecting mechanisms with many deployments in practice in several countries, especially for spectrum and other public sector auctions, see Day and Cramton (2012).

2.5.2 A Class of Non-Decreasing Quadratic Payment Rules

We now consider the following desirable property for payment rules.

Definition 2.5.2. *A payment rule is called non-decreasing, if for every profile* $\mathbf{b}$, *every bidder* $i \in N$ *and every* $b_i' > b_i$ *it holds that*

$$p_i(b_i', \mathbf{b}_{-i}) \geq p_i(\mathbf{b}). \tag{2.19}$$

This notion has been defined independently in Erdil and Klemperer (2010) and Bosshard et al. (2018), with a different motivation in mind. In Erdil and Klemperer (2010), it is argued that payment rules satisfying this property[6] weakly dominate all other payment rules in terms of the so called *marginal incentive to deviate*. Hence, even though such mechanisms may not be truthful, they possess very desirable incentive guarantees. In Bosshard et al. (2018), another advantage of this property is highlighted, which is of computational nature: limiting our attention to non-decreasing payment rules makes the daunting task of computing Bayes Nash equilibria much simpler.

Hence, it becomes important to understand which mechanisms satisfy this property. It can be seen that the VCG mechanism and the pay-your-bid auction do satisfy (2.19).

[6] Erdil and Klemperer originally called these rules monotonic.

In the context of MRCS rules, it is shown by Bosshard et al. (2018), that $\mathbf{p}^{VCG}$-nearest is *not* non-decreasing. To our knowledge, it has remained an open question whether there exist[7] MRCS rules that satisfy (2.19).

We answer this question in the affirmative, by providing a class of quadratic rules that are non-decreasing. To proceed, given a vector $\mathbf{b}$, for all $i \in N$, define $f_i(b_i)$ to be any non-decreasing function of b_i. Let $\mathbf{f}(\mathbf{b}) = (f_1(b_1), \dots, f_n(b_n))$. The following is the main result of this section.

Theorem 2.5.1. *For every $\mathbf{f} = (f_1(\cdot), \dots, f_n(\cdot))$, where each $f_i(\cdot)$ is a non-decreasing function of b_i, the $\mathbf{f}(\mathbf{b})$-nearest payment rule is non-decreasing for binary single-parameter auction domains.*

Proof. For a profile $\mathbf{b}$, let $\mathbf{p}(\mathbf{b})$ be the payment vector of the $\mathbf{f}(\mathbf{b})$-nearest rule. Suppose for a contradiction that (2.19) is not satisfied, i.e., there exists a profile $\mathbf{b}$, a bidder $i \in X^*(\mathbf{b})$, and a bid $b_i' > b_i$ for which

$$p_i(b_i', \mathbf{b}_{-i}) < p_i(\mathbf{b}). \tag{2.20}$$

Note that due to the monotonicity of X^*, it holds that $X^*(\mathbf{b}) = X^*(b_i', \mathbf{b}_{-i})$ for $b_i' > b_i$. Since $\mathbf{f}(\mathbf{b})$-nearest is a Pareto-efficient payment rule, by Lemma 2.3.2, for the deviating bidder i it holds that $p_i(\mathbf{b}) \leq b_i$. By combining this fact along with (2.20), we obtain that $p_i(b_i', \mathbf{b}_{-i}) < b_i$. In its turn, by Theorem 2.3.2 we have that

$$\mathbf{p}(b_i', \mathbf{b}_{-i}) \in CORE(\mathbf{b}). \tag{2.21}$$

Equation (2.21) implies that the optimal solution for the profile $(b_i', \mathbf{b} - i)$ is actually a member of the initial core polyhedron defined for the vector $\mathbf{b}$. We distinguish that each of the following cases leads to a contradiction:

1. $p_i(\mathbf{b}) = p_i^{VCG}(\mathbf{b})$: By (2.20), we know that $p_i(b_i', \mathbf{b}_{-i}) < p_i(\mathbf{b})$, hence $p_i(b_i', \mathbf{b}_{-i}) < p_i^{VCG}(\mathbf{b})$. But under the profile $(b_i', \mathbf{b}_{-i})$, bidder i cannot be asked to pay a payment that is less than $p_i^{VCG}(\mathbf{b})$, as this would violate the constraint for coalition $N \setminus \{i\}$ in (2.3), implying that $\mathbf{p}(b_i', \mathbf{b}_{-i}) \notin CORE(b_i', \mathbf{b}_{-i})$.

2. $p_i(\mathbf{b}) > p_i^{VCG}(\mathbf{b})$ *and* $\mathsf{MREV}(b_i', \mathbf{b}_{-i}) < \mathsf{MREV}(\mathbf{b})$: In this case, the solution $\mathbf{p}(b_i', \mathbf{b}_{-i})$ achieves a strictly lower minimum revenue when compared to $\mathbf{p}(\mathbf{b})$. However, by (2.21), $\mathbf{p}(b_i', \mathbf{b}_{-i}) \in CORE(\mathbf{b})$, which is a contradiction since it implies that the solution $\mathbf{p}(\mathbf{b})$ is not an MRCS solution of $CORE(\mathbf{b})$.

[7] In Bosshard et al. (2018) it was initially claimed that two non-MRCS rules are non-decreasing, however this was later retracted in their corrigendum in Bosshard et al. (2022).

3. $p_i(\mathbf{b}) > p_i^{VCG}(\mathbf{b})$ *and* $\mathsf{MREV}(b_i', \mathbf{b}_{-i}) = \mathsf{MREV}(\mathbf{b})$: To analyze this case, let us first define for every vector $\mathbf{r}$, the function $D(\mathbf{b}, \mathbf{r}) := \sum_{j \in X^*(\mathbf{b})} (p_j(\mathbf{b}) - r_j)^2$. For a given profile $\mathbf{b}$, $D(\mathbf{b}, \mathbf{f}(\mathbf{b}))$ is the optimal value of the objective function of $\mathbf{f}(\mathbf{b})$-nearest. We now have the following implications:

$$\begin{aligned} D((b_i', \mathbf{b}_{-i}), \mathbf{f}(b_i', \mathbf{b}_{-i})) &< D(\mathbf{b}, \mathbf{f}(b_i', \mathbf{b}_{-i})) \\ &= D(\mathbf{b}, \mathbf{f}(\mathbf{b})) + \big(f_i(b_i') - f_i(b_i)\big)\big(f_i(b_i') + f_i(b_i) - 2p_i(\mathbf{b})\big) \\ &< D((b_i', \mathbf{b}_{-i}), \mathbf{f}(\mathbf{b})) + \big(f_i(b_i') - f_i(b_i)\big)\big(f_i(b_i') + f_i(b_i) - 2p_i(\mathbf{b})\big) \end{aligned} \tag{2.22}$$

The first inequality follows from the fact that since $\mathbf{p}(b_i', \mathbf{b}_{-i})$ is the unique optimal solution for $CORE(b_i', \mathbf{b}_{-i})$, and since, by Theorem 2.3.1, $\mathbf{p}(\mathbf{b})$ is also a feasible payment in $CORE(b_i', \mathbf{b}_{-i})$, the value of the objective function $D(\mathbf{b}, \mathbf{f}(b_i', \mathbf{b}_{-i}))$ must be strictly larger. We apply the same argument and obtain the last inequality (Equation (2.22)) for the $CORE(\mathbf{b})$ polyhedron, since by Equation (2.21), $\mathbf{p}(b_i', \mathbf{b}_{-i}) \in CORE(\mathbf{b})$. By rearranging terms, Equation (2.22) implies

$$(f_i(b_i') - f_i(b_i))\,(2p_i(\mathbf{b}) - 2p_i(b_i', \mathbf{b}_{-i})) < 0, \tag{2.23}$$

a contradiction, since for $b_i' > b_i$ we have that $f_i(b_i') \geq f_i(b_i)$ by the monotonicity of $f_i(\cdot)$ and, by assumption $p_i(\mathbf{b}) > p_i(b_i', \mathbf{b}_{-i})$.

□

Notice that the class of $\mathbf{f}(\mathbf{b})$-nearest rules captures both the well known $\mathbf{0}$-nearest and $\mathbf{b}$-nearest mechanisms that were advocated by Erdil and Klemperer (2010) and Ausubel and Baranov (2020) respectively.

Corollary 2.5.1. *The* $\mathbf{0}$*-nearest and the* $\mathbf{b}$*-nearest payment rules are both non-decreasing for binary single-parameter auction domains.*

Chapter 3

The Discriminatory Price Auction: Equilibria and Inefficiency

3.1 Introduction

Multi-unit auctions form a popular transaction means for selling multiple units of a single good. They have been in use for a long time, and there are by now several practical implementations across many countries. Some of the most prominent applications involve government sales of treasury securities to investors, (e.g., see Brenner et al. (2009)), as well as electricity auctions, for distributing electrical energy (e.g. see Rio (2017)). Apart from governmental use, they are also run in other financial markets, and they are being deployed by various online brokers (Ockenfels et al., 2006). In the economics literature, multi-unit auctions have been a subject of study ever since the seminal work of Vickrey (1961), and some formats were conceived even earlier, by Friedman (1960).

Interestingly, according to the theory of core-selecting auctions discussed in Chapter 2, the VCG mechanism should be used for multi-unit auctions, as this is an auction environment in which goods are substitutes[1]. However, despite these theoretical advantages, simpler auction formats are often preferred in practice due to their ease of implementation and familiarity among bidders. One of these formats is the *discriminatory price auction*, also referred to as pay-your-bid auction and its welfare performance is the focus of this chapter[2]. In particular, we study the *uniform bidding interface*, which is the format most often employed in practice. Under this format, each bidder submits

[1] In a multi-unit auction, the set of feasible allocations forms a (poly)matroid.

[2] A conference version with the results of this chapter appeared in WINE '21 (Markakis et al., 2022b).

two parameters, a monetary per-unit bid, along with an upper bound on the number of units desired. Hence, each bidder is essentially asked to declare a *capped-additive* curve (a special case of submodular functions). The auctioneer then allocates the units by satisfying first the demand of the bidder with the highest monetary bid, then moving to the second highest bid, and so on, until there are no units left. As a price, each winning bidder pays his bid multiplied by the number of units received.

It is easy to see that the discriminatory price auction is not a truthful mechanism and the same holds for other auction formats used in practice. In recent years, a series of works have studied the social welfare guarantees that can be obtained at equilibrium. The outcome of these works is quite encouraging for the discriminatory price auction. Namely, pure Nash equilibria are always efficient, whereas for mixed and Bayes-Nash equilibria, the Price of Anarchy is bounded by 1.58 due to De Keijzer et al. (2013) for submodular valuations. These results suggest that simple auction formats can attain desirable guarantees and provide theoretical grounds for the overall success in practice.

Despite these positive findings, there has been no progress on further improving the current Price of Anarchy bounds. The known lower bound of 1.109 by Christodoulou et al. (2016b) is quite far from the upper bounds derived by the commonly used smoothness-based approaches of De Keijzer et al. (2013); Syrgkanis and Tardos (2013), which however do not seem applicable for producing further improvements. We believe the main difficulty in getting tighter results is that one needs to delve more deeply into the properties of Nash equilibria. But obtaining any form of characterization results for mixed or Bayesian equilibria is a notoriously hard problem. Even with two bidders it is often difficult to describe how the set of equilibria looks like. This is precisely the focus of our work, where we manage to either partially or fully characterize equilibrium profiles towards obtaining improved Price of Anarchy bounds, as we outline below.

3.1.1 Contribution

Motivated by the previous discussion, in Section 3.3 we initiate an equilibrium analysis for mixed equilibria. We consider bidders with capped-additive valuations, which is a subclass of submodular valuations, consistent with the underlying bidding format. Our results can be seen as a partial characterization of inefficient mixed equilibria, and our major highlights include both structural properties on the demand profile (see Theorem 3.3.3), as well as properties on the distributions of the mixed strategies (see Corollary 3.3.2, Theorem 3.3.4 and Lemma 3.3.12).

In Section 3.4, we use these results to derive new lower and upper bounds on the Price of Anarchy for mixed equilibria. For two bidders, we arrive at a complete characterization of inefficient equilibria and show an upper bound of 1.1095, which is tight.[3] For multiple bidders, we show that the Price of Anarchy is strictly worse, which also improves the best known lower bound for submodular valuations by Christodoulou et al. (2016b). We further present an improved upper bound of 4/3 for the special case where there exists a "high" demand bidder. We believe these latter instances are representative of the worst-case inefficiency that may arise, and refer to the relevant discussion in Section 3.4.2. To summarize, our results show that in several cases, the Price of Anarchy is even lower than the previous bound of De Keijzer et al. (2013) and strengthen the perception that such auctions can work well in practice.

Finally, in Section 3.5, we also study Bayes-Nash equilibria, and we exhibit a separation result that had been elusive so far: already with two bidders, the Price of Anarchy for Bayes-Nash equilibria is strictly worse than for mixed equilibria. Such separation results, though intuitive, do not hold for all auction formats. For example, it has been shown by Christodoulou et al. (2016a) that in simultaneous second price auctions with submodular valuations, the known tight bounds for mixed equilibria extend to the Bayesian model via the smoothness framework of Roughgarden (2012). This reveals that the Bayesian model in our setting introduces a further source of inefficiency. Note that to obtain this result, we transform the underlying optimization of social welfare at equilibrium to a well-posed variational calculus problem. This technique may be of independent interest and have other applications in mechanism design.

3.1.2 Related Work

The work of Ausubel and Milgrom (2002) was among the first ones that studied the sources of inefficiency in multi-unit auctions. For the discriminatory price auction, the Price of Anarchy was later studied in Syrgkanis and Tardos (2013), and for bidders with submodular valuations, the currently best upper bound of $e/(e-1) \approx 1.58$ has been obtained by De Keijzer et al. (2013) (both for mixed and for Bayes-Nash equilibria). These results exploit the smoothness-based techniques, developed by Roughgarden (2012); Syrgkanis and Tardos (2013). One can also obtain slightly worse upper bounds for subadditive valuations, by using a different methodology, based on Feldman et al. (2020). As for lower bounds, the only construction known for submodular valuations

[3] In Christodoulou et al. (2016b), there is a lower bound of 1.109 that applies to our setting with two bidders and three units. The lower bound we provide here is just slightly better, but most importantly, it is tight and can be seen as a generalization of the instance in Christodoulou et al. (2016b) to many units.

is by Christodoulou et al. (2016b), yielding a bound of at least 1.109. In parallel to these results, there has been a series of works on the inefficiency of many other auction formats, ranging from multi-unit to combinatorial auctions, see among others, Bhawalkar and Roughgarden (2011); Birmpas et al. (2019); Christodoulou et al. (2016a); Feldman et al. (2020).

Apart from social welfare guarantees, several other aspects or properties of equilibrium behavior have been studied. Recently in Pycia and Woodward (2021), a characterization of equilibria is given for a model where the supply of units can be drawn from a distribution. In the past, several works have focused on revenue equivalence results between the discriminatory price and the uniform price auction, see e.g. Ausubel et al. (2014); Swinkels (2001). On a different direction, comparisons from the perspective of the bidders are carried out in Baisa and Burkett (2018).

For a more detailed exposition on multi-unit auctions and their earlier applications, we refer the reader to the books of Krishna (2002) and Milgrom (2004). For more recent applications, we refer to Brenner et al. (2009); Goldner et al. (2020); Rio (2017), for treasury bonds, carbon licence auctions, and electricity auctions, respectively.

3.2 Notation and Definitions

We consider a discriminatory price multi-unit auction, involving the allocation of k identical units of a single item, to a set $\mathcal{N} = \{1, \dots, n\}$ of bidders. Each bidder $i \in \mathcal{N}$ has a private value $v_i > 0$, which reflects her value per unit and a private demand $d_i \in \mathbb{Z}_+$ which reflects the maximum number of units bidder i requires. Therefore, if the auction allocates $x_i \le k$ units to bidder i, her total value will be $\min\{x_i, d_i\} \cdot v_i$. We note that this class of valuations is a subclass of submodular valuations, and includes all additive vectors (when $d_i = k$). We will refer to them as *capped-additive* valuations.

We focus on the following simple format for the discriminatory price auction, which is known as the *uniform bidding* interface. The auctioneer asks each bidder $i \in \mathcal{N}$ to submit a tuple (b_i, q_i), where $b_i \ge 0$, is her monetary bid per unit (not necessarily equal to v_i), and q_i is her demand bid (not necessarily equal to d_i). We denote by $\mathbf{b} = (b_1, \dots, b_n)$ the monetary bidding vector, and similarly $\mathbf{q}$ will be the declared demand vector. For a bidding profile $(\mathbf{b}, \mathbf{q})$, the auctioneer allocates the units by satisfying first the demand of the bidder with the highest monetary bid, then moving to the second highest bid, and so on, until there are no units left. Hence, all the winners have their reported demand satisfied, except possibly for the one selected last, who may be partially satisfied. Moreover, we assume that in case of ties, a deterministic tie-breaking rule is used, which does not depend on the input bids submitted by the players to the auctioneer (e.g., a fixed ordering of the players suffices).

For every bidding profile $(\mathbf{b}, \mathbf{q})$, we let $x_i(\mathbf{b}, \mathbf{q})$ be the number of units allocated to bidder i, where obviously $x_i(\mathbf{b}, \mathbf{q}) \leq q_i$. In the discriminatory auction, the auctioneer requires each bidder i to pay b_i per allocated unit, hence a total payment of $b_i \cdot x_i(\mathbf{b}, \mathbf{q})$. The utility function of bidder $i \in \mathcal{N}$, given a bidding profile $(\mathbf{b}, \mathbf{q})$, is: $u_i(\mathbf{b}, \mathbf{q}) = \min\{x_i(\mathbf{b}, \mathbf{q}), d_i\}v_i - x_i(\mathbf{b}, \mathbf{q})b_i$.

Viewed as games, these auctions have an infinite pure strategy space, and we also allow bidders to play mixed strategies, which are probability distributions over their set of pure strategies. When each bidder $i \in \mathcal{N}$ uses a mixed strategy G_i, she *independently* draws a bid (b_i, q_i) from G_i. We refer to $\mathbf{G} = \times_{i=1}^n G_i$ as the product distribution of bids. Under mixed strategies, the expected utility of a bidder i is $\mathbb{E}_{(\mathbf{b},\mathbf{q})\sim\mathbf{G}}[u_i(\mathbf{b}, \mathbf{q})]$.

Definition 3.2.1. *We say that* $\mathbf{G}$ *is a mixed Nash equilibrium when for all* $i \in \mathcal{N}$*, all* $b'_i \geq 0$ *and all* $q'_i \in \mathbb{Z}_+$

$$\mathop{\mathbb{E}}_{(\mathbf{b},\mathbf{q})\sim\mathbf{G}}[u_i(\mathbf{b}, \mathbf{q})] \geq \mathop{\mathbb{E}}_{(\mathbf{b}_{-i},\mathbf{q}_{-i})\sim\mathbf{G}_{-i}}[u_i((b'_i, \mathbf{b}_{-i}), (q'_i, \mathbf{q}_{-i}))] \,.$$

We note that in any equilibrium, if a bidder i declares with positive probability a bid that exceeds v_i, she should not be allocated any unit, since such strategies are strictly dominated by bidding the actual value v_i.

Fact 3.2.1. *Let* $\mathbf{G}$ *be a mixed Nash equilibrium. The probability that a bidder* i *is allocated some units, conditioned that she bids higher than* v_i*, is 0.*

In the sequel, we focus on equilibria, where the monetary bids never exceed the value per unit.

Given a valuation profile $(\mathbf{v}, \mathbf{d})$, we denote by $OPT(\mathbf{v}, \mathbf{d})$ the optimal social welfare (which can be computed very easily by running the allocation algorithm of the auction with the true value and demand vector). We also denote by $SW(\mathbf{G})$ the expected social welfare of a mixed Nash equilibrium $\mathbf{G}$, i.e., equal to $\mathbb{E}_{(\mathbf{b},\mathbf{q})\sim\mathbf{G}}[\sum_i \min\{x_i(\mathbf{b}, \mathbf{q}), d_i\}v_i]$. The Price of Anarchy is the worst-case ratio $\frac{OPT(\mathbf{v},\mathbf{d})}{SW(\mathbf{G})}$, over all valuation profiles $(\mathbf{v}, \mathbf{d})$, and all equilibria $\mathbf{G}$.

We refer to an equilibrium as *inefficient* when its social welfare is strictly less than the optimal.

3.3 Towards a Characterization of Inefficient Mixed Equilibria

In this section, we derive a series of important properties, that help us understand better how can inefficient equilibria arise. These properties will help us analyze the Price of Anarchy in Section 3.4.

3.3.1 Mixed Nash Equilibria with Demand Revelation

Our first result is that it suffices to focus on equilibria where bidders truthfully reveal their demand, resulting therefore in a single-parameter strategy space for the bidders (Theorem 3.3.1). We further argue that the inefficiency in equilibria appears only when the total demand exceeds k (Lemma 3.3.6) and therefore this is what we assume for the rest of the chapter.

Theorem 3.3.1. *Let* $(\mathbf{v}, \mathbf{d})$ *be a valuation profile, and* $\mathbf{G}$ *be a mixed Nash equilibrium. Then, for every* $i \in \mathcal{N}$*, and in every pure strategy profile* $(b_i, q_i) \sim G_i$*, we can replace* q_i *by* d_i*, so that the resulting distribution remains a mixed Nash equilibrium with the same social welfare.*

Proof. We prove this theorem by a series of lemmas. The first step is the next lemma, showing that it suffices to consider only equilibria, where nobody declares a demand bid that is lower than their true demand.

Lemma 3.3.1. *Let* $\mathbf{G}$ *be any mixed Nash equilibrium and* G'_i *be the same as* G_i *after replacing any* $q_i < d_i$ *with* d_i*. Then* $\mathbf{G}'$ *is also a mixed Nash equilibrium with the same social welfare.*

Proof. We use first the following auxiliary lemma.

Lemma 3.3.2. *Let* $\mathbf{G}$ *be any mixed Nash equilibrium where there exists a bidder* i *such that* $Pr[b_i = v_i, q_i < d_i] > 0$ *for* $(\mathbf{b}, \mathbf{q}) \sim \mathbf{G}$*. Let* G'_i *be the same as* G_i *after replacing any* $q_i < d_i$ *with* d_i*. Then* $\mathbf{G}'$ *is also a mixed Nash equilibrium with the same social welfare.*

Proof. First note that $\mathbb{E}_{(\mathbf{b},\mathbf{q})\sim\mathbf{G}}[u_i(\mathbf{b}, \mathbf{q})] = 0$, since bidder i bids v_i with positive probability that results in zero utility (see also Lemma 3.3.2). Let β^k_{-i} be the random variable expressing the k^{th} maximum payment under $\mathbf{G}_{-i}$. Then $Pr[\beta^k_{-i} < v_i] = 0$, because if β^k_{-i} takes a value less than v_i with positive probability, bidder i has an incentive to deviate to a bid less than v_i and receive positive utility.

For any bidder j, with $v_j > v_i$, and any bidding profile $(\mathbf{b}, \mathbf{q}) \sim \mathbf{G}$, such that $x_i(\mathbf{b}, \mathbf{q}) > 0$ and $x_j(\mathbf{b}, \mathbf{q}) < q_j$, it holds that $b_j \leq b_i = v_i$ (apart maybe from cases that appear with zero probability). Then, $Pr[x_i(\mathbf{b}, \mathbf{q}] > 0, x_j(\mathbf{b}, \mathbf{q}) < q_j) = 0$, otherwise there exists a sufficiently small $\varepsilon > 0$, such that bidder j has an incentive to deviate to $v_i + \varepsilon$ and receive more units. Therefore, if bidders i and j bid both v_i with positive probability the tie-breaking rule is in favour of player j. The same tie-breaking rule should be applied when bidder i increases his quantity bid and so, for any bidding profile $(\mathbf{b}, \mathbf{q}) \sim \mathbf{G}$, $x_j(\mathbf{b}, \mathbf{q}) = x_j(\mathbf{b}, (d_i, \mathbf{q}_{-i}))$.

For any bidder j, with $v_j \leq v_i$, bidder j cannot get any unit by paying less that v_i since $Pr[\beta^k_{-i} < v_i] = 0$. Therefore bidder j may receive units with positive probability only if $v_j = v_i$ and his expected utility is zero.

Overall, if bidder i deviates from G_i to G'_i, either the allocation of the players remains the same (so they still have no incentive to deviate) or they have zero utility (and still no incentive to deviate) and the allocation may change between bidders of the same valuation; so the expected social welfare remains the same and the new strategy profile is a mixed Nash equilibrium.

□

We continue now with the proof of Lemma 3.3.1. Starting by $\mathbf{G}$, we recursively show that if one by one the bidders deviate to G'_i, the bidding profile remains an equilibrium with the same social welfare. It is sufficient to show this for $(G'_i, \mathbf{G}_{-i})$.

First note that, according to the tie-breaking rule, if i deviates from G_i to G'_i, he can only get more units as he only declares the same or more demand. Let S_i be the set of bids (b_i, q_i) such that $q_i < d_i$ and

$$\mathop{\mathbb{E}}_{(\mathbf{b}_{-i}, \mathbf{q}_{-i}) \sim \mathbf{G}_{-i}}[x_i((b_i, \mathbf{b}_{-i}), (q_i, \mathbf{q}_{-i}))] < \mathop{\mathbb{E}}_{(\mathbf{b}_{-i}, \mathbf{q}_{-i}) \sim \mathbf{G}_{-i}}[x_i((b_i, \mathbf{b}_{-i}), (d_i, \mathbf{q}_{-i}))]$$

It should be that for $(\mathbf{b}, \mathbf{q}) \sim \mathbf{G}$, $Pr[(b_i, q_i) \in S_i, b_i < v_i] = 0$, otherwise bidder i would increase her utility by deviating from $(b_i, q_i) \in S_i$ to (b_i, d_i). So, the only case that bidder i may increase his allocation by increasing his demand to d_i is when he bids his value, in which case the lemma follows by Lemma 3.3.2.

So far we have shown that for any bid (b_i, q_i) such that $q_i < d_i$ and $b_i < v_i$ the expected allocation to bidder i remains the same if he deviates to (b_i, d_i), apart maybe from cases that appear with zero probability. It remains to show that under $(G'_i, \mathbf{G}_{-i})$ the allocation to all bidders remains the same. Note that i deviating from G_i to G'_i can only cause other bidders to be allocated less or the same number of units due to the tie breaking rule. Given any $(b_i, q_i) \sim G_i$ with $q_i < d_i$, let $S_{-i}(b_i, q_i) = S_{-i}$ be the set of bidding profiles $(\mathbf{b}_{-i}, \mathbf{q}_{-i}) \sim \mathbf{G}_{-i}$ such that there exists a bidder $j \neq i$, receiving less units by the deviation of i, i.e., $x_j(\mathbf{b}, \mathbf{q}) > x_j(\mathbf{b}, (d_i, \mathbf{q}_{-i}))$, where $\mathbf{b} = (b_i, \mathbf{b}_{-i})$, and $\mathbf{q} = (q_i, \mathbf{q}_{-i})$.

For the sake of contradiction suppose that, under $\mathbf{G}_{-i}$, $Pr[(\mathbf{b}_{-i}, \mathbf{q}_{-i}) \in S_{-i}] > 0$. By summing over all bidders but i and taking the expectation over $\mathbf{G}_{-i}$, we have that

$$\mathop{\mathbb{E}}_{(\mathbf{b}_{-i},\mathbf{q}_{-i})\sim\mathbf{G}_{-i}}[\sum_{j\neq i} x_j(\mathbf{b},\mathbf{q})] > \mathop{\mathbb{E}}_{(\mathbf{b}_{-i},\mathbf{q}_{-i})\sim\mathbf{G}_{-i}}[\sum_{j\neq i} x_j(\mathbf{b},(d_i,\mathbf{q}_{-i}))]\,.$$

This means that $\mathbb{E}_{(\mathbf{b}_{-i},\mathbf{q}_{-i})\sim\mathbf{G}_{-i}}[x_i(\mathbf{b},\mathbf{q})] < \mathbb{E}_{(\mathbf{b}_{-i},\mathbf{q}_{-i})\sim\mathbf{G}_{-i}}[x_i(\mathbf{b},(d_i,\mathbf{q}_{-i}))]$ which leads to a contradiction.

□

The next step is to prove (in Lemma 3.3.3) that if $\sum_i d_i > k$, then it is sufficient to consider only Nash equilibria, where nobody declares more demand that their true demand.

Lemma 3.3.3. *Suppose that $\sum_i d_i > k$ and let $\mathbf{G}$ be any mixed Nash equilibrium where nobody declares less demand and G'_i be the same as G_i after replacing any $q_i > d_i$ with d_i. Then $\mathbf{G}'$ is also a mixed Nash equilibrium with the same social welfare.*

Proof. The proof is established by the following two lemmas.

Lemma 3.3.4. *For any Nash equilibrium $\mathbf{G}$ where nobody declares less demand, if $Pr[x_i(\mathbf{b},\mathbf{q}) > d_i] > 0$ for $(\mathbf{b},\mathbf{q}) \sim \mathbf{G}$, then $\mathbb{E}[b_i \mid x_i(\mathbf{b},\mathbf{q}) > d_i] = 0$.*

Proof. Suppose on the contrary that $\mathbb{E}[b_i \mid x_i(\mathbf{b},\mathbf{q}) > d_i] > 0$. We will show that bidder i has an incentive to declare her true demand instead of a higher demand.

$$\begin{aligned}
\mathbb{E}[u_i(\mathbf{b},\mathbf{q})] &= Pr[x_i(\mathbf{b},\mathbf{q}) > d_i]\,\mathbb{E}[u_i(\mathbf{b},\mathbf{q}) \mid x_i(\mathbf{b},\mathbf{q}) > d_i] \\
&+ Pr[x_i(\mathbf{b},\mathbf{q}) \le d_i]\,\mathbb{E}[u_i(\mathbf{b},\mathbf{q}) \mid x_i(\mathbf{b},\mathbf{q}) \le d_i] \\
&= Pr[x_i(\mathbf{b},\mathbf{q}) > d_i]\,\mathbb{E}[d_i(v_i - b_i) - (x_i(\mathbf{b},\mathbf{q}) - d_i)b_i \mid x_i(\mathbf{b},\mathbf{q}) > d_i] \\
&\quad + Pr[x_i(\mathbf{b},\mathbf{q}) \le d_i]\,\mathbb{E}[x_i(\mathbf{b},\mathbf{q})(v_i - b_i) \mid x_i(\mathbf{b},\mathbf{q}) \le d_i] \\
&= Pr[x_i(\mathbf{b},\mathbf{q}) > d_i]\,\mathbb{E}[d_i(v_i - b_i) \mid x_i(\mathbf{b},\mathbf{q}) > d_i] \\
&\quad + Pr[x_i(\mathbf{b},\mathbf{q}) \le d_i]\,\mathbb{E}[x_i(\mathbf{b},\mathbf{q})(v_i - b_i) \mid x_i(\mathbf{b},\mathbf{q}) \le d_i] \\
&\quad - Pr[x_i(\mathbf{b},\mathbf{q}) > d_i]\,\mathbb{E}[(x_i(\mathbf{b},\mathbf{q}) - d_i)b_i \mid x_i(\mathbf{b},\mathbf{q}) > d_i] \\
&= \mathbb{E}[u_i(\mathbf{b}, d_i, \mathbf{q}_{-i})] - Pr[x_i(\mathbf{b},\mathbf{q}) > d_i]\,\mathbb{E}[(x_i(\mathbf{b},\mathbf{q}) - d_i)b_i \mid x_i(\mathbf{b},\mathbf{q}) > d_i] \\
&\le \mathbb{E}[u_i(\mathbf{b}, d_i, \mathbf{q}_{-i})] - Pr[x_i(\mathbf{b},\mathbf{q}) > d_i]\,\mathbb{E}[b_i \mid x_i(\mathbf{b},\mathbf{q}) > d_i] \\
&< \mathbb{E}[u_i(\mathbf{b}, d_i, \mathbf{q}_{-i})]\,,
\end{aligned}$$

where the last strict inequality is due to our assumption that $\mathbb{E}[b_i \mid x_i(\mathbf{b},\mathbf{q}) > d_i] > 0$. The lemma follows by contradiction. □

Lemma 3.3.5. *If $\sum_i d_i > k$ then in any Nash equilibrium $\mathbf{G}$ where nobody declares less demand, $Pr[x_i(\mathbf{b},\mathbf{q}) > d_i] = 0$ for all i where $(\mathbf{b},\mathbf{q}) \sim \mathbf{G}$.*

Proof. Suppose on the contrary that there exists a bidder i such that $Pr[x_i(\mathbf{b},\mathbf{q}) > d_i] > 0$. Then there is also a bidder j such that $Pr[x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j] > 0$, otherwise $Pr[x_j(\mathbf{b},\mathbf{q}) \geq d_j, \forall j] = 1$ which contradicts the fact that $\sum_i d_i > k$.

Given that $x_i(\mathbf{b},\mathbf{q}) > d_i$ and $x_j(\mathbf{b},\mathbf{q}) < d_j$, bidder j bids 0 (apart maybe for cases that appear with zero probability), otherwise he should have received more units and bidder i less units since by Lemma 3.3.4, bidder i bids 0 and receives at least one unit. Then the expected utility of bidder j can be expressed as:

$$\begin{aligned}
\mathbb{E}[u_j(\mathbf{b},\mathbf{q})] &= Pr[x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j]\,\mathbb{E}[x_j(\mathbf{b},\mathbf{q})(v_j - b_j) \mid x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j] \\
&\quad +(1 - Pr[x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j])\,\mathbb{E}[u_j(\mathbf{b},\mathbf{q}) \mid x_i(\mathbf{b},\mathbf{q}) \leq d_i \text{ or } x_j(\mathbf{b},\mathbf{q}) \geq d_j] \\
&\leq Pr[x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j]\,\mathbb{E}[x_j(\mathbf{b},\mathbf{q})v_j \mid x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j] \\
&\quad +(1 - Pr[x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j])\,\mathbb{E}[u_j(\mathbf{b},\mathbf{q}) \mid x_i(\mathbf{b},\mathbf{q}) \leq d_i \text{ or } x_j(\mathbf{b},\mathbf{q}) \geq d_j]\,.
\end{aligned}$$

Consider now the bidding strategy (b'_j, q_j) where $b'_j = \varepsilon > 0$ when $b_j = 0$ and $b'_j = b_j$ otherwise, for some $\varepsilon < Pr[x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j]v_j/k$. If bidder j deviates to this strategy he should receive at least one more unit since he would bid more than bidder i and his expected utility would be $E' := \mathbb{E}[u_j(b'_j, \mathbf{b}_{-j}, \mathbf{q})]$:

$$\begin{aligned}
E' &= Pr[x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j]\,\mathbb{E}[x_j(b'_j, \mathbf{b}_{-j}, \mathbf{q})(v_j - b'_j) \mid x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j] \\
&\quad +(1 - Pr[x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j])\,\mathbb{E}[u_j(b'_j, \mathbf{b}_{-j}, \mathbf{q}) \mid x_i(\mathbf{b},\mathbf{q}) \leq d_i \text{ or } x_j(\mathbf{b},\mathbf{q}) \geq d_j] \\
&\geq Pr[x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j]\,\mathbb{E}[(x_j(\mathbf{b},\mathbf{q}) + 1)(v_j - \varepsilon) \mid x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j] \\
&\quad +(1 - Pr[x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j])\,\mathbb{E}[u_j(\mathbf{b},\mathbf{q}) - k\varepsilon \mid x_i(\mathbf{b},\mathbf{q}) \leq d_i \text{ or } x_j(\mathbf{b},\mathbf{q}) \geq d_j] \\
&\geq \mathbb{E}[u_j(\mathbf{b},\mathbf{q})] + Pr[x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j](v_j - d_j\varepsilon) \\
&\quad -(1 - Pr[x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j])k\varepsilon \\
&\geq \mathbb{E}[u_i(\mathbf{b},\mathbf{q})] + Pr[x_i(\mathbf{b},\mathbf{q}) > d_i, x_j(\mathbf{b},\mathbf{q}) < d_j]v_j - k\varepsilon \\
&> \mathbb{E}[u_i(\mathbf{b},\mathbf{q})]\,,
\end{aligned}$$

where the strict inequality comes from the definition of ε. This leads to a contradiction that concludes the proof. □

□

The remaining case that has not been covered by Lemma 3.3.3, is when the total demand does not exceed k: $\sum_i d_i \leq k$. But as we state in Lemma 3.3.6, these are efficient equilibria.

The proof of Theorem 3.3.1 follows by combining Lemmas 3.3.1, 3.3.3 and 3.3.6. □

Lemma 3.3.6. *If $\sum_i d_i \leq k$ then the social welfare of any mixed Nash equilibrium is optimal.*

Proof. If $\sum_i d_i \leq k$, the optimum allocation appears when every bidder with positive valuation receives a number of units more or equal to their true demand.

For the sake of contradiction suppose that there exists a Nash equilibrium $\mathbf{G}$ and a bidder i such that $Pr[x_i(\mathbf{b},\mathbf{q}) < d_i] > 0$ for $(\mathbf{b},\mathbf{q}) \sim \mathbf{G}$. Since nobody bids less than their true demand, bidder i receives less units than d_i only because there are either bidders bidding more than their true demand or bidders with zero valuation receive units (or both). By Lemma 3.3.4, we have that $Pr[\max_{j\neq i: x_j(\mathbf{b},\mathbf{q})>d_j} b_i = 0] = 0$ and the expected utility of bidder i can be expressed as follows:

$$
\begin{aligned}
\mathbb{E}[u_i(\mathbf{b},\mathbf{q})] &= Pr[x_i(\mathbf{b},\mathbf{q}) < d_i]\,\mathbb{E}[x_i(\mathbf{b},\mathbf{q})(v_i - b_i) \mid x_i(\mathbf{b},\mathbf{q}) < d_i, \max_{j\neq i, x_j(\mathbf{b},\mathbf{q})>d_j} b_j = 0] \\
&\quad + Pr[x_i(\mathbf{b},\mathbf{q}) \geq d_i]\,\mathbb{E}[d_i v_i - x_i(\mathbf{b},\mathbf{q}) b_i \mid x_i(\mathbf{b},\mathbf{q}) \geq d_i] \\
&\leq Pr[x_i(\mathbf{b},\mathbf{q}) < d_i]\,\mathbb{E}[(d_i - 1) v_i \mid x_i(\mathbf{b},\mathbf{q}) < d_i, \max_{j\neq i, x_j(\mathbf{b},\mathbf{q})>d_j} b_j = 0] \\
&\quad + Pr[x_i(\mathbf{b},\mathbf{q}) \geq d_i]\,\mathbb{E}[d_i v_i - x_i(\mathbf{b},\mathbf{q}) b_i \mid x_i(\mathbf{b},\mathbf{q}) \geq d_i]\,,
\end{aligned}
$$

where the inequality is due to the fact that in the first term $x_i(\mathbf{b},\mathbf{q}) \leq d_i - 1$; further note that in the first term bidder i loses units where the maximum of the other bids is 0, so he should have bid 0.

Consider now the bidding strategy (b'_i, q_i) where $b'_i = \varepsilon > 0$ when $b_i = 0$ and $b'_i = b_i$ otherwise, for some $\varepsilon < Pr[x_i(\mathbf{b},\mathbf{q}) < d_i] v_i / k$. Then the expected utility of bidder i after deviating to this strategy is:

$$
\begin{aligned}
\mathbb{E}[u_i(b'_i, \mathbf{b}_{-i}, \mathbf{q})] &= Pr[x_i(\mathbf{b},\mathbf{q}) < d_i]\,\mathbb{E}[x_i(b'_i, \mathbf{b}_{-i}, \mathbf{q})(v_i - b'_i) \mid x_i(\mathbf{b},\mathbf{q}) < d_i, \max_{j\neq i, x_j(\mathbf{b},\mathbf{q})>d_j} b_j = 0] \\
&\quad + Pr[x_i(\mathbf{b},\mathbf{q}) \geq d_i]\,\mathbb{E}[d_i v_i - x_i(b'_i, \mathbf{b}_{-i}, \mathbf{q}) b'_i \mid x_i(\mathbf{b},\mathbf{q}) \geq d_i] \\
&\geq Pr[x_i(\mathbf{b},\mathbf{q}) < d_i]\,\mathbb{E}[d_i(v_i - \varepsilon) \mid x_i(\mathbf{b},\mathbf{q}) < d_i, \max_{j\neq i, x_j(\mathbf{b},\mathbf{q})>d_j} b_j = 0] \\
&\quad + Pr[x_i(\mathbf{b},\mathbf{q}) \geq d_i]\,\mathbb{E}[d_i v_i - x_i(\mathbf{b},\mathbf{q})(b_i + \varepsilon) \mid x_i(\mathbf{b},\mathbf{q}) \geq d_i] \\
&\geq \mathbb{E}[u_i(\mathbf{b},\mathbf{q})] + Pr[x_i(\mathbf{b},\mathbf{q}) < d_i](v_i - d_i\varepsilon) - Pr[x_i(\mathbf{b},\mathbf{q}) \geq d_i] k\varepsilon \\
&\geq \mathbb{E}[u_i(\mathbf{b},\mathbf{q})] + Pr[x_i(\mathbf{b},\mathbf{q}) < d_i] v_i - k\varepsilon \\
&> \mathbb{E}[u_i(\mathbf{b},\mathbf{q})]\,,
\end{aligned}
$$

where the strict inequality comes from the definition of ε. This leads to a contradiction that concludes the proof. □

3.3.2 Existence of Non-empty-handed Bidders

For the remaining chapter, we consider only strategy profiles where the bidders' demand bid matches their true demand. The main goal of this subsection is to derive Theorem 3.3.3, where we show that in any inefficient mixed equilibrium, there always exists a bidder such that the total demand of the other winners is strictly less than k, meaning that at least one item is allocated to him for sure (with probability one). This is a crucial property for understanding the formation of inefficient mixed equilibria. To proceed, we give first some further notation to be used in this and the following sections.

Further notation. Given Theorem 3.3.1, instead of using distributions on tuples (b_i, q_i), we suppose that each bidder $i \in \mathcal{N}$ independently draws only a monetary bid b_i from a distribution B_i and we refer to $\mathbf{B} = \times_{i=1}^n B_i$ as the product distribution of monetary bids or just bids from now on. For a bidding profile $\mathbf{b}$, the utility of a bidder i will simply be denoted as $u_i(\mathbf{b})$, instead of $u_i(\mathbf{b}, \mathbf{d})$. Definition 3.2.1 is also simplified, and we say that $\mathbf{B}$ is an equilibrium if $\mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})] \geq \mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i((b_i', \mathbf{b}_{-i}))]$, for any i and any $b_i' \geq 0$. Similarly, the social welfare of a mixed Nash equilibrium $\mathbf{B}$ is given by just $SW(\mathbf{B})$.

For a mixed strategy bidding profile $\mathbf{B}$, we denote by $W(\mathbf{B})$ the set of bidders with positive expected utility, i.e., $W(\mathbf{B}) = \{j : \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_j(\mathbf{b})] > 0\}$, and let $\mathbf{B}_W = \times_{i\in W(\mathbf{B})} B_i$. Moreover, the support of a bidder i in $\mathbf{B}$ is the domain of the distribution B_i, that i plays under $\mathbf{B}$, denoted by $Supp(B_i)$. We denote by $\ell(B_i), h(B_i)$ the leftmost and rightmost points, respectively, in the support of bidder i. In particular, if the rightmost part of the domain of B_i is a mass point b or an interval in the form $[a, b]$, then $h(B_i) = b$, and similarly for $\ell(B_i)$. In cases of distributions over intervals, we can safely assume that the domain contains only *closed* intervals, because the endpoints are chosen with zero probability. We further denote by $\ell(\mathbf{B}_W), h(\mathbf{B}_W)$ the leftmost and rightmost points, respectively, of the union of the supports of $W(\mathbf{B})$.

For $i = 1, \ldots, n$ we denote by F_i the CDF of B_i and by f_i their PDF. Moreover, given a profile $\mathbf{b}$, it is often useful in the analysis to consider the vector of bids (thresholds) that a bidder i competes against, denoted by $\beta(\mathbf{b})_{-i} = (\beta_1(\mathbf{b}_{-i}), \ldots, \beta_k(\mathbf{b}_{-i}))$. Here, $\beta_j(\mathbf{b}_{-i})$ is the j-th lowest winning bid of the profile $\mathbf{b}_{-i}$, for $j = 1, \ldots, k$, so that $\beta(b)_{-i}$ describes the winning bids if i didn't participate. This implies that, under profile $\mathbf{b}$, bidder i is allocated $j = 1, \ldots, k-1$ units capped by d_i, when $\beta_j(\mathbf{b}_{-i}) < b_i < \beta_{j+1}(\mathbf{b}_{-i})$ and d_i units, when $\beta_k(\mathbf{b}_{-i}) < b_i$. We note that because we focus on the uniform bidding interface, some consecutive β_j values may coincide and be equal to the bid of the same bidder. When $\mathbf{b}_{-i} \sim \mathbf{B}_{-i}$, for $i = 1, \ldots, n$, we denote the CDF of the random variable

$\beta_j(\mathbf{b}_{-i})$ as $\hat{F}_{ij}$, for $j = 1, \ldots, k$. In the next fact, we express the expected allocation of any bidder i for bidding some $\alpha > 0$, in terms of the values $\hat{F}_{ij}(\alpha)$.

Fact 3.3.1. *Let $\mathbf{B}_{-i}$ be a product distribution of bids. Then for all $\alpha \geq 0$, where no bidder other than (possibly) i has a mass point,* $\mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[x_i(\alpha, \mathbf{b}_{-i})] = \sum_{j=1}^{d_i} \hat{F}_{ij}(\alpha)$.

Given a bidding profile $\mathbf{B}$, for any bidder i we define $\hat{F}_i^{avg}(x) = \frac{\sum_{j=1}^{d_i} \hat{F}_{ij}(x)}{d_i}$, to be the average CDF of the winning bids that bidder i competes against. Note that $\hat{F}_i^{avg}$ is a CDF since it is the average of a number of CDFs.

Remark 3.3.1. *The $\hat{F}_{ij}$ functions are right continuous, as they are CDFs, and moreover, if the F_i functions have no mass point, the same holds for the $\hat{F}_{ij}$ functions. Additionally, if for any j, the $\hat{F}_{ij}$ functions are continuous, so is $\hat{F}_i^{avg}$, as the average of continuous functions.*

We start by ruling out certain scenarios that cannot occur at inefficient equilibria. First, we can safely ignore bidders with zero expected utility, since in any *inefficient* mixed Nash equilibrium they do not receive any units.

Lemma 3.3.7. *Any mixed Nash equilibrium $\mathbf{B}$ with at least one bidder with zero expected utility, but positive expected number of allocated units, is efficient.*

Proof. Let i be such a bidder. Since i receives at least one unit with positive probability, it holds that $Pr[x_i(\mathbf{b}) > 0] > 0$ for $\mathbf{b} \sim \mathbf{B}$. There is only one possible case so that bidder i has zero expected utility and this is that he bids his valuation when he receives at least one unit (or more accurately, the probability that he bids less than his value and receives at least one unit is zero).

First note that the payment for any unit is at least v_i (apart maybe from cases that appear with zero probability), otherwise bidder i has an incentive to bid less than v_i and get a positive utility. For any bidder j with $v_j > v_i$, $Pr[x_i(\mathbf{b}) > 0, x_j(\mathbf{b}) < d_j] = 0$, otherwise there exists a sufficiently small $\varepsilon > 0$, such that bidder j has an incentive to deviate from v_i to $v_i + \varepsilon$ and receive more units. Therefore, it holds that

$$Pr[x_j(\mathbf{b}) = d_j, \ \ \forall j \text{ with } v_j > v_i \mid x_i(\mathbf{b}) > 0] = 1,$$

and since $Pr[x_i(\mathbf{b}) > 0] > 0$ it holds that

$$Pr[x_j(\mathbf{b}) = d_j, \ \ \forall j \text{ with } v_j > v_i] > 0.$$

Since there are allocations where all bidders with higher valuation that v_i receive their demand, it is $\sum_{j:v_j>v_i} d_j < k$. Moreover, whenever bidders with valuation at most v_i receive units (these bidders must have zero expected utility since the lower payment is v_i), those bidders with valuation higher than v_i receive their demand. Overall, bidders with valuation higher than v_i receive their demand with probability 1. The rest of the units are given to bidders with valuation v_i (because the payment is at least v_i) which leads to optimal social welfare. □

Next, we show that to have inefficiency at an equilibrium, there must exist at least two bidders with positive expected utility.

Lemma 3.3.8. *Let* $(\mathbf{v}, \mathbf{d})$ *be a valuation profile and* $\mathbf{B}$ *be an inefficient mixed Nash equilibrium. Then,* $|W(\mathbf{B})| \geq 2$.

Proof. Suppose on the contrary that there exists only one single bidder i with $\mathbb{E}_{\mathbf{b}\sim B}[u_i(\mathbf{b})] > 0$ and for any other bidder $j \neq i$, $\mathbb{E}_{\mathbf{b}\sim B}[u_j(\mathbf{b})] = 0$. By Lemma 3.3.7, $\mathbb{E}_{\mathbf{b}\sim B}[x_j(\mathbf{b})] = 0$, for all $j \neq i$ and therefore, $\mathbb{E}_{\mathbf{b}\sim B}[x_i(\mathbf{b})] = k$. Moreover, since $\mathbf{B}$ is inefficient there exists a bidder $i' \neq i$ with $v_{i'} > v_i$. Since i has a positive expected utility, he receives the units in a price less than v_i. If bidder i' bids v_i, he can satisfy his demand which results in a positive expected utility leading to a contradiction. □

The next warm-up properties involve the expected utility of a bidder under an equilibrium $\mathbf{B}$, conditioned that she bids within a certain interval or at a single point. We start with Fact 3.3.2, which is a straightforward implication of the equilibrium definition, and proceed by arguing that no two bidders may bid on the same point with positive probability. Theorem 3.3.2 concludes by stating the main property regarding the utility of bidders when bidding in their support.

Fact 3.3.2. *Let* $\mathbf{B}$ *be an equilibrium. For a bidder* i*, consider a partition of* $Supp(B_i)$ *(or of a subset of it) into smaller disjoint sub-intervals, say* $I_1, \ldots, I_\ell$*, such that* B_i *has a positive probability on each sub-interval (mass points may also be considered as sub-intervals). Then, it should hold that* $\mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b}) \mid b_i \in I_r] = \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})]$*, for every* $r = 1, \ldots, \ell$.

Based on Fact 3.3.2, we can obtain the following point-wise version. Variations of the version below have also appeared in related works, see e.g., Christodoulou et al. (2016b). For completeness, we present its proof in Section A.2 of Appendix A.

Theorem 3.3.2. *Given a mixed Nash equilibrium* $\mathbf{B}$*, bidder* i *and* $z \in Supp(B_i)$*, where no other bidder has a mass point on* z*,* $\mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(z, \mathbf{b}_{-i})] = \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})]$.

We further give the following observation regarding the existence of mass points on $\ell(\mathbf{B}_W)$.

Observation 3.3.1. *In any inefficient mixed Nash equilibrium* $\mathbf{B}$*, there can be no bidders* $i, j \in W(\mathbf{B})$ *such that both* $Pr[b_i = \ell(\mathbf{B}_W)] > 0$ *and* $Pr[b_j = \ell(\mathbf{B}_W)] > 0$.

Proof. Let $\ell = \ell(\mathbf{B}_W)$ and for the sake of contradiction suppose that $Pr[b_i = \ell] > 0$ and $Pr[b_j = \ell] > 0$. First note that $x_i(\ell, \mathbf{b}_{-i}) > 0$, otherwise, by Fact 3.3.2, $\mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})] = 0$ and $i \notin W(\mathbf{B})$ which is a contradiction. The same holds for bidder j.

By Lemma 3.3.6 we can infer that there exists a bidder with mass point on ℓ that doesn't receive his whole demand while bidding ℓ. W.l.o.g. let i be that bidder. Then, we can find a small enough δ so that transferring all the mass from ℓ to $\ell + \delta$ would yield higher utility for i; the reason is that i would receive more units by winning bidder j for whom we argued above that $x_i(\ell, \mathbf{b}_{-i}) > 0$. □

The main theorem of this section follows, stating the existence of a special bidder, who always receives at least one unit, and is referred to as *non-empty-handed*.

Theorem 3.3.3. *Let* $(\mathbf{v}, \mathbf{d})$ *be a valuation profile, and let* $\mathbf{B}$ *be any inefficient mixed Nash equilibrium. Then, there exists a bidder* $i \in W(\mathbf{B})$*, such that*

$$\sum_{j \in W(\mathbf{B}) \setminus \{i\}} d_j \leq k - 1 \,.$$

Proof. On the contrary, suppose that for every $i \in W(\mathbf{B})$, $\sum_{j \in W(\mathbf{B}) \setminus \{i\}} d_j \geq k$. Let i be some bidder with $\ell = \ell(\mathbf{B}_W) \in Supp(B_i)$. We distinguish two cases.

Case 1: There exists an interval in the form $[\ell, \ell + \epsilon]$, on which B_i has a positive probability mass and on which the bidders of $W(\mathbf{B}) \setminus \{i\}$ have a zero mass. We note that we also allow $\epsilon = 0$, i.e., that i has a mass point on ℓ and the other bidders do not. This means that when bidder i bids within $[\ell, \ell + \epsilon]$, all the other bidders from $W(\mathbf{B})$ are above him. Since we assumed that the total demand of $W(\mathbf{B}) \setminus \{i\}$ is at least k, bidder i does not win any units in this case. Since i bids with positive probability in $[\ell, \ell + \epsilon]$, by Fact 3.3.2, we have $\mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})] = 0$, which contradicts the fact that $i \in W(\mathbf{b})$.

Case 2: Note that by Observation 3.3.1, it cannot happen that both bidder i and at least one bidder $j \in W \setminus \{i\}$, have a mass point on ℓ. Hence, the only remaining case to consider is that any mass point that may exist by the bidders is at some $x > \ell$, and there is also no interval starting from ℓ that is used only by bidder i. Thus, there exists an interval I in the form $I = [\ell, \ell + \epsilon]$ for some small enough $\epsilon > 0$, and a bidder

$j \in W(\mathbf{B}) \setminus \{i\}$, such that both B_i and B_j contain I in their support, and have positive probability mass on I without mass points.

By Theorem 3.3.2, we obtain that $\mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(\ell, \mathbf{b}_{-i})] = \mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})] > 0$. This is a contradiction, because by bidding ℓ, bidder i ranks lower than all other bidders of $W(\mathbf{B})$ with probability one. By our assumption that $\sum_{j \in W(\mathbf{B}) \setminus \{i\}} d_j \geq k$, there are no units left for i when she ranks last among $W(\mathbf{B})$, and therefore, $\mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(\ell, \mathbf{b}_{-i})] = 0$. □ □

The property above already implies the following interesting corollary, that if all bidders have unit demand, any mixed Nash equilibrium is efficient.

Corollary 3.3.1. *Let $(\mathbf{v}, \mathbf{d})$ be a valuation profile with only unit-demand bidders, i.e., $d_i = 1$ for all i. Then any mixed Nash equilibrium $\mathbf{B}$ is efficient.*

Proof. If $|W(\mathbf{B})| < 2$ then, by Lemma 3.3.8, $\mathbf{B}$ is efficient. If $|W(\mathbf{B})| \geq 2$, by Theorem 3.3.3 there exist at most k bidders with positive expected utility. The rest of the bidders have zero expected allocation by Lemma 3.3.7, otherwise $\mathbf{B}$ is efficient. The only way that $\mathbf{B}$ is inefficient is if there exist bidders i and j with $v_i < v_j$ and $i \in W(\mathbf{B})$ whereas $j \notin W(\mathbf{B})$. In such a case, $\mathbb{E}[x_i(\mathbf{b})] > 0$ and if bidder j bids $v_i + \varepsilon < v_j$, $\mathbb{E}[x_j(v_i + \varepsilon, \mathbf{b}_{-j})] > \mathbb{E}[x_i(\mathbf{b})] > 0$, which results in a positive expected utility for bidder j contradicting the fact that $\mathbf{B}$ is a Nash equilibrium. □

3.3.3 The support and the CDFs of Mixed Nash Equilibria

The existence of a non-empty-handed bidder (Theorem 3.3.3) helps us to establish further properties that characterize the structure of inefficient mixed Nash equilibria. These properties (and especially Theorem 3.3.4) will be important to establish the inefficiency results that follow. We start with an observation regarding the highest bid of any bidder $i \in W(\mathbf{B})$, which should be strictly less than v_i.

Observation 3.3.2. *For any bidder $i \in W(\mathbf{B})$, $h(B_i) < v_i$.*

Proof. Suppose on the contrary that for some $i \in W(\mathbf{B})$, $h(B_i) = v_i$. Then, by Theorem 3.3.2, $\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})] = \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(v_i, \mathbf{b}_{-i})] = 0$, which contradicts the fact that $i \in W(\mathbf{B})$. □

The next lemma shows that at any equilibrium $\mathbf{B}$, bidders who are guaranteed units cannot have higher bids in their support than the support of the non-empty-handed bidders. Moreover, any bidder who is non-empty-handed does not have a reason to use bids that are higher than the maximum bid of all other winning bidders. The reason is that if such differences existed, then there would be incentives to win the same number of units by lowering one's bid. Then, Lemma 3.3.10 shows that no bidder will bid alone at any point or interval, and Lemma 3.3.11 specifies that no mass points may exist apart from one case.

Lemma 3.3.9. *Let* $(\mathbf{v}, \mathbf{d})$ *be a valuation profile and* $\mathbf{B}$ *be any inefficient mixed Nash equilibrium. Then, for any non-empty-handed bidder* i*, it holds that* $h(B_i) = h(\mathbf{B}_{W\setminus\{i\}}) = h(\mathbf{B}_W)$*.*

Proof. Suppose for contradiction that i is a non-empty-handed bidder, and there exists a bidder $j \in W(\mathbf{B}) \setminus \{i\}$ (non-empty-handed or not), such that $h(B_j) > h(B_i)$. Since $j \in W(\mathbf{B})$, it must be that $v_j \geq h(B_j)$ and bidder j obtains positive utility when she bids in $Supp(B_j) \cap (h(B_i), h(B_j)]$ (otherwise j would have an incentive not to bid above $h(B_i)$). Moreover, $\mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[x_j(\mathbf{b}) \mid b_i \in (h(B_i), h(B_j)]] = d_j$, since outbidding a non-empty-handed bidder guarantees the allocation of a bidder's entire demand by the auction. However, bidder j can then benefit from transferring probability mass from $(h(B_i), h(B_j)]$ to a point $h(B_i) + \delta$, for some small enough $\delta > 0$, since it still guarantees the allocation of her entire demand but for a strictly better price and thus strictly better expected utility. Hence, we have proved that $h(B_j) \leq h(B_i)$.

We will now also prove that $h(B_i) \leq \max_{j\in W(\mathbf{B})\setminus\{i\}} h(B_j)$. Consider a bidder $j \in W(\mathbf{B}) \setminus \{i\}$. If j is also non-empty-handed, then we can just repeat the argument above by switching the places of i and j, and we are done. Otherwise, i is the only non-empty-handed bidder and suppose $h(B_i) > h(B_j)$ for every $j \in W(\mathbf{B}) \setminus \{i\}$. Then, whenever bidder i bids above every $h(B_j)$, she ranks first, and hence she is granted all her demand. But then, she has incentives to reduce her bid so that she is still above every $h(B_j)$ and win the same units at a lower price, which is a contradiction. So, $h(B_i) = h(\mathbf{B}_{W\setminus\{i\}})$. Then it is straightforward to see that $h(\mathbf{B}_{W\setminus\{i\}}) = h(\mathbf{B}_W)$. □

Lemma 3.3.10. *Let* $(\mathbf{v}, \mathbf{d})$ *be any valuation profile and* $\mathbf{B}$ *be any mixed Nash equilibrium. For all* $i \in W(\mathbf{B})$*, it holds that* $Supp(B_i) \subseteq \bigcup_{j\in W(\mathbf{B})\setminus\{i\}} Supp(B_j)$*.*

Proof. Fix a bidder $i \in W(\mathbf{B})$ and let $I \subseteq Supp(B_i)$ be any interval such that $I \not\subseteq \bigcup_{j\in W(\mathbf{B})\setminus\{i\}} Supp(B_j)$. We distinguish two cases: either $I = [\ell, h]$ for $\ell < h$, or I is an isolated point.

For the first case we can establish that i would have an incentive to bid only on ℓ and still win the same units at a lower price. Indeed,

$$\begin{aligned}\mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b}) \mid b_i \in I] &= \mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[x_i(b_i, \mathbf{b}_{-i})(v_i - b_i) \mid b_i \in I] \\ &= \mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[x_i(\ell, \mathbf{b}_{-i})(v_i - b_i) \mid b_i \in I] \\ &< \mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[x_i(\ell, \mathbf{b}_{-i})(v_i - \ell)]\end{aligned}$$

The second equality is due to the fact that no other bidder in $W(\mathbf{B})$ bids in I with positive probability, whereas the strict inequality follows by the construction of I. This inequality yields a contradiction, by Fact 3.3.2, since there exists a profitable transfer of the probability mass of I to the point ℓ.

For the second case suppose that I is some isolated point z that i bids with positive probability; z is isolated because we have assumed WLOG that all intervals defining the domain of a distribution are closed. Let $h = \max_{j\in W(\mathbf{B})\setminus\{i\}} h(B_j)$. If $z > h$, then i would be benefited by transferring the probability of bidding z to any point between h and z. If $z < h$, let z' be the maximum point such that $[z, z') \not\subseteq \bigcup_{j\in W(\mathbf{B})\setminus\{i\}} Supp(B_j)$ (note that $z \neq z'$). Then, there exists a bidder $i' \neq i$ such that, by Theorem 3.3.2, $\mathbb{E}_{\mathbf{b}_{-i'}\sim B_{-i'}}[u_{i'}(z', \mathbf{b}_{-i'})] = \mathbb{E}_{\mathbf{b}\sim B}[u_{i'}(\mathbf{b})]$. Bidding any bid between z and z' would result to a higher expected utility for bidder i' than $\mathbb{E}_{\mathbf{b}_{-i'}\sim B_{-i'}}[u_{i'}(z', \mathbf{b}_{-i'})]$, which is a contradiction to the fact that $\mathbf{B}$ is a Nash equilibrium. □

Lemma 3.3.11. *Let* $(\mathbf{v}, \mathbf{d})$ *be a valuation profile and* $\mathbf{B}$ *be any inefficient mixed Nash equilibrium.*
1) There exists no bidder $i \in W(\mathbf{B})$ *and no point* $z \in Supp(B_i) \setminus \{\ell(\mathbf{B}_W)\}$, *with* $F_i(z) > \lim_{z\to z^-} F_i(z)$, *i.e., there are no mass points among the bidders of* $W(\mathbf{B})$, *except possibly the leftmost endpoint of all bidders' distributions.*
2) At most one bidder $i \in W(\mathbf{B})$ *may have a mass point on* $\ell(\mathbf{B}_W)$, *in which case,* i *is a non-empty-handed bidder.*

Proof. Regarding the first part of the statement, suppose, for contradiction, that there exists a bidder $i \in W(\mathbf{b})$, and a point $z \in Supp(B_i) \setminus \{\ell(\mathbf{B}_W)\}$ with $F_i(z) > \lim_{z\to z^-} F_i(z)$. We next distinguish the following cases:
Case 1: There does not exist any bidder $j \in W(\mathbf{B}\setminus\{i\})$ with an interval $I = [z-\delta, z] \subset Supp(B_j)$ for some small enough $\delta > 0$, with $z - \delta > \ell(\mathbf{B}_W)$. Then by Lemma 3.3.10 it must also be that the interval $[z - \delta, z)$ is not in the support of bidder i. Furthermore, I is not in the support of any other bidder, who do not belong to $W(\mathbf{B})$ (since the

interval I is to the right of $\ell(\mathbf{B}_W)$, bidders not in $W(\mathbf{B})$ cannot use an interval of this form, because then they would have a positive probability of winning).

Thus, the only way that bidder i would not prefer to choose any bid $\xi \in I$, with $\xi < z$, is when he doesn't win the same number of units as when bidding z; this can only happen when there exists another bidder $i' \in W(\mathbf{B})$ with a mass point on z and the tie breaking rule favors bidder i. But in that case bidder i' can transfer his mass from z to a slightly higher bid [4] (similar to the proof of Observation 3.3.1) and receive higher utility. This results to a contradiction of $\mathbf{B}$ being an equilibrium.

Case 2: There exists a bidder $j \in W(\mathbf{B})\setminus\{i\}$ with an interval $I = [z-\delta, z] \subset Supp(B_j)$ for some small enough $\delta > 0$. But then

$$\begin{aligned}
\mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[u_j(b_j, \mathbf{b}_{-j}) \mid b_j \in [z-\delta, z]] &= \mathop{\mathbb{E}}_{b_j\sim B_j}\big[\mathop{\mathbb{E}}_{\mathbf{b}_{-j}\sim\mathbf{B}_{-j}}[u_j(b_j, \mathbf{b}_{-j})] \mid b_j \in [z-\delta, z]\big] \\
&= \lim_{\xi\to z^-} \mathop{\mathbb{E}}_{\mathbf{b}_{-j}\sim\mathbf{B}_{-j}}[u_j(\xi, \mathbf{b}_{-j})] \\
&= \lim_{\xi\to z^-} \mathop{\mathbb{E}}_{\mathbf{b}_{-j}\sim\mathbf{B}_{-j}}[x_j(\xi, \mathbf{b}_{-j})] \lim_{\xi\to z}(v_j - \xi) \\
&= d_j \lim_{\xi\to z^-} \hat{F}_j^{avg}(\xi)(v_j - z) \\
&< d_j \lim_{\xi\to z^+} \hat{F}_j^{avg}(\xi)(v_j - z)
\end{aligned}$$

The last equality in the above expressions is by Fact 3.3.1. The last inequality holds because $\hat{F}_j^{avg}$ has a discontinuity at z due to the fact that i assigns positive probability to z.

To conclude, the above series of equations imply that there exists a small enough ϵ such that

$$\mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[u_j(b_j, \mathbf{b}_{-j}) \mid b_j \in [z-\delta, z]] < \mathop{\mathbb{E}}_{\mathbf{b}_{-j}\sim\mathbf{B}_{-j}}[u_j(z+\epsilon, \mathbf{b}_{-j})]$$

which contradicts $\mathbf{B}$ being an equilibrium.

Regarding the second part of the statement, by Observation 3.3.1 it cannot be that two bidders have a mass point on $\ell(\mathbf{B}_W)$. For the sake of contradiction, suppose that there exists a bidder i with $Pr[b_i = \ell(\mathbf{B}_W)] > 0$ and i is not a non-empty-handed bidder. Then, the rest of the bidders in $W(\mathbf{B})$ bid higher than $\ell(\mathbf{B}_W)$ with probability one and therefore, bidder i doesn't win any unit by bidding $\ell(\mathbf{B}_W)$. By Theorem 3.3.2, $\mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})] = \mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(\ell(\mathbf{B}_W), \mathbf{b}_{-i})] = 0$, which contradicts the fact that $i \in W(\mathbf{B})$. □

[4]Note that $z < v_{i'}$, otherwise, by Fact 3.3.2, $\mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_{i'}(\mathbf{b})] = 0$ and $i' \notin W(\mathbf{B})$ which is a contradiction.

By combining Theorem 3.3.2 and Lemma 3.3.11 we get the following Corollary.

Corollary 3.3.2. *For any inefficient mixed Nash equilibrium* $\mathbf{B}$*, the following hold:*
1) For any bidder i *and* $z \in Supp(B_i) \setminus \{\ell(\mathbf{B}_W)\}$, $\mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(z, \mathbf{b}_{-i})] = \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})]$.
2) If there exists a bidder i *with* $Pr[b_i = \ell(\mathbf{B}_W)] > 0$*, then* i *is a non-empty-handed bidder and* $\mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(\ell(\mathbf{B}_W), \mathbf{b}_{-i})] = \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})]$.
3) If no non-empty-handed bidder exists with mass point on $\ell(\mathbf{B}_W)$*, for any bidder* i *with* $\ell(\mathbf{B}_W) \in Supp(B_i)$, $\mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(\ell(\mathbf{B}_W), \mathbf{b}_{-i})] = \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})]$.

Observation 3.3.3. *For any inefficient mixed Nash equilibrium* $\mathbf{B}$*, either there exists a non-empty-handed bidder* $i \in W(\mathbf{B})$ *with a mass point on* $\ell(\mathbf{B}_W)$*, or there are at least two non-empty-handed bidders with* $\ell(\mathbf{B}_W)$ *in their support.*

Proof. If there exists a bidder i with mass point on $\ell(\mathbf{B}_W)$, then by Lemma 3.3.11 i is a non-empty-handed bidder. If there is no such bidder then we argue that no bidder $j \in W(\mathbf{B})$ that is not non-empty-handed has $\ell(\mathbf{B}_W)$ in their support.

Suppose on the contrary that $\ell(\mathbf{B}_W) \in Supp(B_j)$ for some bidder j that is not non-empty-handed. Then since no bidder has a mass point on $\ell(\mathbf{B}_W)$, everybody bids above $\ell(\mathbf{B}_W)$ with probability one, leaving j with no units while bidding $\ell(\mathbf{B}_W)$. By Corollary 3.3.2, $\mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_j(\mathbf{b})] = \mathbb{E}_{\mathbf{b}_{-j}\sim\mathbf{B}_{-j}}[u_j(\ell(\mathbf{B}_W), \mathbf{b}_{-j})] = 0$ contradicting the fact that $j \in W(\mathbf{B})$.

By Lemma 3.3.10 there are at least two bidders bidding on $\ell(\mathbf{B}_W)$, which concludes the proof. □

Given any (inefficient) equilibrium, the next theorem specifies the average CDF of the winning bids that bidder i competes against, i.e., $\hat{F}_i^{avg}$, in i's support.

Theorem 3.3.4. *Let* $(\mathbf{v}, \mathbf{d})$ *be any valuation profile and* $\mathbf{B}$ *be any inefficient mixed Nash equilibrium. Then, for* $i \in W(\mathbf{B})$*, the CDF* $\hat{F}_i^{avg}$ *satisfies*

$$\hat{F}_i^{avg}(z) = \frac{u_i}{d_i(v_i - z)}\,, \quad \forall z \in Supp(B_i)\,,$$

where $u_i = \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})] > 0$.

Proof. Fix a bidder $i \in W(\mathbf{B})$. For all intervals $I \subseteq Supp(B_i)$, by Corollary 3.3.2 it must be that for all $z \in I \setminus \ell(\mathbf{B}_W)$

$$\mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(z, \mathbf{b}_{-i})] = \mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})] = u_i > 0\,.$$

The above equality is equivalent to

$$\mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[x_i(z,\mathbf{b}_{-i})] = \frac{u_i}{v_i - z} \Leftrightarrow \hat{F}_i^{avg}(z) = \frac{u_i}{d_i(v_i - z)}\,,$$

for all $z \in Supp(B_i)$. The last equivalence is due to Fact 3.3.1. The theorem follows since $\hat{F}_i^{avg}$ is right continuous. □

A corollary of Theorem 3.3.4 is that the union of the support of the winners is an interval.

Corollary 3.3.3. *Let* $(\mathbf{v},\mathbf{d})$ *be any valuation profile and* $\mathbf{B}$ *be any inefficient mixed Nash equilibrium. Then, for every bidder* $i \in W(\mathbf{B})$, $\bigcup_{j\in W(\mathbf{B})\setminus\{i\}} Supp(B_j) = [\ell(\mathbf{B}_W), h(\mathbf{B}_W)]$.

Proof. Suppose for contradiction that there exists an interval

$$(\ell', h') \nsubseteq \bigcup_{j\in W(\mathbf{B})\setminus\{i\}} Supp(B_j)$$

with $\ell' > \ell(\mathbf{B}_W)$, $h' < h(\mathbf{B}_W)$ and this is maximal. Then, let i be a bidder with h' in their support. By Theorem 3.3.4, $\hat{F}_i^{avg}(h') = \frac{u_i}{d_i(v_i - h')}$ and $\mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})] = \mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(h',\mathbf{b}_{-i})] = u_i$ by Corollary 3.3.2.

For any $x \in (\ell', h')$, $\hat{F}_i^{avg}(x) = \frac{u_i}{d_i(v_i-h')}$, since $(\ell', h') \nsubseteq \bigcup_{j\in W(\mathbf{B})\setminus\{i\}} Supp(B_j)$. Clearly, $\hat{F}_i^{avg}(x) > \frac{u_i}{d_i(v_i - x)}$ and by bidding x,

$$\mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(x,\mathbf{b}_{-i})] = d_i\hat{F}_i^{avg}(x)(v_i - x) > d_i\frac{u_i}{d_i(v_i - x)}(v_i - x) = u_i = \mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})]\,,$$

which is a contradiction to $\mathbf{B}$ being an equilibrium. □

The final lemma of this section shows that the rightmost point in the support of $\mathbf{B}$ is a function of the parameters of certain non-empty-handed bidders.

Lemma 3.3.12. *Let* $(\mathbf{v},\mathbf{d})$ *be any valuation profile and* $\mathbf{B}$ *be any inefficient mixed Nash equilibrium. Let* $i \in W(\mathbf{B})$ *be the non-empty-handed bidder such that* $Pr[b_i = \ell(\mathbf{B}_W)] > 0$*, or if no such bidder exists, then let* i *be any non-empty-handed bidder with* $\ell(\mathbf{B}_W)$ *in his support. We have*

$$h(\mathbf{B}_W) = h(B_i) = v_i - (k - \sum_{j\in W(\mathbf{B})\setminus\{i\}} d_j)\frac{v_i - \ell(\mathbf{B}_W)}{d_i}\,.$$

Proof. Let i be the bidder specified by the lemma's statement; note that such a bidder always exists by Observation 3.3.3. By Lemma 3.3.9, $h(\mathbf{B}_W) = h(B_i)$, and by applying

Corollary 3.3.2, it must be that

$$\mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(\ell(B_i),\mathbf{b}_{-i})] = \mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(h(B_i),\mathbf{b}_{-i})]\,. \tag{3.1}$$

Moreover,

$$\mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(h(B_i),\mathbf{b}_{-i})] = \mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[x_i(h(B_i),\mathbf{b}_{-i})](v_i - h(B_i)) = d_i(v_i - h(B_i))\,. \tag{3.2}$$

Equation (3.2) holds since bidding $h(B_i)$ guarantees outbidding every other bidder in the auction and thus grants d_i units to i (recall that there is no mass point on $h(B_i)$ due to Lemma 3.3.11, and therefore the event of losing due to tie-breaking by bidding $h(B_i)$ has probability zero).

On the other hand, note that by the way i has been defined, $\ell(B_i) = \ell(\mathbf{B}_W)$ and therefore

$$\begin{aligned}
\mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim B_{-i}}[u_i(\ell(B_i),\mathbf{b}_{-i})] &= \mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim B_{-i}}[u_i(\ell(\mathbf{B}_W),\mathbf{b}_{-i})] \\
&= \mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim B_{-i}}[x_i(\ell(\mathbf{B}_W),\mathbf{b}_{-i})](v_i - \ell(\mathbf{B}_W)) \\
&= (k - \sum_{j\in W(\mathbf{B})\setminus\{i\}} d_j)(v_i - \ell(\mathbf{B}_W))\,,
\end{aligned} \tag{3.3}$$

where in the last equality above, we have that $k - \sum_{j\in W(\mathbf{B})\setminus\{i\}} d_j > 0$, since i is non-empty-handed. By equating now (3.2) and (3.3), the lemma follows. □

3.4 Price of Anarchy for mixed equilibria

We can now exploit the properties derived so far for mixed equilibria, in order to analyze the inefficiency of the discriminatory price auction. Since we focus on inefficient equilibria, we assume that in any valuation profile considered in this section, there are at least two bidders with a different value per unit.

3.4.1 The case of two bidders

We pay particular attention to the case of $n = 2$. This is a setting where we can fully characterize in closed form the distributions of the inefficient mixed Nash equilibria, and derive valuable intuitions for the worst-case instances with respect to the Price of Anarchy, that are helpful also for auctions with multiple bidders. The main result of this subsection is the following theorem, showing that the inefficiency is quite limited.

Theorem 3.4.1. *For $k \geq 2$, $n = 2$ and capped additive valuation profiles, the Price of Anarchy of mixed equilibria is at most* 1.1095*, and this is tight as k goes to infinity.*

We postpone the proof of Theorem 3.4.1, as we first need to establish some properties regarding the form of inefficient mixed Nash equilibria with two bidders. For $n = 2$, a capped-additive valuation profile can be described as $(\mathbf{v}, \mathbf{d}) = ((v_1, d_1), (v_2, d_2))$. Recall also that it is sufficient to focus our attention only on profiles where $d_1 + d_2 > k$, since otherwise, by Lemma 3.3.6 any mixed equilibrium is efficient. We start our analysis by characterizing the support of inefficient mixed Nash equilibria.

Lemma 3.4.1. *Let $(\mathbf{v}, \mathbf{d}) = ((v_1, d_1), (v_2, d_2))$ be any capped-additive valuation profile of two bidders, and $\mathbf{B} = (B_1, B_2)$ be any inefficient mixed Nash equilibrium. Then:*

1. *$Supp(B_1) = Supp(B_2) = [\ell(B_1), h(B_1)]$, and $\ell(B_1) = 0$.*
2. *$h(B_1)$ takes one of the following values*

$$h(B_1) = v_1 \frac{d_1 + d_2 - k}{d_1} \quad \text{or} \quad h(B_1) = v_2 \frac{d_1 + d_2 - k}{d_2}.$$

Proof. To prove the first statement, by Lemma 3.3.10, we have that $Supp(B_1) = Supp(B_2)$. Also, by Corollary 3.3.3, we have that $Supp(B_1)$ is an interval. To prove that $\ell(B_1) = 0$, we utilize Lemma 3.3.8 which states that $|W(\mathbf{B})| \geq 2$. Thus, with exactly two bidders, we have that $|W(\mathbf{B})| = 2$, and both bidders have positive expected utility. Let $\ell = \ell(B_1) = \ell(B_2)$. If $\ell > 0$, we will argue that one of the players has an incentive to deviate to a lower bid.

Note first that by Observation 3.3.1, it cannot be that ℓ is a mass point for both bidders. WLOG suppose that bidder 1 has a mass point at ℓ, which by Fact 3.3.2, implies she has a positive expected utility, when playing ℓ. At the same time, bidder 2 bids higher than ℓ with probability equal to 1. Hence, there must be some left over units that bidder 1 wins when bidding ℓ in order to have a positive expected utility. But now this means that bidder 1 would have a profitable transfer of probability mass to 0 in order to have a zero payment while obtaining the same number of units. If neither bidder has mass on ℓ, we can use Corollary 3.3.2 to have that the expected utility of bidder 1 at ℓ equals $E[u_1(\mathbf{b})] > 0$. Hence, she wins some units with positive probability when bidding ℓ. But then bidder 1 would win the same number of units by bidding 0 resulting in higher utility, a contradiction to $\mathbf{B}$ being an equilibrium.

To prove the second statement of the lemma, we use Theorem 3.3.3 that states that at least one of the bidders, say bidder i, is non-empty-handed and by Lemma 3.3.12

we obtain

$$h(B_i) = v_i \frac{d_1 + d_2 - k}{d_i}.$$

□

The following theorem specifies the cumulative distribution functions that comprise any inefficient mixed Nash equilibrium, along with a necessary condition for the existence of such equilibria. For a bidder i below, we use the notation v_{-i} and d_{-i} to denote the value and demand of the other bidder.

Theorem 3.4.2. *Let* $(\mathbf{v}, \mathbf{d}) = ((v_1, d_1), (v_2, d_2))$ *be a capped-additive valuation profile of two bidders, and* $\mathbf{B} = (B_1, B_2)$ *be any inefficient mixed Nash equilibrium.*

1. *The cumulative distribution function of bidder* i*, for* $i = 1, 2$*, is*

$$F_i(z) = \frac{1}{d_1 + d_2 - k} \left(\frac{d_{-i}(v_{-i} - h(B_i))}{v_{-i} - z} - (k - d_i) \right). \tag{3.4}$$

2. *Furthermore, for* i *being the non-empty-handed bidder with a mass point at* 0*, or if no such bidder exists, being any non-empty-handed bidder, it holds that* $\frac{v_{-i}}{v_i} \geq \frac{d_{-i}}{d_i}$*,*

Proof. It is convenient to look into the $\hat{F}_i^{avg}$ distribution to derive the claimed formulas. In the two-bidder environment, the sole source of competition is a single bidder. Firstly, for any bidder i, the other bidder always obtains $k - d_i$ units[5] and there is a competition for the remaining $d_i - (k - d_{-i}) = d_1 + d_2 - k$ units. Therefore, for the competitor of bidder i, we have:

$$\hat{F}_{-i}^{avg}(z) = \frac{\sum_{j=1}^{d_{-i}} \hat{F}_{-i,j}(z)}{d_{-i}} = \frac{k - d_i + (d_1 + d_2 - k)F_i(z)}{d_{-i}}.$$

Now by Theorem 3.3.4, for $z \in Supp(B_{-i})$ (which is the same as $Supp(B_i)$ by Lemma 3.4.1), we have

$$\hat{F}_{-i}^{avg}(z) = \frac{u_{-i}}{d_{-i}(v_{-i} - z)},$$

for some constant $u_{-i} > 0$. By combining the last two equations and rearranging terms we obtain

$$F_i(z) = \frac{1}{d_1 + d_2 - k} \left(\frac{u_{-i}}{v_{-i} - z} - (k - d_i) \right).$$

[5] This is consistent with Theorem 3.3.3 since when a bidder is not non-empty-handed, it must be that the demand of the other bidder is k, and hence, she obtains $k - d_i = 0$ free units, as expected.

We now determine the appropriate value for $u_{-i} > 0$ so that $F_i(z)$ is a valid cumulative distribution function in $Supp(B_i)$. It must be that $F_i(h(B_i)) = 1$, since $h(B_i)$ is the rightmost point in her support. Hence,

$$1 = \frac{1}{d_1 + d_2 - k}\left(\frac{u_{-i}}{v_{-i} - h_i(B_i)} - (k - d_i)\right) \Leftrightarrow u_{-i} = d_{-i}(v_{-i} - h_i(B_i))\,.$$

This establishes that $F_i(z)$ satisfies Equation (3.4).

The second part of the theorem comes from the fact that the CDFs should also satisfy non-negativity. For this, let i be the bidder specified by the second part of the theorem's statement. Then, by using Lemma 3.3.12 it is a matter of simple calculations to see that $F_i(0) \geq 0$ is equivalent to $\frac{v_{-i}}{v_i} \geq \frac{d_{-i}}{d_i}$. Since F_i is increasing, we have established that this condition is necessary to enforce that $F_i(z) \geq 0$ for every $z \in Supp(B_i)$. □

Remark 3.4.1. *By Lemma 3.4.1 and Theorem 3.4.2, we can see that there can be at most two inefficient equilibria, depending on how the interval of the support was determined.*

We are now ready to prove Theorem 3.4.1.

Proof of Theorem 3.4.1. The properties established so far imply a full characterization of instances that have inefficient equilibria. To establish Theorem 3.4.1, we will group instances into three appropriate classes and we will solve an appropriately defined optimization problem that approximates the Price of Anarchy for each subclass to arbitrary precision.

WLOG, suppose we are given a value profile $(\mathbf{v}, \mathbf{d}) = ((v_1, d_1), (v_2, d_2))$ of k units such that $d_1 \geq d_2$. We define the following two quantities, which we refer to as the *normalized* demands.

$$\bar{d}_1 = \frac{d_1}{k} > 0 \qquad \bar{d}_2 = \frac{d_2}{k} > 0 \tag{3.5}$$

Essentially, we intend to use $v_1, v_2, \bar{d}_1$ and $\bar{d}_2$ as the variables of the optimization problem mentioned before.

Let $\mathbf{B}$ be any inefficient mixed Nash equilibrium. With a slight abuse of notation we view the term $h(B_i)$ as a function of the valuation profile parameters, as established by Lemma 3.4.1, and define the functions $h_i(\mathbf{v}, \bar{\mathbf{d}}) = v_i \frac{\bar{d}_1 + \bar{d}_2 - 1}{\bar{d}_i}$ for $i = 1, 2$. We pair these functions with two additional expressions $SW_i(\mathbf{v}, \mathbf{d})$ for $i = 1, 2$ which are (scaled) restatements of the social welfare of an equilibrium, solely in terms of the value profile $(\mathbf{v}, \mathbf{d})$ and k, and without dependencies on the underlying equilibrium distributions. The reason we are able to do so, is Theorem 3.4.2, which tells us what the CDFs are, in terms of the valuation profile. The exact form of $SW_i(v_1, v_2, \bar{d}_1, \bar{d}_2)$ for $i = 1, 2$ is

$$\bar{d}_{-i}(v_{-i}-v_i)\left(1-\int_0^{h_i(\mathbf{v},\bar{\mathbf{d}})}\frac{1}{\bar{d}_1+\bar{d}_2-1}\left(\frac{\bar{d}_i(v_i-h_i(\mathbf{v},\bar{\mathbf{d}}))}{v_i-z}-(1-\bar{d}_{-i})\right)\frac{v_{-i}-h_i(\mathbf{v},\bar{\mathbf{d}})}{(v_{-i}-z)^2}dz\right)+v_i.$$

With these expressions in mind, the following lemma allows us to obtain the social welfare in a form that we can later exploit for producing our upper bound. The lemma follows by Theorem 3.4.2, which tells us what the equilibrium CDFs are, in terms of the valuation profile.

Lemma 3.4.2. *Let i be a non-empty handed bidder with a mass point at 0. Then, $SW(\mathbf{B}) = kSW_i(\mathbf{v}, \bar{\mathbf{d}})$. If no such bidder exists, then either $SW(\mathbf{B}) = kSW_1(\mathbf{v}, \bar{\mathbf{d}})$ or $SW(\mathbf{B}) = kSW_2(\mathbf{v}, \bar{\mathbf{d}})$.*

To proceed, we will distinguish the following two cases.

1. If $\mathbf{B} = (B_1, B_2)$ is such that $Supp(B_1) = Supp(B_2) = [0, v_1\frac{d_1+d_2-k}{d_1}]$, then, by Lemma 3.4.2, $SW(\mathbf{B}) = kSW_1(\mathbf{v}, \bar{\mathbf{d}})$ and by the second part of Theorem 3.4.2 it must be that $\frac{v_2}{v_1} \geq \frac{d_2}{d_1}$ or, equivalently, in terms of normalized demands as $\frac{v_2}{v_1} \geq \frac{\bar{d}_2}{\bar{d}_1}$. We split the analysis into the following sub-cases:

    (a) When $v_1 > v_2$, the optimal social welfare is determined by allocating bidder 1 her entire demand and, subsequently allocating bidder 2 the leftover units. Therefore, in this case $OPT(\mathbf{v}, \mathbf{d}) = v_1 d_1 + (k - d_1)v_2 = k(v_1\bar{d}_1 + (1 - \bar{d}_1)v_2)$ and $\frac{OPT(B)}{SW(B)} = \frac{v_1\bar{d}_1+(1-\bar{d}_1)v_2}{SW_1(v_1,v_2,\bar{d}_1,\bar{d}_2)}$. Hence, the Price of Anarchy of mixed Nash equilibria for this subclass is upper bounded by the optimal solution to the following problem

$$\begin{aligned} \max_{v_1,v_2,\bar{d}_1,\bar{d}_2} \quad & \frac{v_1\bar{d}_1 + (1-\bar{d}_1)v_2}{SW_1(v_1, v_2, \bar{d}_1, \bar{d}_2)} \\ \text{subject to} \quad & 1 > \frac{v_2}{v_1} \geq \frac{\bar{d}_2}{\bar{d}_1}. \end{aligned} \tag{3.6}$$

    (b) Similarly, when $v_1 < v_2$ the optimal social welfare is determined by allocating bidder 2 her entire demand and, subsequently allocating bidder 1 the leftover units. Therefore, in this case $OPT(B) = v_1(k - d_2) + d_2v_2 = k(v_1(1 - \bar{d}_2) + v_2\bar{d}_2)$ and the Price of Anarchy for this subclass is upper bounded by the optimal solution to the following problem

$$
\begin{aligned}
\max_{v_1, v_2, \bar{d}_1, \bar{d}_2} \quad & \frac{v_1(1-\bar{d}_2) + v_2\bar{d}_2}{SW_1(v_1, v_2, \bar{d}_1, \bar{d}_2)} \\
\text{subject to} \quad & v_2 > v_1 . \\
& \bar{d}_1 \geq \bar{d}_2 \, . \\
& 1 > \bar{d}_2 \, .
\end{aligned}
\tag{3.7}
$$

Note that in this sub-case, we enforce the last constraint that $\bar{d}_2 < 1$ (implicitly enforced in the first sub-case). Since we assumed $d_1 \geq d_2$, there can be no mixed Nash equilibrium with $\bar{d}_2 = 1$, because then both bidders are additive, violating the condition of Theorem 3.3.3.

2. The final case we need to consider is equilibria $\mathbf{B} = (B_1, B_2)$ such that $Supp(B_1) = Supp(B_2) = [0, v_2 \frac{d_1+d_2-k}{d_2}]$ when $d_1 < k$ (recall if $d_1 = k$ bidder 1 is non-empty-handed and the support will be as in the first case). As in the previous case, by the second part of Theorem 3.4.2, it must be that $\frac{v_1}{v_2} \geq \frac{d_1}{d_2}$. However, unlike the class of equilibria described in the previous paragraph, it is sufficient to consider here only the case when $v_1 > v_2$ since, due to our assumption that $d_1 \geq d_2$, the condition $\frac{v_1}{v_2} \geq \frac{d_1}{d_2}$ implies that there cannot exist mixed Nash equilibria when $v_1 < v_2$. Thus, the Price of Anarchy for this subclass is upper bounded by

$$
\begin{aligned}
\max_{v_1, v_2, \bar{d}_1, \bar{d}_2} \quad & \frac{v_1\bar{d}_1 + (1-\bar{d}_1)v_2}{SW_2(v_1, v_2, \bar{d}_1, \bar{d}_2)} \\
\text{subject to} \quad & \frac{v_1}{v_2} \geq \frac{\bar{d}_1}{\bar{d}_2} \geq 1 \\
& 1 > \bar{d}_2 \, .
\end{aligned}
\tag{3.8}
$$

By solving numerically the optimization problems of Equations (3.6), (3.7) and (3.8), we found out that the worst case instances arise by the sub-case given by (3.6). In particular, the maximum value for the objective function we obtained was approximately 1.1095 and the optimal values for the four variables are $v_1 = 1, v_2 \approx 0.526, \bar{d}_1 = 1, \bar{d}_2 \approx 0.357$. This concludes the proof of the upper bound on the Price of Anarchy. Furthermore, it is not hard to convert the variables to the underlying worst case instance, which we present in the next paragraph.

Tight Example. Consider an instance of the discriminatory auction for $k \geq 4$ units and $n = 2$ bidders. Bidder 1 has value $v_1 = 1$ and $d_1 = k$, whereas bidder 2 has a value $v_2 = 0.526$ and $d_2 = \lceil 0.357k \rceil$ units. Let B_1, B_2 be two distributions supported in $[0, \frac{d_2}{k}]$. Note that $v_2 > \frac{d_2}{k}$. In accordance to Equation (3.4), the cumulative distribution functions of B_1 and B_2 are

$$F_1(z) = \frac{v_2 - \frac{d_2}{k}}{v_2 - z}, \qquad\qquad F_2(z) = \frac{k - d_2}{d_2} \frac{z}{1 - z}.$$

It is easy to verify that $\mathbf{B} = (B_1, B_2)$ is indeed a mixed equilibrium. The optimal allocation is for bidder 1 to obtain all k units and the expected social welfare of $\mathbf{B}$, by Lemma 3.4.2, is $SW(\mathbf{B}) = kSW_1(\mathbf{v}, \bar{\mathbf{d}})$, since $F_1(0) > 0$. The worst case inefficiency ratio occurs as k grows and is approximately 1.1095. $\square$

3.4.2 Multiple Bidders

Inspired by the construction in the previous section, we move to instances with more than two bidders and provide first a lower bound on the Price of Anarchy. This bound shows a separation between $n = 2$ and $n > 2$, in the sense that equilibria can be more inefficient with a higher number of bidders. It also improves the best known lower bound of the discriminatory price auction for the class of submodular valuations, which was 1.109, by Christodoulou et al. (2016b). The improvement however is rather small.

Theorem 3.4.3. *For $n > 2$, and for the class of mixed strategy Nash equilibria, the Price of Anarchy is at least* 1.1204.

Proof. Consider a discriminatory auction instance of $k \geq 2$ units. Let the number of bidders be $n + 1$: one additive bidder, denoted by α, that competes against $n < k$ unit-demand bidders. We assume that ties are in favor of bidder α. The value per unit of the additive bidder is 1, whereas the value of the i-th unit-demand bidder is v_i, for $i = 1, \dots, n$. The values of the unit-demand bidders are sorted in increasing order, i.e. $v_1 \leq v_2 \leq \dots \leq v_n$ and $v_n \leq 1$. For convenience, we define, for $i = 0, \dots, n$, the auxiliary terms $h_i = \frac{i}{k-n+i}$. Moreover, we will later choose the values so that they satisfy the following set of inequalities:

1. For $i = 1, \dots, n-1$, $m = i + 1, \dots, n$, and every $z \in [h_{m-1}, h_m]$:

$$\prod_{j=i+1}^{m-1} \frac{v_j - h_j}{v_j - h_{j-1}} \geq \frac{v_i - z}{v_i - h_i} \frac{v_m - h_{m-1}}{v_m - z}. \tag{3.9}$$

2. For $i = 2, \ldots, n,\ \ m = 1, \ldots, i-1$, and every $z \in [h_{m-1}, h_m]$:

$$\prod_{j=m+1}^{i-1} \frac{v_j - h_j}{v_j - h_{j-1}} \leq \frac{v_m - z}{v_m - h_m} \frac{v_i - h_{i-1}}{v_i - z}. \tag{3.10}$$

Let $\mathbf{B}$ be a product distribution. The additive bidder α bids according to a distribution B_α supported in $[0, h_n]$ with the cumulative distribution function F_α. F_α is a branch function with n branches, where for $i = 1, \ldots, n$, the form of F_α at $[h_{i-1}, h_i]$, denoted by F_α^i, is

$$F_\alpha(z) = F_\alpha^i(z) = \prod_{j=i+1}^{n} \left(\frac{v_j - h_j}{v_j - h_{j-1}} \right) \frac{v_i - h_i}{v_i - z}.$$

The distribution B_i of each unit-demand bidder $i = 1, \ldots, n$, is supported in $[h_{i-1}, h_i]$, and the form of its CDF is

$$F_i(z) = \frac{k-n}{1-z} - (k - n + i - 1)$$

We now show that this construction is indeed a mixed Nash equilibrium, with the following lemma.

Lemma 3.4.3. *The profile* $\mathbf{B}$ *is an equilibrium, provided the values* $v_1, \ldots, v_n$ *satisfy Equations* (3.9) *and* (3.10)*.*

Proof. Firstly, when the additive bidder bids the rightmost point in her support $h_n = \frac{n}{k}$, this grants her an allocation of k (since she outbids all the unit-demand bidders) and an expected utility of $k(1 - \frac{n}{k}) = k - n$. Therefore, bidding above h_n is a dominated strategy for bidder α, since she would still win k units but will be asked to pay more than $h_n = \frac{n}{k}$. On the other hand for $i = 1, \ldots, n$, bidding $z \in [h_{i-1}, h_i)$, grants bidder α an expected utility of

$$\mathop{\mathbb{E}}_{\mathbf{b}_{-\alpha} \sim \mathbf{B}_{-\alpha}} [u_\alpha(z, \mathbf{b}_{-a}) | z \in [h_i - 1, h_i)] = (k - n + i - 1 + F_i(z))\,(1 - z) = k - n.$$

Therefore, by taking the expectation over B_α on both sides of this equation, we obtain that $\mathbb{E}_{\mathbf{b} \sim \mathbf{B}} [u_\alpha(\mathbf{b})] = k - n$ and bidder α has no profitable unilateral deviation.

We now examine the incentives for unilateral deviations of the n unit-demand bidders. For each one of the unit-demand bidders $i = 1, \ldots, n$, their expected utility for bidding in the interval of their support (h_{i-1}, h_i) is

$$\mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}\left[u_i(z,\mathbf{b}_{-i})|z\in(h_{i-1},h_i)\right] = F_\alpha^i(z)(v_i - z) = \prod_{j=i+1}^{n}\left(\frac{v_j - h_j}{v_j - h_{j-1}}\right)(v_i - h_i)$$

and, by taking an expectation on both sides of the above equation, the expected utility of unit-demand bidder i is

$$\mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}\left[u_i(\mathbf{b})\right] = \prod_{j=i+1}^{n}\left(\frac{v_j - h_j}{v_j - h_{j-1}}\right)(v_i - h_i)\,.$$

Similarly to the additive bidder, no unit-demand bidder is willing to bid higher than h_n, since, even though this strategy will result in outbidding all other bidders and thus guaranteeing them their unit, it will result in overpaying. Moreover, bidding 0 would result in losing to the additive bidder α since ties are in favor of the additive bidder. Finally, no bidder would ever bid v_i or above since such a deviation would result in a non-positive expected utility.

To conclude the proof that this construction is a mixed Nash equilibrium, we need to examine whether any unit-demand bidder i has an incentive to bid outside her support without exceeding h_n. For $i = 1, \ldots, n-1$, suppose that the unit-demand bidder i is unilaterally deviating to a point $z \in [h_{m-1}, h_m]$ such that $z < v_i$, and $m \in \{i+1, \ldots, n\}$. But then, since the value vector is such that Equation (3.9) holds, we have that

$$\begin{aligned}\mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}\left[u_i(z,\mathbf{b}_{-i})\right] = F_\alpha^m(z)(v_i - z) &= \prod_{j=m+1}^{n}\left(\frac{v_j - h_j}{v_j - h_{j-1}}\right)\frac{v_m - h_m}{v_m - z}(v_i - z)\\ &\leq \prod_{j=i+1}^{n}\left(\frac{v_j - h_j}{v_j - h_{j-1}}\right)(v_i - h_i) = \mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}\left[u_i(\mathbf{b})\right].\end{aligned}$$

Finally, for $i = 2, \ldots, n$, consider the unilateral deviation of bidder i to an interval $[h_{m-1}, h_m]$ for $m \in \{1, \ldots, i-1\}$. However, due to Equation (3.10) we once again obtain that

$$\mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}\left[u_i(z,\mathbf{b}_{-i})\right] \leq \mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}\left[u_i(\mathbf{b})\right].$$

Hence, no unit-demand bidder has a unilateral deviation and $\mathbf{B}$ is a mixed Nash equilibrium. □

Note that the welfare maximizing allocation is to assign all units to bidder α. Therefore the optimal social welfare is k. To obtain the worst case instance, given a number of units $k \geq 2$, and a number of unit-demand bidders $n < k$, we need to specify a value vector $\mathbf{v}$ so that the expected social welfare is minimized and Equations (3.9) and (3.10) hold. By Lemma 3.4.8, we can easily obtain the expression of the expected social welfare (normalized by k) and, therefore, the optimization problem that yields the most inefficient auction instance attainable with the above construction given a number of units $k \geq 2$ and a number of unit-demand bidders $n < k$ is

$$
\begin{aligned}
\min_{\mathbf{v}\in(\mathbf{0},\mathbf{1})} \quad & 1-(1-\frac{n}{k})\sum_{i=1}^{n}\prod_{j=i+1}^{n}\left(\frac{v_j-h_j(n,k)}{v_j-h_{j-1}(n,k)}\right)\int_{h_{i-1}(n,k)}^{h_i(n,k)}\frac{v_i-h_i(n,k)}{v_i-z}\frac{1-v_i}{(1-z)^2}dz \\
\text{subject to} \quad & h_i(n,k)=\frac{i}{k-n+i}, \qquad \forall i\in\{0,\dots,n\} \\
& v_i > h_i(n,k), \qquad \forall i\in\{0,\dots,n\} \\
& \prod_{j=i+1}^{m-1}\frac{v_j-h_j}{v_j-h_{j-1}} \geq \frac{v_i-z}{v_i-h_i}\frac{v_m-h_{m-1}}{v_m-z}, \qquad \forall i\in\{1,\dots,n-1\}, m\in\{i+1,\dots,n\}, \\
& \qquad\qquad z\in[h_{m-1}(n,k),h_m(n,k)] \\
& \prod_{j=m+1}^{i-1}\frac{v_j-h_j}{v_j-h_{j-1}} \leq \frac{v_m-z}{v_m-h_m}\frac{v_i-h_{i-1}}{v_i-z}, \qquad \forall i\in\{2,\dots,n\}, m\in\{1,\dots,i-1\}, \\
& \qquad\qquad z\in[h_{m-1}(n,k),h_m(n,k)]
\end{aligned}
\tag{3.11}
$$

We were able to numerically solve a series of optimization problems of the above format given integers k, n using global optimization routines of the scientific computing library Scipy of Python. We observed that the worst case instances were those in which the ratio $\frac{n}{k} \approx 37\%$. For instance, when $k = 10$ and $n = 4$, the above optimization problem yields 0.8941 and therefore the worst case ratio is $1/0.8941 \approx 1.1184$, which is already higher than the Price of Anarchy attainable with two bidders that we have discussed in Section 3.4.1. If we increase the number of units further to, say, $k = 100$ and set $n = 37$, the optimization problem yields approximately 0.8925 and therefore the worst case inefficiency becomes $1/0.8925 \approx 1.1204$. By experimenting with very large values of k and n, the worst case inefficiency differs only in the 5th decimal digit, hence we have a convergence to 1.1204, using $\frac{n}{k} \approx 0.3732$. □

The above bound is the best lower bound we have been able to establish, even after some extensive experimentation (driven by the results in the remainder of this section). It is natural to wonder if there is a matching upper bound, which would establish that the Price of Anarchy remains very small even for a large number of bidders. Recall that from De Keijzer et al. (2013), we know already a bound of $e/(e-1) \approx 1.58$. Although we have not managed to settle this question, we will provide an improved upper bound for a special case, for which there is evidence that it captures worst-case scenarios of inefficiency. At the same time, we will be able to characterize the format of such worst case equilibria.

To obtain some intuition, it is instructive to look at the proofs of our two lower bounds, in Theorem 3.4.1 and in Theorem 3.4.3. One can notice that the main source of inefficiency is the fact that the auctioneer accepts multi-unit demand declarations. When this does not occur, we have already shown in Corollary 3.3.1 that mixed Nash equilibria attain optimal welfare. When multi-demand bidders are present, Theorem 3.4.1 shows that in the case of two bidders, the most inefficient mixed Nash equilibrium occurs when a participating bidder declares a demand for all the units, whereas the opponent requires a much smaller fraction of the supply. In the proof of Theorem 3.4.3 above, we have extended this paradigm for multiple bidders with an arbitrary demand structure, but under the assumption that one of the bidders requires all the units (the additive bidder). Such a setting, of one large-demand bidder facing competition by multiple small-demand bidders has also been discussed in Baisa and Burkett (2018). Furthermore, there exist other auction formats that also needed such a demand profile at their worst case instances, see e.g., Birmpas et al. (2019) for the uniform price auction. To summarize, it seems unlikely that the worst instances involve only bidders with low demand or small variation on their demands.

Given the above, we will analyze the family of instances where there exists an additive bidder (with demand equal to k), and where she also has the highest value per unit. In fact, the latter assumption is needed only for the Price of Anarchy analysis but not for the characterization of the worst-case demand profile and the equilibrium strategies. We strongly believe that this class is representative of the most inefficient mixed Nash equilibria (which is true already for the case of two bidders).

The main result of this section is the following.

Theorem 3.4.4. *Consider the class of valuation profiles, where there exists an additive bidder α with the highest value, and an equilibrium $\mathbf{B}$, such that $\alpha \in W(\mathbf{B})$. Then, the Price of Anarchy is at most $4/3$.*

The proof of the theorem is by following a series of steps. The existence of the additive bidder helps in the analysis, because a direct corollary of Theorem 3.3.3 is that the additive bidder is the sole non-empty-handed bidder (everyone else faces competition for all the units).

Corollary 3.4.1 (by Theorem 3.3.3)**.** *Consider a valuation profile $(\mathbf{v}, \mathbf{d})$ with an additive bidder α, that admits an equilibrium $\mathbf{B}$, such that $\alpha \in W(\mathbf{B})$. Then, bidder α is the unique non-empty-handed bidder under $\mathbf{B}$, thus, $\sum_{i \in W(\mathbf{B}) \setminus \{\alpha\}} d_i \leq k-1$.*

To proceed, we ensure that for the instances described by Theorem 3.4.4, it suffices to analyze the equilibria where bidder α belongs to $W(\mathbf{B})$, i.e., there cannot exist a more inefficient equilibrium $\mathbf{B}'$ of these instances with $\alpha \notin W(\mathbf{B}')$. This is addressed by the following lemma.

Lemma 3.4.4. *Consider a valuation profile, and suppose that it admits two distinct inefficient equilibria,* $\mathbf{B}$ *and* $\mathbf{B}'$. *If* $i \in W(\mathbf{B})$ *is a non-empty-handed bidder in* $\mathbf{B}$, *then* $i \in W(\mathbf{B}')$ *.*

Proof. Let $i \in W(\mathbf{B})$ be a non-empty-handed bidder in $\mathbf{B}$ and suppose for contradiction that $i \notin W(\mathbf{B}')$. We know that $\sum_{j \in W(\mathbf{B}) \setminus \{i\}} d_j \leq k - 1$. Since $\mathbf{B}'$ is inefficient, and i does not belong to $W(\mathbf{B}')$, by Lemma 3.3.6, there must exist a bidder m such that $m \in W(\mathbf{B}') \setminus W(\mathbf{B})$.

We can now look more closely on the bidding behavior of bidders i and m in $\mathbf{B}'$. Since $i \notin W(\mathbf{B}')$, by Lemma 3.3.7 we know that i ranks lower than all other winning bidders with probability one. From this, we claim that $Pr_{b_m \sim B'_m}[b_m \geq v_i] > 0$. Indeed if this was not the case, then $Pr_{b_m \sim B'_m}[b_m < v_i] = 1$, and bidder i would have an incentive to outbid bidder m by bidding a value lower than v_i and obtain positive utility, which is a contradiction. This implies that $h(B'_m) \geq v_i$. By Observation 3.3.2 on the maximum bid submitted by the players of $W(\mathbf{B}')$, this yields that $v_m > v_i$.

To obtain a contradiction, we come back to the equilibrium $\mathbf{B}$. Again by Observation 3.3.2, $h(B_i) < v_i$, and therefore, $Pr_{b_i \sim B_i}[b_i < v_m] = 1$. But this implies that bidder m has an incentive to outbid bidder i and obtain a positive utility, which completes the proof. □

Using Lemma 3.4.4 and Corollary 3.4.1, from now on, we fix a bidder α and an inefficient equilibrium $\mathbf{B}$, so that α is additive and $\alpha \in W(\mathbf{B})$.

Corollary 3.4.1 already gives us an insight about the competition in such an equilibrium $\mathbf{B}$. While bidder α will have to compete against the other bidders of $W(\mathbf{B})$ to win extra units, in addition to those that she is guaranteed to obtain, each bidder in $W(\mathbf{B}) \setminus \{\alpha\}$ only competes against α. Each of them is not guaranteed any units, unless she outbids α (bidder α is the only cause of externality for bidders in $W(\mathbf{B}) \setminus \{\alpha\}$, and anyone bidding lower than α cannot get any units). If bidder α did not exist, the other winners could be automatically granted the demand they are requesting since, in total, it is smaller than k and hence, there is no competition among them.

Observation 3.4.1. $\hat{F}_i^{avg}(z) = F_\alpha(z)$, *for every* $i \in W(\mathbf{B}) \setminus \{\alpha\}$, *where* F_α *is the CDF of bidder* α.

We continue with further properties on the support of the mixed strategies.

Lemma 3.4.5. *For the equilibrium* $\mathbf{B}$ *under consideration, it is true that:*

1. $Supp(B_\alpha) = [\ell(\mathbf{B}_W), h(\mathbf{B}_W)]$.
2. *For any two bidders* $i, j \in W(\mathbf{B}) \setminus \{\alpha\}$ *such that* $v_i \neq v_j$*, the set* $Supp(B_i) \cap Supp(B_j)$ *is of measure* 0 *(intersection points can occur only at endpoints of intervals).*

Proof. For the first statement, note that it cannot be the case that the set difference between $[\ell(\mathbf{B}_W), h(\mathbf{B}_W)]$ and $Supp(B_\alpha)$ is a collection of isolated points, since the distributions utilize closed intervals. Suppose now that there exists an interval $I \notin Supp(B_\alpha)$, with $I \subseteq [\ell(\mathbf{B}_W), h(\mathbf{B}_W)]$. We can choose I to be sufficiently small, so that there exists a bidder $i \in W(\mathbf{B}) \setminus \{\alpha\}$ such that $I \subseteq Supp(B_i)$. This is feasible, since by Corollary 3.3.3 the union of the supports of bidders in $W(\mathbf{B}) \setminus \{\alpha\}$ is an interval. Assuming $I = [\ell', h']$, where we can also enforce that $\ell' > \ell(\mathbf{B}_W)$, we obtain

$$\begin{aligned}\mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})] &= \mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(h', \mathbf{b}_{-i})] = (v_i - h')d_i\hat{F}_i^{avg}(h') = (v_i - h')d_i\hat{F}_i^{avg}(\ell') \\ &< \mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(\ell', \mathbf{b}_{-i})] = \mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})].\end{aligned}$$

The first and last equalities above are due to Corollary 3.3.2 (since $\ell' > \ell(\mathbf{B}_W)$), the second equality holds by Fact 3.3.1, and the third equality follows by Observation 3.4.1 and the fact that $F_\alpha(h') = F_\alpha(\ell')$ (because $I \notin Supp(B_\alpha)$) and therefore $\hat{F}_i^{avg}(h') = \hat{F}_i^{avg}(\ell')$. The contradiction we get establishes the first statement of the theorem.

To prove the second statement, suppose for contradiction that there exist two bidders $i, j \in W(\mathbf{B}) \setminus \{\alpha\}$ such that $v_i \neq v_j$ and an interval $I \subseteq Supp(B_i) \cap Supp(B_j)$. By Theorem 3.3.4, we obtain $\hat{F}_i^{avg}(z) = \frac{u_i}{di(v_i - z)}$ and since again by Observation 3.4.1, $\hat{F}_i^{avg}(z) = F_\alpha(z)$, we conclude that $F_\alpha(z) = \frac{u_i}{di(v_i - z)}$ for $z \in I$.

Now for bidder j, and every $z \in I$ we obtain

$$\mathop{\mathbb{E}}_{\mathbf{b}_{-j}\sim\mathbf{B}_{-j}}[u_j(z, \mathbf{b}_{-i})] = d_j F_\alpha(z)(v_j - z) = d_j \frac{u_i}{di(v_i - z)}(v_j - z).$$

The right hand side must be the same for all $z \in I$ by Corollary 3.3.2. However, this is a contradiction since this can only be true for an infinite set of values for z, only when $v_i = v_j$. □

Lemma 3.4.5 suggests that we can group the bidders according to their values (since only bidders with the same value can overlap in their support). Let $r \leq |W(\mathbf{B}) \setminus \{\alpha\}|$

represent the number of distinct values $v_1, \ldots, v_r$, that bidders in $W(\mathbf{B}) \setminus \{\alpha\}$ have. We can partition the bidders of $W(\mathbf{B}) \setminus \{\alpha\}$ into r groups $W_1(\mathbf{B}), \ldots, W_r(\mathbf{B})$, such that, for $j = 1, \ldots, r$, the bidders in group $W_j(\mathbf{B})$ have value v_j. Similarly, we split the support of the winning bidders $[\ell(\mathbf{B}_W), h(\mathbf{B}_W)]$ into r intervals, i.e., $[\ell(\mathbf{B}_W), h(\mathbf{B}_W)] = \bigcup_{j=1}^{r} I_j(\mathbf{B})$, where each interval $j \in \{1, \ldots, r\}$ is formed as $I_j(\mathbf{B}) = \bigcup_{i \in W_j(\mathbf{B})} Supp(B_i)$. The following is a direct corollary of Lemma 3.4.5.

Corollary 3.4.2. *For every $s, t \in \{1, \ldots, r\}$ with $s \neq t$, the set $I_s(\mathbf{B}) \cap I_t(\mathbf{B})$ is of measure* 0*.*

When all bidders in $W(\mathbf{B}) \setminus \{\alpha\}$ have distinct values there are precisely $|W(\mathbf{B}) \setminus \{\alpha\}|$ intervals, whereas when they all have a common value, they must be bidding on the entire interval $[\ell(W(\mathbf{B})), h(W(\mathbf{B}))]$ (the equilibrium in the 2-bidder case when $d_1 = k$, in Section 3.4.1, is one such example). We sometimes denote as $I_0(\mathbf{B})$ the interval of losing bidders $[0, \ell(\mathbf{B}_W)]$, i.e., for the bidders in $\mathcal{N} \setminus W(\mathbf{B})$. Note that given $\mathbf{B}$, the only criterion for the membership of the support of a bidder i in an interval $I_s(\mathbf{B})$ is their value.

The next step is quite crucial in simplifying the extraction of our upper bound. We show that the worst case demand structure for the bidders in $W(\mathbf{B}) \setminus \{\alpha\}$ is when they all have unit demand.

Theorem 3.4.5. *For the value profile $(\mathbf{v}, \mathbf{d})$ and the equilibrium $\mathbf{B}$ under consideration, there exists another value profile $(\mathbf{v}', \mathbf{d}')$ and a product distribution $\mathbf{B}'$ such that*

1. *$\alpha \in W(\mathbf{B}')$ is an additive bidder and for every bidder $i \in W(\mathbf{B}') \setminus \{\alpha\}$, it holds that $d'_i = 1$.*
2. *$\mathbf{B}'$ is a mixed Nash equilibrium for $(\mathbf{v}', \mathbf{d}')$.*
3. $\frac{OPT(\mathbf{v},\mathbf{d})}{SW(\mathbf{B})} = \frac{OPT(\mathbf{v}',\mathbf{d}')}{SW(\mathbf{B}')}$.

Proof. We first construct the value profile $(\mathbf{v}', \mathbf{d}')$ and the product distribution $\mathbf{B}'$. We then argue that they follow the three properties in the statement of the theorem. Firstly, we construct the valuation profile $(\mathbf{v}', \mathbf{d}')$ by modifying the profile $(\mathbf{v}, \mathbf{d})$ as follows: we *replace* every bidder $i \in W(\mathbf{B}) \setminus \{\alpha\}$ with d_i unit-demand bidders, each of them having value v_i. All other bidders (the additive bidder and the losing bidders) retain their values and demands.

We construct the product distribution $\mathbf{B}'$ from $\mathbf{B}$ as follows. We let bidders α and $\mathcal{N} \setminus W(\mathbf{B})$ to bid as before. This leaves us only the newly generated unit-demand bidders. These bidders use the same CDF, and on the same support, as the bidder they were derived from. This completes the description of $\mathbf{B}'$.

Now that we have defined $\mathbf{B}'$ the first property follows easily by observing that the bidders with positive expected utility are precisely all the newly generated unit-demand bidders as well as bidder α.

To see that $\mathbf{B}'$ is an equilibrium, note that the losing bidders have no incentives to deviate, just as in $\mathbf{B}$. Since the CDF of bidder α has not changed, the unit-demand bidders have no incentive to deviate because they face the same competition from α as the bidders in $\mathbf{B}$. If there was a successful deviation by a unit-demand bidder, this would directly translate to a deviation in $\mathbf{B}$. The same is true for the additive bidder since she also sees the same competition on average, and this does not affect the improvement of her expected utility .

Finally, it is very easy to see that $SW(\mathbf{B}) = SW(\mathbf{B}')$, and also that $OPT(\mathbf{v}, \mathbf{d}) = OPT(\mathbf{v}', \mathbf{d}')$, which establishes the last statement.

□

For the remainder of the section, it suffices to analyze valuation profiles, that possess equilibria where the members of $W(\mathbf{B})$ are either additive or unit-demand. Recall, that due to Corollary 3.4.1, there must be a unique additive bidder. Hence, we fix an instance given by a valuation profile $(\mathbf{v}, \mathbf{d})$, so that at the equilibrium $\mathbf{B}$, the set $W(\mathbf{B})$ consists of n unit-demand bidders plus the additive bidder α, i.e., $n = |W(\mathbf{B}) \setminus \{\alpha\}|$. Moreover, due to the following observation we may assume, without loss of generality, that the support of each unit-demand bidder has no overlapping intervals with other bidders from $W(\mathbf{B}) \setminus \{\alpha\}$.

Lemma 3.4.6. *Let* $(\mathbf{v}, \mathbf{d})$ *be a value profile, and let* $\mathbf{B}$ *be any mixed Nash equilibrium, such that the members of* $W(\mathbf{B})$ *are all unit-demand bidders aside from one additive bidder. Then, there exists a mixed Nash equilibrium* $\mathbf{B}'$ *with disjoint support intervals such that* $SW(\mathbf{B}) = SW(\mathbf{B}')$.

Proof. If $\mathbf{B}$ is such that no support intervals intersect for unit-demand bidders, we are done. Otherwise, there exists an interval in which more than one unit-demand bidders bid. Let $S \subseteq W_{\mathbf{B}} \setminus \{\alpha\}$ be such a set of unit-demand bidders and let $I_j = \bigcup_{i \in S} Supp(B_i)$.

Consider the perspective of bidder α when she bids at $I_j \in [L, R]$. Her average CDF of winning bids when bidding in I_j is by Theorem 3.3.4

$$\hat{F}^{avg}_{\alpha}(z) = \frac{D_j(v_\alpha - R)}{k(v_\alpha - z)} .$$

Suppose that we partition the interval $[L, R]$ to $D_j - D_{j-1}$ disjoint subintervals and assign each bidder to one of them. For $i = 1, \ldots, D_j - D_{j-1}$, the CDF of bidder i will be such that the equation above remains satisfied. This means that

$$F_i(z) = \frac{D_j(v_a - R)}{v_\alpha - z} - (k - D_{j-1} - (i-1)),$$

and for every two p_i, p_{i+1} it must be that

$$(k - n + D_{j-1} + i)(v_\alpha - p_i) = (k - n + D_{j-1} + i + 1)(v_\alpha - p_{i+1})$$

These points are clearly inside the interval $[L, R]$.

We partition the interval I_j into $|S|$ intervals. Then, we assign a unit demand bidder in S to bid in a different subinterval with the CDF function H_s. Call this new product distribution B'. Similarly to Theorem 3.4.5, it is clear that the incentives of the small bidders remain unchanged since bidder α did not change their distribution. The same is true for bidder α due to the transformation we defined above. This concludes the proof. □

Therefore, by Corollary 3.4.2 and the discussion preceding it, the support of each bidder $i = 1, \ldots, n$ is $[\ell(B_i), h(B_i)]$. Note that due to Lemma 3.4.5, the unit-demand bidders must cover the entire interval $[\ell(\mathbf{B}_W), h(\mathbf{B}_W)]$. Hence, for a unit-demand bidder $i = 1, \ldots, n$, it must be that $\ell(B_i) = h(B_{i-1})$, assuming for convenience that $h(B_0) = \ell(\mathbf{B}_W)$.

The next theorem provides a more complete understanding of the support intervals and the distributions of the equilibrium $\mathbf{B}$.

Theorem 3.4.6. *For the value profile* $(\mathbf{v}, \mathbf{d})$ *under consideration, the following properties hold:*

1. *For bidder α, we have $h(B_\alpha) = h(B_n) = h(\mathbf{B}_W) = v_\alpha - (k-n)\frac{v_\alpha - \ell(B_\alpha)}{k}$.*

   Moreover, for every unit-demand bidder $i = 1, \ldots, n-1$ it holds that

   $$\ell(B_{i+1}) = h(B_i) = v_\alpha - \frac{(k-n)(v_\alpha - \ell(B_\alpha))}{k - n + i}.$$

2. *The CDF F_α of bidder α, is a branch function, so that for $i = 1, \ldots, n$, $F_\alpha(z) = F_\alpha^i(z)$ for every $z \in [h(B_{i-1}), h(B_i)]$ with*

   $$F_\alpha^i(z) = \prod_{j=i+1}^{n} \left(\frac{v_j - h(B_j)}{v_j - h(B_{j-1})} \right) \frac{v_i - h(B_i)}{v_i - z}.$$

Proof. For the first part of the theorem, we can easily obtain the expression for $h(B_\alpha)$, for the additive bidder α, since she is the sole non-empty-handed bidder, by applying Lemma 3.3.12. To obtain the expressions for the rightmost points of the unit-demand bidders, we study the utility function of bidder α focusing on the points $h(B_1), h(B_2), \ldots, h(B_n)$. In fact, by Corollary 3.3.2 it must be that the expected utility at all these points is equal. Since these are the rightmost endpoints of the support of the unit-demand bidders (and none of them is a mass point for any of them), bidder α is guaranteed $i+k-n$ units when she bids $h(B_i)$. This means that for $i = 1, \ldots, n-1$,

$$\mathop{\mathbb{E}}_{\mathbf{b}_{-\alpha}\sim\mathbf{B}_{-\alpha}}[u_\alpha(h(B_i),\mathbf{b}_{-\alpha})] = \mathop{\mathbb{E}}_{\mathbf{b}_{-\alpha}\sim\mathbf{B}_{-\alpha}}[u_\alpha(h(B_{i+1}),\mathbf{b}_{-\alpha})] \Leftrightarrow$$
$$(k-n+i)(v_a - h(B_i)) = (k-n+i+1)(v_a - h(B(i+1)))\,.$$

The above yields a recursive relation, where $h(B_i)$ can be obtained as a function of $h(B_{i+1})$. Since $h(B_n) = h(B_a)$ is known, we can use induction and establish the desired equation.

For the second part of the theorem, we use that for $i = 1, \ldots, n$, and $z \in Supp(B_i) = [h(B_{i-1}), h(B_i)]$, we have

$$\mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(z,\mathbf{b}_{-i})] = \mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(h(B_i),\mathbf{b}_{-i})] = F_\alpha(h(B_i))(v_i - h(B_i))\,,$$

where the last equality is by Observation 3.4.1, that $\hat{F}_i^{avg}(z) = F_\alpha(z)$. By using Theorem 3.3.4 for $\hat{F}_i^{avg}(z)$, we have

$$F_\alpha(z) = \frac{F_\alpha(h(B_i))(v_i - h(B_i))}{v_i - z} \quad \forall z \in Supp(B_i)\,. \tag{3.12}$$

To proceed, it is convenient to think of F_α as a branch function, with a different branch corresponding to each support interval of the unit-demand bidders. In particular, we let $F_\alpha^i(z) = F_\alpha(z)$ for $z \in Supp(B_i)$. Moreover, by Lemma 3.3.11 the distribution F_α must have no mass points in its support at any intermediate point $h(B_i)$ for $i = 1, \ldots, n-1$. Therefore, since $h(B_i)$ belongs both to $Supp(B_i)$ and to $Supp(B_{i+1})$, in order to have continuity, it must hold that

$$F_\alpha^i(h(B_i)) = F_\alpha^{i+1}(h(B_i)) \;\; \forall i = 1, \ldots, n-1\,.$$

By combining the last two equalities, Equation (3.12) can be rewritten as

$$F_\alpha^i(z) = \frac{F_\alpha^{i+1}(h(B_i))(v_i - h(B_i))}{v_i - z}\,.$$

Hence, we have expressed F_α^i as dependent on the term $F_\alpha^{i+1}(h(B_i))$. Finally, since we know that $F_\alpha^n(h(B_n)) = 1$, we can work inductively and obtain the closed form of each branch F_{α_i}, which completes the proof. □

Before proving our upper bound, we present two additional lemmas. The first is a straightforward inequality, that is a direct consequence of the definition of a mixed equilibrium, and the second is an expression for the social welfare. Both of these are useful for obtaining our final Price of Anarchy upper bound.

Lemma 3.4.7. *Consider a value profile* $(\mathbf{v}, \mathbf{d})$*, and any inefficient mixed Nash equilibrium* $\mathbf{B}$*, with* $W(\mathbf{B})$ *consisting only of additive or unit-demand bidders. Then, for* $i = 2, \ldots, n$*,* $m = 1, \ldots, i-1$*, and every* $z \in [h(B_{m-1}), h(B_m)]$*,*

$$\prod_{j=m+1}^{i-1} \frac{v_j - h(B_j)}{v_i - h(B_{j-1})} \leq \frac{v_m - z}{v_m - h(B_m)} \frac{v_i - h(B_{i-1})}{v_i - z}\,. \tag{3.13}$$

Proof. For bidder $i = 1, \ldots, n-1$, and $m = i+1, \ldots, n$ consider a unilateral deviation $z \in [h(B_{m-1}), h(B_m)]$ of bidder i. Then by the definition of a mixed Nash equilibrium, Definition 3.2.1, it holds that

$$\mathop{\mathbb{E}}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(z, \mathbf{b}_{-i})] \leq \mathop{\mathbb{E}}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})] \Leftrightarrow F_{\alpha^m}(z)(v_i - z) \leq F_{\alpha^m}(h(B_i))(v_i - h(B_i))\,.$$

By substituting the appropriate branches of F_α, the first inequality follow. The reasoning is identical for the second inequality, although this time a bidder $i = 2, \ldots, n$ examines a deviation to a lower interval $m = 1, \ldots, i-1$. □

Lemma 3.4.8. *Consider a value profile* $(\mathbf{v}, \mathbf{d})$*, and any inefficient mixed Nash equilibrium* $\mathbf{B}$*, with* $W(\mathbf{B})$ *consisting only of additive or unit-demand bidders. The expected social welfare is*

$$kv_\alpha - (k-n)(v_\alpha - \ell(B_\alpha)) \sum_{i=1}^{n} \prod_{j=i+1}^{n} \left(\frac{v_j - h(B_j)}{v_j - h(B_{j-1})} \right) \int_{h(B_{i-1})}^{h(B_i)} \frac{v_i - h(B_i)}{v_i - z} \frac{v_\alpha - v_i}{(v_a - z)^2} dz\,.$$

Proof. For brevity we denote $\ell(B_a)$ as ℓ and, for $j = 1, \ldots, n$, we denote $h(B_j)$ as h_j.

We have that

$$
\begin{aligned}
SW(\mathbf{B}) &= \mathop{\mathbb{E}}_{b_\alpha \sim B_\alpha}\left[\mathop{\mathbb{E}}_{\mathbf{b}_{-\alpha}\sim \mathbf{B}_{-\alpha}}\left[\sum_{i=1}^{n} x_i(\mathbf{b})v_i\right]\right]\\
&= F_\alpha(\ell)(k-n+\sum_{j=1}^{n} v_j) + \sum_{i=1}^{n}\int_{h_{i-1}}^{h_i} f_{\alpha^i(z)}\left(F_i(z)(v_\alpha - v_i) + v_\alpha(k-n+i-1) + \sum_{j=i}^{n} v_j\right)dz\\
&= F_\alpha(\ell)(k-n+\sum_{j=1}^{n} v_j) + \sum_{i=1}^{n}\int_{h_{i-1}}^{h_i} f_{\alpha^i(z)}\left(F_i(z)(v_\alpha - v_i) + v_\alpha(k-n+i-1) + \sum_{j=i}^{n} v_j\right)dz\\
&= F_\alpha(\ell)(k-n+\sum_{j=1}^{n} v_j) + \sum_{i=1}^{n}\int_{h_{i-1}}^{h_i} f_{\alpha^i(z)}\left(\frac{k(v_\alpha - h_n)(v_\alpha - v_i)}{v_a - z} + v_i(k-n+i) + \sum_{j=i+1}^{n} v_j\right)dz
\end{aligned}
$$

By integrating the integral by parts, rearranging terms and substituting F_{α^i} by its definition the lemma follows. □

Proof of Theorem 3.4.4. For brevity, we denote $\ell(B_a)$ as ℓ and for $j = 1, \ldots, n$, we denote $h(B_j)$ as h_j. Moreover, by assumption $v_a \geq v_n$. To simplify the calculations, we assume that $v_a = 1$ by rescaling all values in the instance.

Given a mixed Nash equilibrium $\mathbf{B}$, we lower bound the expected social welfare $SW(\mathbf{B})$ described in the equation of Lemma 3.4.8 as

$$
\begin{aligned}
SW(\mathbf{B}) &= k - (k-n)(1-\ell)\sum_{i=1}^{n}\prod_{j=i+1}^{n}\left(\frac{v_j - h_j}{v_j - h_{j-1}}\right)\int_{h_{i-1}}^{h_i}\frac{v_i - h_i}{v_i - z}\frac{1-v_i}{(1-z)^2}dz\\
&= k - (k-n)(1-\ell)\sum_{i=1}^{n}\prod_{j=i+1}^{n}\left(\frac{v_j - h_j}{v_j - h_{j-1}}\right)\\
&\qquad\left(\int_{h_{i-1}}^{h_i}\frac{v_i - h_i}{v_i - z}\frac{1}{(1-z)}dz - \int_{h_{i-1}}^{h_i}\frac{v_i - h_i}{(1-z)^2}dz\right)\\
&> k - (k-n)(1-\ell)\sum_{i=1}^{n}\prod_{j=i+1}^{n}\left(\frac{v_j - h_j}{v_j - h_{j-1}}\right)\int_{h_{i-1}}^{h_i}\frac{v_i - h_i}{v_i - z}\frac{1}{(1-z)}dz\\
&\geq k - (k-n)(1-\ell)\int_{\ell}^{h_n}\frac{v_n - h_n}{(v_n - z)(1-z)}dz\\
&\geq k - (k-n)(1-\ell)\int_{\ell}^{h_n}\frac{1-h_n}{(1-z)^2}dz \geq k - (k-n)(1-\ell)\\
&= k - (k-n)(h_n - \ell) = k - (k-n)\left(\frac{n}{k}(1-\ell)\right) \geq k - \frac{(k-n)n}{k} \geq \frac{3}{4}k\,.
\end{aligned}
$$

The first inequality is true since for all bidders $i = 1, \ldots, n$, it holds that $v_i > h_i$ by Observation 3.3.2. The second one is an application of the mixed Nash equilibrium property encoded by Equation (3.13) of Lemma 3.4.7. The next two inequalities occur by observing that the respective functions are increasing in terms of v_n (which, by assumption, we upper bound with $v_n \leq 1$) and ℓ (which we lower bound with $\ell \geq 0$). The last inequality follows by setting $x = \frac{n}{k}$ and minimizing the function $s(x) = 1 - x + x^2$ for $x \in (0, 1)$. The theorem follows by observing that the optimal welfare is k, since the additive bidder has the highest value. $\square$

3.5 A Separation between Mixed and Bayesian Cases

In this section we explore the more general solution concept of Bayes Nash equilibrium. We consider the following incomplete information setting. Let (v_i, d_i) be the type of bidder $i \in \mathcal{N}$. We suppose that the private value v_i of a bidder i is drawn independently from a distribution V_i. The second part of bidder i's type is his demand d_i; for the purposes of this section (we only construct a lower bound instance), we assume d_i to be deterministic private information.

Each bidder i is aware of her own value per unit v_i and the product distribution formed by the V_j's, and decides a strategy $(b_i, q_i) \sim G_i(v_i)$ for each value $v_i \sim V_i$. The bidding strategy is in general a mixed strategy. In the special case that bidder i chooses a single bid $(b_i(v_i), q_i)$ for each drawn value v_i, he submits a pure strategy, where q_i is not necessarily d_i.

Definition 3.5.1. *Given $\mathbf{V} = \times_{i=1}^{n} V_i$ and $\mathbf{d}$, a profile $\mathbf{G}(\mathbf{v})$ is a Bayes Nash equilibrium if for all $i \in \mathcal{N}$, v_i in V_i's domain, $b_i' \geq 0$ and $q_i' \in \mathbb{Z}_+$ it holds that*

$$\mathop{\mathbb{E}}_{\mathbf{v}_{-i} \sim \mathbf{V}_{-i}} \left[\mathop{\mathbb{E}}_{(\mathbf{b},\mathbf{q}) \sim \mathbf{G}(\mathbf{v})} [u_i^{v_i}(\mathbf{b}, \mathbf{q})] \right] \geq \mathop{\mathbb{E}}_{\mathbf{v}_{-i} \sim \mathbf{V}_{-i}} \left[\mathop{\mathbb{E}}_{(\mathbf{b}_{-i},\mathbf{q}_{-i}) \sim \mathbf{G}_{-i}(\mathbf{v}_{-i})} [u_i^{v_i}((b_i', q_i'), (\mathbf{b}_{-i}, \mathbf{q}_{-i}))] \right],$$

where $u_i^{v_i}(\cdot)$ stands for bidder i's utility when his value is v_i.

We can define the Bayesian Price of Anarchy in the same way as before, by comparing against the expected optimal welfare, over the value distributions.

Although in a few other auction formats, the inefficiency does not get worse when one moves to incomplete information games, we exhibit that this is not the case here. We present a lower bound on the Bayesian Price of Anarchy of 1.1204, with two bidders.

For mixed equilibria and two bidders, Theorem 3.4.1 showed that the Price of Anarchy is at most 1.1095. Although this difference is small, it shows that the Bayesian model is more expressive and can thus create more inefficiency. In particular, we stress that the bound obtained here for two bidders is inspired by the same bound of 1.1204 for mixed equilibria in Theorem 3.4.3, where we had to use a large number of bidders.

Theorem 3.5.1. *For $n = 2$, $k \geq 2$, and capped additive valuation profiles, the Price of Anarchy of Bayes Nash equilibria is at least* 1.1204*.*

Proof. Consider an instance of the discriminatory auction for $k \geq 2$ units and $n = 2$ bidders. Let $d \in \{1, \dots, k-1\}$ be an integer to be specified later and let $h = \frac{d}{k} < 1$. The type of bidder 1 is $(1, k)$ (i.e., deterministically additive with a value of 1 for each of the k units). The type of bidder 2 is (v_2, d) where v_2 is drawn from a continuous distribution V_2. Both bidders reveal their demands. Bidder 1 bids a mixed strategy according to the distribution B_1 supported in $[0, h]$. We denote the (continuous) CDF of B_1 as F_1 and will present its formula in the sequel. Moreover, we denote by f_1 its probability density function.

Given a value v_2 drawn from V_2, the second bidder bids according to some bidding function $b_2(v_2)$ that maps the drawn value v_2 to a bid. Therefore, for each value v_2, bidder 2 specifies a pure strategy bid. Nevertheless, due to the randomness of the value distribution V_2, bidder 1 competes against mixed strategies and observes a CDF

$$F_2(x) = \frac{x}{1-x}\frac{1-h}{h}$$

that describes the distribution of the random variable $b_2(v_2) \in [0, h]$ which we denote as $B_2(V_2)$.

Optimality of bidding function $b_2(\cdot)$ Consider the utility maximization problem of bidder 2. For all values v_2 drawn from distribution V_2, bidder 2 must, at a Bayes Nash equilibrium, specify a pure bid $b_2 \in [0, h]$ that maximizes her expected utility, i.e.

$$\max_{b_2 \in [0,h]} \mathop{\mathbb{E}}_{b_1 \sim B_1}[x_2(b_1, b_2)](v_2 - b_2) \Leftrightarrow \max_{b_2 \in [0,h]} dF_1(b_2)(v_2 - b_2),$$

and the equivalence is due to the fact that the sole source of competition for bidder 2 is bidder 1. Working backwards and viewing the value of bidder 2 as a function of her bid (which is the inverse of the bidding function), and taking first order conditions with respect to b_2, we obtain that

$$v_2(b_2) = b_2 + \frac{F_1(b_2)}{f_1(b_2)}, \qquad b_2 \in (0, h]. \tag{3.14}$$

Hence, a utility-maximizing bidding function satisfies Equation (3.14).

We now argue that this auction instance is a Bayes Nash equilibrium. Firstly, neither bidder has an incentive to bid above h, as bidding h already guarantees them their entire demand. Moreover, lying about one's demand is also a weakly dominated strategy for both bidders.

When bidder 1 declares a bid $z \in [0, h]$, her utility is

$$\mathop{\mathbb{E}}_{v_2 \sim V_2}[u_1(z, b_2(v_2))] = (k - d + dF_2(z))(1 - z) = k - d$$

Therefore, since the expected utility of bidder 1 is $k - d$ and is a constant at every subinterval of her support, bidder 1 has no incentive to deviate unilaterally.

In the case of bidder 2, we have chosen her bidding function $b_2(v_2)$ to be one that satisfies Equation (3.14). Since this bidding function maximizes her utility, it holds that, given a type v_2

$$\mathop{\mathbb{E}}_{b_1 \sim B_1}[x_2(b_1, b_2(v_2))](v_2 - b_2(v_2)) \geq \mathop{\mathbb{E}}_{b_1 \sim B_1}[x_2(b_1, b')](v_2 - b')$$

for all $b' \in [0, h]$. Therefore, this instance is a Bayes Nash equilibrium.

The expected social welfare of this BNE is

$$SW := \mathop{\mathbb{E}}_{b_2 \sim B_2(V_2)} \left[\mathop{\mathbb{E}}_{b_1 \sim B_1} [x_1(b_1, b_2) + x_2(b_1, b_2) v_2(b_2)] \right]$$

. We continue as follows:

$$\begin{aligned} SW &= \int_0^h f_2(z) \left(\mathop{\mathbb{E}}_{b_1 \sim B_1} [x_1(b_1, z) + x_2(b_1, z) v_2(z)] \right) dz \\ &= \int_0^h f_2(z) \left(k - \mathop{\mathbb{E}}_{b_1 \sim B_1} [x_2(b_1, z)](1 - v_2(z)) \right) dz \\ &= k \int_0^h f_2(z) \left(1 - hF_1(z) \left(1 - z - \frac{F_1(z)}{f_1(z)} \right) \right) dz \\ &= \frac{k(1-h)}{h} \int_0^h \frac{1}{(1-z)^2} \left(1 - hF_1(z) \left(1 - z - \frac{F_1(z)}{f_1(z)} \right) \right) dz \\ &= k \left(1 + (1-h) \int_0^h \frac{F_1(z)^2}{(1-z)^2 f_1(z)} - \frac{F_1(z)}{1-z} dz \right) \end{aligned} \tag{3.15}$$

$$\tag{3.16}$$

The third equality in the above derivation is due to Equation (3.14). In Equation (3.15) the social welfare of this instance is written in terms of $h = \frac{d}{k}$ and the functions F_1 and f_1.

We have already shown that this instance is in fact a Bayes Nash equilibrium for any continuous function F_1 supported in $[0, h]$. The question remains which function F_1 (and consequently f_1 as its derivative) should we choose.

Let us now focus on the integral in the right hand side Equation (3.15) i.e.,

$$I := \int_0^h \frac{F_1(z)^2}{(1-z)^2 f_1(z)} - \frac{F_1(z)}{1-z} dz$$

Ideally, we would want to pick the continuous and increasing function F_1 that minimizes I as long as $F_1(h) = 1$. This is possible, as we are dealing with a well-posed problem of variational calculus, a field of mathematics with a goal of finding functional maxima and minima. Using such an approach (solving the Euler-Lagrange equation of the problem), we were able to determine that the function

$$F_1(x) = \frac{\mathcal{W}\left(-e^{-1}(1-h)^2\right)}{\mathcal{W}(-e^{-1}(1-x)^2)}$$

is such a functional minimum for Equation (3.15). Here, $\mathcal{W}(x)$ for $x \geq -\frac{1}{e}$ is the principal branch (commonly denoted by $\mathcal{W}_0$) of the *Lambert* function, the multi-valued inverse of $we^w = x$, see the work of Corless et al. (1996) for a complete reference. The derivative of $\mathcal{W}(x)$ is $\frac{d\mathcal{W}}{dx} = \frac{\mathcal{W}(x)}{x(1+\mathcal{W}(x))}$. By applying this rule, we can obtain that

$$f_1(z) = \frac{dF_1}{dz} = \frac{2F_1(z)}{(1-z)\left(1+\mathcal{W}(-e^{-1}(1-z)^2)\right)}. \tag{3.17}$$

We observe that $F_1(z)$ is indeed a valid CDF.

Fact 3.5.1. *For $x \in [0,h]$, it holds that:*

1. $F_1(0) \geq 0$
2. *$F_1(x)$ is increasing in $[0,h]$.*
3. $F_1(h) = 1$.

Therefore, by replacing F_1 and f_1 into Equation (3.15) we obtain

$$\begin{aligned} \frac{SW}{k} &= 1 + (1-h)\int_0^h \frac{F_1(z)^2}{(1-z)^2 f_1(z)} - \frac{F_1(z)}{1-z} dz \\ &= 1 + \frac{(1-h)\mathcal{W}(-e^{-1}(1-h)^2)}{2}\int_0^h \frac{\mathcal{W}(-e^{-1}(1-z)^2)-1}{\mathcal{W}(-e^{-1}(1-z)^2)(1-z)} dz. \end{aligned}$$

The second equality follows by Equation 3.17. What remains now is to find the global minimum of the right hand side of the above with respect to $h \in (0,1)$. Indeed, we can numerically verify that for $h \approx 0.36$, we have that $\frac{SW}{k} \approx 0.8925$.

Finally we observe that for all $z \in [0,h]$ it holds that

$$v(z) = z + \frac{F(z)}{f(z)} = z + (1-z)(1+\mathcal{W}(-e^{-1}(1-z)^2)) \leq 1$$

and therefore, the optimal social welfare is exactly k, since the optimal allocation is to always give k units to bidder 1. Thus, for this instance, the ratio of the optimal social welfare over the expected social welfare for this equilibrium is $\frac{OPT}{SW} \approx \frac{1}{0.8925} \approx 1.1204$ and the proof follows. □

Remark 3.5.1. *When $k=1$, there is a lower bound of 1.156 in Jin and Lu (2022) for the first price auction. However this requires a very large number of bidders. There is a simpler construction with two bidders in Syrgkanis (2014) but it only yields a lower bound of 1.06.*

Chapter 4

Interval Cover

4.1 Introduction

For the remainder of the dissertation, we shift our attention to procurement auctions. In this chapter[1], we consider a mechanism design problem, that emerges under certain spatial models for crowdsourcing and labour matching markets. It is instructive to start with an example, so as to introduce the main aspects of the model. Imagine a set of cities, located geographically in a consecutive order, and consider a company that has opened a store in each of these cities. In Figure 4.1, we see an instance with 5 cities, named A to E. The company needs to meet a demand constraint, i.e., a lower bound on the volume of goods that need to be transported to each store, based on consumption and future planning. To achieve this goal, it chooses to hire other firms, or single individuals, that can make deliveries, via a reverse (procurement) auction (for an exposition on auctions for transportation routes, see Cramton et al. (2006)). Suppose that every participating entity, referred to as a bidder or a worker, can only cover a certain interval of contiguous cities, at a cost that she specifies, and furthermore, she can only accommodate a certain volume of goods, assumed to be the same for each city in her declared interval (among others, dependent on the transportation means that she owns). The problem boils down to selecting a set of winning bidders who can jointly cover all the demand constraints at minimum cost, and in a way that prevents the workers from misreporting their true preferences.

In the crowdsourcing jargon, we can view the store in each city as a task with a demand requirement. For instance, in labeling/classification tasks, the demand could correspond to the number of people who should execute the task in order to acquire

[1]A conference version with the results of this chapter appeared in SAGT '22 (Markakis et al., 2022a).

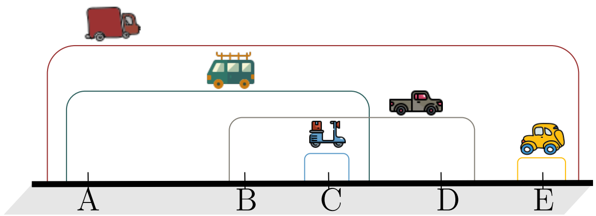

Figure 4.1 An example with 5 tasks (A to E) and 5 workers.

a higher confidence on the outcome. In other cases, it could be interpreted as a (not necessarily integral) volume coverage requirement.

vehicle	contribution	cost
truck	8	10
van	6	5
pickup	5	4
motorbike	1	2
car	3	4

Table 4.1 Contribution and costs of the workers of Figure 4.1.

Coming back to the example of Figure 4.1, say that the demand of cities A to E are given by the vector $(6, 2, 9, 1, 3)$. The contribution per task, as well as the cost of the workers are given by Table 4.1. The interval covered by each worker is visible in Figure 4.1. Obviously, hiring all the workers is a feasible solution. Selecting the workers with the van, the pickup and the car also forms a feasible solution since the demands of all cities are indeed satisfied. Notice that the optimal solution is to hire the workers with the truck and the motorbike at a total cost of 12.

From an algorithmic viewpoint, there has been already significant progress on the problem. As it can be easily seen to be NP-hard, the main results on this front include constant factor approximation algorithms, with further improvements for special cases. However, in the context of mechanism design, one needs to consider strategic aspects as well. Bidders may attempt to report higher costs in order to achieve higher payments, or they can lie about the subset of tasks they can actually fulfil. It would be therefore desirable to have mechanisms that, on the one hand, achieve competitive approximation guarantees, and on the other, deter bidders from misreporting. To our knowledge, the currently best result on this direction is by Dayama et al. (2015), where however the approximation ratio of their truthful mechanism is unbounded in terms of the number of tasks and workers (and depends on numerical parameters of each instance). It has remained open since then, if truthfulness can be compatible with bounded approximation guarantees.

Contribution

We focus on truthful mechanisms and their approximation guarantees against the optimal cost. To this end, we provide two main results. The first one, in Section 4.3, is a truthful Δ-approximation mechanism, where Δ is the maximum number of workers that are willing to work on the same task (the maximum being taken over all tasks). This mechanism, improves significantly the state of the art, coming to a guarantee that is polynomially bounded in terms of the input size. Apart from the improvement, we note also that our result is based on the local-ratio technique from approximation algorithms (see Bar-Yehuda et al. (2005) for an extensive reference), a technique that has not been used very often for building truthful mechanisms (for an exception, see Elkind et al. (2007)). Moving on, our second result, in Section 4.4, concerns the class of instances with a constant number of tasks, which generalizes MIN-KNAPSACK, an NP-hard variant of KNAPSACK that corresponds to the single-task case of our problem. For this class we provide a truthful FPTAS, by mainly exploiting and adapting the framework of Briest et al. (2011). In doing so, we also identify a flaw in a previous attempt for designing a truthful FPTAS for MIN-KNAPSACK by Chen et al. (2019). Finally, in Section 4.5, we discuss some further implications and extensions.

Related Work

As already discussed, the work most related to ours is Dayama et al. (2015), which introduced the model in the context of crowdsourcing. They provide a truthful optimal mechanism when the workers have an identical contribution, whereas for the general case, they present a truthful approximate mechanism, the ratio of which is dependent on the contribution parameters of the workers. Furthermore, Xu et al. (2015) study the case of unit-demand tasks and unit-contribution workers and identifies a truthful optimal mechanism. For the same setting, they also propose a mechanism when workers can submit multiple bids, for which they attain a logarithmic factor. Additionally, in Zhang et al. (2015) the authors studied the prominent special case of a single task (MIN-KNAPSACK) and provided a randomized, truthful-in-expectation mechanism that achieves an approximation factor of 2.

Regarding the purely algorithmic problem, without the constraint of truthfulness, it has appeared under the name of (0-1)RESOURCE ALLOCATION, and a 4-approximation was presented in Chakaravarthy et al. (2011). The currently best known algorithm achieves a factor of 3, in Mondal (2018). For the MIN-KNAPSACK problem, a PTAS is implied by Csirik (1991) and a FPTAS is given in Kothari et al. (2005).

There are quite a few problems that can be viewed as generalizations of what we study here, such as *general scheduling problem* (Bansal and Pruhs, 2014), *multidimensional min-knapsack* (Pritchard, 2010), *column restricted covering integer programs* (Chakrabarty et al., 2010). Moreover, several problems in discrete optimization can be seen as related variants, but are neither extensions nor special cases of ours. Indicatively: *bandwidth allocation* (Chen et al., 2002), *multiset multicover*(Bredereck et al., 2020; Rajagopalan and Vazirani, 1993), *geometric knapsack* (Gálvez et al., 2021), *capacitated network design* (Carr et al., 2000), *admission control* (Phillips et al., 2000).

Finally, for general spatio-temporal models appearing in the crowdsourcing literature, we refer to two recent surveys by Gummidi et al. (2019) and Tong et al. (2020), which cover to a big extent the relevant results.

4.2 Preliminaries

In this section, we first define formally the problem that we study, together with some additional necessary notation. In the sequel, we discuss the relevant definitions for the design of truthful mechanisms.

Problem Statement

We are interested in the optimization problem defined below. For motivating applications, we refer the reader to Section 1.2 in Dayama et al. (2015).

Cost Minimization Interval Cover (CMIC): Consider a set of tasks, say $\{1, \ldots, m\}$, that are ordered in a line, and a set of available bidders $\mathcal{N} = \{1, \ldots, n\}$. We will interchangeably use the term bidder or worker in the sequel. An instance of CMIC is determined by a tuple $(\mathbf{b}, \mathbf{q}, \mathbf{d})$, where:

- The vector $\mathbf{b} = (b_1, \ldots, b_n)$ is the bidding profile. For each bidder $j \in \mathcal{N}$, $b_j = (c_j, [s_j, f_j])$, where $[s_j, f_j] = I_j$ is the interval of contiguous tasks $\{s_j, s_j + 1, \ldots, f_j\} \subseteq [m]$, that j is able to contribute to, and $c_j \in \mathbb{R}_{\geq 0}$ is the cost incurred to her, if she is selected to contribute. We can assume positive costs since workers of zero cost are trivially included in the solution. We often denote $\mathbf{b}$ as $(\mathbf{c}, \mathbf{I})$ where $\mathbf{c}$ is the cost vector and $\mathbf{I}$ is the vector of the intervals.

- The vector $\mathbf{q} = (q_1, \ldots, q_n)$ is the contribution vector, so that $q_j \in \mathbb{R}_{>0}$ denotes the contribution that bidder j can make to the tasks that belong to $[s_j, f_j]$.

- The vector $\mathbf{d} = (d_1, \ldots, d_m)$ specifies the demand $d_j \in \mathbb{R}_{>0}$ of a task j.

The goal is to select a set of bidders $S \subseteq \mathcal{N}$ of minimum cost, that satisfies

$$\sum_{i \in \mathcal{N}_j(\mathbf{I}) \cap S} q_i \geq d_j, \ \forall \ j \in [m], \tag{4.1}$$

where for $j \in [m]$, $\mathcal{N}_j(\mathbf{I}) := \{i \in \mathcal{N} \mid j \in [s_i, f_i]\}$, i.e. is the set of bidders who can contribute to task j.

The problem belongs to the broad family of problems described by covering integer programs. It may also seem reminiscent of multicover variants of the SET COVER problem on an interval universe. We note however that in CMIC, each worker is allowed to be picked at most once, and moreover, the coverage requirement is not necessarily integral, which is a substantial difference.

Given a feasible solution $S \subseteq \mathcal{N}$, we refer to its total cost, $\sum_{i \in S} c_i$, as the derived *social cost*. Throughout this chapter, we focus on deterministic allocation algorithms that use a consistent deterministic tie-breaking rule. Given an instance $P = (\mathbf{b}, \mathbf{q}, \mathbf{d})$, and an allocation algorithm A, we denote by $\mathcal{W}_A(\mathbf{b})$ the set of bidders selected by A, when $\mathbf{q}, \mathbf{d}$ are clear from the context[2]. In the same spirit, we let also $C(A, \mathbf{b}) := \sum_{i \in \mathcal{W}_A(\mathbf{b})} c_i$ be the *social cost* derived by algorithm A on input P. Finally, we use $OPT(\mathbf{b})$ to denote the cost of an optimal solution.

Truthful Mechanisms

We move now to the strategic scenario, where bids can be private information. A mechanism for a reverse auction, like the ones we study here, is a tuple $M = (A, \boldsymbol{\pi})$, consisting of an *allocation algorithm* A and a *payment rule* $\boldsymbol{\pi}$. Initially, each bidder $j \in \mathcal{N}$ is asked to submit a bid $b_j := (c_j, [s_j, f_j])$, which may differ from her actual cost and interval. Then, given a bidding profile $\mathbf{b} = (\mathbf{c}, \mathbf{I})$, the allocation algorithm $A(\mathbf{b})$ selects the set of winning bidders, i.e., this is a binary setting where the allocation decision for each bidder is whether she is included in the solution or not. Finally, the mechanism computes a vector of payments $\boldsymbol{\pi}(\mathbf{b})$, so that $\pi_i(\mathbf{b})$ is the amount to be paid to bidder i by the auctioneer. Naturally, we will consider that non-winning bidders do not receive any payment.

Note that we consider the contribution q_j, for $j \in \mathcal{N}$, to be public information, available to the auctioneer. The reason is that such a parameter could be estimated

[2]It is convenient to highlight the dependence on $\mathbf{b}$, especially when arguing about truthful mechanisms in the remaining sections.

by past statistical information on the performance or capacity of the worker. As an example, we refer to Dayama et al. (2015), where for labeling tasks, it is explained how q_j can be computed as a function of a worker's quality (i.e., the probability for a worker to label correctly), that can be available in a crowdsourcing platform with rating scores or reviews for the workers.

Our setting corresponds to what is usually referred to as a pseudo-2-parameter environment (or almost-single-parameter as referred to by Blumrosen and Nisan (2007)), since each bidder has two types of private information, the monetary cost and the interval. And in particular, our model can be seen as a special case of *single-minded* bidders, who are interested in a single subset each, but for reverse auctions. As in the paper of Dayama et al. (2015), when the true type of a worker i is $(c_i, [s_i, t_i])$, and she declares her true interval or any non-empty subset of it, then she enjoys a utility of $\pi_i - c_i$, if she is selected as a winner by the mechanism (with π_i being her payment). If she declares any other interval that contains any task $j \notin [s_i, t_i]$, and she is selected as a winner, then the bidder has a utility of $-\infty$, or equivalently has an infinite cost. This simply models the fact that the worker may not be capable of executing or does not desire to be assigned to any task outside her true interval (and therefore would have no incentive for such deviations).

The previous discussion allows us to exploit the sufficient conditions proposed by Lehmann et al. (2002) (for forward auctions), to obtain truthful mechanisms, as an extension of the seminal result by Myerson (1981), which has been stated in Chapter 2 as Lemma 2.2.1. For reverse auctions, the same framework is also applicable, implying that as long as an allocation algorithm is exact (each selected bidder is assigned her declared interval), the crucial property we need to enforce is *monotonicity*. Monotonicity of an algorithm means that if a winning bidder declares a more competitive bid, she should still remain a winner. To be more precise, we define first the following partial order on possible bids.

Definition 4.2.1. *Let $b_i = (c_i, [s_i, f_i])$ and $b'_i = (c'_i, [s'_i, f'_i])$ be two bids of bidder $i \in \mathcal{N}$. We say that $b'_i \succeq b_i$, if $c_i \geq c'_i$ and $[s_i, f_i] \subseteq [s'_i, f'_i]$.*

Definition 4.2.2. *An allocation algorithm A is monotone if for every bidding profile $\mathbf{b}$, for any bidder $i \in \mathcal{W}_A(\mathbf{b})$ and any bid $b'_i \succeq b_i$, it holds that $i \in \mathcal{W}_A(b'_i, \mathbf{b}_{-i})$.*

Theorem 4.2.1 (cf. Lehmann et al. (2002), Theorem 9.6)**.** *In settings with single-minded bidders, given a monotone and exact algorithm A, there exists an efficiently computable payment rule $\boldsymbol{\pi}$, such that $\mathcal{M} = (A, \boldsymbol{\pi})$ is a truthful mechanism.*

Finally, we stress that all our proposed algorithms are exact by construction, which means that bidders will only be assigned to the set of tasks which they asked for, as in Briest et al. (2011), and hence we only need to care about monotonicity.

4.3 An Improved Truthful Mechanism

As already stated in Section 4.1, the currently best algorithm for CMIC has an approximation ratio of 3 due to Mondal (2018), based on refining the 4-approximation algorithm by Chakaravarthy et al. (2011). However, as the next proposition shows, these algorithms are not monotone[3]. More generally, it has been noted by Dayama et al. (2015) that primal-dual algorithms with a "delete phase" at the end, are typically non-monotone, without, however, providing an example. For the sake of completeness, we provide a concrete example, in Section B.1 of the Appendix B, that proves the following statement:

Proposition 4.3.1. *The current state of the art constant factor approximation algorithms for CMIC by Chakaravarthy et al. (2011); Mondal (2018) are not monotone.*

So far, the only truthful mechanism that has been identified by Dayama et al. (2015), achieves an approximation ratio of $2 \cdot \max_{i \in \mathcal{N}}\{q_i\}$. Note that this approximation ratio is dependent on the contribution parameters of the workers, which can become arbitrarily large, and not bounded by any function of n and m. The main result of this section is the following theorem, established via a greedy, local-ratio algorithm, which reduces the gap between truthful and non-truthful mechanisms. In particular, this gap is related to the maximum number of workers that contribute to any given task, which for an instance $((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$ of CMIC, is $\Delta(\mathbf{I}) := \max_{j=1,\ldots,m} |\mathcal{N}_j(\mathbf{I})|$. We denote it simply by Δ when $\mathbf{I}$ is clear from context. Obviously, Δ is always upper bounded by the number of workers, n.

Theorem 4.3.1. *There exists a truthful, polynomial-time mechanism, that achieves a Δ-approximation for the CMIC problem.*

The rest of the section is devoted to the proof of Theorem 4.3.1. The main component of the proof is an approximation-preserving reduction to a particular job scheduling problem for a single machine (Bar-Noy et al., 2001), defined as follows:

[3] For a similar reason, the 40-approximation for CMIC by Chakrabarty et al. (2010), which uses as a subroutine a primal-dual algorithm involving a "delete phase", is non-monotone as well.

Loss Minimization Interval Scheduling (LMIS): We are given a limited resource whose amount may vary over a time period, which WLOG, is defined by the integral time instants $\{1, \ldots, m\}$. We are also given a set of activities $\mathcal{J} = \{1, \ldots, n\}$, each of which requires the utilization of the resource, for an interval of time instants. An instance of LMIS is determined by a tuple $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$, where:

- The vector $\mathbf{p} = (p_1, \ldots, p_n)$ specifies a penalty $p_j \in \mathbb{R}_{>0}$, for each activity $j \in \mathcal{J}$, reflecting the cost that is incurred by not scheduling the activity.
- For each activity $j \in \mathcal{J}$, we are given an interval[4] $T_j = [s_j, f_j]$, such that $s_j, f_j \in \{1, \ldots, m\}$ are the start and finish times of j respectively. Let $\mathbf{T} = (T_1, \ldots, T_n)$, be the vector of all activity intervals.
- The vector $\mathbf{r}$ contains, for each activity $j \in \mathcal{J}$, the width $r_j \in \mathbb{R}_{>0}$, reflecting how much resource the activity requires, i.e., this means that activity j requires r_j units of resource at every integral time instant of its interval T_j.
- The vector $\mathbf{D} = (D_1, \ldots, D_m)$ specifies the amount of available resource $D_i \in \mathbb{R}_{>0}$, at each integral time instant $i \in [m]$.

Let $\mathcal{J}_i(\mathbf{T}) := \{j \in \mathcal{J} \mid i \in T_j\}$. The goal in LMIS is to select a set of activities $S \subseteq \mathcal{J}$ to schedule, that meet the resource constraint

$$\sum_{j \in S \cap \mathcal{J}_i(\mathbf{T})} r_j \leq D_i, \quad i = 1, \ldots, m, \tag{4.2}$$

and such that $\sum_{j \in \mathcal{J} \setminus S} p_j$ is minimized, i.e., we want to minimize the sum of the penalties for the non-scheduled activities.

Our work highlights an interesting connection between LMIS and CMIC. This can be seen via the reduction provided by algorithm $\hat{A}$ below, where for an instance $((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$, we let $\mathbf{Q}(\mathbf{I}) = (Q_1(\mathbf{I}), \ldots, Q_m(\mathbf{I}))$ and $Q_j(\mathbf{I}) := \sum_{i \in \mathcal{N}_j(\mathbf{I})} q_i, \forall j \in [m]$.

Algorithm 4.1: $\hat{A}(\mathbf{b})$

▷ **Input:** A bidding profile $\mathbf{b} = (\mathbf{c}, \mathbf{I})$ of a CMIC instance $((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$

Construct the LMIS instance $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D}) = (\mathbf{c}, \mathbf{I}, \mathbf{q}, \mathbf{Q}(\mathbf{I}) - \mathbf{d})$, with $\mathcal{J} = \mathcal{N}$.

Run an approximation algorithm for the LMIS instance, and let S be the set of scheduled activities.

return $\mathcal{N} \setminus S$

[4] Originally, the problem was defined using a semi-closed interval for each activity, but it is easy to see that defining it using a closed one instead, is equivalent and more convenient for our purposes.

Theorem 4.3.2. *Algorithm $\hat{A}$ converts any α-approximation algorithm for* LMIS *to an α-approximation algorithm for* CMIC.

Theorem 4.3.2 is based on the lemma below, which shows the connection between the feasible solutions of the two problems.

Lemma 4.3.1. *Consider a* CMIC *instance $P = ((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$. Let also P' be the* LMIS *instance defined by $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D}) = (\mathbf{c}, \mathbf{I}, \mathbf{q}, \mathbf{Q}(\mathbf{I}) - \mathbf{d})$, with $\mathcal{J} = \mathcal{N}$. Then, for every feasible solution S of P, it holds that $\mathcal{J} \setminus S$ is a feasible solution for P' with the same cost, and vice versa.*

Proof. Observe first that for $\ell = 1, \ldots, m$, the set $\mathcal{J}_\ell(\mathbf{I})$ in P' coincides with $\mathcal{N}_\ell(\mathbf{I})$ in P. For the first direction, fix a feasible schedule S for P'. For every $\ell = 1, \ldots, m$, using Equation (4.2) we obtain

$$\sum_{j \in S \cap \mathcal{N}_\ell(\mathbf{I})} q_j \leq Q_\ell(\mathbf{I}) - d_\ell = \sum_{j \in \mathcal{N}_\ell(\mathbf{I})} q_j - d_\ell,$$

where the equality follows from the definition of $\mathbf{Q}$. By rearranging terms, we have that $\sum_{j \in \mathcal{N}_\ell(\mathbf{I}) \setminus S} q_j \geq d_\ell$, which is exactly the feasibility constraint of Equation (4.1), when we take $\mathcal{N} \setminus S$ as the CMIC solution, since $\mathcal{N}_\ell(\mathbf{I}) \setminus S = \mathcal{N}_\ell(\mathbf{I}) \cap (\mathcal{N} \setminus S)$. Note also that the cost of the solution S in P' is $\sum_{i \in \mathcal{N} \setminus S} c_i$, which is the same as the cost of the solution $(\mathcal{N} \setminus S)$ for P.

For the reverse direction, fix a feasible solution W of P. For all $\ell = 1, \ldots, m$, using Equation (4.1), we obtain

$$d_\ell \leq \sum_{j \in W \cap \mathcal{N}_\ell(\mathbf{I})} q_j = \sum_{j \in \mathcal{N}_\ell(\mathbf{I})} q_j - \sum_{j \in \mathcal{N}_\ell(\mathbf{I}) \cap (\mathcal{N} \setminus W)} q_j,$$

and by rearranging terms and using the fact that $\mathcal{N}_\ell(\mathbf{I}) = \mathcal{J}_\ell(\mathbf{I})$, we obtain

$$\sum_{j \in \mathcal{J}_\ell(\mathbf{I}) \cap (\mathcal{N} \setminus W)} q_j \leq Q_\ell(\mathbf{I}) - d_\ell,$$

which completes the proof, since $\mathcal{N} \setminus W$ satisfies the feasibility constraint of Equation (4.2) for P'. Again, the costs of the two solutions coincide. □

Proof of Theorem 4.3.2. Lemma 4.3.1 already suggests how to derive an algorithm for CMIC, using as a black box, an algorithm for LMIS. Suppose we run Algorithm $\hat{A}$ on a CMIC instance $P = ((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$. Let P' be the LMIS instance defined by $(\mathbf{c}, \mathbf{I}, \mathbf{q}, \mathbf{Q}(\mathbf{I}) - \mathbf{d})$ and suppose we run an α-approximation algorithm for P'. Note first of all that by Lemma 4.3.1, the solution returned is feasible for CMIC. We will also prove that Algorithm $\hat{A}$ achieves an α-approximation for P.

To see this, let S be the solution returned by the algorithm for LMIS. Let also $OPT_{LMIS}(P')$ be the cost of an optimal solution in P'. Since, we have used an α-approximation algorithm for LMIS, it returned a solution of cost $\sum_{i \in \mathcal{N} \setminus S} c_i \leq \alpha OPT_{LMIS}(P')$. By the definition of $\hat{A}$, the solution returned for P is $\mathcal{N} \setminus S$, and hence the cost of our solution for CMIC equals $\sum_{i \in \mathcal{N} \setminus S} c_i$. It remains now to observe that as a consequence of Lemma 4.3.1, it follows that $OPT(P) = OPT_{LMIS}(P')$, where $OPT(P)$ is the cost of an optimal solution for P. □

Approximation Guarantee and Monotonicity of $\hat{A}$

If we only cared about the approximation ratio of $\hat{A}$, it would suffice to use as a black box any algorithm for LMIS. And in fact, the best known algorithm for LMIS achieves a 4-approximation, and was obtained by Bar-Noy et al. (2001), using the local-ratio framework. Plugging in this algorithm however, does not ensure that we will end up with a truthful mechanism for CMIC. Instead, we will consider an appropriate modification of the algorithm by Bar-Noy et al. (2001), which enforces monotonicity of $\hat{A}$, but at the price of a higher approximation ratio. This is presented as Algorithm 4.2 below.

We introduce first a notion that will be useful both for the statement of Algorithm 4.2 and for our analysis. Consider an instance $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$ of LMIS. Given a set of jobs $S \subseteq \mathcal{J}$, and a time instant $i = 1, \ldots, m$, we define

$$R_i(S, \mathbf{T}, \mathbf{D}) := \sum_{\ell \in S \cap \mathcal{J}_i(\mathbf{T})} r_\ell - D_i.$$

The quantity $R_i(S, \mathbf{T}, \mathbf{D})$ measures how much (if at all), the resource constraint of Equation (4.2), for the i-th time instant is violated when scheduling all the activities in S. Accordingly, define $R^*(S, \mathbf{T}, \mathbf{D}) := \max_{i=1,\ldots,m} R_i(S, \mathbf{T}, \mathbf{D})$. Note that a schedule S is feasible if and only if $R^*(S, \mathbf{T}, \mathbf{D}) \leq 0$.

Algorithm 4.2 constructs a feasible schedule S as follows: Initially, it checks if the entire set of activities $S = \mathcal{J}$ constitutes a feasible schedule. If not, the algorithm iteratively removes one activity per iteration from S, in a greedy fashion, until S becomes feasible. The algorithm determines the time instant t^* with the most violated feasibility constraint, by computing $R^*(S, \mathbf{T}, \mathbf{D})$, and considers all activities from S whose interval contains t^*. Then, it removes from S one of these activities that minimizes a certain ratio, dependent on the current penalties and resource requirements, while it simultaneously decreases the penalty of all other activities that contain t^*.

Algorithm 4.2 LMIS-LR$(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$

▷ **Input:** An instance $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$ of LMIS

Initialize $S = \mathcal{J}, k = 0,$ and $\mathbf{p}_k = (p_{i,k})_{i\in[n]} = \mathbf{p}$.

while $R^*(S, \mathbf{T}, \mathbf{D}) > 0$ **do**

  Let $t^* \in [m]$ be a maximizer of $R^*(S, \mathbf{T}, \mathbf{D})$.

  $S_k = S \cap \mathcal{J}_{t^*}(\mathbf{T})$

  $\varepsilon_k = \min\limits_{i\in S_k} \dfrac{p_{i,k}}{\min\{R^*(S, \mathbf{T}, \mathbf{D}), r_i\}}$

  For $i = 1, \ldots, m$ let

$$p_{i,k+1} = \begin{cases} p_{i,k} - \varepsilon_k \min\{R^*(S, \mathbf{T}, \mathbf{D}), r_i\}, & \text{if } i \in S_k, \\ p_{i,k}, & \text{o/w.} \end{cases}$$

  Let $j^* \in S_k$ be a minimizer of ε_k.

  Set $S = S \setminus \{j^*\}$, and $k = k + 1$.

return S

Remark 4.3.1. *A variation of Algorithm 4.2 for* LMIS *is stated in Bar-Noy et al. (2001). The main difference is that the algorithm of Bar-Noy et al. (2001) has an additional step to ensure that the solution returned is maximal. The extra step helps in improving the approximation ratio, but it destroys the hope for monotonicity of $\hat{A}$. This can be demonstrated using the same example that we used for the primal-dual algorithms of Proposition 4.3.1. Furthermore, we note that the algorithm of Bar-Noy et al. (2001) is presented using the local-ratio jargon. We have chosen to present Algorithm 4.2 here in a self-contained way for ease of exposition but for its analysis (Section B.2 of Appendix B), we do make use of the local-ratio framework.*

Theorem 4.3.3. *Algorithm 4.2 achieves a Δ-approximation for the* LMIS *problem, where $\Delta = \max_{j=1,\ldots,m} |\mathcal{N}_j(\mathbf{I})|$, and the analysis of its approximation is tight.*

It remains to be shown that $\hat{A}$, with Algorithm 4.2 as a subroutine, becomes a monotone allocation algorithm for CMIC. For establishing the monotonicity, according to Definition 4.2.2, we have to examine the ways in which a winning worker i can deviate from the truth with a bid $b_i' \succeq b_i$, where b_i is her initial bid. By Definition 4.2.1, this means that under b_i', a lower or equal cost and a larger or the same interval are declared, compared to b_i. But to argue about $\hat{A}$, we first have to understand how such deviations from the truth affect the outcome of Algorithm 4.2, when it is called by $\hat{A}$.

Lowering the cost at a CMIC instance corresponds to lowering the penalty of an activity at the LMIS instance that $\hat{A}$ constructs. The lemma below examines precisely what happens when we lower the penalty of a non-scheduled activity in a LMIS instance.

Lemma 4.3.2. *For an instance* $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$ *of* LMIS*, let S be the schedule returned by Algorithm 2, and $j \in \mathcal{J} \setminus S$. Then, for any $p'_j \leq p_j$, it holds that $j \in \mathcal{J} \setminus S'$, where S' is the schedule returned by Algorithm 2 for* $((p'_j, \mathbf{p}_{-j}), \mathbf{T}, \mathbf{r}, \mathbf{D})$.

Proof. Let k^* be the iteration of Algorithm 4.2, in which activity j was removed from the solution set, when run on $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$. If the unilateral change of p_j to a $p'_j < p_j$ causes activity j to be removed at an iteration $k < k^*$ by being the chosen minimizer of ε_k, then the lemma follows.

When this is not the case, we claim that the only alternative is for activity j to be removed, once again, at iteration k^*. To see this, we will show that all previous iterations run exactly as in the original instance, if j is not removed in one of them. Let $p^{new}_{i,k}$ be the value of $p_{i,k}$ when Algorithm 4.2 is run on the new instance. The crucial observation here (enforcing that there will be at least k^* iterations) is that the quantity $R_\ell(S, \mathbf{T}, \mathbf{D})$ is independent of $\mathbf{p}$ for every ℓ. Hence $R^*(S, \mathbf{T}, \mathbf{D})$ remains unaffected by the change of p_j to p'_j (nothing else changes in the new instance). This in turn implies that as long as Algorithm 4.2 uses a deterministic tie-breaking rule, the maximizer of $R^*(S, \mathbf{T}, \mathbf{D})$ remains the same, and so does the set S_k at each iteration k. Since we assumed that activity j was not removed at iterations $k = 1, \dots, k^* - 1$, the minimizers of each ε_k are exactly the same as before, and different from j. Putting everything together, all iterations before k^* will run in exactly the same manner as before, and when we reach iteration k^*, for all activities $\ell \neq j$, we have that $p^{new}_{\ell,k^*} = p_{\ell,k^*}$, whereas for j we have that $p^{new}_{j,k^*} \leq p_{j,k^*}$ (since we started with a lower value for p_j). Hence,

$$\frac{p^{new}_{j,k^*}}{\min\{R^*(S, \mathbf{T}, \mathbf{D}), r_j\}} \leq \frac{p_{j,k^*}}{\min\{R^*(S, \mathbf{T}, \mathbf{D}), r_j\}}.$$

This implies that the activity j is still a minimizer of ε_{k^*} at iteration k^* and j will once again be left out from the schedule at this iteration. □

The next lemma examines enlarging the interval of an activity. This needs a different argument from the previous lemma because the deviation causes an activity to participate in more time instants.

Lemma 4.3.3. *For an instance* $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$ *of* LMIS*, let S be a schedule returned by Algorithm 2 and $j \in \mathcal{J} \setminus S$. For any interval $T'_j \supseteq T_j$, consider the instance*

$P' = (\mathbf{p}, (T'_j, \mathbf{T}_{-j}), \mathbf{r}, \mathbf{D}')$, where $\mathbf{D}' = (D'_1, \ldots, D'_m)$ such that:

$$D'_\ell = \begin{cases} D_\ell + r_j, & \textit{if } \ell \in T'_j \setminus T_j, \\ D_\ell, & o/w. \end{cases}$$

Then, $j \in \mathcal{J} \setminus S'$, where S' is the schedule returned by Algorithm 2 for P'.

Proof. Let k^* be the iteration of the algorithm in which activity j was removed when run on the $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$ instance. If the simultaneous change of T_j to a superset T'_j and of $\mathbf{D}$ to $\mathbf{D}'$, as defined in the statement of the lemma, causes activity j to be removed at an iteration $k < k^*$ by being the chosen minimizer of ε_k, then the lemma follows.

When this is not the case, we claim that the only alternative is for activity j to be removed again at iteration k^*. Let $p_{i,k}^{new}$ be the value of $p_{i,k}$ when Algorithm 4.2 is run on the new instance. The crucial observation is that as long as j is not removed, the quantity $R_\ell(S, \mathbf{T}, \mathbf{D})$ will remain the same as in the original instance for every $\ell \in [m]$. This enforces that the condition of the while loop will not be violated before iteration k^*.

Claim 1. *For all $S \subseteq \mathcal{J}$ such that $j \in S$, and for every $\ell \in [m]$, it holds that $R_\ell(S, \mathbf{T}, \mathbf{D}) = R_\ell(S, (T'_j, \mathbf{T}_{-j}), \mathbf{D}')$.*

Proof. Fix a time instant $\ell \in [m]$. We distinguish the cases:
Case 1: $\ell \in T'_j \setminus T_j$. We have that,

$$\begin{aligned} R_\ell(S, \mathbf{T}, \mathbf{D}) &= \sum_{i \in S \cap \mathcal{J}_\ell(\mathbf{T})} r_i - D_\ell \\ &= \sum_{i \in S \cap \mathcal{J}_\ell(\mathbf{T})} r_i + r_j - (D_\ell + r_j) \\ &= \sum_{i \in S \cap \mathcal{J}_\ell(T'_j, \mathbf{T}_{-j})} r_i - D'_\ell \\ &= R_\ell(S, (T'_j, \mathbf{T}_{-j}), \mathbf{D}'). \end{aligned}$$

The third equality follows by the facts that, for $\ell \in T'_j \setminus T_j$, it holds that $D'_\ell = D_\ell + r_j$ and also that $j \in S$ and $\{j\} = \mathcal{J}_\ell(T'_j, \mathbf{T}_{-j}) \setminus \mathcal{J}_\ell(\mathbf{T})$.

Case 2: $\ell \notin T'_j \setminus T_j$. We have that,

$$\begin{aligned} R_\ell(S, \mathbf{T}, \mathbf{D}) &= \sum_{i \in S \cap \mathcal{J}_\ell(\mathbf{T})} r_i - D_\ell \\ &= \sum_{i \in S \cap \mathcal{J}_\ell(T'_j, \mathbf{T}_{-j})} r_i - D'_\ell \\ &= R_\ell(S, (T'_j, \mathbf{T}_{-j}), \mathbf{D}'). \end{aligned}$$

The second equality follows by the facts that, for $\ell \notin T'_j \setminus T_j$, it holds that $D'_\ell = D_\ell$ and also that $\mathcal{J}_\ell(T'_j, \mathbf{T}_{-j}) = \mathcal{J}_\ell(\mathbf{T})$. □

By using Claim 1 for iterations $k = 1, \ldots, k^* - 1$ (we can do that since j has not been removed by the algorithm yet, therefore $j \in S$ in all these iterations), we notice that the time instants that maximize R^* in each iteration are exactly the same as before. Furthermore, since the interval of activity j now includes more time instants, there may be more iterations where j is in the set S_k, compared to the original instance (this happens when the time instant maximizing R^* at iteration k is in $T'_j \setminus T_j$). But this could only cause $p^{new}_{j,k}$ to further decrease, and therefore we will eventually have that $p^{new}_{j,k^*} \leq p_{j,k^*}$ or (using Claim 1) that

$$\frac{p^{new}_{j,k^*}}{\min\{R^*(S, (T'_j, \mathbf{T}_{-j}), \mathbf{D}), r_j\}} \leq \frac{p_{j,k^*}}{\min\{R^*(S, \mathbf{T}, \mathbf{D}), r_j\}}.$$

This implies that activity j is once again the minimizer of ε_{k^*}, as in the original instance, and the proof is complete. □

Combining Lemma 4.3.2 and Lemma 4.3.3 we get the following:

Theorem 4.3.4. *Algorithm $\hat{A}$ is monotone, when using Algorithm 4.2 as a black box for solving* LMIS*.*

Proof. Fix a bidding profile $\mathbf{b}$ and a winning worker $i \in \mathcal{W}_{\hat{A}}(\mathbf{b})$. Let $b_i = (c_i, [s_i, f_i])$, and consider an arbitrary deviation of i, say $b''_i = (c'_i, [s'_i, f'_i])$, such that $b''_i \succeq b_i$. We need to show that $i \in \mathcal{W}_{\hat{A}}(b''_i, \mathbf{b}_{-i})$ and we will do this in two steps. First, we consider the deviation $b'_i = (c'_i, [s_i, f_i])$. Since b'_i differs from b_i only with respect to the declared cost and the bids of the remaining workers remain the same, we can directly use Lemma 4.3.2 to conclude that $i \in \mathcal{W}_{\hat{A}}(b'_i, \mathbf{b}_{-i})$. Having established that i is still a winner under $(b'_i, \mathbf{b}_{-i})$, consider now the deviation from b'_i to b''_i. Note that $(b''_i, \mathbf{b}_{-i})$ differs from $(b'_i, \mathbf{b}_{-i})$ only with respect to the declared interval of bidder i. Recall also that Algorithm $\hat{A}$ calls Algorithm 4.2 with input the tuple $(\mathbf{c}, \mathbf{I}, \mathbf{q}, \mathbf{Q}(\mathbf{I}) - \mathbf{d})$. This means that under b_i'', the vector $\mathbf{Q}(\mathbf{I}) - \mathbf{d}$ in the constructed LMIS instance changes only for time instants that belong to $[s_i', f_i'] \setminus [s_i, f_i]$ (where we simply add q_i). But then, this can be handled by Lemma 4.3.3, and obtain that $i \in \mathcal{W}_{\hat{A}}(b_i'', \mathbf{b}_{-i})$. □

To conclude, it is trivial that Algorithm $\hat{A}$ runs in polynomial time, when using Algorithm 4.2 for solving LMIS, and hence, the monotonicity of $\hat{A}$, together with Theorems 4.3.2 and 4.3.3 complete the proof of Theorem 4.3.1.

4.4 A Truthful FPTAS for a Small Number of Tasks

At what follows, we investigate whether we can have truthful mechanisms with a better approximation ratio for special cases of restricted problem size. An instance of CMIC with a constant number of workers can be optimally solved in polynomial time by a brute force algorithm, which, together with the VCG payment scheme, results in a truthful mechanism. On the other hand, the story is different when we have a small number of tasks, since CMIC is NP-hard even for one task (Dayama et al., 2015). Building upon this negative result, we provide a truthful mechanism that achieves the best possible approximation factor, for the case of a constant number of tasks, and our main result of this section is the following:

Theorem 4.4.1. *There exists a truthful FPTAS for CMIC, when the number of tasks is constant.*

We would like first to pay attention to the special case of a single task, which corresponds to the MIN-KNAPSACK problem, the minimization version of KNAPSACK where items have costs (instead of values) and there is a covering requirement (instead of a capacity constraint) for the selected items. The work of Briest et al. (2011), which proposes a truthful FPTAS for the classic maximization version of KNAPSACK, claims (without providing a proof) that the analogous result holds for MIN-KNAPSACK too. To our knowledge, the only published work that explicitly attempts to extend the work of Briest et al. (2011) and describe a truthful FPTAS for MIN-KNAPSACK is by Chen et al. (2019). However, we found that the analysis of truthfulness there is flawed and we refer to Section B.3 of Appendix B for a counterexample.

Proposition 4.4.1. *The FPTAS for MIN-KNAPSACK, proposed by Chen et al. (2019), is not monotone.*

Therefore, our Theorem 4.4.1 helps to resolve any potential ambiguities for MIN-KNAPSACK. Finally, for two or more tasks, we are not aware of any truthful mechanism attaining any bound better than the one provided by Theorem 4.3.1.

Remark 4.4.1. *As in other approaches for constructing a truthful FPTAS, e.g. Briest et al. (2011), we also make the assumption from here onwards, that* $c_i \geq 1$ *for every worker* i*, i.e., the workers will not be allowed to declare a cost lower than* 1*. At Section B.4 of Appendix B, we prove that we can adjust the assumption to* $c_i \geq \delta$ *for any arbitrarily small* δ*, but for convenience, here we stick to* $\delta = 1$*.*

Furthermore, we bring in some additional notation that we use in this section. Given a bidding profile $\mathbf{b}$, let $c_{sum}(\mathbf{b}) := \sum_{i \in \mathcal{N}} c_i$. Similarly let $c_{min}(\mathbf{b})$ (resp. $c_{max}(\mathbf{b})$), be the minimum (resp. maximum) cost by the bidders.

4.4.1 A Pseudopolynomial Dynamic Programming Algorithm

The first step towards designing the FPTAS, is a pseudopolynomial dynamic programming algorithm that returns the optimal solution for the case of a constant number of tasks. For simplicity, we focus on describing the algorithm for the case of two tasks. The generalization is rather obvious (and discussed briefly at the end of this subsection).

Given an instance with two tasks, let d_1, d_2 be the demand requirements of the tasks. Note that $\mathcal{N}$ can be partitioned into three sets, W_0, W_1, W_2, since we can have at most three types of workers: W_0 is the set of workers who can contribute to both tasks, and for $\ell \in \{1, 2\}$, W_ℓ is the set of workers who are capable of contributing only to task ℓ.

We define a 3-dimensional matrix $Q[\ell, i, c]$, where for $\ell = 0, 1, 2$, for $i = 0, 1, \dots n$, and for $c = 0, 1, \dots, c_{sum}(\mathbf{b})$, $Q[\ell, i, c]$ denotes the maximum possible contribution that can be jointly achieved by any set of workers in $W_\ell \cap \{1, \dots, i\}$ with a total cost of exactly c. For our purposes, we assume[5] that for $i \in [n]$, each c_i is an integer, so that c also takes only integral values. Our algorithm is based on computing the values of the cells of Q and we claim that this can be done by exploiting the following recursive relation:

$$Q[\ell, i, c] = \begin{cases} 0, & \text{if } i = 0 \\ Q[\ell, i-1, c], & \text{if } i > 0 \text{ and either } i \notin W_\ell \text{ or } c_i > c \\ \max\{Q[\ell, i-1, c], Q[\ell, i-1, c-c_i] + q_i]\}, & \text{o/w} \end{cases} \tag{4.3}$$

Observe that for a feasible solution S, the workers who contribute to the demand of task 1 (resp. 2) are those from $S \cap W_0$ and $S \cap W_1$ (resp. $S \cap W_0$ and $S \cap W_2$). Hence, for

[5] It becomes clear in the next subsection, that the dynamic programming procedure is only needed for integral cost values.

$\ell \in \{1, 2\}$, it should hold that $\sum_{j \in S \cap W_0} q_j + \sum_{j \in S \cap W_\ell} q_j \geq d_\ell$. Our algorithm then can work as follows: After computing the values of Q, according to Equation (4.3), return the set of workers that minimize $c^{(0)} + c^{(1)} + c^{(2)}$, subject to $Q[0, n, c^{(0)}] + Q[\ell, n, c^{(\ell)}] \geq d_\ell$, for $\ell \in \{1, 2\}$. This can be done by enumerating all possible options, for breaking down the final cost as a sum of 3 values, $c^{(0)}$, $c^{(1)}$ and $c^{(2)}$. The formal statement can be found below.

Algorithm 4.3: DP($\mathbf{b}$) (presented for two tasks)

```
▷ Input: A bidding profile b = (c, I) of a CMIC instance (b, q, d) with m = 2
for ℓ ∈ {0, 1, 2} do
    for i ∈ {0, 1, ..., n} do
        for c ∈ {0, 1, ..., c_sum(b)} do
            Compute Q[ℓ, i, c] using Equation (4.3)
return the set of workers that minimize c^(0) + c^(1) + c^(2), s.t.
    Q[0, n, c^(0)] + Q[ℓ, n, c^(ℓ)] ≥ d_ℓ, ∀ℓ ∈ {1, 2} (or +∞, if (b, q, d) has no solution)
```

The optimality of the DP algorithm is straightforward from the preceding discussion. Furthermore, its running time is pseudopolynomial, since the size of the table Q is $3 \cdot (|\mathcal{N}| + 1) \cdot (c_{sum}(\mathbf{b}) + 1)$ and to find the optimal solution we need to check at most $\binom{c_{sum}(\mathbf{b})}{3}$ different combinations for the decomposition of the total cost in three terms, as described earlier. It is easy to extend these ideas, for more tasks given the interval structure of the problem, i.e., the first dimension of Q will have a range of $O(m^2)$ and the enumeration part of the algorithm will require an order of $\binom{c_{sum}}{m^2}$ steps. Finally, we note that since this is an optimal mechanism and we use a deterministic, consistent tie-breaking rule, it will trivially be monotone.

Henceforth, we will be referring to this pseudopolynomial dynamic programming algorithm for any constant number of tasks, as the DP algorithm.

Theorem 4.4.2. *Given an instance of* CMIC *on a profile* $\mathbf{b}$*, with a constant number of tasks and integer costs, Algorithm DP(*$\mathbf{b}$*) is optimal, monotone, and runs in pseudopolynomial time, i.e. polynomial in the input size and in* c_{sum}.

4.4.2 The FPTAS

In order to convert the DP algorithm to a truthful FPTAS, we adapt the framework of Briest et al. (2011). To that end, we define, for every integer k, an algorithm $A_k(\mathbf{b}, \epsilon)$, that uses the DP algorithm as a subroutine, on a subset of the initial set of bidders, with rounded costs, as follows:

Algorithm 4.4: $A_k(\mathbf{b}, \epsilon)$

▷ **Input:** A bidding profile $\mathbf{b} = (\mathbf{c}, \mathbf{I})$ of a CMIC instance $(\mathbf{b}, \mathbf{q}, \mathbf{d})$, $\epsilon \in (0, 1)$

Let $\mathcal{L}_k(\mathbf{c}) = \{i \in \mathcal{N} : c_i \leq 2^{k+1}\}$

$a_k = \frac{n}{\epsilon 2^k}$.

for $i \in \mathcal{L}_k(\mathbf{c})$ **do**

  $\bar{c}_i = \lceil a_k \cdot c_i \rceil$

$\bar{\mathbf{b}} = (\bar{c}_i, [s_i, f_i])_{i \in \mathcal{L}_k(\mathbf{c})}$

return DP($\bar{\mathbf{b}}$)

Lemma 4.4.1. *Let* $0 < \epsilon < 1$. *For a bidding profile* $\mathbf{b}$ *and* $k \geq 0$, *the algorithm* $A_k(\mathbf{b}, \epsilon)$ *runs in time polynomial in the input size and in* $\frac{1}{\epsilon}$, *and if* $2^k \leq OPT(\mathbf{b}) < 2^{k+1}$, *it computes a solution of cost at most* $(1 + \epsilon) OPT(\mathbf{b})$.

Proof. Starting with the time complexity of $A_k(\mathbf{b}, \epsilon)$, in order to prove that it terminates after a polynomial amount of time, we need to examine the running time of $DP(\bar{\mathbf{b}})$, since all other steps are clearly polynomial. Each component of the cost vector that is given as input to the algorithm DP($\bar{\mathbf{b}}$) is bounded by $\frac{2^{k+1}n}{\epsilon 2^k} + 1$, by the definitions of the set $\mathcal{L}_k(\mathbf{c})$ and a_k and hence $c_{sum}(\bar{\mathbf{b}})$ is polynomial in n and $\frac{1}{\epsilon}$. Since all other parameters of the instance remain the same, we conclude that DP($\bar{\mathbf{b}}$) is polynomial in the original input size of the algorithm $A_k(\mathbf{b}, \epsilon)$ and in $1/\epsilon$. Hence, $A_k(\mathbf{b}, \epsilon)$ is also polynomial in the input size and in $\frac{1}{\epsilon}$.

For every worker $i \in \mathcal{L}_k(\mathbf{c})$, let $c'_i = \frac{\bar{c}_i}{a_k}$. For brevity, let $\mathcal{W}^*$ be the optimal set of workers w.r.t. the bidding vector $\mathbf{b}$. We have the following implications:

$$\sum_{i \in \mathcal{W}^*} c'_i - \mathrm{OPT}(\mathbf{b}) = \sum_{i \in \mathcal{W}^*} c'_i - \sum_{i \in \mathcal{W}^*} c_i \leq$$
$$\sum_{i \in \mathcal{W}^*} \left(\frac{a_k c_i + 1}{a_k} - c_i\right) = \sum_{i \in \mathcal{W}^*} \frac{1}{a_k} \leq \frac{|\mathcal{W}^*|}{a_k} \leq$$
$$\frac{n}{a_k} = \epsilon 2^k \leq \epsilon \mathrm{OPT}(\mathbf{b})$$

Now let W_k be the set of selected workers by the Algorithm $A_k(\mathbf{b}, \epsilon)$, so that the cost of the returned solution is SOL $= \sum_{i \in \mathcal{W}_k} c_i$. Since for every worker i, we have that $c_i = \frac{c_i a_k}{a_k} \leq \frac{\lceil c_i a_k \rceil}{a_k} = \frac{\bar{c}_i}{a_k} = c'_i$, it holds that

$$\text{SOL} \le \sum_{i \in \mathcal{W}_k} c'_i = \frac{1}{a_k} \sum_{i \in \mathcal{W}_k} \bar{c}_i \le \frac{1}{a_k} \sum_{i \in \mathcal{W}^*} \bar{c}_i = \sum_{i \in \mathcal{W}_k} c'_i.$$

The rightmost inequality in the above implications holds because of the following claim together with the fact that, due to the optimality of the dynamic programming procedure, the set $\mathcal{W}_k$ is optimal w.r.t. to the bidding vector $\bar{\mathbf{b}}$.

Claim 2. *The set $\mathcal{W}^*$ is a feasible solution for the instance with the scaled bidding vector $\bar{\mathbf{b}}$ of $|\mathcal{L}_k(\mathbf{c})|$ bidders, produced in Algorithm $A_k(\mathbf{b}, \epsilon)$, i.e., the constraints of Equation (4.1) are satisfied by $\mathcal{W}^*$, and $\mathcal{W}^* \subseteq \mathcal{L}_k(\mathbf{c})$.*

Proof. We recall that feasibility of a solution is not dependent on the bidding vector but only on the contribution and demand parameters, according to Equation (4.1). Since $A_k(\mathbf{b}, \epsilon)$ leaves unaltered both the qualities of the workers and the demands of the tasks, any subset of workers from $\mathcal{L}_k(\mathbf{c})$ that is a feasible solution for the bidding profile $\mathbf{b}$ is also a feasible solution for $\bar{\mathbf{b}}$. Consequently, it suffices to show that $\mathcal{W}^* \subseteq \mathcal{L}_k(\mathbf{c})$. But note that any worker i who does not belong to $\mathcal{L}_k(\mathbf{c})$, satisfies $c_i > 2^{k+1}$. Hence, i cannot be part of the optimal solution $\mathcal{W}^*$, since we have assumed that $\text{OPT}(\mathbf{b}) < 2^{k+1}$. □

Combining the established inequalities, we get that $\text{SOL} \le (1+\epsilon)\text{OPT}(\mathbf{b})$. □

Using Lemma 4.4.1, we can achieve the desired approximation by checking all possible values for k. However, to provide a polynomial-time mechanism, we can only test polynomially many such values, and hope that we compute the same outcome as if we were able to test all such algorithms. We will show that Algorithm 4.5, that tests all values up to a certain threshold, is what we need.

Algorithm 4.5: $A_{\text{FPTAS}}(\mathbf{b}, \epsilon)$

▷ **Input:** A bidding profile $\mathbf{b} = (\mathbf{c}, \mathbf{I})$ of a CMIC instance $((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$, $\epsilon \in (0, 1)$

for $k = 0, \ldots, \lceil \log\left(\frac{n c_{max}(\mathbf{b})}{\epsilon}\right) \rceil$ **do**
  Run $A_k(\mathbf{b}, \epsilon)$ and store the winning set and its cost.
return the set of workers that achieve the minimum cost among the above, breaking ties in favor of the algorithm with the lowest index

Let $k^*(\mathbf{b}) := \lceil \log\left(\frac{n c_{max}(\mathbf{b})}{\epsilon}\right) \rceil$ (or simply k^* when the bidding profile is clear from the context). To see that the algorithm is well-defined, recall that $c_{max}(\mathbf{b}) \ge 1$ and since also $\frac{n}{\epsilon} \ge 1$, we have that $k^* \ge 0$. To establish that this is indeed a FPTAS, we prove in the following lemma that one cannot find a better solution by running an algorithm A_k for a value of k higher than k^*. In combination with Lemma 4.4.1, this directly establishes that A_{FPTAS} is a FPTAS for CMIC.

Lemma 4.4.2. *Given an instance of* CMIC *and an* $\epsilon \in (0,1)$*, it holds that* $\mathcal{W}_{A_k(\mathbf{b})} = \mathcal{W}_{A_{k^*}(\mathbf{b})}$*, for every* $k > k^*(\mathbf{b})$*.*

Monotonicity of A_{FPTAS}

To establish monotonicity, we will make use of the following operator:

Definition 4.4.1. *Let* $\mathcal{A} = \{A_0, A_1, \ldots\}$ *be the set of all allocation algorithms* A_k*. For a profile* $\mathbf{b}$ *and a finite collection of algorithms* $S \subseteq \mathcal{A}$*, let* $MIN(S, \mathbf{b}) := \arg\min_{A \in S} C(A, \mathbf{b})$*, with ties broken in favor of the lowest index.*

Given a profile $\mathbf{b}$, the algorithm A_{FPTAS} can be expressed as $\text{MIN}\{A_0, \ldots, A_{k^*(\mathbf{b})}\}$. Hence, the next step is to determine when is the MIN operator monotone. The framework of Briest et al. (2011) defines a set of sufficient conditions, for maximization objectives. We adapt these properties below, and we note that they are sufficient conditions for minimization problems as well.

Definition 4.4.2. *A monotone allocation algorithm* A *is bitonic w.r.t. the social cost function* C *if for any bidding profile* $\mathbf{b}$ *and any worker* i*, the following hold:*

1. $i \in \mathcal{W}_A(\mathbf{b}) \Rightarrow C(A, \mathbf{b}) \geq C(A, (b'_i, \mathbf{b}_{-i})) \quad \forall b'_i \succeq b_i$
2. $i \notin \mathcal{W}_A(\mathbf{b}) \Rightarrow C(A, \mathbf{b}) \geq C(A, (b'_i, \mathbf{b}_{-i})) \quad \forall b'_i \preceq b_i$

Lemma 4.4.3. *For* $k \geq 0$*, Algorithm 4.4 is monotone and bitonic w.r.t. the social cost function* C*.*

Proof. To prove first the monotonicity of A_k, fix a bid vector $\mathbf{b}$ and a bidder $i \in \mathcal{W}_{A_k}(\mathbf{b})$. We will prove that when bidder i issues a bid $b'_i \succeq b_i$ it will hold that $i \in \mathcal{W}_{A_k}(b'_i, \mathbf{b}_{-i})$. Suppose that $b'_i = (c'_i, [s'_i, f'_i])$. Firstly, it is clear that $i \in W_{A_k}(\mathbf{b})$ implies $i \in \mathcal{L}_k(\mathbf{c})$ and since $b'_i \succeq b_i$, this means that $c'_i \leq c_i$, therefore, it trivially holds that $i \in \mathcal{L}_k(c'_i, \mathbf{c}_{-i})$ and in fact $\mathcal{L}_k(\mathbf{c}) = \mathcal{L}_k(c'_i, \mathbf{c}_{-i})$. Moreover, $\bar{c}'_i \leq \bar{c}_i$.

The last step of Algorithm A_k is to run the dynamic programming procedure DP. Since, DP is an optimal algorithm and with a consistent, deterministic, tie-breaking rule, it is easy to verify that it is also monotone. Hence, if $i \in \mathcal{W}_{\text{DP}}(\bar{\mathbf{b}})$, then $i \in \mathcal{W}_{\text{DP}}(\bar{b}'_i, \bar{b}_{-i})$, because $\bar{b}'_i \succeq \bar{b}_i$, and therefore

$$i \in \mathcal{W}_{A_k}(\mathbf{b}) \implies i \in \mathcal{W}_{A_k}(b'_i, \mathbf{b}_{-i}). \tag{4.4}$$

We continue with proving the bitonicity of A_k w.r.t. the social cost function C. For the first statement in Definition 4.4.2, given k, a bidding profile $\mathbf{b}$, a bidder $i \in \mathcal{W}_{A_k}(\mathbf{b})$ and a bid $b'_i \succeq b_i$, we have

$$\begin{aligned} C(A_k, \mathbf{b}) &= \sum_{j \in \mathcal{W}_{A_k}(\mathbf{b}) \setminus \{i\}} c_j + c_i \geq \sum_{j \in \mathcal{W}_{A_k}(\mathbf{b}) \setminus \{i\}} c_j + c'_i \\ &= \sum_{j \in \mathcal{W}_{A_k}(b'_i, \mathbf{b}_{-i}) \setminus \{i\}} c_j + c'_i = C(A_k, (b'_i, \mathbf{b}_{-i})). \end{aligned}$$

The inequality holds since $b'_i \succeq b_i$ implies that $c_i \geq c'_i$. Also, regarding the second equality, we note that it may not necessarily be the case that $\mathcal{W}_{A_k}(\mathbf{b}) \setminus \{i\} = \mathcal{W}_{A_k}(b'_i, \mathbf{b}_{-i}) \setminus \{i\}$. However, for all the optimal solutions containing i, of the instance for which the dynamic programming algorithm DP is run by A_k, the total cost of the selected workers other than i must be the same. These solutions will remain optimal when i bids c'_i and this justifies the second equality above.

For the second statement of Definition 4.4.2, consider an integer k, a bidding profile $\mathbf{b}$, a bidder $i \notin \mathcal{W}_{A_k}(\mathbf{b})$, and a bid b'_i with $b'_i \preceq b_i$. Clearly, If $i \notin \mathcal{L}_k(\mathbf{c})$ then $i \notin \mathcal{L}_k(c'_i, \mathbf{c}_{-i})$, since $b'_i \preceq b_i$ implies that $c'_i \geq c_i$. Therefore bidder $i \notin \mathcal{W}_{A_k}(b'_i, \mathbf{b}_{-i})$, which means that $C(A_k, \mathbf{b}) = C(A_k, (b'_i, \mathbf{b}_{-i}))$. On the other hand, when $i \in \mathcal{L}_k(\mathbf{c})$ but $i \notin \mathcal{W}_{\mathrm{DP}}((\bar{\mathbf{b}}))$, it will also hold that $i \notin \mathcal{W}_{\mathrm{DP}}(\bar{b}'_i, \bar{\mathbf{b}}_{-i})$, since DP is monotone. Again, this implies that $C(A_k, \mathbf{b}) = C(A_k, (b'_i, \mathbf{b}_{-i}))$. □

Lemma 4.4.3 will be used in conjunction with the following one, which is implied by Briest et al. (2011). For completeness, we present a proof of it.

Lemma 4.4.4. *(implied by Briest et al. (2011)) For a profile* $\mathbf{b}$*, MIN*$(\{A_0, \ldots, A_\ell\}, \mathbf{b})$ *is monotone if* $A_0, \ldots, A_\ell$ *are monotone algorithms that are additionally bitonic w.r.t. the social cost function* C*.*

Proof. Denote by $\mathcal{W}_{\mathrm{MIN}}(\mathbf{b})$ the set of winners returned by $\mathrm{MIN}(\{A_0, \ldots, A_\ell\}, \mathbf{b})$, for a profile $\mathbf{b}$ and assume that $\mathrm{MIN}(\{A_0, \ldots, A_\ell\})$ is not monotone. Then, there exists a specific bidding profile, say $\mathbf{b}$, a bidder $i \in \mathcal{W}_{\mathrm{MIN}}(\mathbf{b})$ and a bid $b'_i \succeq b_i$ such that $i \notin \mathcal{W}_{\mathrm{MIN}}(b'_i, \mathbf{b}_{-i})$. Let A_j be the algorithm that returns the solution $\mathcal{W}_{\mathrm{MIN}}(\mathbf{b})$ with $k \in \{0, \ldots, \ell\}$. This means that

$$C(A_j, \mathbf{b}) \leq C(A_k, \mathbf{b}), \text{ for } k = 0, \ldots, \ell,$$

and moreover, from the bitonicity of A_j it holds that

$$C(A_j, \mathbf{b}) \geq C(A_j, (b'_i, \mathbf{b}_{-i})).$$

Let A_m be the algorithm that returns the solution $W_{\min}(b'_i, \mathbf{b}_{-i})$ (which by assumption does not include bidder i) with $m \in \{0, \ldots, \ell\}$. Note that $j \neq m$, otherwise, we would immediately obtain a contradiction to the fact that A_j is monotone. For A_m we have that

$$C(A_m, (b'_i, \mathbf{b}_{-i})) \leq C(A_k, (b'_i, \mathbf{b}_{-i})), \text{ for } k = 0, \ldots, \ell.$$

Now from the bitonicity of A_m (second sentence of Definition 4.4.2, and using $\mathbf{b}$ as a deviation from $(b'_i, \mathbf{b}_{-i})$), we have

$$C(A_m, (b'_i, \mathbf{b}_{-i})) \geq C(A_m, \mathbf{b}).$$

By combining the inequalities above we have that for all $k = 0, \ldots, \ell$, the following hold:

$$C(A_k, \mathbf{b}) \geq C(A_j, (b'_i, \mathbf{b}_{-i})) \tag{4.5}$$

$$C(A_k, (b'_i, \mathbf{b}_{-i})) \geq C(A_m, \mathbf{b}) \tag{4.6}$$

If we set $k = m$ in Equation (4.5) and $k = j$ in Equation (4.6), then we obtain

$$C(A_m, \mathbf{b}) = C(A_j, (b'_i, \mathbf{b}_{-i})).$$

But we have already argued that $C(A_m, \mathbf{b}) \geq C(A_j, \mathbf{b}) \geq C(A_j, (b'_i, \mathbf{b}_{-i}))$. Hence $C(A_j, \mathbf{b}) = C(A_m, \mathbf{b})$. Since $j \neq m$, it must be that $j < m$ due to our tie-breaking rule. But note that $C(A_j, (b'_i, \mathbf{b}_{-i})) \geq C(A_m, (b'_i, \mathbf{b}_{-i})) \geq C(A_m, \mathbf{b}) = C(A_j, (b'_i, \mathbf{b}_{-i}))$. Therefore, $C(A_j, (b'_i, \mathbf{b}_{-i})) = C(A_m, (b'_i, \mathbf{b}_{-i}))$. This means that again, due to the tie-breaking rule, we should have picked the solution of A_j and not of A_m, under the profile $(b'_i, \mathbf{b}_{-i})$, reaching a contradiction. □

Before we proceed, we would like to comment on a subtle point regarding the MIN operator. We stress that the set of algorithms A_k that are called by A_{FPTAS}, depends on the input profile $\mathbf{b}$. There is no a priori fixed set of algorithms that are run in every profile, but instead, this is determined by the quantity $k^*(\mathbf{b})$. As a result, Lemma 4.4.4 does not suffice on its own. For the monotonicity of A_{FPTAS}, we also need to consider the case that a winning worker declares a lower cost that changes $c_{max}(\mathbf{b})$, and decreases the number of algorithms that A_{FPTAS} runs. This is where Lemma 4.4.2 comes to rescue, as explained in the proof of Theorem 4.4.3, and this is where the flaw in Chen et al. (2019) is located (i.e., the algorithm of Chen et al. (2019) does not perform enough iterations, as explained further in Section B.3 of the Appendix).

Given the above discussion, by performing a suitable case analysis and by using Lemmas 6 and 7, we can prove the following theorem, which also completes the proof of Theorem 4.4.1 and concludes this section.

Theorem 4.4.3. *For the domain of bidding profiles* $\mathbf{b}$*, such that* $c_{min}(\mathbf{b}) \geq 1$*, the algorithm* A_{FPTAS} *is monotone.*

Proof. Consider a bidder $i \in \mathcal{W}_{A_{\text{FPTAS}}}(\mathbf{b})$, and consider a deviation from b_i to $b_i' = (c_i', [s_i', f_i]')$, with $b_i' \succeq b_i$. For brevity, let $(b_i', \mathbf{b}_{-i}) = \mathbf{b}'$. We distinguish the following cases:

Case 1: Suppose that the deviation to b_i' is such that the quantity $k^*(\mathbf{b}')$ remains unaffected, and equal to $k^*(\mathbf{b})$, as before the deviation (e.g. if $c_{max}(\mathbf{b}) = c_{max}(\mathbf{b}')$ or if $c_{max}(\mathbf{b}')$ did not change significantly so as to affect $k^*(\mathbf{b})$). Then, A_{FPTAS} runs exactly the same set of algorithms under both $\mathbf{b}$ and $\mathbf{b}'$ and in both cases corresponds to $\text{MIN}(\{A_0, A_1, \ldots, A_{k^*(\mathbf{b})}\})$. But then, by applying Lemma 4.4.4, we have that $i \in \mathcal{W}_{A_{\text{FPTAS}}}(\mathbf{b}')$, and we are done.

Case 2: Suppose that the deviation is such that $k^*(\mathbf{b}') < k^*(\mathbf{b})$. This implies that $c_i = c_{max}(\mathbf{b})$ and $c_i' < c_i$. It should also hold that $c_i > c_{max}(\mathbf{b}_{-i}) = \max_{j \in \mathcal{N} \setminus \{i\}} c_j$. Thus, worker i declares a lower cost and causes a decrease in the number of algorithms that A_{FPTAS} will use under $\mathbf{b}'$. In particular, $A_{\text{FPTAS}}(\mathbf{b}')$ corresponds to $\text{MIN}(\{A_0, \ldots, A_{k^*(\mathbf{b}')}\}, \mathbf{b}')$ and, again, we need to show that i remains a winner.

Case 2a: Suppose that $c_i' \geq c_{max}(\mathbf{b}_{-i})$. Then, by applying Lemma 4.4.2 for profile $\mathbf{b}'$ and $k = k^*(\mathbf{b}')+1, \ldots, k^*(\mathbf{b})$, it holds that $\mathcal{W}_{A_{k^*(\mathbf{b}')}}(\mathbf{b}') = \mathcal{W}_{A_k}(\mathbf{b}')$. But this implies that calling $\text{MIN}(\{A_0, \ldots, A_{k^*(\mathbf{b}')}\}, \mathbf{b}')$ is equivalent to $\text{MIN}(\{A_0, \ldots, A_{k^*(\mathbf{b})}\}, \mathbf{b}')$. By applying Lemma 4.4.4 the theorem follows.

Case 2b: Suppose that $c_i' < c_{max}(\mathbf{b}_{-i})$. For convenience, we view b_i' as a deviation consisting of two steps: a deviation from b_i to $b_i'' = (c_{max}(\mathbf{b}_{-i}), [s_i', f_i'])$, and a deviation from b_i'' to b_i'. Note that it holds that both $b_i'' \succeq b_i$ and $b_i' \succeq b_i''$. In the first deviation, from b_i to b_i'', the maximum cost changes from c_i to $c_{max}(\mathbf{b}_{-i})$. Let $\mathbf{b}'' = (b_i'', \mathbf{b}_{-i})$. We will have that $k^*(\mathbf{b}'') < k^*(\mathbf{b})$, and we can use again the same argument, as in Case 2a, to conclude that $i \in A_{\text{FPTAS}}(\mathbf{b}'')$. Consider the second deviation from b_i'' to b_i'. By lowering further the cost, worker i will not affect the maximum cost, since $c_{max}(\mathbf{b}') = c_{max}(\mathbf{b}'') = c_{max}(\mathbf{b}_{-i})$. But this means that $k^*(\mathbf{b}'') = k^*(\mathbf{b}')$, and the argument of Case 1 applies, so that we can conclude that $i \in A_{\text{FPTAS}}(\mathbf{b}')$. □

4.5 Generalization to Non-Interval Structures

The focus of our work in this chapter has been largely on highlighting the difference on the approximation ratio between truthful and non-truthful algorithms. We conclude this chapter by showing that one can get tighter results, when moving to more general scenarios. A direct generalization of CMIC is to drop the linear arrangement of the tasks, and allow each worker $j \in \mathcal{N}$ to declare an arbitrary subset of tasks $I_j \subseteq [m]$. The rest of the input remains the same (costs, contributions and demands), and we refer to this problem as Cost Minimization Demand Cover (CMDC). Notice that CMDC with $q_i = 1$ for every worker i, and $d_j = 1$ for every task j, is nothing but the famous SET COVER problem.

Further extensions of CMDC have been studied under various names in a series of works regarding covering integer programs, see e.g., Kolliopoulos and Young (2005); Rajagopalan and Vazirani (1993), and approximation algorithms that match the factor of Algorithm 4.1 exist (see e.g. Carr et al. (2000); Fujito and Yabuta (2004); Koufogiannakis and Young (2013); Pritchard and Chakrabarty (2011)). However, the focus of these works was not about the monotonicity of their algorithms. It is unclear if any of these algorithms are monotone (and in fact some of them are certainly non-monotone, as they are based on the construction of certain approximate separation oracles).

Towards obtaining a monotone algorithm, a careful inspection of the proofs of Section 4.3, suffices to deduce that Algorithm 4.1 can be used for this more general setting as well (after defining first the appropriate generalization of LMIS) and it continues to yield the Δ factor for CMDC. Furthermore, under this setting, Algorithm 4.1 yields essentially a tight result, according to the following Proposition. The proof of Proposition 4.5.1 is straightforward due to the hardness results for k-UNIFORM HYPERGRAPH VERTEX COVER Dinur et al. (2005); Khot and Regev (2008), which is a special case of SET COVER, and therefore a special case of CMDC.

Proposition 4.5.1. *For the class of instances where Δ is constant, CMDC is $\Delta - 1 - \epsilon$ inapproximable, unless $P = NP$, and $\Delta - \epsilon$ inapproximable assuming the Unique Games Conjecture.*

Proof. Firstly we will give the definition of the problem k-UNIFORM HYPERGRAPH VERTEX COVER. We remind the reader that a k-uniform hypergraph $H = (V, E)$ consists of a set of vertices V, and a collection E, of k-element subsets of V called hyperedges. A vertex cover of H is a subset of vertices S such that every hyperedge in E has a non-empty intersection with S. The k-UNIFORM HYPERGRAPH VERTEX

COVER problem is the problem of finding a minimum size vertex cover in a k-uniform hypergraph. Let k be a constant. Assuming the unique games conjecture, Khot and Regev (2008) proved a $k - \epsilon$ inapproximability result for this problem and furthermore, Dinur et al. (2005) proved that it is NP-hard to approximate k-UNIFORM HYPERGRAPH VERTEX COVER within a factor of $(k - 1 - \epsilon)$, for any small ϵ.

To prove the proposition, we establish that this problem is a special case of CMIC. Indeed, given a hypergraph H, we can view the vertices as workers and the hyperedges as tasks. We say that a worker w is capable of executing a task t, if the hyperedge that corresponds to t contains the vertex that corresponds to w. Let every worker have unit bid and contribution. Additionally, let the demand of every task be equal to 1. Evidently, $\Delta = k$. Then, the existence of a vertex cover of H of a certain size is equivalent to the existence of a set of workers, that are able to execute all tasks at a cost that equals the size of the cover.

□

Finally, we note that the algorithms of Section 4.4 can in principle also be applied for CMDC with no loss in the approximation factor, but at the expense of a much higher running time (doubly exponential in m, which still remains polynomial, as long as the number of tasks is a constant).

Chapter 5

Budget-feasible Mechanisms for Divisible Agents and Multiple Levels of Service

5.1 Introduction

In Chapter 4, we focused on procurement auctions in which the primary objective of the auctioneer was to hire bidders in a cost-effective manner, with the implicit aim of minimizing the payments to the bidders. However, in that scenario, there was no strict budget constraint for the auctioneer and the feasibility of the allocation scheme depended solely on the natural constraints of the underlying combinatorial optimization problem.

In this chapter we examine a different class of single-parameter procurement auctions where the auctioneer is constrained by a strict budget, meaning that the total *payments* made by the auctioneer to the bidders cannot exceed a certain threshold. In this context, the auctioneer is seeking to hire services from a group of bidders, each of whom has a private cost parameter that represents their cost for providing the service. It is assumed that the bidders may strategically misreport their costs. Additionally, each hiring scheme selected by the auctioneer has a known value, which is not subject to deception by the bidders. The goal of the auctioneer is to select a hiring scheme that attains a good chunk of the optimal value compared to the objective of the underlying combinatorial optimization problem (a variant of Knapsack), while also ensuring that the total payments made by the auctioneer do not exceed the predetermined budget threshold.

Given the strategic behavior by the bidders, our focus in this mechanism design problem is to devise mechanisms that incentivize the bidders to report their true costs. Since Singer (2010) proposed the original problem, concerning an additive value function and a simple binary environment, where bidders could either be hired or not, a large body of work has been devoted to obtaining improved results on the original model, as well as proposing a number of extensions. These extensions include additional feasibility constraints, richer objectives, and additional assumptions. In this chapter, we concentrate on two settings that have received relatively less attention in the literature. In both settings below, the auctioneer is allowed to partially hire the services offered by each bidder.

1. **Divisible agents:** In this scenario, the auctioneer is dealing with bidders who offer a divisible service, and each bidder's value is known to the auctioneer. The auctioneer has the flexibility to hire each bidder for any fraction of the service. The primary goal is to maximize the total value while ensuring that the total payments do not exceed the budget constraint. It is worth noting that this problem is the fractional relaxation of the one introduced by Singer (2010). Anari et al. (2014) were the first to study the fractional problem. However, in their work they employed a *large market* assumption, which, in the context of budget-feasible mechanism design, roughly means that the cost of each bidder for their entire service is insignificant compared to the budget of the auctioneer. Very recently, Klumper and Schäfer (2022) revisited this problem without the "large market" assumption but rather under the standard assumption of the literature that the cost of each bidder for the entire service is below the budget. They came up with a new truthful and budget-feasible mechanism that is tailored to this specific scenario with an improved approximation guarantee compared to the state of the art in the literature.[1]

2. **Agents with multiple levels of service:** In this scenario, each bidder offers a service with a range of multiple levels, and the auctioneer can choose to hire some, but not all, of the available levels of service. Furthermore, each bidder's value function is concave, meaning that the marginal value of each level of service is non-increasing. The main objective is once again to maximize the total value of the auctioneer while staying within the budget. This type of scenario was first

[1] All previous works addressing the problem of indivisible agents achieve an approximation guarantee for the underlying fractional knapsack instance, which in turn guarantees an approximation to the knapsack problem. As such, these guarantees apply to the setting examined by Klumper and Schäfer (2022) and to the work presented in this chapter.

introduced by Chan and Chen (2014), but in their study, they assumed that the costs of individual levels are arbitrary, meaning that it could be the case that the auctioneer only affords to hire a single level of service of a single bidder. Chan and Chen (2014) were able to obtain randomized, truthful, and budget-feasible mechanisms for this setting, with approximation guarantees that depend on the number of bidders. The difference with our setting is that we assume that the auctioneer has enough available budget to hire each individual bidder for all levels of service offered, if she wishes to do so.

The two classes of procurement auctions have a number of practical applications in various domains. For example, in the context of crowdsourcing, the divisible setting would be useful to model the time availability of a worker as the fraction which the auctioneer can hire. Moreover, these types of auctions can also be applied to other industries, such as transportation and logistics, where the delivery of goods and services can be broken down into multiple levels of service. For instance, in the transportation industry, the first level of service can represent the basic delivery service, while the higher levels can represent more premium and specialized services, such as express delivery or temperature-controlled shipping. The auctioneer can then choose to hire each bidder up to an available level of service, not necessarily the best offered, based on the budget constraint and the value of the services provided.

Contribution In this chapter, we propose truthful and budget-feasible mechanisms both for the divisible agents and the multiple-level scenarios. Specifically, in Section 5.3, we present a deterministic mechanism for the divisible setting, that is truthful, budget-feasible, and achieves a $1 + \sqrt{2} \approx 2.41$ approximation. This mechanism represents an adaptation of the 3-approximate posted-price mechanism proposed by Gravin et al. (2020) for the case of indivisible agents to our fractional procurement auction model. Our proposed mechanism improves upon the existing state-of-the-art mechanism of Klumper and Schäfer (2022), which achieves a $1 + \phi \approx 2.62$ approximation. Next, in Section 5.4, we introduce a deterministic mechanism that is truthful, budget-feasible, and achieves a $2 + \sqrt{3} \approx 3.73$ approximation for the multiple levels of service scenario. For this setting, no constant-factor approximation mechanism was previously known[2].

[2]Note that this mechanism is not comparable to the results of Chan and Chen (2014), since, in their setting, they assume that the auctioneer can only afford to hire a single bidder on their own for a single level.

5.1.1 Related Work

The design of truthful budget-feasible mechanisms was introduced by Singer (2010), who gave a deterministic mechanism for additive valuation functions with an approximation guarantee of 5, along with a lower bound of 2 for deterministic mechanisms. This guarantee was subsequently improved to $2+\sqrt{2} \approx 3.41$ by Chen et al. (2011), who also provided a lower bound of 2 for randomized mechanisms and a lower bound of $1+\sqrt{2} \approx 2.41$ for the deterministic case. Gravin et al. (2020) gave a 3-approximate deterministic mechanism, which is the best guarantee for deterministic mechanisms, known to date. Regarding randomized mechanisms, Gravin et al. (2020) settled the question by providing a 2-approximate randomized mechanism, matching the lower bound of Chen et al. (2011). Finally, the question has also been settled under the large market assumption by Anari et al. (2014), who extended their $\frac{e}{e-1} \approx 1.58$ mechanism for the fractional setting we have mentioned above to the indivisible agents case.

For indivisible agents, the problem has also been extended to richer valuation functions. This line of inquiry also started by Singer (2010), who gave a randomized algorithm with an approximation guarantee of 112 for a monotone submodular objective. Once again, this result was improved by Chen et al. (2011) to a 7.91 guarantee, and the same authors devised a deterministic mechanism with a 8.34 approximation. Subsequently, the bound for randomized mechanisms was improved by Jalaly and Tardos (2018) to 5. Very recently, Balkanski et al. (2022) proposed a new method of designing mechanisms that goes beyond the sealed-bid auction paradigm we are following throughout this dissertation. Instead, Balkanski et al. (2022) presented mechanisms in the form of deterministic clock auctions[3] and, for the monotone submodular case, present a deterministic clock auction which achieves a 4.75 guarantee.

Beyond monotone submodular valuations, it becomes much harder to obtain truthful mechanisms with small constants as approximation guarantees. Namely, for non-monotone submodular objectives the best sealed-bid randomized mechanism is due to Amanatidis et al. (2019) and its approximation guarantee is 505. While there is no known deterministic sealed-bid mechanism with a constant approximation for this class of valuations, Balkanski et al. (2022) have provided a deterministic clock auction mechanism which is 64-approximate. Richer valuations that have been studied are XOS valuation functions (see Amanatidis et al. (2017); Bei et al. (2017)) and subadditive valuation functions (see Balkanski et al. (2022); Bei et al. (2017); Dobzinski et al. (2011)). For subadditive valuation functions, no mechanism achieving a constant approximation is known. However, Bei et al. (2017) have proved that such a mechanism

[3] We give the definition of a clock auction and briefly discuss clock auctions in Section 6.2.3.

should exist, using a non-constructive argument. Finding such as mechanism is an intriguing open question.

Finally, other settings that have been studied include environments with underlying feasibility constraints, such as downward-closed environments (Amanatidis et al., 2016) and matroid constraints (Leonardi et al., 2017) and auction environments in which the auctioneer wants to get a set of heterogeneous tasks done. In this setting, each task requires that the hired agent has a certain skill, see Goel et al. (2014).

5.2 Model and Preliminaries

We consider a scenario where a procurement auction is being held by an auctioneer who has a budget $B > 0$. The auction is for a service that will be provided by a set of bidders $N = \{1, \ldots, n\}$. Each bidder $i \in N$ has a private cost parameter $c_i > 0$, that represents their true cost for providing the service. In this extended model of budget-feasible mechanism design, the auctioneer is allowed to hire bidders partially and the value he derives from each hiring scheme is public information. Bidders cannot deceive the auctioneer about this value, and it can be considered as data available on a website or a public forum.

A deterministic auction mechanism $\mathcal{M}$ in this setting consists of an allocation algorithm $\mathbf{x} : \mathbb{R}^n_{\geq 0} \mapsto \mathbb{R}^n_{\geq 0}$, coupled with a payment rule $\mathbf{p} : \mathbb{R}^n_{\geq 0} \mapsto \mathbb{R}^n_{\geq 0}$. To begin with, the auctioneer collects a vector of bids from the bidders, denoted by $\mathbf{b} = (b_i)_{i \in N}$. Here, b_i denotes the cost declared by bidder $i \in N$, which may differ from their true cost c_i. The auctioneer then determines an allocation (hiring scheme) $\mathbf{x}(\mathbf{b}) = (x_1(\mathbf{b}, \ldots, x_n(\mathbf{b})))$, where $x_i(\mathbf{b}) \in \mathbb{R}_{\geq 0}$ is the allocation decision for bidder i. Then, the auctioneer determines a vector of payments $\mathbf{p}(\mathbf{b}) = (p_1(\mathbf{b}), \ldots, p_n(\mathbf{b}))$, where $p_i(\mathbf{b})$ is the payment bidder i will receive for their service.

We assume that bidders have quasi-linear utilities, i.e., for a deterministic procurement auction mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$, the utility of bidder $i \in N$ for a profile $\mathbf{b}$ is $u_i(\mathbf{b}) = p_i(\mathbf{b}) - c_i x_i(\mathbf{b})$. We are interested in designing mechanisms that satisfy three properties for any true profile $\mathbf{c}$ and any declared profile $\mathbf{b}$:

- *Individual rationality:* All bidders $i \in N$ should have a utility of at least 0, i.e., $u_i(\mathbf{b}) \geq 0$.

- *Budget Feasibility:* The sum of all payments made by the auctioneer, $\sum_{i \in N} p_i(\mathbf{b})$, should not exceed the budget B.

- *Truthfulness:* All bidders $i \in N$ should have no incentive to lie about their true cost, i.e., $u_i(c_i, \mathbf{b}_{-i}) \geq u_i(\mathbf{b})$.

As in Chapters 2 and 4, this is a single-parameter environment and the characterization of Myerson (1981) applies. Therefore it is sufficient to focus on the class of mechanisms with monotone allocation algorithms. Similarly to the definition of Chapter 4, an allocation algorithm is monotone if for every bidder $i \in N$, every profile $\mathbf{b}$, and all $b'_i \leq b_i$, we have $x_i(\mathbf{b}) \leq x_i(b'_i, \mathbf{b}_{-i})$. However it is a slightly more involved task to compute the payments of the bidders in this case. The reason behind this is the fact that the procurement auction settings we study in this chapter, are not *binary* single-parameter domains i.e., bidders may be selected in an allocation partially. Fortunately, this case is also handled in the work of Myerson (1981). Below we present Myerson's Lemma in its complete form, which is commonly referred to in the literature as Myerson's *payment identity.*

Lemma 5.2.1. *Given a monotone allocation algorithm* $\mathbf{x}$*, there is a unique payment rule* $\mathbf{p}$*, such that* $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ *is a truthful and individually rational mechanism. For every profile* $\mathbf{b}$*, the payment to bidder* $i \in N$ *is*

$$p_i(\mathbf{b}) = b_i x_i(\mathbf{b}) + \int_{b_i}^{\infty} x_i(y, \mathbf{b}_{-i}) dy. \tag{5.1}$$

Given an allocation $\mathbf{x}$, the total value the auctioneer obtains is denoted by $V(\mathbf{x})$. The exact form of this function depends on the setting we are studying and will be formalized in the paragraphs that follow.

In the sequel, we will consistently refer to the true valuation profile of bidders as $\mathbf{c}$, rather than $\mathbf{b}$. This is because we will be exclusively presenting provably monotone allocation algorithms and payments to bidders will be determined by Equation (5.1). In the remainder of this chapter, we present the two settings we study.

5.2.1 Divisible Agents

In the "divisible agents" setting, we are allowed to hire each bidder for any arbitrary percentage of their full service. To be precise, given a profile $\mathbf{c}$, we have that $\mathbf{x}(\mathbf{c}) \in [0, 1]^n$. Moreover, for each bidder $i \in N$, the auctioneer has access to a *publicly known* parameter $v_i > 0$, which represents how valuable the bidder is, should she get hired entirely. Finally, we assume that $\mathbf{c}$ is such that each bidder can be hired fully on their own, i.e., for all $i \in N$, it holds that $c_i \leq B$, as in the work of Klumper and Schäfer

(2022). Note that this assumption is much weaker than the large market assumption of Anari et al. (2014).

In this scenario, $V(\mathbf{x}(\mathbf{c})) = \sum_{i \in N} v_i x_i(\mathbf{c})$, is the total value the auctioneer enjoys under the allocation $\mathbf{x}(\mathbf{c})$. To evaluate the performance of a monotone and budget-feasible mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$, we compare its respective $V(\mathbf{x}(\mathbf{c}))$ with

$$OPT_F(\mathbf{c}) := \max \left\{ \sum_{i=1}^{n} v_i x_i \mid (x_i)_{i=1}^n \in [0,1]^n, \quad \sum_{i=1}^{n} c_i x_i \leq B \right\},$$

which is precisely the fractional relaxation of Knapsack. Oftentimes, this problem is also mentioned in the literature as the setting with *fractional additive* valuations. Thus, for the divisible agents scenario, we say that a mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ is α-approximate if, given an $\alpha > 1$, and for every profile $\mathbf{c}$, it holds that $V(\mathbf{x}(\mathbf{c})) \geq \frac{1}{\alpha} OPT_F(\mathbf{c})$.

5.2.2 Multiple Levels of Service

In this alternative extension of the classical budget-feasible mechanism design model, each bidder $i \in N$ is offering k levels of service. Analogously to the "divisible agents" case, we assume that the auctioneer can afford hiring each bidder entirely on their own, i.e., given a profile $\mathbf{c}$, it holds that $c_i k \leq B$ for all $i \in N$. Moreover, for each bidder $i \in N$, we associate a value function $v_i : \{0, \dots, k\} \mapsto \mathbb{R}_+$ with $v_i(0) = 0$. Each function $v_i(\cdot)$ is *concave* i.e., for $j = 1, \dots, k-1$, it holds that $v_i(j) - v_i(j-1) \geq v_i(j+1) - v_i(j)$. For notational convenience, let $m_i(j) := v_i(j) - v_i(j-1)$. Thus, given a profile $\mathbf{c}$, the total value the auctioneer enjoys under an allocation is $V(\mathbf{x}(\mathbf{c})) = \sum_{i \in N} v_i(x_i(\mathbf{c}))$, a *concave separable* function. We will be measuring the performance of a mechanism comparing $V(\mathbf{x}(\mathbf{c}))$ with the underlying non-strategic combinatorial optimization problem, which in this case is

$$OPT_I^k(\mathbf{c}) := \max \left\{ \sum_{i=1}^{n} v_i(x_i) \mid (x_i)_{i=1}^n \in \{0, \dots, k\}^n, \quad \sum_{i=1}^{n} c_i x_i \leq B \right\},$$

a problem commonly referred to as Bounded Knapsack[4] (e.g. see Martello and Toth (1990) for a classification of knapsack problems). For the "multiple levels of service" scenario, we say that a mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ is α-approximate if, given an $\alpha > 1$ and for every profile $\mathbf{c}$, it holds that $V(\mathbf{x}(\mathbf{c})) \geq \frac{1}{\alpha} OPT_I^k(\mathbf{c})$.

It is often the case in the literature for the fractional relaxation of an integral knapsack problem (a quantity that can be computed in polynomial or pseudopolynomial

[4] Note that for $k = 1$, this becomes the well-known 0-1 Knapsack problem.

time) to be used as a *proxy* in order to approximate the integral problem (which is NP-Hard). To be precise, given a profile $\mathbf{c}$, we define the *fractional relaxation* of the Bounded Knapsack problem for k units as

$$OPT_F^k(\mathbf{c}) := \max\left\{\sum_{i=1}^{n}(v_i(\lfloor x_i\rfloor) + m_i(\lceil x_i\rceil)(x_i - \lfloor x_i\rfloor)) \mid (x_i)_{i=1}^n \in [0,\dots,k]^n, \quad \sum_{i=1}^{n} c_i x_i \le B\right\},$$

Note that $OPT_F^1(\mathbf{c})$ is the fractional relaxation of knapsack, defined as $OPT_F(\mathbf{c})$ in Section 5.2.1. Thus, $OPT_F^k(\mathbf{c})$ inherits the well-known properties of its one-dimensional analogue.

Fact 5.2.1. *Let $(\mathbf{c}, (v_1(\cdot), \dots, v_n(\cdot)), B)$ be an instance of the fractional relaxation of a Bounded Knapsack instance for k units. Then, Algorithm 5.1 returns a vector $\mathbf{x}^*$ which is a solution to $OPT_F^k(\mathbf{c})$ in time $O(kn\log(kn))$.*

Algorithm 5.1: An Algorithm for Fractional k-Bounded Knapsack

```
▷ Input: An instance c, (v_1(·), ..., v_n(·)), B.
Initialize an empty list of kn elements.
for i ∈ n do
    for j = 1, ..., k do
        Add m_i(j)/c_i to the list.
Sort the elements of the list in decreasing order.
Let x* = 0 and j = 1.
while ∑_{i∈N} c_i x*_i ≤ B do
    Let ℓ ∈ {1, ..., n} be the index corresponding to the j-th
      marginal-value-per-cost in the list.
    Set x*_ℓ = x*_ℓ + min((B − ∑_{i∈N} c_i x*_i)/c_ℓ, 1).
    Set j = j + 1.
Return x*.
```

Observe that Algorithm 5.1 assigns a fractional x_i to at most one $i \in \{1, \dots, n\}$. Naturally, it also holds that $OPT_F^k(\mathbf{c}) \ge OPT_I^k(\mathbf{c})$.

5.3 A Truthful and Budget-Feasible Mechanism for Divisible Agents

In this section, we present a novel mechanism for divisible agents that is both truthful and individually rational while also being budget-feasible. Our mechanism achieves a $1+\sqrt{2}$ approximation ratio, which improves upon the state-of-the-art mechanism proposed by Klumper and Schäfer (2022), which achieved a $(1+\phi)$ approximation ratio. Our mechanism is a posted-price mechanism, where appropriate take-it-or-leave-it prices are offered to bidders in the final step. A similar approach was taken by Gravin et al. (2020) to obtain a 3-approximate mechanism in the case when agents are indivisible, as in the classical setting of Singer (2010). The added flexibility of being able to allocate fractionally allows our mechanism to achieve even better results.

In order to determine the appropriate posted prices to be offered, it is important to first exclude bidders with a low value-per-cost ratio. To achieve this, we propose an adapted pruning allocation algorithm inspired by the work of Gravin et al. (2020) in Section 5.3.1. This mechanism is used to obtain a provisional allocation that can be further modified based on information elicited by posted prices, as detailed in Section 5.3.2.

We establish that the proposed composition of mechanisms is truthful, individually rational, and budget-feasible, with an improved performance guarantee.

5.3.1 A Pruning Algorithm for Divisible Agents

We first present the pruning algorithm of Gravin et al. (2020). This algorithm is part of the 3-approximate mechanism Gravin et al. (2020) present for the case of indivisible agents. However, it will be a useful starting point for the setting of divisible agents as well.

Given a profile $\mathbf{c}$, this algorithm computes an allocation $\bar{\mathbf{x}}(\mathbf{c})$ and a positive quantity $r(\mathbf{c})$, which we refer to as the *rate*. We assume that the bidders are initially relabeled by their decreasing value-per-cost ratio, i.e. $\frac{v_1}{c_1} \geq \frac{v_2}{c_2} \geq \cdots \geq \frac{v_n}{c_n}$.

Algorithm 5.2: The Pruning Algorithm of Gravin et al. (2020)

▷ **Input:** A profile $\mathbf{c}$ such that $\frac{v_1}{c_1} \geq \frac{v_2}{c_2} \geq \cdots \geq \frac{v_n}{c_n}$

Let $r := \frac{1}{B}\max\{v_i \mid i \in N\}$.
For $i \in N$, set $\bar{x}_i = 1$ if $\frac{v_i}{c_i} \geq r$ and $\bar{x}_i = 0$ otherwise.
Let $\ell := \arg\max\{i \mid \bar{x}_i = 1\}$.
while $rB < \sum_{i=1}^{\ell} v_i - \max\{v_i \mid i = 1, \ldots, \ell\}$ **do**
    Continuously increase rate r.
    If $\frac{v_\ell}{c_\ell} \leq r$, set $\bar{x}_\ell = 0$ and $\ell = \ell - 1$.
return $(r, \bar{\mathbf{x}})$

Notice that $\bar{\mathbf{x}} \neq \mathbf{0}$, since the condition of the while-loop is violated should $\ell = 1$ and, therefore, the provisional allocation is never empty.

Suppose now that we consider Algorithm 5.2 to be the allocation algorithm of a mechanism. It is not hard to observe that this algorithm is monotone and this property is naturally proven by Gravin et al. (2020). In fact, an even stronger property holds and we present it in Lemma 5.3.1.

Lemma 5.3.1 (Implied by Lemma 3.1 of Gravin et al. (2020))**.** *Let* $\mathbf{c}$ *be a profile. Fix a bidder* $i \in N$ *such that* $\bar{x}_i(\mathbf{c}) = 1$*. Then, for all* c'_i *such that* $\bar{x}_i(c'_i, \mathbf{c}_{-i}) = 1$*, it holds that*

1. $\bar{\mathbf{x}}(c'_i, \mathbf{c}_{-i}) = \bar{\mathbf{x}}(\mathbf{c})$.

2. $r(c'_i, \mathbf{c}_{-i}) = r(\mathbf{c})$.

Lemma 5.3.1 asserts that a bidder who has been selected in the provisional allocation cannot alter the outcome of Algorithm 5.2 by declaring a signal (cost) unilaterally that keeps her in the set of hired bidders. As a result, Algorithm 5.2 can be used as a preliminary filtering step to eliminate unappealing bidders in mechanism composition schemes. If the subsequent allocation algorithm, which given a profile $\mathbf{c}$, takes $(r(\mathbf{c}), \bar{\mathbf{x}}(\mathbf{c}))$ as input, is monotone, then the resulting allocation algorithm is monotone as well.

In the remaining part of this section, for the sake of brevity, we shall omit mentioning the dependence on profile $\mathbf{c}$. Furthermore, we shall frequently make references to the bidder with the highest value in the provisional allocation. To facilitate our analysis, we define $i^* = \arg\max_{i \in N}\{v_i \mid \bar{x}_i = 1\}$. The properties that follow, proved by Gravin et al. (2020), will be valuable in our analysis.

Lemma 5.3.2 (Lemma 3.2 of Gravin et al. (2020))**.** *Let* $(r, \bar{\mathbf{x}})$ *be the output of Algorithm 5.2 for profile* $\mathbf{c}$*. Then*

1. $c_i \leq \frac{v_i}{r} \leq B$*, for all* $i \in N$ *with* $\bar{x}_i = 1$.

2. $V(\bar{\mathbf{x}}) - v_{i^*} \leq rB < V(\bar{\mathbf{x}})$.

3. $OPT_F \leq V(\bar{\mathbf{x}}) + r \cdot (B - \sum_{i \in N} c_i \bar{x}_i)$.

Property 1 essentially affirms that any bidder included in the provisional allocation will accept a posted price of $\frac{v_i}{r}$. Property 2 of Lemma 5.3.2, not only establishes an upper and lower bound for r but also directly implies that, for a provisional allocation $\bar{x}$, there exists a $z \in [0, 1)$ such that

$$V(z, \bar{\mathbf{x}}_{-i^*}) = rB. \tag{5.2}$$

Therefore, we can introduce an additional step after Algorithm 5.2, allowing us to allocate fractionally by reducing the provisional allocation of the highest-valued bidder i^* from 1 to z. This additional flexibility yields Equation (5.2) which, in turn, leads us to a restatement of Property 3 of Lemma 5.3.2.

Observation 5.3.1. *Let* $(r, \bar{\mathbf{x}})$ *be the output of Algorithm 5.2 for profile* $\mathbf{c}$, *and let* $z \in [0, 1)$ *be such that Equation (5.2) holds. Then,*

$$OPT_F \leq (1 - z)v_{i^*} + 2V(z, \bar{\mathbf{x}}_{-i^*}) - r \sum_{i \in N} c_i \bar{x}_i. \tag{5.3}$$

Proof. We can rewrite Property 3 of Lemma 5.3.2 as

$$\begin{aligned} OPT_F &\leq V(\bar{\mathbf{x}}) + r \cdot \left(B - \sum_{i \in N} c_i \bar{x}_i \right) \\ &= V(\bar{\mathbf{x}}) + rB - r \sum_{i \in N} c_i \bar{x}_i \\ &= (1 - z)v_{i^*} + V(z, \bar{\mathbf{x}}_{-i^*}) + V(z, \bar{\mathbf{x}}_{-i^*}) - r \sum_{i \in N} c_i \bar{x}_i. \end{aligned}$$

The second equality is due to the definition of $V(\mathbf{x})$ for a vector $\mathbf{x} \in [0, 1]^n$ and due to Equation (5.2). □

Throughout the rest of our analysis, Equation (5.3) will serve as the central point and guide our selection of appropriate posted prices in the mechanism of Section 5.3.2.

5.3.2 An Adaptive Posted-price Mechanism

In this section we a truthful, individually rational and budget-feasible mechanism that achieves a $(1 + \sqrt{2})$ approximation ratio when agents are divisible.

The mechanism first applies the pruning algorithm of Gravin et al. (2020) and potentially singles out the highest-valued bidder in the provisional allocation. Then, depending on the comparison between the value of this bidder and the total value achieved by the other provisionally allocated bidders, a different allocation scheme is selected.

Intuitively, if the value of the highest-valued bidder is high enough, then this bidder is the only one selected. Alternatively, if the value of the highest-valued bidder is low compared to the total value of the others, then the provisional allocation is implemented with further pruning, as presented in Section 5.3.1. Importantly, these two cases are not mutually exclusive. In fact, there is a "middle" case in which the highest-valued bidder is offered a posted price and, depending on his answer, may be allocated entirely in the best case. The final component of this mechanism is that the other bidders in the provisional allocation may also be asked to accept a posted price when the highest-valued bidder accepts, and they may be allocated a tailored fraction. The usefulness of this posted price scheme is due to the fact that, for a profile $\mathbf{c}$, a posted price rejection implies a lower bound on the cost $c_i > 0$ of a bidder $i \in N$, a crucial piece of information that we can exploit using Equation (5.3).

Let $\alpha \in (0, 1)$ be a parameter, the value of which we will specify later. Moreover, for $x \in [0, 1]$ define

$$\lambda(x) := \frac{2}{1 - x + \alpha} - \frac{\alpha}{1 - x \cdot (1 + \alpha)}$$

In Mechanism 5.3 this scheme is described. The main result of this section is Theorem 5.3.1.

Theorem 5.3.1. *Mechanism 5.3 is truthful, individually rational and budget-feasible. Moreover, For $\alpha = \sqrt{2} - 1$ it is $(1 + \sqrt{2})$-approximate.*

Note that 5.3 clearly runs in polynomial time. We will prove Theorem 5.3.1 using the three Lemmas that follow. We start by proving that the allocation algorithm of Mechanism 5.3 is monotone.

Lemma 5.3.3. *The allocation algorithm of Mechanism 5.3 is monotone, i.e., for every profile $\mathbf{c}$, every bidder $i \in N$ and every $c'_i \leq c_i$ it holds that $x_i(c'_i, \mathbf{c}_{-i}) \geq x_i(\mathbf{c})$.*

Proof. Fix a bidder $i \in N$ with $x_i > 0$ and a $c'_i \leq c_i$. Recall that, By Lemma 5.3.1, a unilateral deviation $c'_i \leq c_i$ of any bidder does not alter r or the provisional allocation $\bar{\mathbf{x}}$ of Algorithm 5.2, and as a result, it does not affect the quantity z either, in step 2. We analyze each one of the cases named in lines 4, 6, 10 and 16 separately.

(a) In this case bidder i^* is the only bidder with $x_i > 0$ and no improved unilateral deviation of i^* can change that.

(b) Here, the output of Algorithm 5.2 with further pruning $(z, \bar{\mathbf{x}}_{-i^*})$ is selected and, once again, no bidder can affect this outcome unilaterally.

Algorithm 5.3: A Posted-price Mechanism for Divisible Agents

▷ **Input:** A profile $\mathbf{c}$ such that $\frac{v_1}{c_1} \geq \frac{v_2}{c_2} \geq \cdots \geq \frac{v_n}{c_n}$ and a parameter $\alpha \in (0,1)$

Obtain $(r, \bar{\mathbf{x}})$ by running Algorithm 5.2 for profile $\mathbf{c}$.
Let $z \in [0,1)$ be s.t. Equation (5.2) holds.
Let $i^* = \arg\max_{i \in N}\{v_i \mid \bar{x}_i = 1\}$ and $T = \{i \in N \setminus \{i^*\} \mid \bar{x}_i = 1\}$.
if $v_i^* \geq \frac{2}{1+z+\alpha} V(z, \bar{\mathbf{x}}_{-i^*})$ **then**
  `// Case (a): High-valued` i^*
  Set $x_{i^*} = 1$ and $x_i = 0 \quad \forall i \in N \setminus \{i^*\}$.
else if $v_i^* \leq \frac{a}{1-z} V(z, \bar{\mathbf{x}}_{-i^*})$ **then**
  `// Case (b): Low-valued` i^*
  Set $\mathbf{x} = (z, \bar{\mathbf{x}}_{-i^*})$.
else
  Let $B_{i^*} = \min\left\{\frac{v_{i^*}}{r}, \frac{v_{i^*}(1-z)-\alpha V(z,\bar{\mathbf{x}}_{-i^*})}{r(1-z)\lambda(z)(1-z(1+\alpha))}\right\}$.
  if $c_{i^*} \leq B_{i^*}$ **then**
    `// Case (c):` i^* `accepts posted price`
    Set $x_{i^*} = 1$.
    for $i \in T$ **do**
      Let $B_i = \min\left\{\frac{v_i}{r}, \sqrt{2+\alpha} \cdot \frac{v_i}{\sum_{j \in T} v_j}\left(B - (1-z)B_{i^*} - z\frac{v_{i^*}}{r}\right)\right\}$.
      if $c_i \leq B_i$ **then**
        Set $x_i = \frac{1}{\sqrt{2+\alpha}}$.
      else
        Set $x_i = 0$.
  else
    `// Case d:` i^* `rejects posted price`
    Set $\mathbf{x} = (z, \bar{\mathbf{x}}_{-i^*})$.
Allocate $\mathbf{x}$ and set payments according to Equation (5.1).

(c) Since i^* has met her posted price B_{i^*}, she is assigned an allocation of $x_i = 1$ and this remains true for all $c_i' \le c_i$. Similarly, a subset of bidders in T are allocated $\frac{1}{\sqrt{2+\alpha}}$ and this remains true for every $c_i' \le c_i$.

(d) For bidders in T the argument of Case (b) applies. For i^* we observe that since $c_{i^*} > B_{i^*}$, they are (potentially) allocated $z \in [0, 1)$ and should they unilaterally declare a $c_{i^*}' \le B_{i^*}$ their allocation will increase to 1, as per Case 3.

□

Since the final payments can be described by Equation (5.1) we conclude that Mechanism 5.3 is truthful and individually rational.

We now prove that Mechanism 5.3 is also budget-feasible.

Lemma 5.3.4. *Mechanism 5.3 is budget-feasible, i.e. for every profile* $\mathbf{c}$ *it holds that* $\sum_{i \in N} p_i(\mathbf{c}) \le B$.

Proof. Payments are assigned according to Equation (5.1). We analyze each one of the four cases (corresponding to lines 4, 6, 10 and 18 of the allocation algorithm) separately.

(a) Bidder i^* is the only bidder being fully allocated and thus asked to pay. The payment identity implies that, bidder i^* should pay the highest $c_{i^*}' \ge c_i$ so that she remains a part of the provisional allocation, which in any case is at most B by using the assumption that the auctioneer can always hire one bidder entirely.

(b) For every bidder $i \in N$ with $\bar{x}_i > 0$ let $c_i' \ge c_i$ be the highest cost she can declare and remain a provisional winner under $(c_i', \mathbf{c}_{-i})$. Thus, by Property 1 of Lemma 5.3.2 we have that $c_i' \le \frac{v_i}{r}$, for all i and by using the fact that the allocation is $(z, \bar{\mathbf{x}}_{-i^*})$ and since, by Lemma 5.3.1 no bidder can unilaterally change the provisional allocation, we obtain

$$\begin{aligned}
\sum_{i \in N} p_i(\mathbf{c}) &= p_{i^*}(\mathbf{c}) + \sum_{i \in T} p_i(\mathbf{c}) \\
&= z c_{i^*} + \int_{c_{i^*}}^{\frac{v_{i^*}}{r}} z dy + \sum_{i \in T} \left(c_i + \int_{c_i}^{\frac{v_i}{r}} dy \right) \\
&= z \frac{v_{i^*}}{r} + \sum_{i \in T} \frac{v_i}{r} = \frac{V(z, \bar{\mathbf{x}}_{-i^*})}{r} = B.
\end{aligned}$$

and the last equality follows by Equation (5.2).

(c) For bidders $i \in T$, the critical bid is exactly B_i and this is the price they pay for their fixed fractional allocation since

$$p_i(\mathbf{c}) = \frac{c_i}{\sqrt{2+\alpha}} + \int_{c_i}^{B_i} \frac{1}{\sqrt{2+\alpha}} dy = \frac{B_i}{\sqrt{2+\alpha}}. \tag{5.4}$$

Moreover, by the definition of the case, i^* accepts the posted price of B_{i^*} and is assigned an allocation of 1. To compute the payment of i^*, note that if bidder i^* declares a cost $c'_{i^*} > B_{i^*}$, her allocation will become equal to z (as she will fall into Case d), as long as $c'_{i^*} \leq \frac{v_{i^*}}{r}$ (by Property 1 of Lemma 5.3.2). Therefore, we can breakdown the payment of i^* as follows:

$$p_{i^*}(\mathbf{c}) = c_{i^*} + \int_{c_{i^*}}^{B_{i^*}} dy + \int_{B_{i^*}}^{\frac{v_{i^*}}{r}} z dy = (1-z)B_{i^*} + z\frac{v_{i^*}}{r}. \tag{5.5}$$

We now proceed with the upper bound on total payments. Let U be the (potentially empty) subset of bidders in T that have accepted their posted price, i.e. for all $i \in U$ it holds that $c_i \leq B_i$. Then

$$\begin{aligned}
\sum_{i \in N} p_i(\mathbf{c}) &= p_{i^*}(\mathbf{c}) + \sum_{i \in U} p_i(\mathbf{c}) \\
&= p_{i^*}(\mathbf{c}) + \frac{1}{\sqrt{2+\alpha}} \sum_{i \in U} B_i \\
&\leq p_{i^*}(\mathbf{c}) + \sum_{i \in U} \cdot \frac{v_i}{\sum_{j \in T} v_j} \left(B - (1-z)B_{i^*} - z\frac{v_{i^*}}{r} \right) \\
&\leq p_{i^*}(\mathbf{c}) + B - (1-z)B_{i^*} - z\frac{v_{i^*}}{r} = B.
\end{aligned}$$

The second equality follows by Equation (5.4) and the first inequality is true because of the definition of B_i. Finally, the last equality is due to Equation (5.5).

(d) The analysis here is identical to Case (b).

□

The final component for the proof of Theorem 5.3.1 is its approximation guarantee.

Lemma 5.3.5. *Let* $\mathbf{c}$ *be a profile and* $\mathbf{x}(\mathbf{c})$ *be the final allocation of Mechanism 5.3. Then for* $\alpha \in [\sqrt{2}-1, 1)$ *it holds that* $V(\mathbf{x}(\mathbf{c})) \geq \frac{1}{2+\alpha} OPT_F(\mathbf{c})$.

Before presenting the proof of Lemma 5.3.5, we state and prove an auxiliary technical lemma which will be useful for our analysis.

Lemma 5.3.6. *Let* $\mathbf{c}$ *be a profile and let* $T = \{i \in N \setminus \{i^*\} \mid \bar{x}_i(\mathbf{c}) = 1\}$. *Then:*

$$rB - (1-z)B_{i^*}r - zv_{i^*} \geq \frac{\frac{2}{1-z+\alpha}}{\lambda(z)} \sum_{i \in T} v_i - \frac{v_{i^*}}{\lambda(z)}$$

Proof. By the definition of B_{i^*} we have

$$\begin{aligned} rB - (1-z)B_{i^*}r - zv_{i^*} &\geq rB - \frac{v_{i^*}(1-z) - \alpha V(z, \bar{\mathbf{x}}_{-i^*})}{\lambda(z)(1-z(1+\alpha))} - zv_{i^*} \\ &= V(z, \bar{\mathbf{x}}_{-i^*}) - \frac{v_{i^*}(1-z) - \alpha V(z, \bar{\mathbf{x}}_{-i^*})}{\lambda(z)(1-z(1+\alpha))} - zv_{i^*} \\ &= \sum_{i \in T} v_i - \frac{v_{i^*}(1-z) - \alpha(\sum_{i \in T} v_i + zv_{i^*})}{\lambda(z)(1-z(1+\alpha))} \\ &= \left(1 + \frac{\frac{\alpha}{1-z(1+\alpha)}}{\lambda(z)}\right) \sum_{i \in T} v_i - \frac{v_{i^*}}{\lambda(z)}. \end{aligned}$$

The first equality follows by Equation (5.2). The second equality follows by the fact that $V(z, \bar{\mathbf{x}}_{-i^*}) = \sum_{i \in T} v_i + zv_i^*$. Then, the last equality is just a rearrangement of terms. □

We now present the proof of Theorem 5.3.1.

Proof of Lemma 5.3.5. Recall that by Equation (5.3) we have

$$OPT_F \leq (1-z)v_{i^*} + 2V(z, \bar{\mathbf{x}}_{-i^*}) - r \sum_{i \in N} c_i \bar{x}_i.$$

(a)

$$\begin{aligned} OPT_F &\leq (1-z)v_{i^*} + 2V(z, \bar{\mathbf{x}}_{-i^*}) - r \sum_{i \in N} c_i \bar{x}_i \\ &\leq (1-z)v_{i^*} + 2v_{i^*} \frac{1+z+a}{2} = (2+\alpha)v_{i^*} = (2+\alpha)V(\mathbf{x}). \end{aligned}$$

The second inequality follows by the definition of the case we are considering, and the equality follows by observing that $x_{i^*} = 1$ and $x_i = 0$ for all $i \neq i^*$.

(b) In this case $\mathbf{x} = (z, \bar{\mathbf{x}}_{-i^*})$. Similarly to case (a)

$$\begin{aligned} OPT_F &\leq (1-z)v_{i^*} + 2V(z, \bar{\mathbf{x}}_{-i^*}) - r\sum_{i \in N} c_i \bar{x}_i \\ &\leq (1-z)\frac{\alpha}{1-z}V(z, \bar{\mathbf{x}}_{-i^*}) + 2V(z, \bar{\mathbf{x}}_{-i^*}) \\ &= (2+\alpha)V(z, \bar{\mathbf{x}}_{-i^*}) = (2+\alpha)V(\mathbf{x}). \end{aligned}$$

(c) Let U be the subset of T such that $c_i \leq B_i$. Note that

$$x_i = \begin{cases} 1 & i = i^* \\ \frac{1}{\sqrt{2+\alpha}} & i \in U \\ 0 & \text{otherwise} \end{cases}$$

The definition of the case, i.e. that $\frac{\alpha}{1-z}V(z, \bar{\mathbf{x}}_{-i^*}) < v_{i^*} < \frac{2}{1+z+\alpha}V(z, \bar{\mathbf{x}}_{-i^*})$ has a few implications that will be useful in our analysis which we now prove

Claim 5.3.1. *Under Case (c), the following are true:*

1) $\frac{\alpha}{1-z(1+\alpha)}\sum_{i \in T} v_i < v_{i^*} < \frac{2}{1-z+\alpha}\sum_{i \in T} v_i$.

2) $z < 1-\alpha$.

3) $T \neq \emptyset$.

4) $\lambda(z) > 0$.

Proof. 1) For the lower bound we have

$$\begin{aligned} v_{i^*} > \frac{\alpha \sum_{i \in T} v_i}{1-z(1+\alpha)} &\Leftrightarrow v_{i^*}\left(1 + \frac{\alpha z}{1-z(1+\alpha)}\right) > \frac{\alpha\left(\sum_{i \in T} v_i + z v_{i^*}\right)}{1-z(1+\alpha)} \\ &\Leftrightarrow v_{i^*}\left(\frac{1-z}{1-z(1+\alpha)}\right) > \frac{\alpha V(z, \bar{\mathbf{x}}_{-i^*})}{1-z(1+\alpha)} \\ &\Leftrightarrow v_{i^*} > \frac{\alpha}{1-z}V(z, \bar{\mathbf{x}}_{-i^*}) \end{aligned}$$

which holds by definition. Similarly, for the upper bound we have

$$\begin{aligned} v_{i^*} < \frac{2\sum_{i \in T} v_i}{1-z+\alpha} &\Leftrightarrow v_{i^*}\left(1 + \frac{2z}{1-z+\alpha}\right) < \frac{2\left(\sum_{i \in T} v_i + z v_{i^*}\right)}{1-z+\alpha} \\ &\Leftrightarrow v_{i^*}\left(\frac{1+z+\alpha}{1-z+\alpha}\right) < \frac{2V(z, \mathbf{x}_{-i^*})}{1-z+\alpha} \\ &\Leftrightarrow v_{i^*} < \frac{2}{1+z+\alpha}V(z, \mathbf{x}_{-i^*}). \end{aligned}$$

2) By the first statement we infer that

$$\frac{2}{1-z+\alpha} > \frac{\alpha}{1-z(1+\alpha)}$$

and this can be rearraranged to $z < 1-\alpha$.

3) Suppose not, i.e $T = \emptyset$. Then, by $v_{i^*} < \frac{2}{1-z+\alpha} \sum_{i \in \emptyset} v_i = 0$, a contradiction.

4) Again, by the first property, $\frac{2}{1-z+\alpha} > \frac{\alpha}{1-z(1+\alpha)}$ and this is equivalent to $\lambda(z) > 0$.

□

Consider now the set $T \setminus U$, which is the set of bidders that did not accept their personalized posted price B_i and yet were part of the provisional allocation $\bar{\mathbf{x}}$ of Algorithm 5.2. Therefore, by Property 1 of Lemma 5.3.2, it is also true that $c_i \leq \frac{v_i}{r}$ for all members of $T \setminus U$. By combining the two inequalities we obtain that for these bidders

$$\frac{v_i}{r} \geq c_i > B_i = \min\left\{\frac{v_i}{r}, \sqrt{2+\alpha} \cdot \frac{v_i}{\sum\limits_{j \in T} v_j}\left(B - (1-z)B_{i^*} - z\frac{v_{i^*}}{r}\right)\right\}.$$

Observe that a direct implication of this is that

$$B_i = \sqrt{2+\alpha} \cdot \frac{v_i}{\sum\limits_{j \in T} v_j}\left(B - (1-z)B_{i^*} - z\frac{v_{i^*}}{r}\right), \quad \forall i \in T \setminus U \tag{5.6}$$

We now turn back to obtaining an upper bound on OPT_F. We can rewrite Equation (5.3) as:

$$
\begin{aligned}
OPT_F &\le (1+z)v_{i^*} + 2\sum_{i\in T} v_i - r\sum_{i\in N} c_i\bar{x}_i \\
&\le (1+z)v_{i^*} + 2\sum_{i\in T} v_i - r\sum_{i\in T\setminus U} B_i \\
&= (1+z)v_{i^*} + 2\sum_{i\in T} v_i - \sqrt{2+\alpha}\sum_{i\in T\setminus U}\frac{v_i}{\sum_{j\in T} v_j}\left(Br-(1-z)B_{i^*}r - zv_{i^*}\right) \\
&\le (1+z)v_{i^*} + 2\sum_{i\in T} v_i - \sqrt{2+\alpha}\cdot\frac{\sum_{i\in T\setminus U} v_i}{\sum_{j\in T} v_j}\left(\frac{\frac{2}{1-z+\alpha}}{\lambda(z)}\sum_{j\in T} v_j - \frac{v_{i^*}}{\lambda(z)}\right) \\
&= \left(1+z+\frac{\sqrt{2+\alpha}}{\lambda(z)}\right)v_{i^*} + \left(2-\frac{\frac{2\sqrt{2+\alpha}}{1-z+\alpha}}{\lambda(z)}\right)\sum_{i\in T} v_i + \frac{\sqrt{2+\alpha}}{\lambda(z)}\left(\frac{2}{1-z+a} - \frac{v_{i^*}}{\sum_{i\in T} v_i}\right)\sum_{i\in U} v_i \\
&\le \left(1+z+\frac{\sqrt{2+\alpha}}{\lambda(z)}\right)v_{i^*} + \left(2-\frac{\frac{2\sqrt{2+\alpha}}{1-z+\alpha}}{\lambda(z)}\right)\sum_{i\in T} v_i + \sqrt{2+\alpha}\sum_{i\in U} v_i \\
&\le \left(1+z+\frac{\sqrt{2+\alpha}}{\lambda(z)}\right)v_{i^*} + \left(2-\frac{\frac{2\sqrt{2+\alpha}}{1-z+\alpha}}{\lambda(z)}\right)v_{i^*}\frac{1-z+\alpha}{2} + \sqrt{2+\alpha}\sum_{i\in U} v_i \\
&= (2+\alpha)\left(v_{i^*} + \frac{1}{\sqrt{2+\alpha}}\sum_{i\in U} v_i\right) = (2+\alpha)V(\mathbf{x}).
\end{aligned}
$$

The second equality follows by Equation (5.6), and the inequality that follows is by Lemma 5.3.6. The remaining two inequalities follow by the properties of Claim 5.3.1 and the fact that $\left(2-\frac{\frac{2\sqrt{2+\alpha}}{1-z+\alpha}}{\lambda(z)}\right) < 0$ for all $z \le 1-\alpha$.

(d) As in Case (b), we have that $\mathbf{x} = (z, \bar{\mathbf{x}}_{-i^*})$ Moreover, bidder i^* rejects her personalized posted price B_{i^*}, i.e $c_{i^*} > B_{i^*}$. However, as in Case (c), by Property 1 of Lemma 5.3.2, it also holds that $c_{i^*} \le \frac{v_{i^*}}{r}$ and therefore we can infer that

$$B_{i^*} = \frac{v_{i^*}(1-z) - \alpha V(z, \bar{\mathbf{x}}_{-i^*})}{r(1-z)\lambda(z)(1-z(1+\alpha))}. \tag{5.7}$$

Therefore, by Equation (5.3) we have

$$
\begin{aligned}
OPT_F &\le (1-z)v_{i^*} + 2V(z, \bar{\mathbf{x}}_{-i^*}) - r\sum_{i\in N} c_i \bar{x}_i \\
&\le (1-z)v_{i^*} + 2V(z, \bar{\mathbf{x}}_{-i^*}) - rB_{i^*} \\
&= (1-z)v_{i^*} + 2V(z, \bar{\mathbf{x}}_{-i^*}) - \frac{v_{i^*}(1-z) - \alpha V(z, \bar{\mathbf{x}}_{-i^*})}{(1-z)\lambda(z)(1-z(1+\alpha))} \\
&= \left(1 - z - \frac{1}{\lambda(z)(1-z(1+\alpha))}\right) v_{i^*} + \left(2 + \frac{\alpha}{(1-z)\lambda(z)(1-z(1+\alpha)}\right) V(z, \bar{\mathbf{x}}_{-i^*}) \\
&\le \left(\left(1 - z - \frac{1}{\lambda(z)(1-z(1+\alpha))}\right) \frac{\alpha}{1-z} + 2 + \frac{\alpha}{(1-z)\lambda(z)(1-z(1+\alpha)}\right) V(z, \bar{\mathbf{x}}_{-i^*}) \\
&= (2+\alpha)V(z, \bar{\mathbf{x}}_{-i^*}) = (2+\alpha)V(\mathbf{x}).
\end{aligned}
$$

The second equality follows by Equation (5.7) and the last inequality by the properties of Claim 5.3.1 and the fact that $\left(1 - z - \frac{1}{\lambda(z)(1-z(1+\alpha))}\right) \le 0$, for $z \le 1-\alpha$.

□

5.4 A Truthful and Budget-Feasible Mechanism for Multiple Levels of Service

In this section, we propose a mechanism that is truthful, individually rational, and budget-feasible for an auction setting where bidders can offer up to k levels of service.

We will need the following auxiliary notation to describe the mechanism. For an allocation $\mathbf{x}$, we denote by $W(\mathbf{x}) = \{i \in N \mid x_i > 0\}$, the set of bidders who have been hired under $\mathbf{x}$ (either for at least one level of service in the integral case or fractionally when we consider the relaxation of the Bounded Knapsack problem). Finally, we denote by $\ell(\mathbf{x})$ the bidder whose $x_{\ell(\mathbf{x})}$-th level of service is the least valuable in $\mathbf{x}$, in terms of her marginal-value-per-cost ratio. Notice that due to the fact that the valuation functions are concave, the worst case marginal-value-per-cost ratio indeed corresponds to the $x_{\ell(\mathbf{x})}$-th ratio of bidder $\ell(\mathbf{x})$. For the sake of brevity, for the allocation $\mathbf{x}(\mathbf{c})$ of Mechanism 5.4, we denote this bidder as ℓ, when it is clear from context.

The main result of this section is the next theorem.

Theorem 5.4.1. *Mechanism 5.4 is a truthful, individually rational, budget-feasible and $(2+\sqrt{3})$-approximate mechanism.*

Algorithm 5.4: A Mechanism for k Levels of Service

▷ **Input:** A profile $\mathbf{c}$.

Set $i^* = \arg\max_{i \in N} \frac{v_i(k)}{OPT_F^k(\mathbf{c}_{-i})}$.
if $v_{i^*}(k) \geq \frac{1}{1+\sqrt{3}} \cdot OPT_F^k(\mathbf{c}_{-i^*})$ **then**
    Set $x_{i^*} = k$ and $x_i = 0 \quad \forall i \in N \setminus \{i^*\}$.
else
    Denote by $\mathbf{x}^*(\mathbf{c})$ the optimal allocation of $OPT_F^k(\mathbf{c})$.
    Initialize an empty list of $\sum_{i=1}^n \lfloor x_i^*(\mathbf{c}) \rfloor$ elements.
    for $i \in W(\mathbf{x}^*(\mathbf{c}))$ **do**
        for $j = 1, \ldots, \lfloor x_i^*(\mathbf{c}) \rfloor$ **do**
            Add $\frac{m_i(j)}{c_i}$ to the list.
    Sort the elements of the list in *decreasing* order.
    Initialize $\mathbf{x} = (\lfloor x_1^*(\mathbf{c}) \rfloor, \ldots, \lfloor x_n^*(\mathbf{c}) \rfloor)$ and let ℓ be the bidder who corresponds to the last element of the list.
    while $V(\mathbf{x}) - m_\ell(x_\ell) \geq \frac{1}{2+\sqrt{3}} OPT_F^k(\mathbf{c})$ **do**
        Set $x_\ell = x_\ell - 1$.
        Remove the last element from the list and update ℓ.
Allocate $\mathbf{x}$ and set $\mathbf{p}$ according to Equation (5.1).

In the rest of this section, we present a set of lemmas that demonstrate the properties described in Theorem 5.4.1. We start with the following fact, which is an obvious property for the solution returned by the mechanism.

Fact 5.4.1. *By construction, the allocation $\mathbf{x}$ returned by Mechanism 5.4 satisfies $x_i \leq x_i^*$ for every $i \in N$.*

We now prove that the allocation algorithm of Mechanism 5.4 is monotone.

Lemma 5.4.1. *The allocation algorithm of Mechanism 5.4 is monotone.*

Proof. Let $\mathbf{c}$ be a bidding profile. We distinguish the following two cases.

1. $v_{i^*}(k) \geq \frac{1}{1+\sqrt{3}} \cdot OPT_F^k(\mathbf{c}_{-i^*})$. In this case, i^* is hired for k levels of service. Suppose that i^* unilaterally decreases their cost to a $c'_{i^*} \leq c_{i^*}$. Note that such a unilateral deviation does not alter the condition of the case, since the quantity $OPT_F^k(\mathbf{c}_{-i^*})$ does not depend on their bid. Therefore, for any such deviation, $i^*(\mathbf{c})$ will remain the sole winner and will keep being hired for k levels of service, i.e., $x_{i^*}(c'_{i^*}, \mathbf{c}_{-i^*}) = x_{i^*}(\mathbf{c}) = k$. No other bidder was winning under this case, hence there is no need to examine deviations by other bidders.

2. $v_{i^*}(k) < \frac{1}{1+\sqrt{3}} \cdot OPT_F^k(\mathbf{c}_{-i^*}(\mathbf{c}))$. Here, the allocation algorithm of Mechanism 5.4 is allocating to a set $W(\mathbf{x}(\mathbf{c}))$. Fix a bidder i in $W(\mathbf{x}(\mathbf{c}))$ and suppose

she unilaterally deviates and declares $c'_i \leq c_i$. First of all, note that for every $j \in N \setminus \{i\}$ it holds that

$$\frac{v_j(k)}{OPT_F^k(\mathbf{c}_{-\{j,i\}}, c'_i)} \leq \frac{v_j(k)}{OPT_F^k(\mathbf{c}_j)} \leq \frac{v_{i^*}(k)}{OPT_F^k(\mathbf{c}_{-i^*})}$$

Therefore, even if i^* changes under profile $(c'_i, \mathbf{c}_{-i})$, we will always remain in the **else** case on line 4 and the new set $W(\mathbf{x}(c'_i, \mathbf{c}_{-i}))$ will not lose to the new i^* under $(c'_i, \mathbf{c}_{-i})$.

Consider now what happens to the **while** condition on line 11. On the one hand, for the optimal allocation, it can only be that $OPT_F^k(c'_i, \mathbf{c}_{-i}) \geq OPT_F^k(\mathbf{c})$. On the other hand, the marginal-value-per-cost ratios of bidder i under profile $(c'_i, \mathbf{c}_{-i})$ have a better position in the ordering constructed by Mechanism 5.4. At the same time, the unilateral deviation of bidder i may, in fact, increase the allocation of other bidders. Crucially though, since $\mathbf{x}^*(c'_i, \mathbf{c}_{-i})$ is constructed by Algorithm 5.1, it is guaranteed that the new marginal-value-per-cost ratios belonging to other bidders that enter the solution $\mathbf{x}^*(c'_i, \mathbf{c}_{-i})$ will not move in front of the ratios of i that guarantee i an allocation of $x_i(\mathbf{c})$. Note that this is true as long as we employ a deterministic tie-breaking rule. Therefore, it holds that $x_i(c'_i, \mathbf{c}_i) \geq x_i(\mathbf{c})$ and monotonicity holds.

□

Since the payments follow the payment identity of Equation (5.1), we conclude that the mechanism is truthful and individually rational.

We continue by proving that Mechanism 5.4 achieves the claimed approximation guarantee.

Lemma 5.4.2. *Let* $\mathbf{c}$ *be a profile and* $\mathbf{x}(\mathbf{c})$ *be the final allocation of Mechanism 5.4. Then, it holds that* $V(\mathbf{x}(\mathbf{c})) \geq \frac{1}{2+\sqrt{3}} OPT_I^k(\mathbf{c})$.

Proof. For a profile $\mathbf{c}$ we will prove the claimed guarantee against the optimal solution to the fractional relaxation of the bounded knapsack instance, i.e., we will show that $V(\mathbf{x}(\mathbf{c})) \geq \frac{1}{2+\sqrt{3}} OPT_F^k(\mathbf{c})$. As we have mentioned in Section 5.2.2, this establishes our guarantee since $OPT_F^k(\mathbf{c}) \geq OPT_I^k(\mathbf{c})$. We will prove that Mechanism 5.4 achieves this guarantee in each of the two cases that follow.

1. $v_{i^*}(k) \geq \frac{1}{1+\sqrt{3}} \cdot OPT_F^k(\mathbf{c}_{-i^*})$. We directly have

$$v_{i^*}(k) \geq \frac{1}{1+\sqrt{3}} \cdot OPT_F^k(\mathbf{c}_{-i^*}) \geq \frac{1}{1+\sqrt{3}} \left(OPT_F^k(\mathbf{c}) - v_{i^*}(k)\right)$$

The second inequality follows by the fact that $OPT_F^k(\mathbf{c}_{-i}) + v_i(k) \geq OPT_F^k(\mathbf{c})$, for all $i \in N$. By rearranging terms, we obtain

$$v_{i^*}(k) \geq \frac{1}{2+\sqrt{3}} OPT_F^k(\mathbf{c})$$

and the proof for the case follows since $v_{i^*)}(k) = V(\mathbf{x}(\mathbf{c}))$.

2. $v_{i^*}(k) < \frac{1}{1+\sqrt{3}} \cdot OPT_F^k(\mathbf{c}_{-i^*})$. In this case, per line 11, the allocation algorithm of Mechanism 5.4 computes a solution $\mathbf{x}(\mathbf{c})$ such that $V(\mathbf{x}(\mathbf{c})) \geq \frac{1}{2+\sqrt{3}} OPT_F^k(\mathbf{c})$. Note that for the final solution $\mathbf{x}(\mathbf{c})$, we have that $\mathbf{x}(\mathbf{c}) \neq \mathbf{0}$, since the **while** condition evaluates to False when a single bidder is hired for a single level of service.

□

Finally, we prove that Mechanism 5.4 is budget-feasible.

Lemma 5.4.3. *Mechanism 5.4 is budget-feasible.*

Before presenting the proof of Lemma 5.4.3, we present a series of auxiliary statements, which will prove to be useful in our analysis. The purpose of these statements is to characterise and give upper bounds on the individual payments of winning bidders, whenever we are in the **else** part of the mechanism.

We begin with Lemma 5.4.4 in which we derive an upper bound on the costs of winning bidders. Recall that ℓ is the index of the bidder with the least attractive value-per-cost ratio in the solution output by the mechanism.

Lemma 5.4.4. *Let* $\mathbf{c}$ *be a bidding profile for an instance with* $v_{i^*}(k) < \frac{1}{1+\sqrt{3}} OPT_F^k(\mathbf{c}_{-i^*})$. *It holds that,*

$$c_\ell \leq \frac{2+\sqrt{3}}{1+\sqrt{3}} \cdot B \frac{m_\ell(x_\ell(\mathbf{c}))}{OPT_F^k(\mathbf{c})}. \tag{5.8}$$

Proof. For brevity, let $\mathbf{x} := \mathbf{x}(\mathbf{c})$ and $\mathbf{x}^* := \mathbf{x}^*(\mathbf{c})$. Observe that, since, $v_{i^*}(k) < \frac{1}{1+\sqrt{3}} OPT_F^k(\mathbf{c}_{-i^*})$, the algorithm constructs a solution such that $V(\mathbf{x}) - m_\ell(x_\ell) \leq \frac{1}{2+\sqrt{3}} OPT_F^k(\mathbf{c})$ which implies that

$$OPT_F^k(\mathbf{c}) - V(\mathbf{x}) + m_\ell(x_\ell) \geq \frac{1+\sqrt{3}}{2+\sqrt{3}} OPT_F^k(\mathbf{c}). \tag{5.9}$$

Recall, that in $\mathbf{x}^*$, computed by Algorithm 5.1, there exists at most one bidder in $W(\mathbf{x}^*)$ with a non-integer allocation. Denote that bidder as f. We prove Equation

(5.8) as follows:

$$
\begin{aligned}
B &\geq \sum_{i\in W(\mathbf{x}^*)} c_i x_i^* \geq c_\ell(x_\ell^* - x_\ell + 1) + \sum_{i\in W(\mathbf{x}^*)\setminus\{\ell\}} c_i(x_i^* - x_i) \\
&= \sum_{j=x_\ell}^{\lfloor x_\ell^*\rfloor} \frac{c_\ell}{m_\ell(j)} m_\ell(j) + \sum_{i\in W(\mathbf{x}^*)\setminus\{\ell\}} \sum_{j=x_i+1}^{\lfloor x_i^*\rfloor} \frac{c_i}{m_i(j)} m_i(j) + \frac{c_f}{m_f(\lceil x_f^*\rceil)}(x_f^* - \lfloor x_f^*\rfloor) m_f(\lceil x_f^*\rceil) \\
&\geq \frac{c_\ell}{m_\ell(x_\ell)} \left(\sum_{j=x_\ell}^{\lfloor x_\ell^*\rfloor} m_\ell(j) + \sum_{i\in W(\mathbf{x}^*)\setminus\{\ell\}} \sum_{j=x_i+1}^{\lfloor x_i^*\rfloor} m_i(j) + (x_f^* - \lfloor x_f^*\rfloor) m_f(\lceil x_f^*\rceil) \right) \\
&= \frac{c_\ell}{m_\ell(x_\ell)} \left(OPT_F^k(\mathbf{c}) - V(\mathbf{x}) + m_\ell(x_\ell) \right) \geq \frac{c_\ell}{m_\ell(x_\ell)} \frac{1+\sqrt{3}}{2+\sqrt{3}} OPT_F^k(\mathbf{c}).
\end{aligned}
$$

The first inequality follows by the feasibility of $\mathbf{x}^*$ whereas the second inequality (in the first line) by the fact that $x_\ell \geq 1$. Then, the penultimate inequality is due to Fact 5.2.1 and, specifically, the marginal-value-per-cost ordering Algorithm 5.1 performs. Finally, the last inequality is due to Equation (5.9). By rearranging terms the Lemma follows. □

Let $\mathbf{c}$ be a bidding profile and let $i \in W(\mathbf{x}(\mathbf{c}))$. For $j = x_i(\mathbf{c}), x_i(\mathbf{c}) - 1, \ldots, 1$, we define by $p_{ij}(\mathbf{c}_{-i}) = \sup\{z \geq c_i \mid x_i(z, \mathbf{c}_{-i}) = j\}$, the critical cost for the j-th level of service for bidder i, should the supremum exist. If the supremum does not exist, we define $p_{ij}(\mathbf{c}_{-i}) = p_{i(j+1)}(\mathbf{c}_{-i})$. Notice that $p_{ix(\mathbf{c})}(\mathbf{c}_{-i})$ is always equal to $\sup\{z \geq c_i \mid x_i(z, \mathbf{c}_{-i}) = x_i(\mathbf{c})\}$ i.e., this supremum always exists since, by assumption, $c_i \leq \frac{B}{k}$. We now show that we can describe the payment of bidder i as the sum of these critical payments for the levels of service she was hired for.

Observation 5.4.1. *Let $\mathbf{c}$ be a bidding profile and let $i \in W(\mathbf{x}(\mathbf{c}))$. It holds that*

$$
p_i(\mathbf{c}) = \sum_{j=1}^{x_i(\mathbf{c})} p_{ij}(\mathbf{c}_{-i}). \tag{5.10}
$$

Proof. Indeed, by Equation (5.1) we obtain that

$$
\begin{aligned}
p_i(\mathbf{c}) &= c_i x_i(\mathbf{c}) + \int_{c_i}^{\infty} x_i(z, \mathbf{c}_{-i}) dz \\
&= c_i x_i(\mathbf{c}) + \sum_{j=1}^{x_i(\mathbf{c})-1} \int_{p_{i(j+1)(\mathbf{c}_{-i})}}^{p_{ij}(\mathbf{c}_{-i})} j dz + \int_{c_i}^{p_{ix_i(\mathbf{c}_{-i})}(\mathbf{c})} x_i(\mathbf{c}) dz = \sum_{j=1}^{x_i(\mathbf{c})} p_{ij}(\mathbf{c}_{-i}).
\end{aligned}
$$

□

We now proceed to obtaining an upper bound on the payments each bidder receives for each level of service.

Lemma 5.4.5. *Let* $\mathbf{c}$ *be a bidding profile such that* $v_{i^*}(k) < \frac{1}{1+\sqrt{3}} OPT_F^k(\mathbf{c}_{-i^*})$*. Moreover, let* $i \in W(\mathbf{x}(\mathbf{c}))$*. For* $j = 1, \dots, x_i(\mathbf{c})$*, it holds that*

$$p_{ij}(\mathbf{c}_{-i}) \leq \frac{2+\sqrt{3}}{1+\sqrt{3}} \cdot B \frac{m_i(j)}{OPT_F^k(\mathbf{c}_{-i})}. \tag{5.11}$$

Proof. For the sake of brevity, let $\ell' := \ell(\mathbf{x}(p_{ij}(\mathbf{c}_{-i}), (\mathbf{c}_{-i}))$ and $p_{ij} := p_{ij}(\mathbf{c}_{-i})$. Since p_{ij} is a bid that guarantees bidder i at least j levels of service under the profile $(p_{ij}, \mathbf{c}_{-i})$, it holds that

$$p_{ij} \leq \frac{c_{\ell'} m_i(j)}{m_{\ell'}(x_{\ell'}(p_{ij}, \mathbf{c}_{-i}))} \leq \frac{2+\sqrt{3}}{1+\sqrt{3}} \cdot B \frac{m_i(j)}{OPT_F^k(p_{ij}, \mathbf{c}_{-i})} \leq \frac{2+\sqrt{3}}{1+\sqrt{3}} \cdot B \frac{m_i(j)}{OPT_F^k(\mathbf{c}_{-i})}.$$

The second inequality follows by applying the inequality in Equation (5.8) for the profile $(p_{ij}, \mathbf{c}_{-i})$, whereas the third inequality is due to the fact that $OPT_F^k(\mathbf{c}) \geq OPT_F^k(\mathbf{c}_{-i})$, for every profile $\mathbf{c}$ and every bidder $i \in N$. □

The final component needed for the proof of Lemma 5.4.3 is a lower bound on the optimal fractional objective, when one bidder is excluded.

Lemma 5.4.6. *Let* $\mathbf{c}$ *be a bidding profile such that* $v_{i^*}(k) < \frac{1}{1+\sqrt{3}} OPT_F^k(\mathbf{c}_{-i^*})$*. For every bidder* $i \in N$ *it holds that*

$$OPT_F^k(\mathbf{c}_{-i}) \geq (1+\sqrt{3})(V(\mathbf{x}(\mathbf{c})) - m_\ell(x_\ell)). \tag{5.12}$$

Proof. By the stopping condition of the **while** loop, we have:

$$\begin{aligned}
V(\mathbf{x}(\mathbf{c})) - m_\ell(x_\ell) &\leq \frac{1}{2+\sqrt{3}} OPT_F^k(\mathbf{c}) \\
&\leq \frac{1}{2+\sqrt{3}} \left(OPT_F^k(\mathbf{c}_{-i}) + v_i(k)\right) \\
&\leq \frac{OPT_F^k(\mathbf{c}_{-i})}{2+\sqrt{3}} \left(1 + \frac{v_{i^*}(k)}{OPT_F^k(\mathbf{c}_{-i^*})}\right) \\
&\leq \frac{OPT_F^k(\mathbf{c}_{-i})}{2+\sqrt{3}} \left(1 + \frac{1}{1+\sqrt{3}}\right) = \frac{OPT_F^k(\mathbf{c}_{-i})}{1+\sqrt{3}}
\end{aligned}$$

.

The third inequality follows by the definition of i^*, whereas the last inequality follows directly by the assumption of the lemma. □

We can now present the proof of Lemma 5.4.3.

Proof of Lemma 5.4.3. Let $\mathbf{c}$ be a bidding profile. If $v_{i^*}(k) \geq \frac{1}{1+\sqrt{3}} OPT_F^k(\mathbf{c}_{-i^*})$, bidder i^* is the only bidder hired and she is hired for k levels of services. By Equation (5.10), we obtain

$$p_{i^*}(\mathbf{c}) = \sum_{j=1}^{k} p_{i^* j}(\mathbf{c}_{-i^*}) = k p_{i^* k}(\mathbf{c})_{-i^*} = k\frac{B}{k} = B$$

The second equality, follows by the definition of the critical payments whereas the third equality by the fact that, for this case, $\sup\{z \geq c_{i^*} \mid x_{i^*}(z, \mathbf{c}_{-i^*}) = k)\} = \frac{B}{k}$.

Consider now the other case, i.e., a profile $\mathbf{c}$ such that $v_{i^*}(k) < \frac{1}{1+\sqrt{3}} OPT_F^k(\mathbf{c}_{-i^*})$. We can upper bound the total payments made to bidders as follows:

$$\begin{aligned}
\sum_{i \in W(\mathbf{x}(\mathbf{c}))} p_i(\mathbf{c}) &= p_\ell(\mathbf{c}) + \sum_{i \in W(\mathbf{x}(\mathbf{c})) \setminus \{\ell\}} p_i(\mathbf{c}) \\
&= \sum_{j=1,\dots x_\ell(\mathbf{c})} p_{\ell j}(\mathbf{c}_{-\ell}) + \sum_{i \in W(\mathbf{x}(\mathbf{c})) \setminus \{\ell\}} \sum_{j=1,\dots,x_i(\mathbf{c})} p_{ij}(\mathbf{c}_{-i}) \\
&\leq B\frac{2+\sqrt{3}}{1+\sqrt{3}} \left(\frac{v_\ell(x_\ell(\mathbf{c}))}{OPT_F^k(\mathbf{c}_{-\ell})} + \sum_{i \in W(\mathbf{x}(\mathbf{c})) \setminus \{\ell\}} \frac{v_i(x_i(\mathbf{c}))}{OPT_F^k(\mathbf{c}_{-i})} \right) \\
&\leq B\frac{2+\sqrt{3}}{1+\sqrt{3}} \left(\frac{v_\ell(k)}{OPT_F^k(\mathbf{c}_{-\ell})} + \sum_{i \in W(\mathbf{x}(\mathbf{c})) \setminus \{\ell\}} \frac{v_i(x_i(\mathbf{c}))}{OPT_F^k(\mathbf{c}_{-i})} \right) \\
&\leq B\frac{2+\sqrt{3}}{1+\sqrt{3}} \left(\frac{v_{i^*}(k)}{OPT_F^k(\mathbf{c}_{-i^*})} + \sum_{i \in W(\mathbf{x}(\mathbf{c})) \setminus \{\ell\}} \frac{v_i(x_i(\mathbf{c}))}{(1+\sqrt{3})(V(\mathbf{x}(\mathbf{c})) - m_\ell(x_\ell(\mathbf{c})))} \right) \\
&\leq B\frac{2+\sqrt{3}}{1+\sqrt{3}} \left(\frac{1}{1+\sqrt{3}} + \frac{1}{1+\sqrt{3}} \right) = B.
\end{aligned}$$

The second equality is due to Equation (5.10) and the inequality that follows by applying Equation (5.11) for every bidder $i \in W(\mathbf{x}(\mathbf{c}))$, and every $j = 1\dots, x_i(\mathbf{c})$. Observe also that $\sum_{j=1}^{x_i(\mathbf{c})} m_i(j) = v_i(x_i(\mathbf{c}))$. Then, the third inequality is due to Equation (5.12) and to the definition of i^*. Finally, the last inequality is by the fact that $\sum_{i \in W(\mathbf{x}(\mathbf{c})) \setminus \{\ell\}} v_i(x_i(\mathbf{c})) = V(\mathbf{x}(\mathbf{c})) - v_\ell(x_\ell(\mathbf{c})) \leq V(\mathbf{x}(\mathbf{c})) - m_\ell(x_\ell(\mathbf{c}))$, and by the fact that we analyze the **else** case of the mechanism.

□

Chapter 6

Conclusions and Future Directions

Throughout this dissertation, we have delved into the design and analysis of algorithmic mechanisms in both forward and reverse auction environments. Our work has tackled classic themes in algorithmic game theory, addressing computational and incentive concerns. In this final chapter, we provide an extended discussion of our findings, highlight open problems for future research, and propose several modern and relatively unexplored directions in the analysis of auctions under computational lens.

6.1 Discussion and Open Problems

6.1.1 Forward Auctions

In the first part of the thesis, we focused on the design and analysis of forward auctions. In Chapter 2 we focused on the notion of the core in auctions whereas in Chapter 3 we focused on the mixed and Bayes Nash equilibria of the discriminatory auction.

In Chapter 2, our findings provide insight into core-selecting and core-competitive mechanisms, and we believe that the established properties of the core polyhedron in Section 2.3 and the analysis of quadratic rules in Section 2.5 have broader applicability and independent interest. However, there are still interesting avenues for further investigation. The recent experimental evaluation of Bünz et al. (2022) has sparked a debate on identifying the most appropriate MRCS rules. We find that the notion of non-decreasing payment rules, defined in Section 2.5, is a suitable refinement of MRCS rules towards this direction. It would be interesting to better understand or even characterize which MRCS rules can satisfy this property. Moreover, the literature on designing truthful mechanisms that are core-competitive is still scarce, and it would be desirable to identify special cases of single-parameter domains, where better than

$O(\log n)$-competitiveness can be achieved, as in the Text and Image setting of Goel et al. (2015). Generalizations to multi-parameter domains would also be enlightening for further understanding the structure of the core constraints and their broader applicability.

In Chapter 3, a complete characterization was performed for the mixed equilibria of the discriminatory price auction in the case of two bidders under the uniform bidding interface, assuming capped-additive valuations. However, for n bidders, with $n > 2$, this characterization is only partial, which raises the question of obtaining a complete characterization for this case. Recently, Jin and Lu (2022) achieved an impressive feat by obtaining a complete characterization for the Bayes Nash equilibria of the First Price Auction, (which can be viewed as a special case of $n \geq 2$ and $k = 1$ in our model, albeit for a richer class). They also established that the Price of Anarchy (PoA) for this class of equilibria is $\frac{e^2}{e^2-1} \approx 1.1565$, and in order to improve the previously known lower bound of Hartline (2013), they formalized a calculus of variations problem using techniques similar to those we use in Section 3.5. This suggests that our technique may be of wider applicability, even for unrelated combinatorial auctions environments. In general, our work in this chapter and other emerging approaches do not rely on the smoothness framework of Syrgkanis and Tardos (2013). Roughgarden et al. (2017) have provided a survey of the smoothness framework's successes and limitations in this regard. While this framework has enabled the community to obtain many important PoA results, some of which are tight, its usefulness may not extend to all scenarios and it is important to understand these limitations in a principled way. It is a particularly challenging research direction to understand whether one needs to develop a refined framework beyond smoothness or whether a complete characterization of equilibria is unavoidable for obtaining tight PoA results (for problems where smoothness techniques have not been able to yield tight bounds).

6.1.2 Procurement Auctions

In the second part of this thesis, our focus was on mechanism design for procurement auctions. In Chapter 4, we delved into a mechanism design problem with covering constraints that has relevance for crowdsourcing applications. In Chapter 5, we extended our focus to the design of budget-feasible auction mechanisms for two extensions of the original model that was initiated by Singer (2010).

Regarding the CMIC problem of Chapter 4, from a mechanism design viewpoint, the most important question for future research is to design truthful mechanisms with better approximation guarantees, as there is still a large gap between the non-truthful

3-approximation and our truthful Δ-approximation of Section 4.3. Moreover, exploring the approximability of well-motivated special cases of CMIC, other than the restriction on a constant number of tasks that we examined in Section 4.4, is also an intriguing topic. In the context of crowdsourcing, some of these special cases are also meaningful to study in a 2-dimensional model, where workers can cover circular areas of a given radius, or other geometric shapes. Apart from positive results, we believe it is very interesting to investigate the existence of lower bounds on the worst case performance of polynomial time truthful mechanisms. After all, it is conceivable that there may be a strict separation on the approximability by truthful and non-truthful algorithms, as, e.g., for the combinatorial public project problems (Papadimitriou et al., 2008). Finally, from a purely algorithmic viewpoint, if we set truthfulness aside, it would be quite interesting to obtain an algorithm with a better guarantee than the 3-approximation one due to Mondal (2018).

Finally, in Chapter 5 we studied two procurement auction settings in which the auctioneer may hire bidders partially, and obtained novel mechanisms for both settings. For the divisible agents setting, Mechanism 5.3 achieves a $1+\sqrt{2}$ approximation, retrieving truthfulness and budget-feasibility. As we have already mentioned in Section 5, this mechanism improves upon the previously known state-of-the-art approximation that the mechanism of Klumper and Schäfer (2022) achieved for the same setting. Intriguingly, the approximation guarantee of our mechanism matches the best known lower bound for the *indivisible* agents case, due to Chen et al. (2011). This means that, we are on the verge of obtaining a separation result between the divisible and the indivisible agents scenario: a mechanism that is truthful, budget-feasible and has a $1+\sqrt{2}-\varepsilon$ approximation, for any small constant $\varepsilon > 0$, would settle this question in the affirmative. Moreover, regarding the multiple levels of service scenario, Mechanism 5.4 is the first constant factor approximation mechanism for this setting. It would be interesting to understand whether we can obtain mechanisms with approximation guarantees closer to those possible for single-level type settings, or alternatively, determine whether allowing multiple levels of service is an inherently harder problem. Finally, as far as simple valuation functions are concerned, the most important open problem is still the indivisible agents case with additive valuations, for which the approximation gap is $[1+\sqrt{2}, 3]$ (due to Chen et al. (2011); Gravin et al. (2020)). Any progress on that front may give rise to novel techniques, which may subsequently be also used for problems in richer environments.

6.2 Avenues for Further Research

In this section, we will highlight a few contemporary directions in algorithmic mechanism design, specifically in the context of auctions. The following three themes represent modern areas of interest for members of the Economics and Computation community.

6.2.1 Auction Design with Predictions

Recently, there has been a growing interest in algorithms that go beyond the conventional worst-case analysis paradigm. These algorithms utilize prediction information about the input or the optimal solution's structure to achieve better efficiency. Even if the prediction is incorrect, the goal is to design algorithms that can still fall back on a worst-case guarantee. This approach aims to combine accuracy (consistency), which roughly means that the algorithm should utilize the prediction in case it is of value, and worst-case performance guarantees (robustness) in algorithm design, as in classical algorithms. For a survey of algorithmic results in this direction, we refer to Mitzenmacher and Vassilvitskii (2022).

In the context of mechanism design, this approach presents additional challenges because the input is provided by strategic agents. In a forward auction, bidders may report lower bids to win at a lower cost and pay less. Naturally, we seek incentive-compatible or truthful mechanisms where agents do not have an incentive to misreport their preferences. However, in the presence of predictions, this is largely unexplored territory.

A recent work by Xu and Lu (2022) suggests that combining improved performance and truthfulness is achievable in learning-augmented settings, which should be applicable to auction design as the authors, in fact, study such a setting. Moreover, Gkatzelis et al. (2022) show that the quality of Nash equilibria can improve in games where truthful reporting is not guaranteed, under the assumption that the mechanism has access to predictions. They study a class of scheduling games and a class of network formation games. It would be quite interesting to analyze a non-truthful auction format, such as the setting of Chapter 3 under this prism.

6.2.2 Data-driven Auctions

The design of auctions depends on the valuation function of each bidder, which shows their willingness to pay for different combinations of goods. Classic Bayesian auction models assume that the auctioneer knows the probability distribution from which valuations are drawn. This has led to a significant amount of research on mechanisms for revenue or social welfare maximization, among others, the celebrated theorem of Myerson (1981), which make explicit use of the distribution in the rules of the mechanism.

However, it is not always feasible for an auctioneer to know or approximate the distribution of valuations, let alone use it a priori in the mechanism's description. Thus, there has been a growing interest in learning-based approaches for auctions. One such approach is batch learning and data-driven models in auctions, as studied in Huang et al. (2015) and Mohri and Medina (2014). These models assume that the auction designer has access to data from past auctions, and a supervised learning approach is used to estimate reserve prices for future auctions. We refer also to Roughgarden and Wang (2019) and Derakhshan et al. (2019, 2021), for LP-based approaches regarding the offline version of the problem.

6.2.3 Clock Auctions

A clock auction is a type of auction where the bidding starts at a low price and gradually increases over time, hence the name "clock". The auctioneer initiates the auction with a low price and progressively raises it until a bidder agrees to the current price or until a predetermined end time is reached.

In Chapter 5, we referenced the work of Balkanski et al. (2022), who introduced a promising direction for research in designing budget-feasible mechanisms, which is a field that has mainly consisted of sealed-bid auctions, as in this dissertation. According to Milgrom and Segal (2020), clock auctions offer several advantages over sealed bid auctions, such as simpler implementations that are easier to comprehend for the bidders as well as transparency. The latter means that bidders are able to trust that the auctioneer will handle their private information with care and adhere to the auction protocol. Additionally, clock auctions are in most cases *obviously strategyproof* (Li, 2017), meaning that bidders can easily understand that the protocol is indeed strategyproof. Relevant works in this area include those of Dütting et al. (2017), Gkatzelis et al. (2017), and Feldman et al. (2022).

Furthermore, from an algorithmic perspective, these auction protocols possess an interesting feature - unlike sealed-bid auctions, they are often truthful by design. As a result, the mechanism design problem is essentially transformed into an algorithm design problem since there is no need to argue about truthfulness, see Milgrom and Segal (2020).

Considering the aforementioned factors, exploring the development of clock auction mechanisms as an alternative to sealed-bid auctions may offer several benefits in terms of performance and real-life implementation of protocols. As previously mentioned, Balkanski et al. (2022) demonstrated the superiority of clock auctions over state-of-the-art sealed-bid auctions in achieving budget-feasible mechanisms for a number of valuation functions. It would be intriguing to investigate the performance of such protocols in other domains, including the design of core-selecting and core-competitive mechanisms, as well as in environments involving predictions.

Bibliography

Aggarwal, G. and Hartline, J. D. (2006). Knapsack auctions. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1083–1092.

Amanatidis, G., Birmpas, G., and Markakis, E. (2016). Coverage, matching, and beyond: new results on budgeted mechanism design. In *Web and Internet Economics: 12th International Conference, WINE 2016, Proceedings*, pages 414–428.

Amanatidis, G., Birmpas, G., and Markakis, E. (2017). On budget-feasible mechanism design for symmetric submodular objectives. In *Web and Internet Economics: 13th International Conference, WINE 2017, Proceedings*, pages 1–15. Springer.

Amanatidis, G., Kleer, P., and Schäfer, G. (2019). Budget-feasible mechanism design for non-monotone submodular objectives: Offline and online. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019*, pages 901–919.

Anari, N., Goel, G., and Nikzad, A. (2014). Mechanism design for crowdsourcing: An optimal 1-1/e competitive budget-feasible mechanism for large markets. In *IEEE 55th Annual Symposium on Foundations of Computer Science, FOCS 2014, Proceedings*, pages 266–275. IEEE.

Ausubel, L. and Baranov, O. (2020). Core-selecting auctions with incomplete information. *International Journal of Game Theory*, 49(1):251–273.

Ausubel, L., Cramton, P., Pycia, M., Rostek, M., and Weretka, M. (2014). Demand Reduction and Inefficiency in Multi-Unit Auctions. *The Review of Economic Studies*, 81:1366–1400.

Ausubel, L. and Milgrom, P. (2002). Ascending auctions with package bidding. *The BE Journal of Theoretical Economics*, 1(1).

Ausubel, L. and Milgrom, P. (2006). The lovely but lonely vickrey auction. P. Cramton, Y. Shoham, R. Steinberg, eds., Combinatorial Auctions.

Baisa, B. and Burkett, J. (2018). Large Multi-Unit Auctions With a Large Bidder. *Journal of Economic Theory*, 174:1–15.

Balkanski, E., Garimidi, P., Gkatzelis, V., Schoepflin, D., and Tan, X. (2022). Deterministic budget-feasible clock auctions. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2940–2963.

Bansal, N. and Pruhs, K. (2014). The geometry of scheduling. *SIAM Journal on Computing*, 43(5):1684–1698.

Bar-Noy, A., Bar-Yehuda, R., Freund, A., Naor, J., and Schieber, B. (2001). A unified approach to approximating resource allocation and scheduling. *Journal of the ACM*, 48(5):1069–1090.

Bar-Yehuda, R., Bendel, K., Freund, A., and Rawitz, D. (2005). The local ratio technique and its application to scheduling and resource allocation problems. In *Graph Theory, Combinatorics and Algorithms*. Springer.

Beck, M. and Ott, M. (2010). Revenue monotonicity in core-selecting auctions. Technical report, Stanford University.

Bei, X., Chen, N., Gravin, N., and Lu, P. (2017). Worst-case mechanism design via bayesian analysis. *SIAM Journal on Computing*, 46(4):1428–1448.

Bhawalkar, K. and Roughgarden, T. (2011). Welfare Guarantees for Combinatorial Auctions with Item Bidding. In *ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, Proceedings*, pages 700–709.

Bikhchandani, S. and Ostroy, J. (2002). The package assignment model. *Journal of Economic theory*, 107(2):377–406.

Birmpas, G., Markakis, E., Telelis, O., and Tsikiridis, A. (2019). Tight Welfare Guarantees for Pure Nash Equilibria of the Uniform Price Auction. *Theory of Computing Systems*, 63(7):1451–1469.

Blumrosen, L. and Nisan, N. (2007). Algorithmic Game Theory. *Introduction to mechanism design, Cambridge University Press.*

Bosshard, V., Wang, Y., and Seuken, S. (2018). Non-decreasing payment rules for combinatorial auctions. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI 2018*, pages 105–113.

Bosshard, V., Wang, Y., and Seuken, S. (2022). Non-decreasing payment rules for combinatorial auctions. *arXiv preprint arXiv:2206.03857*.

Bredereck, R., Faliszewski, P., Niedermeier, R., Skowron, P., and Talmon, N. (2020). Mixed integer programming with convex/concave constraints: Fixed-parameter tractability and applications to multicovering and voting. *Theoretical Computer Science*, 814:86–105.

Brenner, M., Galai, D., and Sade, O. (2009). Sovereign Debt Auctions: Uniform or Discriminatory? *Journal of Monetary Economics*, 56(2):267–274.

Briest, P., Krysta, P., and Vöcking, B. (2011). Approximation techniques for utilitarian mechanism design. *SIAM Journal on Computing*, 40(6):1587–1622.

Bünz, B., Lubin, B., and Seuken, S. (2022). Designing core-selecting payment rules: A computational search approach. *Information Systems Research*, 33(4):1157–1173.

Bünz, B., Seuken, S., and Lubin, B. (2015). A faster core constraint generation algorithm for combinatorial auctions. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Carr, R. D., Fleischer, L. K., Leung, V. J., and Phillips, C. A. (2000). Strengthening integrality gaps for capacitated network design and covering problems. In *Proceedings of the 11th ACM-SIAM Symposium on Discrete Algorithms, SODA 2000*, page 106–115.

Chakaravarthy, V. T., Kumar, A., Roy, S., and Sabharwal, Y. (2011). Resource allocation for covering time varying demands. In *Algorithms–ESA 2011: 19th Annual European Symposium, ESA 2011, Proceedings*, pages 543–554.

Chakrabarty, D., Grant, E., and Könemann, J. (2010). On column-restricted and priority covering integer programs. In *International Conference on Integer Programming and Combinatorial Optimization, IPCO 2010, Proceedings*, pages 355–368.

Chan, H. and Chen, J. (2014). Truthful multi-unit procurements with budgets. In *Web and Internet Economics: 10th International Conference, WINE 2014, Proceedings*, pages 89–105.

Chen, B., Hassin, R., and Tzur, M. (2002). Allocation of bandwidth and storage. *IIE Transactions*, 34(5):501–507.

Chen, J., Ye, D., Ji, S., He, Q., Xiang, Y., and Liu, Z. (2019). A truthful fptas mechanism for emergency demand response in colocation data centers. In *IEEE Conference on Computer Communications, INFOCOM 2019, Proceedings*, pages 2557–2565.

Chen, N., Gravin, N., and Lu, P. (2011). On the approximability of budget feasible mechanisms. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms, SODA 2011*, pages 685–699.

Christodoulou, G., Kovács, A., and Schapira, M. (2016a). Bayesian combinatorial auctions. *Journal of ACM*, 63(2):11:1–11:19.

Christodoulou, G., Kovács, A., Sgouritsa, A., and Tang, B. (2016b). Tight Bounds for the Price of Anarchy of Simultaneous First-Price Auctions. *ACM TEAC*, 4(2):9:1–9:33.

Clarke, E. H. (1971). Multipart pricing of public goods. *Public choice*, 11(1):17–33.

Corless, R. M., Gonnet, G. H., Hare, D. E. G., Jeffrey, D. J., and Knuth, D. E. (1996). On the Lambert W Function. *Advances in Computational Mathematics*, 5:329–359.

Cramton, P., Shoham, Y., Steinberg, R., et al. (2006). Combinatorial auctions. Technical report, University of Maryland.

Csirik, J. (1991). Heuristics for the 0-1 min-knapsack problem. *Acta Cybernetica*.

Day, R. and Cramton, P. (2012). Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research*, 60(3):588–603.

Day, R. and Milgrom, P. (2008). Core-selecting package auctions. *International Journal of Game Theory*, 36(3-4):393–407.

Day, R. and Raghavan, S. (2007). Fair payments for efficient allocations in public sector combinatorial auctions. *Management Science*, 53(9):1389–1406.

Dayama, P., Narayanaswamy, B., Garg, D., and Narahari, Y. (2015). Truthful interval cover mechanisms for crowdsourcing applications. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015*.

De Keijzer, B., Markakis, E., Schäfer, G., and Telelis, O. (2013). Inefficiency of standard multi-unit auctions. In *Algorithms–ESA 2013: 21st Annual European Symposium, Proceedings*, pages 385–396.

Derakhshan, M., Golrezaei, N., and Paes Leme, R. (2019). LP-based approximation for personalized reserve prices. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*, pages 589–589.

Derakhshan, M., Pennock, D. M., and Slivkins, A. (2021). Beating greedy for approximating reserve prices in multi-unit vcg auctions. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1099–1118.

Dinur, I., Guruswami, V., Khot, S., and Regev, O. (2005). A new multilayered PCP and the hardness of hypergraph vertex cover. *SIAM Journal on Computing*, 34(5):1129–1146.

Dobzinski, S., Papadimitriou, C. H., and Singer, Y. (2011). Mechanisms for complement-free procurement. In *Proceedings of the 12th ACM conference on Electronic Commerce, EC 2011*, pages 273–282.

Dütting, P., Gkatzelis, V., and Roughgarden, T. (2017). The performance of deferred-acceptance auctions. *Mathematics of Operations Research*, 42(4):897–914.

Elkind, E., Goldberg, L. A., and Goldberg, P. W. (2007). Frugality ratios and improved truthful mechanisms for vertex cover. In *Proceedings of the 8th ACM Conference on Electronic Commerce, EC 2007*.

Erdil, A. and Klemperer, P. (2010). A new payment rule for core-selecting package auctions. *Journal of the European Economic Association*, 8(2-3):537–547.

Feldman, M., Fu, H., Gravin, N., and Lucier, B. (2020). Simultaneous Auctions Without Complements are (almost) Efficient. *Games and Economic Behavior*, 123:327–341.

Feldman, M., Gkatzelis, V., Gravin, N., and Schoepflin, D. (2022). Bayesian and randomized clock auctions. In *Proceedings of the 23rd ACM Conference on Economics and Computation, EC 2022*, pages 820–845.

Fiat, A., Goldberg, A., Hartline, J., and Karlin, A. (2002). Competitive generalized auctions. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC)*, pages 72–81.

Friedman, M. (1960). *A Program for Monetary Stability*. Fordham University Press, New York, NY.

Fujito, T. and Yabuta, T. (2004). Submodular integer cover and its application to production planning. In *International Workshop on Approximation and Online Algorithms, WAOA 2004, Proceedings*, pages 154–166.

Gálvez, W., Grandoni, F., Ingala, S., Heydrich, S., Khan, A., and Wiese, A. (2021). Approximating geometric knapsack via L-packings. *ACM Transactions on Algorithms*, 17(4):1–67.

Gkatzelis, V., Kollias, K., Sgouritsa, A., and Tan, X. (2022). Improved price of anarchy via predictions. In *Proceedings of the 23rd ACM Conference on Economics and Computation, EC 2022*, pages 529–557.

Gkatzelis, V., Markakis, E., and Roughgarden, T. (2017). Deferred-acceptance auctions for multiple levels of service. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC 2017*, pages 21–38.

Goel, G., Khani, M. R., and Leme, R. P. (2015). Core-competitive auctions. In *Proceedings of the 16th ACM Conference on Economics and Computation, EC 2015*, pages 149–166.

Goel, G., Nikzad, A., and Singla, A. (2014). Allocating tasks to workers with matching constraints: truthful mechanisms for crowdsourcing markets. In *Proceedings of the 23rd International Conference on World Wide Web, WWW 2014*, pages 279–280.

Goeree, J. and Lien, Y. (2016). On the impossibility of core-selecting auctions. *Theoretical Economics*, 11(1):41–52.

Goldberg, A. and Hartline, J. (2003). Competitiveness via consensus. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2003*, pages 215–222.

Goldberg, A., Hartline, J., Karlin, A., Saks, M., and Wright, A. (2006). Competitive auctions. *Games and Economic Behavior*, 55(2):242–269.

Goldberg, A., Hartline, J., and Wright, A. (2001). Competitive auctions and digital goods. In *Proceedings of the 12th ACM-SIAM Symposium on Discrete Algorithms, SODA 2001*, pages 735–744.

Goldner, K., Immorlica, N., and Lucier, B. (2020). Reducing Inefficiency in Carbon Auctions with Imperfect Competition. In *Innovations in Theoretical Computer Science, ITCS 2020, Proceedings*, pages 15:1–15:21.

Gravin, N., Jin, Y., Lu, P., and Zhang, C. (2020). Optimal budget-feasible mechanisms for additive valuations. *ACM Transactions on Economics and Computation (TEAC)*, 8(4):1–15.

Groves, T. (1973). Incentives in teams. *Econometrica*, 41(4):617–631.

Gummidi, S. R. B., Xie, X., and Pedersen, T. B. (2019). A survey of spatial crowdsourcing. *ACM Transactions on Database Systems*, 44(2):1–46.

Guruswami, V., Hartline, J., Karlin, A., Kempe, D., Kenyon, C., and McSherry, F. (2005). On profit-maximizing envy-free pricing. In *Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms, SODA 2005*, pages 1164–1173.

Harsanyi, J. C. (1967). Games with incomplete information played by "bayesian" players, i–iii part i. the basic model. *Management science*, 14(3):159–182.

Hartline, J. (2013). Mechanism design and approximation. *Book draft. October*, 122.

Hartline, J. D. et al. (2013). Bayesian mechanism design. *Foundations and Trends® in Theoretical Computer Science*, 8(3):143–263.

Huang, Z., Mansour, Y., and Roughgarden, T. (2015). Making the most of your samples. In *Proceedings of the 16th ACM Conference on Economics and Computation, EC 2015*, pages 45–60.

Jalaly, P. and Tardos, É. (2018). Simple and efficient budget feasible mechanisms for monotone submodular valuations. In *Web and Internet Economics: 14th International Conference, WINE 2018, Proceedings*, pages 246–263.

Jin, Y. and Lu, P. (2022). First price auction is 1–1/e 2 efficient. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 179–187.

Khot, S. and Regev, O. (2008). Vertex cover might be hard to approximate to within $2 - \varepsilon$. *Journal of Computer and System Sciences*, 74(3):335–349.

Klumper, S. and Schäfer, G. (2022). Budget feasible mechanisms for procurement auctions with divisible agents. In *Algorithmic Game Theory: 15th International Symposium, SAGT 2022, Proceedings*, pages 78–93.

Kolliopoulos, S. and Young, N. (2005). Approximation algorithms for covering/packing integer programs. *Journal of Computer and System Sciences*, 71(4):495–505.

Kothari, A., Parkes, D., and Suri, S. (2005). Approximately-strategyproof and tractable multiunit auctions. *Decision Support Systems*, 39(1):105–121.

Koufogiannakis, C. and Young, N. E. (2013). Greedy δ-approximation algorithm for covering with arbitrary constraints and submodular cost. *Algorithmica*, 66:113–152.

Koutsoupias, E. and Papadimitriou, C. (1999). Worst-case equilibria. In *Annual Symposium on Theoretical Aspects of Computer Science, STACS 1999*, pages 404–413.

Krishna, V. (2002). *Auction Theory*. Academic Press.

Lamy, L. (2010). Core-selecting package auctions: a comment on revenue-monotonicity. *International Journal of Game Theory*, 39(3):503–510.

Lehmann, B., Lehmann, D., and Nisan, N. (2006). Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296.

Lehmann, D., Oćallaghan, L. I., and Shoham, Y. (2002). Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM (JACM)*, 49(5):577–602.

Leonardi, S., Monaco, G., Sankowski, P., and Zhang, Q. (2017). Budget feasible mechanisms on matroids. In *Integer Programming and Combinatorial Optimization: 19th International Conference, IPCO 2017, Proceedings*, pages 368–379.

Li, S. (2017). Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–3287.

Markakis, E., Papasotiropoulos, G., and Tsikiridis, A. (2022a). On improved interval cover mechanisms for crowdsourcing markets. In *Algorithmic Game Theory: 15th International Symposium, SAGT 2022, Proceedings*, pages 94–112.

Markakis, E., Sgouritsa, A., and Tsikiridis, A. (2022b). Towards a characterization of worst case equilibria in the discriminatory price auction. In *Web and Internet Economics: 17th International Conference, WINE 2021, Proceedings*, pages 186–204.

Markakis, E. and Tsikiridis, A. (2019). On core-selecting and core-competitive mechanisms for binary single-parameter auctions. In *Web and Internet Economics: 15th International Conference, WINE 2019, Proceedings*, pages 271–285.

Martello, S. and Toth, P. (1990). *Knapsack problems: algorithms and computer implementations*. John Wiley & Sons, Inc.

Micali, S. and Valiant, P. (2007). Collusion-resilient revenue in combinatorial auctions. Technical report, Technical Report MIT-CSAIL-TR-2007-052, MIT-CSAIL.

Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press.

Milgrom, P. and Segal, I. (2020). Clock auctions and radio spectrum reallocation. *Journal of Political Economy*, 128(1):1–31.

Mitzenmacher, M. and Vassilvitskii, S. (2022). Algorithms with predictions. *Communications of the ACM*, 65(7):33–35.

Mohri, M. and Medina, A. M. (2014). Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *International Conference on Machine Learning, ICML 2014, Proceedings*, pages 262–270.

Mondal, S. (2018). Improved algorithm for resource allocation problems. *Asia-Pacific Journal of Operational Research*, 35(01):1–23.

Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.

Nguyen, T. T., Roos, M., and Rothe, J. (2013). A survey of approximability and inapproximability results for social welfare optimization in multiagent resource allocation. *Annals of Mathematics and Artificial Intelligence*, 68(1-3):65–90.

Niazadeh, R., Hartline, J., Immorlica, N., Khani, M. R., and Lucier, B. (2022). Fast core pricing for rich advertising auctions. *Operations Research*, 70(1):223–240.

Nisan, N. and Ronen, A. (1999). Algorithmic mechanism design. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing, STOC 1999*, pages 129–140.

Ockenfels, A., Reiley Jr, D. H., and Sadrieh, A. (2006). Online auctions.

Ott, M. and Beck, M. (2013). Incentives for overbidding in minimum-revenue core-selecting auctions. Technical report, Working paper, Stanford University.

Papadimitriou, C., Schapira, M., and Singer, Y. (2008). On the hardness of being truthful. In *Proceedings of the 49th Foundations of Computer Science, FOCS 2008*, pages 250–259.

Parkes, D. (2002). On indirect and direct implementations of core outcomes in combinatorial auctions. *Harvard University Technical Report*.

Parkes, D., Kalagnanam, J., and Eso, M. (2001). Achieving budget-balance with Vickrey-based payment schemes in exchanges. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence, IJCAI 2001*, pages 1161–1168.

Phillips, C. A., Uma, R., and Wein, J. (2000). Off-line admission control for general scheduling problems. *Journal of Scheduling*, 3(6):365–381.

Pritchard, D. (2010). An LP with integrality gap $1 + \epsilon$ for multidimensional knapsack. *arXiv preprint arXiv:1005.3324*.

Pritchard, D. and Chakrabarty, D. (2011). Approximability of sparse integer programs. *Algorithmica*, 61(1):75–93.

Pycia, M. and Woodward, K. (2021). Auctions of homogeneous goods: A case for pay-as-bid. In *Proceedings of the 22nd ACM Conference on Economics and Computation, EC 2021*, pages 787–787.

Rajagopalan, S. and Vazirani, V. (1993). Primal-dual RNC approximation algorithms for (multi)-set (multi)-cover and covering integer programs. In *Proceedings of the 34th Foundations of Computer Science, FOCS 1993*.

Rio, P. D. (2017). Designing Auctions for Renewable Electricity Support. Best Practices From Around the World. *Energy for Sustainable Development*, 41:1–13.

Roughgarden, T. (2012). The Price of Anarchy in Games of Incomplete Information. In *ACM Conference on Economics and Computation, EC 2012*, pages 862–879.

Roughgarden, T., Syrgkanis, V., and Tardos, E. (2017). The Price of Anarchy in Auctions. *Journal of Artificial Intelligence Research*, 59:59–101.

Roughgarden, T. and Wang, J. R. (2019). Minimizing regret with multiple reserves. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–18.

Sano, R. (2011). Incentives in core-selecting auctions with single-minded bidders. *Games and Economic Behavior*, 72(2):602–606.

Singer, Y. (2010). Budget feasible mechanisms. In *51st Annual Symposium on Foundations of Computer Science, FOCS 2010*, pages 765–774.

Swinkels, J. (2001). Efficiency of Large Private Value Auctions. *J. Econ. Theory*, 69(1):37–68.

Syrgkanis, V. (2014). *Efficiency of Mechanisms in Complex Markets*. PhD thesis, Cornell University.

Syrgkanis, V. and Tardos, E. (2013). Composable and Efficient Mechanisms. In *ACM Symposium on Theory of Computing, STOC 2013*, pages 211–220.

Tong, Y., Zhou, Z., Zeng, Y., Chen, L., and Shahabi, C. (2020). Spatial crowdsourcing: a survey. *The VLDB Journal*, 29(1):217–250.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37.

Xu, C. and Lu, P. (2022). Mechanism design with predictions. In Raedt, L. D., editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 571–577.

Xu, J., Xiang, J., and Yang, D. (2015). Incentive mechanisms for time window dependent tasks in mobile crowdsensing. *IEEE Transactions on Wireless Communications*, 14(11):6353–6364.

Zhang, L., Ren, S., Wu, C., and Li, Z. (2015). A truthful incentive mechanism for emergency demand response in colocation data centers. In *Proceedings of the IEEE International Conference on Computer Communications, INFOCOM 2015, Proceedings.*

Appendix A

Missing Material from Forward Auctions

A.1 Missing Material from Chapter 2

Proof of Proposition 2.3.1. Consider the following combinatorial auction with 6 single-minded bidders and 3 items for sale, $M = \{A, B, C\}$. For $i = 1, \ldots, 6$ we denote by $b_i(T)$ the bid of i for the set of items $T \subseteq \{A, B, C\}$. Since bidders are single-minded, each bidder declares a single bid of this form. The bids are summarized below:

$$\begin{array}{ll} b_1(\{A\}) = 9.5 & b_4(\{A, B\}) = 15 \\ b_2(\{B\}) = 6 & b_5(\{A, C\}) = 15 \\ b_3(\{C\}) = 6 & b_6(\{B, C\}) = 10 \end{array}$$

An allocation is feasible when each of the three items is assigned to a unique bidder. For instance, bidders 1 and 4 cannot be a part of a feasible allocation. For the vector $\mathbf{b}$ above, it is easy to see that the welfare-maximizing allocation algorithm assigns $\{A, B, C\}$ to bidders $X^*(\mathbf{b}) = \{1, 2, 3\}$. Per Equation (2.6), the MRCS linear program with variables p_1, p_2, p_3 is:

$$
\begin{array}{lll}
\text{minimize} & p_1 + p_2 + p_3 & \\
 & & \\
\text{subject to} & p_1 & \geq \; 9 \\
 & p_2 & \geq \; 5.5 \\
 & p_3 & \geq \; 5.5 \\
 & p_1 + p_2 & \geq \; 15 \\
 & p_1 + p_3 & \geq \; 15 \\
 & p_2 + p_3 & \geq \; 10 \\
 & p_1 + p_2 + p_3 & \geq \; 15
\end{array}
$$

Recall that each of the constraints is defined in Equation (2.3) of Section 2.2. For the instance above, it is easy to compute the minimum revenue of the auctioneer (the value of the objective function) by taking all the combinations of the MRCS constraints and observing that one set of constraints that must be satisfied with equality are $p_2 \geq 5.5$ and $p_1 + p_3 \geq 15$ which, combined yield that

$$(p_1 + p_3) + (p_2) \geq 20.5.$$

Hence, $\mathsf{MREV}(\mathbf{b}) = 20.5$.

Now suppose that bidder 1 unilaterally declares $b'_i(\{A\}) = 10 > 9.5$. The optimal allocation remains $X^*(b_1, \mathbf{p}_{-1}) = \{1, 2, 3\}$. However, the new MRCS LP becomes:

$$
\begin{array}{lll}
\text{minimize} & p_1 + p_2 + p_3 & \\
 & & \\
\text{subject to} & p_1 & \geq \; 9 \\
 & p_2 & \geq \; 5 \\
 & p_3 & \geq \; 5 \\
 & p_1 + p_2 & \geq \; 15 \\
 & p_1 + p_3 & \geq \; 15 \\
 & p_2 + p_3 & \geq \; 10 \\
 & p_1 + p_2 + p_3 & \geq \; 15
\end{array}
$$

Once again, we consider all combinations of constraints. Observe that since this time the VCG constraints of bidders 2 and 3 have been relaxed, the "blocking" constraints are all greater than all or equal to 20 or strictly weaker. Hence, $\mathsf{MREV}(b'_i, \mathbf{b}_{-i}) = 20 < 20.5 = \mathsf{MREV}(\mathbf{b})$ and the proof follows. Note that this example also implies that the core polyhedron has strictly increased (Theorem 2.3.1). For example, the solution $(p_1, p_2, p_3) = 10, 5, 5$ is now a feasible core point, whereas when bidder 1 was bidding 9.5 this was not possible. □

A.2 Missing Material from Chapter 3

Proof of Theorem 3.3.2. If z is a mass point for bidder i, then we are done by Fact 3.3.2. If not, then consider an interval $I \subseteq Supp(B_i)$ with $z \in I$ where nobody has a mass point in it (recall that the other bidders have no mass point on z, so such I exists). We analyze first the expected utility of a bidder i, given that she bids in I:

$$\begin{aligned}
\mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b}) \mid b_i \in I] &= \mathop{\mathbb{E}}_{b_i\sim B_i}\big[\mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(b_i, \mathbf{b}_{-i})] \mid b_i \in I\big] \\
&= \int_{z\in I} f_{b_i|b_i\in I}(z) \mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(z, \mathbf{b}_{-i})]dz \\
&= \frac{1}{Pr[b_i \in I]} \int_{z\in I} f_i(z) \mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(z, \mathbf{b}_{-i})]dz \qquad \text{(A.1)}
\end{aligned}$$

where f_i is the pdf of B_i, and $f_{b_i|b_i\in I}$ is the conditional pdf when $b_i \in I$. Note that for all $z \in I$ it holds that $f_i(z) \geq 0$ and f_i is continuous. Since no bidder has a mass point in I, by Fact 3.3.1 and Remark 3.3.1 it holds that $\mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(z, \mathbf{b}_{-i})]$ is also continuous in I as a function of z.

We now use a standard fact from calculus, commonly referred to as the integral version of the mean value theorem.

Fact A.2.1. *Let f, g be continuous functions on $[a, b]$ such that f is non-negative. Then there exists a $c \in [a, b]$ such that*

$$\int_a^b f(x)g(x)dx = g(c)\int_a^b f(x)dx.$$

Using this fact, we get that there exists $\xi \in I$, so that we can write Equation (A.1) as

$$\mathop{\mathbb{E}}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b}) \mid b_i \in I] = \mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(\xi, \mathbf{b}_{-i})]\frac{\int_{z\in I} f_i(z)dz}{Pr[b_i \in I]} = \mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(\xi, \mathbf{b}_{-i})]$$

Let $u = \mathbb{E}_{\mathbf{b}\sim B}[u_i(\mathbf{b})]$. By Fact 3.3.2, what we have established so far is that there exists a $\xi \in I$ for which

$$\mathop{\mathbb{E}}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(\xi, \mathbf{b}_{-i})] = u.$$

Consider now making the interval I smaller and smaller, by taking a sequence $I_1, I_2, \dots$ such that in the limit, I_k collapses to z as $k \to \infty$. By the previous arguments, for every I_k, there exists a ξ_k such that $\mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(\xi_k, \mathbf{b}_{-i})] = u$. In the limit, $\xi_k \to z$ and we obtain that $\mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(z, \mathbf{b}_{-i})] = u$. □

Appendix B

Missing Material from Chapter 4

B.1 Proof of Proposition 4.3.1

In this section we show that the algorithms of Chakaravarthy et al. (2011) (4-approximate) and, consequently, of Mondal (2018) (3-approximate) for CMIC are not monotone. This establishes that our Algorithm 4.1 in Section 4.3 is the only currently known monotone algorithm with bounded approximation guarantee, as we prove in Theorem 4.3.1. To proceed, we provide first a description of the algorithms of Chakaravarthy et al. (2011); Mondal (2018).

An ILP formulation for cmic: The problem can be described with a a rather simple integer program. However, the algorithms we present here work with a more involved ILP. Suppose we are given an instance $(\mathbf{b}, \mathbf{q}, \mathbf{d})$ of CMIC with $\mathbf{b} = (\mathbf{c}, \mathbf{I})$. Given a set $S \subseteq \mathcal{N}$, the *residual demand* of a task $j = 1, \ldots, m$ w.r.t. S is denoted by $d_j(S) = d_j - \sum_{i \in \mathcal{N}_j(\mathbf{I}) \cap S} q_i$. This is precisely the demand that remains to be covered by workers outside of S, had we taken S as part of our solution. Moreover, for each worker $i = 1, \ldots, n$ and each $j = 1, \ldots, m$ let $q_i(S, j) = \min\{q_i, d_j(S)\}$. We can consider now the following ILP formulation, where the feasibility constraint says that for any task j, and any set S, if we take any workers from $\mathcal{N} \setminus S$ in our solution, they should collectively cover the residual demand $d_j(S)$.

$$\begin{aligned}
&\text{minimize} \sum_{i=1}^{n} c_i x_i \\
&\text{subject to} \sum_{i \in \mathcal{N}_j(\mathbf{I}) \setminus S} q_i(S,j) x_i \geq d_j(S), \quad \forall j \in [m], \forall S \subseteq \mathcal{N} \\
&\qquad x_i \in \{0,1\}, \quad \forall i \in [n]
\end{aligned}$$

Note that for $S = \emptyset$ the related ILP constraints represent exactly the problem of finding a feasible solution for $(\mathbf{b}, \mathbf{q}, \mathbf{d})$. Then, for every other subset $S \subseteq \mathcal{N}$, the relevant constraints represent the residual problem to be solved assuming we have decided that workers in S are to be included.

Primal-dual algorithm Chakaravarthy et al. (2011); Mondal (2018): The corresponding dual of the LP-relaxation of the above ILP, with dual variables $z(S,j)$ for every pair of j and S, is:

$$\begin{aligned}
&\text{maximize} \sum_{(S,j)} d_j(S) z(S,j) \\
&\text{subject to} \sum_{(S,j): j \in [s_i, f_i]; i \notin S} q_i(S,j) z(S,j) \leq c_i, \quad \forall i \in \mathcal{N} \\
&\qquad z(S,j) \geq 0, \quad \forall S \subseteq \mathcal{N}, \forall j \in [m]
\end{aligned}$$

The primal-dual algorithm runs in two phases: the forward phase and the delete phase.

In the forward phase, initially no bidder is selected and $S = \emptyset$. The algorithm iteratively constructs a feasible solution for the primal as follows: Firstly, it selects the task j with the maximum value for $d_j(S)$. Then, it increases the dual variable $z(S,j)$ until some dual constraint holds with equality. Finally, the algorithm adds bidder i whose dual-constraint became tight to S. This procedure is repeated until S becomes a feasible solution, i.e. $\max_j d_j(S) \leq 0$.

After the forward phase, the delete phase begins. Here, the algorithm considers the bidders in S in the reverse order[1] with which they were added to S in the forward phase. A bidder is removed when the remaining bidders can constitute a feasible

[1]This is prescribed by Mondal (2018) but it is also captured by the algorithm of Chakaravarthy et al. (2011), in which it is not being specified a particular order for deletion consideration.

solution. Thus, the solution returned is minimal feasible.

Counterexample: To argue that the described algorithm of Chakaravarthy et al. (2011); Mondal (2018) is not monotone, we construct an instance where one bidder is initially selected, but after a unilateral decrease in her cost, she is removed from the solution. Consider the following instance with 2 tasks with demands $\mathbf{d} = (4, 1)$ and 4 workers:

worker	c_i	q_i	task 1	task 2
A	4	4	✓	–
B	3	3	✓	–
C	4	1	–	✓
D	4	4	✓	✓

The interval covered by each worker can be inferred from the last two columns, e.g., only worker D can cover both tasks, whereas the first three workers cover a single task each. We also specify a deterministic, tie-breaking rule, determined by the following three conditions:

- TB1: Task 1 is prioritized before Task 2.
- TB2: Ties between worker B and any other worker are broken in favor of B.
- TB3: If worker B has been selected in the solution, then ties between D and any other worker are broken against worker D but if B has not been selected in the solution, ties are broken in favor of D.

Let us now run the primal-dual algorithm. Initially, $S = \emptyset$, and let $\mathbf{d}(S) = (d_1(S), d_2(S))$ be the residual demand vector. Therefore, initially, $\mathbf{d}(\emptyset) = (4, 1)$. The dual constraints for $S = \emptyset$ are the following (where z_j below corresponds to the dual variable $z(\emptyset, j)$):

$$\begin{aligned} \text{constraint of A:} \quad & 4z_1 \leq 4 \\ \text{constraint of B:} \quad & 3z_1 \leq 3 \\ \text{constraint of C:} \quad & z_2 \leq 4 \\ \text{constraint of D:} \quad & 4z_1 + z_2 \leq 4 \end{aligned}$$

Since $d_1(\emptyset) > d_2(\emptyset)$ we start by increasing z_1 and we observe that (by using TB2), the constraint of worker B will be the first to hold with equality. This results to adding B to S, and now $\mathbf{d}(\{B\}) = (1, 1)$. The relevant dual constraints for the remaining workers (where z_j below corresponds to $z(\{B\}, j)$ are now:

$$\begin{aligned} \text{constraint of A:} \quad & z_1 \leq 4 \\ \text{constraint of C:} \quad & z_2 \leq 4 \\ \text{constraint of D:} \quad & z_1 + z_2 \leq 4 \end{aligned}$$

Using TB1, we start again increasing z_1. We observe (by using TB3), that the constraint of A will be the first to become tight which adds worker A to the solution. Now, $\mathbf{d}(\{B, A\}) = (-3, 1)$ and the dual constraints for $S = \{B, A\}$ are:

$$\begin{aligned} \text{constraint of C:} \quad & z_2 \leq 4 \\ \text{constraint of D:} \quad & -3z_1 + z_2 \leq 4 \end{aligned}$$

The only remaining uncovered task is task 2 and hence we increase variable z_2. Using once again TB3, we add C in the solution and the forward phase of the algorithm terminates with the solution $S = \{B, A, C\}$.

The next step is to run the delete phase, which eliminates B, since worker B is not necessary for feasibility, having C and A present. Hence, the final solution will be $\{A, C\}$. Note that bidders C, A are considered for deletion first, but they are indispensable to maintain feasibility and cannot be deleted.

Consider now a unilateral decrease of the cost of worker A from 4 to $4 - \epsilon$, for a small $\epsilon > 0$. Again, with $S = \emptyset$, the initial dual constraints are:

$$\begin{aligned} \text{constraint of A:} \quad & 4z_1 \leq 4 - \epsilon \\ \text{constraint of B:} \quad & 3z_1 \leq 3 \\ \text{constraint of C:} \quad & z_2 \leq 4 \\ \text{constraint of D:} \quad & 4z_1 + 1z_2 \leq 4 \end{aligned}$$

This will result in the initial selection of A on the first iteration, since A is no longer tied with B. The updated residual demands are $\mathbf{d}(\{A\}) = (0, 1)$. With task 2 being the only task left to be covered, the updated constraints are:

$$\text{constraint of C:} \quad z_2 \leq 4$$
$$\text{constraint of D:} \quad z_2 \leq 4$$

Here, worker D will be selected (since worker B is not in the solution as per TB3) and we obtain the solution $\{A, D\}$. This completes the forward phase.

But now, the delete phase will remove worker A since worker D, who is the first candidate for deletion, can cover both tasks on her own. This concludes the proof, since an initially winning bidder A has decreased her cost and this deviation resulted in her removal. Therefore, the algorithm is not monotone.

B.2 Proof of Theorem 4.3.3

To prove the theorem, we will first introduce the relevant terminology of the local-ratio framework in the design of approximation algorithms. For a more detailed exposition on local-ratio algorithms, we refer the reader to Bar-Yehuda et al. (2005). Roughly speaking, local-ratio algorithms are iterative procedures that terminate when a feasible solution has been constructed, and at each step perform a “weight" decomposition. At each iteration, the algorithm determines a vector of subtractions from the current instance. A subtraction here means that certain weight parameters (typically, the ones that contribute to the cost of the objective function) are reduced, implying that we obtain a new instance with a lower optimal value, but at the same time, this incurs a cost that we will account for in the final solution (determined by the amount subtracted from the weights). The ratio between the incurred cost and the gain in the optimal value (which is the difference between the initial optimal value and the new optimal value in the reduced instance), is called “effectiveness” of the weight decomposition. The approximation factor of a local-ratio algorithm depends on the effectiveness of the proposed decomposition.

We introduce some auxiliary notation. Let z be the total number of iterations that Algorithm 2 performs, and let $\bar{S}_k$ be the set of remaining activities right before iteration k begins, for $k = 1, \ldots, z$. Note that $\bar{S}_1 = \mathcal{J}$, and by slightly abusing notation, we can view the returned feasible schedule S, as the set $\bar{S}_{z+1}$.

It is not hard to notice that Algorithm 2 does indeed implement a local-ratio scheme by focusing on step 6 at each iteration $k = 1, \ldots, z$. In fact, for every activity $i = 1, \ldots, n$, and iteration k, step 6 can be rewritten as $p_{i,k} = w_{1,i}(k) + w_{2,i}(k)$, where

$$w_{1,i}(k) = \begin{cases} \varepsilon_k \min\{R^*(\bar{S}_k, \mathbf{T}, \mathbf{D}), r_i\}, & \text{if } i \in \bar{S}_k \cap \mathcal{J}_{t^*}(\mathbf{T}), \\ 0, & \text{o/w}, \end{cases}$$
$$w_{2,i}(k) = p_{i,k+1}.$$

The following theorem is the adjustment for LMIS, of the local-ratio Theorem for minimization problems of Bar-Yehuda et al. (2005).

Theorem B.2.1. *(implied by Bar-Yehuda et al. (2005)) Consider an instance* ($\mathbf{p}$, $\mathbf{T}$, $\mathbf{r}$, $\mathbf{D}$) *of* LMIS *and let* $\mathbf{w_1}, \mathbf{w_2}$ *be penalty vectors such that* $\mathbf{p} = \mathbf{w_1} + \mathbf{w_2}$. *Let* $S \subseteq \mathcal{J}$ *be a solution that is* α*-approximate for the instances* ($\mathbf{w_1}$, $\mathbf{T}$, $\mathbf{r}$, $\mathbf{D}$) *and* ($\mathbf{w_2}$, $\mathbf{T}$, $\mathbf{r}$, $\mathbf{D}$). *Then,* S *is* α*-approximate for* ($\mathbf{p}$, $\mathbf{T}$, $\mathbf{r}$, $\mathbf{D}$).

Hence, to prove that Algorithm 2 is α-approximate we only need to satisfy the conditions of Theorem B.2.1. To that end, it suffices to consider the first iteration, where the input vector $\mathbf{p}$ is decomposed into two vectors. In fact, we establish something stronger, regarding the decompositions in all iterations, and show that for the returned solution S of the algorithm, the following hold at every iteration $k = 1, \ldots, z$:

$$\sum_{j \in \mathcal{J} \setminus S} w_{1,j}(k) \leq \alpha \sum_{j \in \mathcal{J} \setminus OPT(\mathbf{w_1}(k))} w_{1,j}(k) \tag{B.1}$$

$$\sum_{j \in \mathcal{J} \setminus S} w_{2,j}(k) \leq \alpha \sum_{j \in \mathcal{J} \setminus OPT(\mathbf{w_2}(k))} w_{2,j}(k) \tag{B.2}$$

Here, S is the feasible schedule returned by Algorithm 2 and $OPT(\mathbf{w_1}(k))$, $OPT(\mathbf{w_2}(k))$ are the optimal solutions for the respective instances of the decomposition. Before continuing, it is important to notice an important feature of this decomposition: at any iteration k, the constant ε_k is such that for the index j^* selected at iteration k, it holds that $w_{2,j^*}(k) = 0$, whereas $w_{2,i}(k) \geq 0$, for all $i \neq j^*$. Furthermore, since $w_{2,j}(z) = 0$ for all $j \in \mathcal{J} \setminus S$, and hence the total cost of S w.r.t. w_2 is zero, we trivially obtain that $S = OPT(\mathbf{w_2}(z))$. This fact will be crucial for the proof of the following claim.

Claim 3. *Suppose that* $S \subseteq \mathcal{J}$ *is such that Equation* (B.1) *holds for* $k = 1, \ldots, z$, *and some value of* α. *Then, for* $k = 1, \ldots, z$, *Equation* (B.2) *holds as well.*

Proof. For $k = z$ the claim follows trivially by the discussion right before. For any job $j \in \mathcal{J} \setminus S$ and any $k = 1, \ldots, z - 1$, by the definition of the weight decomposition, we have that

$$\begin{aligned} &w_{2,j}(k) = p_{j,k+1} = w_{1,j}(k+1) + w_{2,j}(k+1) = \\ &w_{1,j}(k+1) + w_{1,j}(k+2) + w_{2,j}(k+2) = \cdots = \textstyle\sum_{\ell=k+1}^{z} w_{1,j}(\ell) + w_{2,j}(z). \end{aligned} \tag{B.3}$$

Since for $j \in \mathcal{J} \setminus S$, we have $w_{2,j}(z) = 0$, by summing for all jobs j in $\mathcal{J} \setminus S$ and by our assumption that Equation (B.1) holds, we have,

$$\begin{aligned} \sum_{j \in \mathcal{J}\setminus S} w_{2,j}(k) &= \sum_{j \in \mathcal{J}\setminus S} \sum_{\ell=k+1}^{z} w_{1,j}(\ell) \\ &= \sum_{\ell=k+1}^{z} \sum_{j \in \mathcal{J}\setminus S} w_{1,j}(\ell) \\ &\le \alpha \sum_{\ell=k+1}^{z} \sum_{j \in \mathcal{J}\setminus OPT(\mathbf{w_1}(\ell))} w_{1,j}(\ell) \\ &\le \alpha \sum_{j \in \mathcal{J}\setminus OPT(\mathbf{w_2}(k))} \sum_{\ell=k+1}^{z} w_{1,j}(\ell) \\ &\le \alpha \sum_{j \in \mathcal{J}\setminus OPT(\mathbf{w_2}(k))} w_{2,j}(k), \end{aligned}$$

where the first equality and the last inequality are due to Equation (B.3) and the fact that $w_{2,j}(k) \ge 0$, and the second inequality holds because $OPT(\mathbf{w_1})$ is optimal w.r.t. $\mathbf{w_1}$, but $OPT(\mathbf{w_2})$ is just feasible for the instance ($\mathbf{w_1}$, $\mathbf{T}$, $\mathbf{r}$, $\mathbf{D}$), (note that feasibility is not affected by the weights).

□

With the above claim, it suffices to focus on proving Equation (B.1) for all k, and for some value of α. We prove below that indeed, Equation (B.1) holds for $\alpha = \Delta$.

Lemma B.2.1. *Let $S \subseteq \mathcal{J}$ be a solution of Algorithm 2 for the* LMIS *instance* $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$. *Then, for any k,*

$$\sum_{j \in \mathcal{J}\setminus S} w_{1,j}(k) \le \Delta \cdot \sum_{j \in \mathcal{J}\setminus OPT(\mathbf{w_1}(k))} w_{1,j}(k)$$

Proof. Let $R_k^* = R^*(\bar{S}_k, \mathbf{T}, \mathbf{D})$ and let t^* be the time instant considered at iteration k. By the definition of $w_{1,j}(k)$, and the fact that the number of jobs, whose interval contains t^*, are at most Δ, we have:

$$\sum_{j \in \mathcal{J}\setminus S} w_{1,j}(k) = \sum_{j \in (\mathcal{J}\setminus S)\cap(\bar{S}_k \cap \mathcal{J}_{t^*}(\mathbf{I}))} \varepsilon_k \min\{R_k^*, r_j\} \le \Delta \varepsilon_k R_k^*.$$

To complete the proof, we need to show the following:

$$\sum_{j \in \mathcal{J} \setminus OPT(\mathbf{w_1}(k))} w_{1,j}(k) \geq \varepsilon_k R_k^*. \tag{B.4}$$

For brevity, let $OPT(\mathbf{w_1}(k)) = OPT_1$. Observe that, since OPT_1 is a feasible solution it holds that $\sum_{j \in \mathcal{J}_{t^*}(\mathbf{T}) \cap OPT_1} r_j \leq D_{t^*}$, or equivalently, that

$$\sum_{j \in \bar{S}_k \cap \mathcal{J}_{t^*}(\mathbf{T})} r_j - \sum_{j \in \mathcal{J}_{t^*}(\mathbf{T}) \cap OPT_1} r_j \geq \sum_{j \in \bar{S}_k \cap \mathcal{J}_{t^*}(\mathbf{T})} r_j - D_{t^*} = R_k^*,$$

and since $(\bar{S}_k \cap \mathcal{J}_{t^*}(\mathbf{T})) \setminus (\mathcal{J}_{t^*}(\mathbf{T}) \cap OPT_1) = (\bar{S}_k \cap \mathcal{J}_{t^*}(\mathbf{T})) \setminus OPT_1$, this is equivalent to

$$\sum_{j \in (\bar{S}_k \cap \mathcal{J}_{t^*}(\mathbf{T})) \setminus OPT_1} r_j \geq R_k^*. \tag{B.5}$$

Now to prove that Equation B.4 holds, we have

$$\begin{aligned} \sum_{j \in \mathcal{J} \setminus OPT(\mathbf{w_1}(k))} w_{1,j}(k) &= \sum_{j \in (\bar{S}_k \cap \mathcal{J}_{t^*}(\mathbf{T})) \setminus OPT_1} \varepsilon_k \min\{R_k^*, r_j\} \\ &= \varepsilon_k \left(\sum_{\substack{j \in (\bar{S}_k \cap \mathcal{J}_{t^*}(\mathbf{T})) \setminus OPT_1 \\ s.t. \quad r_j \leq R_k^*}} r_j + \sum_{\substack{j \in (\bar{S}_k \cap \mathcal{J}_{t^*}(\mathbf{T})) \setminus OPT_1 \\ s.t. \quad r_j > R_k^*}} R_k^* \right). \end{aligned}$$

When there is a job $j \in \mathcal{J}$ with $r_j > R_k^*$ in $(\bar{S}_k \cap \mathcal{J}_{t^*}(\mathbf{T})) \setminus OPT_1$, the claim follows. When no such job exists, we use Equation (B.5) to complete the proof.

□

Tightness of the approximation ratio: Consider the LMIS instance with n jobs of unit width, namely $\{1, 2, \ldots, n\}$, to be scheduled on a machine active on n time-instants, namely $\{t_1, t_2, \ldots, t_n\}$. Throughout the example, we assume ties break lexicographically. At time instant t_i, the amount of available resource is $n - i$, except for instant t_n, where it is $1 - \epsilon$, for some $\epsilon \in (0, 1)$. Let the interval in which job i is active be $T_i = \{1, \ldots, t_i\}$ for $i = 1, 2 \ldots, n$. The penalties of the jobs are given by the vector $\mathbf{p}$ such that $p_i = 1 - y(n - i + 1)$ for $i = 1, 2 \ldots, n$, and we fix $y = \frac{2}{(n+1)(n^2-2n+2)}$. Observe that $\Delta = n$.

The optimal solution is to schedule all jobs except j_n at a cost of $1-y=$

$$1-\frac{2}{(n+1)(n^2-2n+2)}=\frac{(n+1)(n^2-2n+2)-2}{(n+1)(n^2-2n+2)}=\frac{n^2(n-1)}{(n+1)(n^2-2n+2)}$$

The first iteration of Algorithm 2 will select $t^*=t_1$, due to the tie-breaking rule, since $R^*=\max\limits_{i=1,\ldots,n-1}\{R_i\}=\max\{1,\ldots,1\}=1$. Then, it will reduce the cost of all jobs by ε_1, where ε_1 is defined as the unique quantity that makes the penalty of the first job (which has the lowest penalty) to drop down to zero (in which case it equals p_1), after the update in step 6 of the algorithm. Hence the first job will be removed from the schedule. The current solution remains infeasible, and continuing analogously, in the second iteration, time instant t_2 will be selected and the second job will be the one to be removed. It is easy to verify that the current solution will remain infeasible after every iteration, and Algorithm 2 will gradually remove from the schedule all jobs. Hence, the output of the algorithm is an empty schedule, which is feasible, and the solution cost equals the sum of all job penalties, which is:

$$\sum_{i=1}^{n}[1-y(n-i+1)]=n-yn^2-yn+y\sum_{i=1}^{n}i=n-\frac{yn(n+1)}{2}=\frac{n(n-1)^2}{n^2-2n+2}.$$

But then, the ratio between the solution cost and the optimal cost is $\frac{(n-1)(n+1)}{n}=n-\frac{1}{n}$, which approaches Δ as n grows larger.

B.3 Proof of Proposition 4.4.1

The authors in Chen et al. (2019) study a problem that can be seen as a generalization of the Min-Knapsack problem and in our model corresponds to the case of having a single task. It is claimed that Algorithm 3 in Chen et al. (2019) is monotone, based on the framework of Briest et al. (2011) and leads to a truthful FPTAS. In this section we show that this was not a correct application of the technique by Briest et al. (2011), and their algorithm is in fact non-monotone.

For the precise statement of the algorithm, we refer to Chen et al. (2019). We provide here a brief overview, adapted to our notation, and we focus on the special case of their problem that corresponds to Min-Knapsack. Their algorithm is similar in spirit to our Algorithm A_{FPTAS} in Section 4.4, which in turn uses the algorithms A_k (described by Algorithm 4.4). The differences with our work for each algorithm A_k, are as follows:

- The set $\mathcal{L}_k(\mathbf{c})$ is defined as: $\mathcal{L}_k(\mathbf{c}) = \{i \in \mathcal{N} : c_i \leq 2^k\}$,
- the parameter a_k is defined as: $a_k = \frac{\epsilon 2^k}{n+1}$,
- the rounding is different, namely, $\bar{c}_i = \lfloor \frac{c_i}{a_k} \rfloor$.

Most importantly, the for loop in their algorithm A_{FPTAS} is examining the algorithms A_k only for $k = 1, \dots, \log c_{max}(\mathbf{b})$. This is a crucial point as the number of iterations do not suffice for monotonicity (in contrast to the number of iterations that we use in Section 4.4, which is equal to $\log(nc_{max}(\mathbf{b})/\epsilon + 1)$. We note that it is not specified in Chen et al. (2019) if $\log c_{max}(\mathbf{b})$ is rounded to its floor or ceiling. Hence we are going to consider both cases. We will prove that for each case there exists an instance in which the decrease of the cost of a winning worker could result in her exclusion from the final solution.

Case 1. Say that A_{FPTAS} runs algorithm A_k for $k = 1, \dots, \lceil \log c_{max}(\mathbf{b}) \rceil$. We fix $\epsilon = 0.097$ and we create an instance with a single task of demand 1, and 3 workers, namely A, B and C, the costs and the qualities of which are shown in the following table:

i	c_i	q_i
A	2.01	1
B	1.01	0.5
C	1.01	0.5

Note that maximum cost is slightly above 2, and hence $\lceil \log c_{max} \rceil = 2$. Thus, the number of algorithms that are going to be executed by A_{FPTAS} equals 2. Both algorithms A_1 and A_2 compute a scaling factor a_1 and a_2 respectively. Their values for the selected ϵ and since $n = 3$, are $a_1 = 0.0485, a_2 = 0.097$. The rounded cost $\bar{c}_i^{(k)}$ of element i concerning algorithm A_k is shown below ($\bar{c}_i^{(k)} = \lfloor c_i/a_k \rfloor$, (this is denoted as c_i' in their work).

i	$\bar{c}_i^{(1)}$	$\bar{c}_i^{(2)}$	q_i
A	41	20	1
B	20	10	0.5
C	20	10	0.5

It is now easy to see that A_1 will return elements B and C having a total cost of 2.02, since element A will not be examined by the algorithm due to its high cost. On the other hand, A_2 will return the solution $\{A\}$, of total cost 2.01 and thus A_{FPTAS} will return $\{A\}$, given that it is the solution of minimum cost computed by the executed algorithms.

Let us now examine what will happen if agent A declares a lower cost, in particular, suppose he reports a cost of 2 instead of 2.01. We claim that such a decision will leave

A outside of the solution. In that case, just a single algorithm A_1 (for $k = 1$) will be executed by A_{FPTAS} since $\lceil \log 2 \rceil = 1$. We observe that the scaled cost of element A will not change after that deviation (since $\lfloor \frac{2}{a_1} \rfloor = \lfloor \frac{2.01}{a_1} \rfloor$), and of course the scaled costs of the remaining elements will remain the same as well. Thus, element A will not be in the solution of A_{FPTAS} this time because $40 = 2\lfloor \frac{1.01}{a_1} \rfloor < \lfloor \frac{2}{a_1} \rfloor = 41$ and thus A_1 will return $\{B, C\}$ once again. Hence, by lowering his cost, A stopped being a winner.
Case 2. Say that A_{FPTAS} runs algorithm A_k for $k = 1, \ldots, \lfloor \log c_{max}(\mathbf{b}) \rfloor$. We create an instance with a single task of unit demand, and n workers, whose costs and qualities are shown below:

worker	cost	q_1
A	4	1
B_1	1	$\frac{1}{n-1}$
...	...	...
B_{n-1}	1	$\frac{1}{n-1}$

It is easy to see that exactly 2 algorithms will be evaluated from A_{FPTAS}. In A_1, all workers except A will be considered and hence the only feasible solution is to select them all. In A_2, worker A will also be considered. In that algorithm $a_2 = \frac{4\epsilon}{n+1}$ and $\bar{c}_i = \lfloor \frac{c_i(n+1)}{\epsilon} \rfloor$. The dynamic programming procedure runs optimally for the rounded costs and hence it will produce the feasible solution $\{A\}$, the cost of which is less than the cost of the solution returned by A_1. Thus, A_{FPTAS} will return A as a single winner.

Suppose now that A is changing her cost to $4-\epsilon$. That negligible deviation will have a significant impact on the number of algorithms that will be called by A_{FPTAS} since $\lfloor \log c_{max}(\mathbf{b}) \rfloor = 1$. Given that $c_{max}(\mathbf{b}) < 2^1$, worker A will again not be considered in A_1, by the definition of $\mathcal{L}_k(\mathbf{c})$ and hence, the only feasible solution is to select all workers but A. This means that A will no longer be part of the solution returned by A_{FPTAS}.

We have concluded that the FPTAS by Chen et al. (2019) is not monotone as claimed.

B.4 Handling the case of $c_{min}(\mathbf{b}) < 1$

We remind the reader that in Section 4.4 we assumed that $c_i \geq 1$ for all $i \in \mathcal{N}$. Essentially, one of the reasons to do that is to be able to use Lemma 4.4.1 for $k = 0$ (our lowest index) and this requires that, for a bidding profile $\mathbf{b}$, $OPT(\mathbf{b}) \geq 1$. The second reason is to enforce that $c_{max}(\mathbf{b}) \geq 1$, which means that the number of iterations performed by A_{FPTAS}, is well-defined. These two conditions suffice in order to be able to run A_{FPTAS} and have the desired guarantee.

Claim 4. *For a bidding profile* $\mathbf{b}$*, if* $c_{max(\mathbf{b})} \geq 1$*, and* $OPT(\mathbf{b}) \geq 1$*, then we can run* $A_{FPTAS}(\mathbf{b})$ *and retain all the properties established in Section 4.4.*

Hence, the assumption that $c_i \geq 1$ for every $i \in \mathcal{N}$ helps us enforce what we need in the above claim. In this subsection, we present a way to adjust the assumption to: $c_i \geq \delta$ for all $i \in \mathcal{N}$ for any arbitrary value δ such that $1 > \delta > 0$.

Given a $\delta > 0$, we will first provide a reduction from an instance of CMIC with $c_i \geq \delta$ for all $i \in \mathcal{N}$ to an instance with $\text{OPT}(\mathbf{b}) \geq 1$ and $c_{max}(\mathbf{b}) \geq 1$, and then we will argue that this is without any loss in the approximation factor.

Suppose that we want to solve an instance of CMIC $I = (\mathbf{b}, \mathbf{q}, \mathbf{d})$, with n workers and m tasks such that $c_i \geq \delta$. Obviously, if it also holds that $c_i \geq 1$ for every i, we just run A_{FPTAS}. If not, we suggest running A_{FPTAS} for an instance I' of $n+1$ workers (the n workers of I and an extra dummy worker say x), and $m+1$ tasks (the m tasks of I and an extra dummy task say $m+1$, so that we have the interval from 1 to $m+1$). In this modified instance, the dummy worker x is only interested in contributing to task $m+1$, and at the same time no other worker is interested in contributing to $m+1$. Let both the contribution of x and the demand of $m+1$ be equal to 1. Let also the cost of the new worker be equal to 1, whereas the workers and the tasks of I, remain exactly the same in I'.

Notice that in the instance I' it necessarily holds that $c_{max} \geq 1$, no matter what the initial bidding profile in I is. Also, every optimal solution of the new instance I' has a cost that is greater than 1 since x must be included in every feasible solution and her cost is 1. Therefore, A_{FPTAS} for I' can run without a problem, since the conditions of Claim 4 are satisfied. Furthermore, due to the monotonicity of A_{FPTAS}, we can easily verify that the new algorithm is also monotone. At what concerns the approximation ratio we claim that we can still obtain a $(1+\epsilon)$-approximation, if we run $A_{\text{FPTAS}}(\mathbf{b}, \epsilon')$ for $\epsilon' = \epsilon\delta/(\delta+1)$. By doing so, if $SOL(I')$ is the solution returned for I', and $SOL(I)$ is the solution we get by removing the dummy worker, we obtain

$$SOL(I) + 1 = SOL(I') \leq (1+\epsilon')OPT(I') =$$
$$(1+\epsilon')(OPT(I)+1) = (1+\epsilon')OPT(I) + 1 + \epsilon',$$

where the inequality is due to the fact that A_{FPTAS} is a FPTAS. By rearranging terms we continue as follows:

$$SOL(I) \leq (1+\epsilon')OPT(I) + \epsilon' = (1+\epsilon')OPT(I) + \frac{\delta\epsilon'}{\delta}$$
$$\leq (1+\epsilon')OPT(I) + \frac{\epsilon'}{\delta}OPT(I) = (1+\epsilon)OPT(I)$$

Here, the last inequality holds since $\delta \leq c_{min} \leq OPT(I)$. With this transformation, the resulting algorithm is once again an FPTAS. Crucially, the monotonicity is not destroyed since the value we use in A_{FPTAS} for ϵ' does not in any way depend on the profile $\mathbf{b}$, since no bidder can affect it by any unilateral deviation.