The Gap and the Gain:

Improving the Approximate Mechanism Design Frontier in Constrained Environments

Sophie Klumper

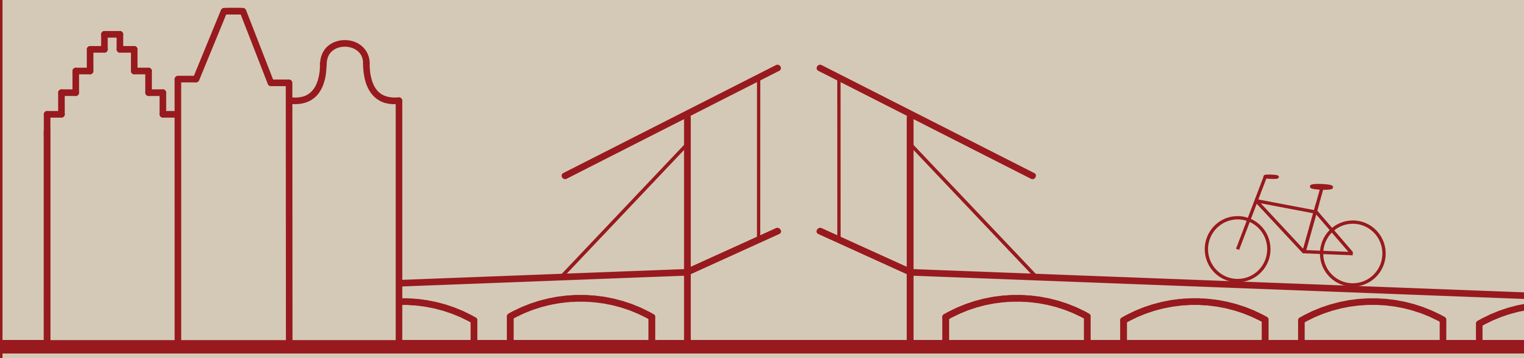

The Gap and the Gain:

Improving the Approximate Mechanism Design Frontier in Constrained Environments

ILLC Dissertation Series DS-2025-08

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

For further information about ILLC-publications, please contact

Institute for Logic, Language and Computation
Universiteit van Amsterdam
Science Park 107
1098 XG Amsterdam
phone: +31-20-525 6051
e-mail: illc@uva.nl
homepage: http://www.illc.uva.nl/

This research has been carried out in the Networks and Optimization group at the Centrum Wiskunde & Informatica in Amsterdam, supported by the Open Technology Program of the Dutch Research Council (NWO), project number 18938.

Cover design by Sophie Klumper.
Printed and bound by Ipskamp Printing.

ISBN: 978-90-6196-431-5

The Gap and the Gain:

Improving the Approximate Mechanism Design Frontier in Constrained Environments

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. P.P.C.C. Verbeek
ten overstaan van een door het College voor Promoties ingestelde commissie,
in het openbaar te verdedigen in de Aula der Universiteit
op donderdag 23 oktober 2025, te 14.00 uur

door

Sophie Joan Klumper

geboren te Amsterdam

Promotiecommissie

Promotores:	prof. dr. G. Schäfer	Universiteit van Amsterdam
	prof. dr. R.D. van der Mei	Vrije Universiteit Amsterdam
Copromotor:	prof. dr. S. Bhulai	Vrije Universiteit Amsterdam
Overige leden:	prof. dr. M. Feldman	Tel Aviv University
	prof. dr. E. Markakis	Athens University of Economics and Business
	prof. dr. C. Ventre	King's College London
	dr. D. Grossi	Universiteit van Amsterdam
	dr. rer. nat. R.E.M. Reiffenhäuser	Universiteit van Amsterdam
	prof. dr. ing. R.A.M. van Rooij	Universiteit van Amsterdam

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

Contents

Part Three: Inefficiency of Equilibria

Final Remarks and Appendix

Acknowledgments

I am grateful to have been given the opportunity to pursue my PhD at CWI and to have been a part of the Algorithmic Game Theory community. I would like to thank everyone who has contributed to this memorable journey.

I especially want to thank my promoter and daily supervisor Guido Schäfer. Guido, after you introduced me to AGT during my master's at the VU, I became intrigued by the topic and we discussed the possibility of you supervising my master's thesis. You introduced the problem of budget-feasible mechanism design, and even though mechanism design was my least favorite part of the course (back then!), I agreed on the topic as I considered you an inspiring and knowledgeable lecturer, and I really wanted to work together. At the time I didn't know it would ultimately lead to a PhD thesis! A great decision in hindsight. I still appreciate how approachable you were, and for showing me that doing research could be fun! And so, when looking for a PhD position after a little detour in industry, I was very excited to start a position with you at CWI. After having made seemingly limited progress in the first months, and with limited in-person interactions due to the pandemic, I started doubting whether I had made the right decision. I'm truly grateful for your understanding at that time, for your continued belief in me, and for encouraging me to take some time over Christmas to reflect. Shortly after deciding to continue, I was fortunate to experience some of the special moments of academia when our first paper was accepted at SAGT, and I got to attend my first workshop abroad. Since then, despite the inevitable struggles of pursuing a PhD, it has been a great three-and-a-half years filled with unforgettable memories. I want to thank you for your guidance and support, for introducing me to great people within our field, for making almost every research meeting fun, and for your dedication to everything you do. Especially in the final months of wrapping up the thesis, I truly appreciated that your door was always open.

Mijn promotor Rob van der Mei en copromotor Sandjai Bhulai wil ik bedanken voor hun vertrouwen en enthousiasme, en de vrijheid binnen het project. Rob, bedankt voor al je positieve woorden, vooral na presentaties, dit heb ik altijd erg gewaardeerd!

Furthermore, I would like to express my appreciation to the other members of the doctorate committee. Michal Feldman, Evangelos Markakis, Carmine Ventre, Davide Grossi, Rebecca Reiffenhäuser, and Robert van Rooij, thank you for taking the time to assess my thesis and for participating in the doctoral defense ceremony.

This thesis would not have been possible without my other co-authors, which I would like to thank for the great collaborations. Vangelis, Yorgos and Artem, you were the first people outside of CWI that I worked with, and I couldn't have wished for a better experience. Thank you for your kindness, enthusiasm, and insights. Riccardo and Artem, thank you for not giving up on our project when we initially only managed to get negative results. And of course, thank you for all the fun moments. I didn't know writing a rebuttal could be so enjoyable. Stefano, thank you for agreeing to host Guido, Artem and me at Sapienza in Rome after I bluntly asked you if I could visit. I think we all had an amazing time. This was in large part due to the kindness and warmth of you and your group. I'm grateful for all the time you freed up, not only for discussing research, but also for the occasional aperitivo, dinner, or gelato. Argyris, thank you for proposing a multi-day visit to Royal Holloway. It was my first research visit alone, and besides being excited about it, I was also a bit nervous. But you, Eduard, and the whole group were very welcoming, quickly making me feel part of the group. I appreciate how much time you and Eduard freed up to talk about research and besides being productive, our whiteboard sessions were very enjoyable!

Carmine and Artem, thank you for being partners in crime to get SAGT 2024 to Amsterdam. Surprisingly, it all worked out! It's quite some work to plan a conference on top of everything else, but I had a lot of fun organizing it together with you, Guido, and Susanne. Thank you also for the time spent together at workshops and conferences. Besides enjoying beach volleyball and dinners together, I'm grateful for your support and the fact that I could always count on a few familiar faces during talks or poster sessions.

It has been a pleasure to work at CWI for the last four years, in large part due to many wonderful colleagues. Thank you to all the members of the Networks and Optimization group for the great atmosphere, and thank you to our senior researchers for being kind and approachable while conducting excellent research. I think this is a very valuable aspect of our group.

A special thank you to all the office mates of M236, and in particular Willem and Artem. Willem, bedankt voor alle gezelligheid en steun. Het was ontzettend waardevol om samen hard te kunnen werken, en vervolgens samen een rondje te wandelen. Ik ben blij dat we na onze studietijd collega's zijn geweest, en zowel alle leuke dingen buiten werk als onze PhD struggles konden delen op kantoor. Artem, I'm grateful that you decided to join CWI. I truly believe that my PhD would not have been as fun and fruitful otherwise. Thank you for being kind, open, and honest, and always having time for my questions. There are many memories and jokes that still put a smile on my face when I think about them. I am already looking forward to a future project together, but before that, I am looking forward to you being my paranymph during the defense.

Dank ook aan alle andere collega's op het CWI. Ik kijk met plezier terug op de gesprekjes in de ochtend met Remco, op de gang op de tweede verdieping met Karima of Rick, bij de koffiemachine met Paul of Ton, en de ontspanningsmomenten op donderdag met de CWI yogi's. En in het bijzonder, de wat langere gesprekken met Bikkie, Minnie, Susanne en Nada. Bedankt voor jullie warmte en gezelligheid, de soms nodige afleiding, en de aanmoedigingen tijdens het afronden. Als ik aan het CWI denk dan denk ik aan jullie. Bedankt voor het bijdragen aan een fijne plek om te werken.

Mijn vrienden wil ik ook bedanken voor hun steun en alle tijd die we samen hebben doorgebracht. Ik ben dankbaar voor onze vriendschappen, en koester dat we steeds meer onderdelen van het leven samen mogen meemaken. Michelle, het is moeilijk om goed uit te leggen hoe het echt is om een PhD te doen, en daarom was het heel waardevol om bij jou aan een paar woorden genoeg te hebben; bedankt! Fijn dat je mijn paranimf bent.

Mijn hele familie wil ik bedanken voor al hun aanmoedigingen en steun, niet alleen tijdens mijn PhD, maar in alles wat ik doe. Pap en mam, bedankt voor de liefde die jullie geven, alle avonturen die we samen hebben meegemaakt, alles wat jullie me geleerd hebben, en de altijd onvoorwaardelijke steun. Wouter en Jacky, ik ben dankbaar voor onze band, alles wat we al hebben beleefd, en voor wat we samen met Ziko en Alexander nog gaan beleven.

Alexander, dank je voor de afgelopen drieënhalf jaar en voor wie je bent. Je maakt zo veel gewone dagelijkse momenten mooi. Bedankt voor je oprechte interesse in mijn onderzoek, zoals bewijzen echt willen begrijpen en het meedenken op sommige aspecten. Bedankt dat je me aanmoedigt het beste uit mezelf te halen, je steun tijdens de moeilijke momenten, en je enthousiasme om samen op avontuur te gaan. Ik kijk uit naar al onze volgende avonturen.

Sophie
Amsterdam, 2025

Chapter 1

Introduction

Nowadays, algorithms have become an integral part of our daily lives, with many designed to make life more efficient. A concrete example of this is a navigational tool, that uses algorithms to compute the optimal route for you. Roughly speaking, an algorithm receives some information as input, which it processes by applying a predefined set of rules, resulting in a solution that the algorithm returns as output (see Figure 1.1). Additionally, constraints can be defined which the computed solution must satisfy, as well as an objective. In the example of finding a route, the input given to the algorithm includes the desired origin and destination. The input could also include constraints that you specify, such as the means of transportation or avoiding toll roads. The objective could be to find the fastest or the shortest route. The algorithm takes all of this into account when computing a solution. This all sounds very convenient, but you might only be satisfied if the algorithm meets certain crucial requirements. It might be non-negotiable for the route to only use bicycle paths, and to exactly start and end at the origin and destination that you defined, i.e., the algorithm must compute a *feasible* solution. If you are in a hurry, you might want the computed route to indeed be the fastest route satisfying all your constraints, i.e., the algorithm must compute an *optimal solution*. Otherwise, you might be okay with a route that satisfies all your constraints and that takes 2 minutes longer, i.e., the algorithm must compute an *approximate solution*.

Using a navigational tool is an example in which the user defines the problem they wish to see solved. It's up to you to decide what the objective and constraints are. On the other hand, there are many examples in which you are not the one defining the problem, but the solution does affect you. For example, an algorithm

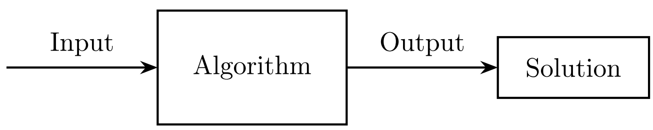

Figure 1.1: Schematic overview of an algorithm.

might be used to set the price of a hotel you are interested in booking. The data used to compute this price might be based on historical demand and not directly depend on you. Moreover, it is likely that the objective of the algorithm, e.g., maximizing the selling price, does not align with your objective.

Let us now consider a problem in which the input directly depends on you. Imagine that your manager has to decide which team member is assigned to which individual task. However, a team member, referred to as an *agent*, might not be able to execute every task due to a certain skill that is required or due to their availability. The set of tasks that an agent is able to execute is referred to as the compatibility, and the manager will have to respect this when constructing an assignment of agents to tasks. Now, suppose that this compatibility is *private* information: only an agent themself knows which tasks they are able to do. And so, the manager asks the agents to declare what their compatibility is in order to construct a solution: an assignment that respects the declared compatibilities and in which each agent is assigned to at most one task. It is common knowledge, i.e., *public* information, how valuable it is when a specific agent executes a specific task. A natural goal of the manager is to compute a solution of maximum value.

Example 1.0.1. Consider an *instance* of the problem described above with two agents and two tasks for which the private compatibilities are depicted by the *bipartite graph* in Figure 1.2a; one agent is able to do both tasks, while the other agent is only able to do one task. The values of the compatible agent-task pairs are also displayed in Figure 1.2a. The manager is looking for a *matching* in the bipartite graph; an assignment in which each agent is assigned to at most one task, and each task is assigned to at most one agent. Suppose that the agents are honest and declare their true compatibilities. We can then use the Hungarian Algorithm [Kuhn, 1955] to compute an optimal matching in polynomial time. The optimal matching is highlighted in Figure 1.2b and has a value of 20.

However, the agents have a goal of their own: they want to maximize the value of the compatible task to which they are assigned. And as the agents are strategic, they will misreport their compatibility in order to achieve this goal. This is also what happens in our example: instead of truthfully declaring the

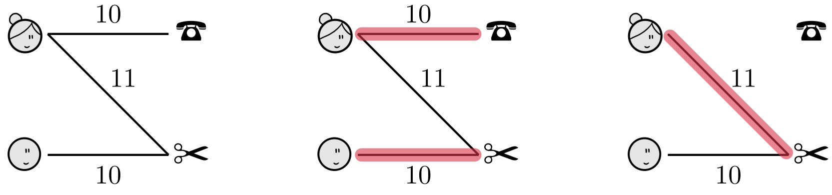

(a) (b) (c)

Figure 1.2: The problem instance used in Example 1.0.1 is depicted in (a). In (b) and (c), the optimal solution is highlighted in red for the given compatibilities.

compatibilities, the agent compatible with both tasks can lie and declare that they are only compatible with the higher-valued task, as depicted in Figure 1.2c. The optimal solution computed by the Hungarian Algorithm will now match this agent with their highest-valued compatible task, and so this agent is better off by lying about their compatibility.

Mechanism Design. The problem faced by the manager in our example falls into the category of *mechanism design problems*. Here, the primary goal is to design a mechanism (see Figure 1.3) for a strategic setting that computes an outcome that is desirable from the designer's perspective. The challenge is that strategic agents hold some private information relevant for computing the outcome and act in their own best interest. There are two standard approaches in the literature to study a mechanism design problem: in the *Bayesian setting* it is assumed that the distributions of the agents' private information are known, whereas in the *prior-free setting* it is assumed that nothing is known about these distributions. In this thesis we only focus on the prior-free setting. A key objective in mechanism design is to incentivize the agents to reveal their private information truthfully. We say that a mechanism is *strategyproof (or truthful)* if it is always in an agent's best interest to reveal their true private information, regardless of the behavior of other agents. The property of strategyproofness is desirable from a societal perspective, as the absence of it may lead to unpredictable behavior and outcomes. And ideally, we want to be able to guarantee something about the quality of the computed outcome. Furthermore, strategyproofness is also desirable from an agent's perspective as there is no need to strategize: they can just declare their private information truthfully as this is in their best interest.

If we revisit our previous example, we can say that the goal is to design a strategyproof mechanism that computes a matching of agents to tasks respecting the compatibilities, and ideally of maximum value. This problem has been solved by Dughmi and Ghosh [2010], who derive a strategyproof mechanism that is *2-approximate*. This means that the ratio between the value of an optimal solution and the value of the solution computed by the mechanism is at most 2 for *any* instance of the problem: it is a *worst-case* guarantee. They also show that this is best possible and prove that there is no strategyproof mechanism that is better than 2-approximate.

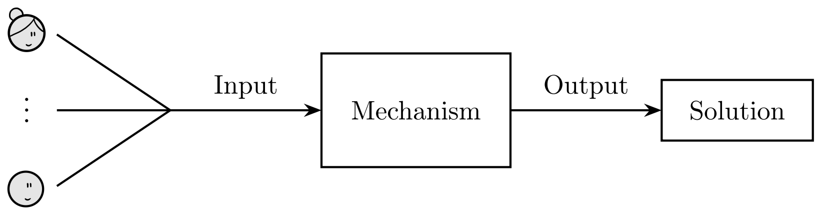

Figure 1.3: Schematic overview of a mechanism.

Learning-Augmented Mechanism Design. In general, and as we have seen in the above example, the requirement of strategyproofness can impose an impossibility result on the approximation guarantee of the mechanism. As a consequence, worst-case guarantees can be rather disappointing from a practical perspective (see, e.g., [Roughgarden, 2021]). A way to circumvent these impossibility results is to consider beyond worst-case analysis. Algorithms with predictions represent one perspective within this paradigm. Here, the goal is to overcome existing worst-case guarantees by augmenting each instance of a problem with a prediction, possibly generated with machine learning. Hence, this line of work is also sometimes referred to as *learning-augmented algorithms*. While this perspective first emerged in the area of online algorithms, it is particularly appealing in the context of mechanism design as the prediction can be informative about the private information of the agents. This resulted into the field of *learning-augmented mechanism design* or *mechanism design with predictions*, introduced by Xu and Lu [2022] and Agrawal et al. [2024]. The nature of the prediction may depend on the problem at hand, e.g., it may predict the private information of the agents or some aggregate information thereof. In this thesis, we will only consider *output predictions*, which are predictions of the optimal outcome.

The goal is to design a strategyproof mechanism that leverages the prediction and provably achieves improved, ideally optimal, approximation guarantees when the prediction is perfect, and ensures that the approximation guarantee does not deteriorate arbitrarily, ideally maintaining the best-known worst-case guarantee when the prediction is erroneous (see [Lykouris and Vassilvitskii, 2021]). If we fully trust the prediction, i.e., use it as a solution, and it turns out that the prediction is perfect, this would lead to an optimal solution, i.e., 1-consistency. However, we do not know if a prediction is perfect as it concerns private information, and so fully trusting the prediction would lead to unbounded robustness as the prediction could be completely erroneous. Therefore, the challenge is to design a mechanism that trusts and uses the prediction to some extent.

We study a generalization of Example 1.0.1 in this framework of learning-augmented mechanism design. Namely, we consider the generalized assignment problem in the private graph model with output predictions. Besides a value, each agent-task pair has a non-negative size, and each task has a positive capacity. The sizes and capacities are also public information, and it is allowed to assign multiple agents to a single task as long as the total size that is assigned does not exceed the capacity of the task. Note that our example is a specific variant of this problem in which all sizes and capacities are equal to 1. The main question that we answer in this thesis is the following:

Can we derive strategyproof mechanisms that leverage output predictions for the generalized assignment problem in the private graph model?

Another example of a mechanism design problem is *facility location* on the real line (see, e.g., [Moulin, 1980]). In this problem, the mechanism designer must determine the location of a facility on a line. The facility will serve a set of agents, which each have a private preferred location for the facility. After the location of the facility is chosen, the agents will incur a cost that is equal to the distance between their private location and the facility, a quantity they wish to minimize. The goal of the mechanism designer is to design a strategyproof mechanism that minimizes some social cost objective. A well-studied objective is minimizing the egalitarian social cost, which is the maximum cost of an agent. For this objective, a strategyproof 2-approximate mechanism is known due to Procaccia and Tennenholtz [2013], who also show that this is best possible. However, in the following example we show that a single agent can have a disproportional influence on the location of the facility, which raises the question of whether the tight 2-approximation is where the story should end. 1

Example 1.0.2. Consider an instance of facility location on the line consisting of 5 agents with private locations depicted as crosses in Figure 1.4. If we use the mechanism in [Procaccia and Tennenholtz, 2013] that chooses the first order statistic (smallest value) as the location of the facility, the facility will be placed at 2 (depicted in red). In this case, the agent with a location at 2 incurs no cost at all, whereas all other agents incur a cost of at least 6. Even for the non-strategic problem, the optimal location of the facility is at 6 (depicted in blue), placing all but one agent on the right side of the facility. Wouldn't it be better to disregard the single, very distant agent on the left? If we would exclude this agent the first order statistic will place the facility at 8 (depicted in green). In this case, the majority of the agents would incur a cost of at most 2, whereas only the agent with a location at 2 would incur a cost of 8.

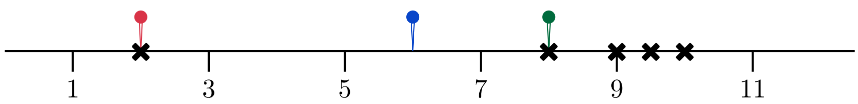

Figure 1.4: Private locations of the agents in Example 1.0.2.

Mechanism Design with Outliers. We introduce a new perspective on mechanism design problems and introduce the notion of *outliers*. Although outliers have been widely studied in algorithm design (see, e.g., [Charikar et al., 2001]), to the best of our knowledge, it has not yet been considered in the context of mechanism design. In the domain of mechanism design with outliers, the social cost objective only accounts for a certain number of agents, instead of all of them. As we have seen in Example 1.0.2, a single agent might greatly influence the location of the facility which might negatively affect the costs of the majority of the agents. This makes this new domain especially relevant for mechanism design problems in which a number of agents have extreme or atypical private information. We use facility location as a first, natural test case and consider the following research question:

Using facility location on the line as a test case, can we derive strategyproof mechanisms that leverage outliers and achieve improved approximation guarantees for this new domain?

Budget-Feasible Mechanism Design. All mechanism design problems that we have considered so far belong to the domain of mechanism design without money introduced by Procaccia and Tennenholtz [2013]. Here, monetary transfers cannot be used to incentivize the agents to reveal their private information truthfully. There are also mechanism design problems in which monetary transfers are allowed, but in which the payments are limited. Namely, consider a *knapsack procurement auction* in which the auctioneer (buyer) is interested in buying goods (or services) from a group of strategic agents (sellers). Each agent specifies a cost for the good that they offer, but their true cost for producing the good is private information. The auctioneer has a value for each good and a maximum budget they can spend, which are both public information, and wants to maximize the total value of the acquired goods while not exceeding their budget. In order to achieve this, the auctioneer wants to use a strategyproof, *budget-feasible* mechanism, prescribing which goods to buy and at what price. Note that the auctioneer only cares that the payments are budget-feasible, i.e., within budget, and does not value left-over budget. Budget-feasible mechanism design was introduced by Singer [2010], who, among other things, studied the knapsack procurement auction for *indivisible* agents, meaning that the good of an agent is either fully acquired or not at all. Motivated by agents offering services, such as in crowdsourcing, Anari et al. [2018] studied this problem for *divisible* agents—meaning that any fraction of the service can be acquired—while assuming that the cost of a single agent is very small compared to the budget of the auctioneer (large market assumption).

In other scenarios, it might make sense to allow the auctioneer to only partially buy what an agent is offering, e.g., an agent might have multiple copies of a good. In this case, we say that *partial allocations* are allowed and as the granularity of the partial allocation increases, this setting converges to the setting of divisible agents. We study the setting of partial allocations under the assumption that the budget of the auctioneer is large enough to acquire everything from each individual agent separately, and focus on the following question:

Can we derive budget-feasible mechanisms for knapsack procurement auctions that are truthful and achieve improved approximation guarantees if partial allocations are allowed?

Price of Anarchy. Up to now, we have only focused on the design of mechanisms in strategic settings, but it could be that the environment in which the strategic agents interact is already defined. Consider for example a single-item first-price auction: agents simultaneously declare a bid for the item that is up for sale, and the agent with the highest bid wins the item and has to pay a price equal to their bid. In this case, we might still be interested in the outcomes that arise. More specifically, we might be interested in the efficiency of *equilibria* resulting from the strategic interactions. In an equilibrium outcome, there is no agent who has an incentive to change their strategy—their bid in the example—given the strategies of all other players. A measure for the efficiency of equilibria is the price of anarchy (POA), introduced by Koutsoupias and Papadimitriou [1999], which is equal to the worst-possible ratio of an optimal outcome to an equilibrium outcome.

In this thesis, we consider simultaneous single-item first-price auctions. In this case, multiple items are auctioned off at the same time, and the bids of the agents consist of a bid for each item separately. Every individual item is sold to the agent with the highest bid for this item, and this agent has to pay a price equal to their bid. In the classical model of utility maximizers, the utility of each agent is equal to their total value for the acquired items minus the total payment they have to make. We study these auctions under the hybrid agent model in which agents have different payment sensitivities, meaning that the total payment in the utility function of an agent is discounted by a parameter between 0 and 1. Additionally, each agent in our model has a return-on-investment constraint. An example of an application of this model is online advertising: advertisers may want to maximize their value for the acquired advertisement slots, i.e., their payment sensitivity is 0, given some budgetary constraint. Our goal is to study the loss in efficiency due to the strategic behavior of the agents:

Can we derive bounds on the price of anarchy of equilibria of simultaneous first-price auctions under the hybrid agent model?

1.1 Overview and Publications

In Chapter 2, we lay the groundwork for the rest of this thesis. We introduce general notation and concepts that will be used throughout the subsequent chapters, such as incentive compatibility and approximation notions. To avoid an overload of notation and improve readability, problem specific notation and concepts will be introduced in the relevant chapter. The subsequent chapters are subdivided into three parts, each highlighting a different aspect of improvements or developments of the efficiency frontier in constrained environments with strategic agents. The author of this thesis is one of the main contributors to the publications and manuscripts mentioned in this subsection.

1.1.1 Part One: Mechanism Design without Money

In the first part of this thesis, we consider two problems in the domain of mechanism design without money. To both problems, we apply the perspective of an emerging new line of research that bridges worst-case analysis and machine learning. Namely, we augment the problems with a prediction[1] of the optimal solution, with the goal of achieving improved approximation guarantees when the prediction is perfect (consistency), while maintaining close to worst-case approximation guarantees when the prediction is imperfect (robustness).

Chapter 3. We study variants of the generalized assignment problem (GAP) in the private graph model, which has been considered in the literature in the non-augmented setting. In this problem, the mechanism designer knows the value and size that an agent has for a resource, but the agent-resource compatibilities are private information. The aim is to find a strategyproof mechanism that computes an assignment of agents to resources of maximum value, while respecting the capacity constraints.

For the GAP variant where all sizes and capacities are equal to 1, i.e., the bipartite matching problem (BMP), we design a deterministic mechanism that is weakly group-strategyproof (WGSP) and achieves the optimal consistency-robustness trade-off. We improve upon this approximation by designing a randomized mechanism that is universally WGSP. For more general variants of GAP, our universally WGSP mechanism randomizes over a greedy mechanism, our mechanism for BMP, and the predicted assignment. Additionally, a confidence parameter can be set for each mechanism which determines the trade-off between the respective consistency and robustness guarantees. Choosing a higher confidence value leads to a better consistency but a worse robustness, and vice versa.

This chapter is based on the following publication:

- R. Colini-Baldeschi, S. Klumper, G. Schäfer and A. Tsikiridis. To trust or not to trust: Assignment mechanisms with predictions in the private graph model. In *Proceedings of the 25th ACM Conference on Economics and Computation,* pages 1134–1154, 2024.

Chapter 4. Before we consider the framework of learning-augmented mechanism design, we first introduce a new perspective on mechanism design problems and introduce the notion of *outliers.* In this new domain of mechanism design with outliers, the social cost objective *only* accounts for a certain number of agents, instead of *all* of them. More specifically, the input of the problem includes an extra integer z that denotes the number of outliers, i.e., the number of agents to be excluded from the social cost objective.

[1]How this prediction is learned, e.g., from historical data, is beyond the scope of this thesis.

We apply this framework to the fundamental problem of single facility location on the real line. Here, n agents declare their private preferred locations on the line, and a mechanism chooses the location of a new facility with the goal of minimizing the social cost. We derive tight bounds for the two well-studied objectives of utilitarian and egalitarian social cost if $z < n/2$, and show that no strategyproof mechanism can achieve a bounded approximation guarantee when the number of outliers exceeds this threshold. For the egalitarian objective, our strategyproof mechanism is 2-approximate. In contrast, for the utilitarian objective we demonstrate that strategyproof mechanisms cannot effectively leverage outliers, leading to the counterintuitive result that the approximation guarantee deteriorates as the number of outliers increases.

We then further enrich our model by augmenting with predictions. Interestingly, we find that for the egalitarian objective the lower bound of 2 remains for various types of predictions—including the commonly studied prediction of the optimal location—contrary to the problem without outliers. For the utilitarian objective, we can leverage predictions of the optimal location and derive a strategyproof mechanism that achieves the best possible trade-off between consistency and robustness for $n \geq 3z$.

This chapter is based on the following manuscript—which will soon be made available—and originated during a multi-day research visit to Royal Holloway, University of London, in November 2024:

- A. Deligkas, E. Eiben, S. Klumper, G. Schäfer and A. Tsikiridis. Mechanism Design with Outliers and Predictions.

1.1.2 Part Two: Budget-Feasible Mechanism Design

In Part Two we consider a mechanism design problem with limited monetary transfers. Namely, we consider the problem of designing budget-feasible mechanisms for knapsack procurement auctions, where the auctioneer has a strict budget constraint for the total payments. The auctioneer is interested in buying services (or goods) from a group of strategic agents, and has a concave and non-decreasing valuation function for each service that is publicly known. The total cost of the service of an agent on the other hand, is private information. The main focus of the literature has been on settings where the service of each agent is either fully acquired or not at all. We focus on settings in which partial allocations are allowed, which ranges from the k-level setting in which the allocation can be any number in $\{0, 1, \dots, k\}$ for $k \geq 1$, to the divisible setting in which the allocation can be any fraction in $[0, 1]$. The parameter k bridges the indivisible and divisible setting as k grows from 1 to ∞, a property which we leverage in several of our results.

We derive truthful and budget-feasible mechanisms for both settings under the assumption that the auctioneer is able to afford each individual agent entirely, circumventing strong inapproximability results. For the k-level setting, we derive a mechanism that is $(2+\sqrt{3})$-approximate, whereas for the special case of linear valuations we derive a mechanism with an approximation that improves (decreases) as the granularity k increases, converging to $1+\varphi \approx 2.62$. For the divisible setting, we improve the first constant factor approximation from $2(2+\sqrt{3})$ to $1+\varphi$. For the special case of linear valuations, we improve the best-known approximation from $1+\varphi$ to 2, establishing a separation between this setting and its indivisible counterpart.

Chapter 5. This is the sole chapter of Part Two and is based on the following publications:

- S. Klumper and G. Schäfer. Budget feasible mechanisms for procurement auctions with divisible agents. In *Proceedings of the 15th International Symposium on Algorithmic Game Theory,* pages 78–93, 2022,

- G. Amanatidis, S. Klumper, E. Markakis, G. Schäfer and A. Tsikiridis. Partial allocations in budget-feasible mechanism design: bridging multiple levels of service and divisible agents. *ACM Transactions on Economics and Computation,* 13(2), pages 1–28, 2025.

1.1.3 Part Three: Inefficiency of Equilibria

In the third part of this thesis, we shift our focus from the design and analysis of strategyproof mechanisms to solely the analysis of equilibria resulting from strategic interactions within an already predefined environment. Contrary to the previous part, it is now the agents who have budgetary constraints, on which we further elaborate below. We study the efficiency of equilibrium outcomes (POA) of simultaneous single-item first-price auctions (FPAs) under the hybrid agent model, in which the payments in the utility function of an agent are discounted by a specific parameter. This parameter can be any value between 0 (for value maximizers) and 1 (for utility maximizers). Additionally, each agent has a return-on-investment constraint, enforcing that their expected payment does not exceed their expected acquired value. Besides the already well-established works on utility maximizers, a new focus of the literature has been on settings solely containing value maximizers (autobidding) and on the mixed agent model, which allows for both value and utility maximizing agents.

We bridge two lines of work by extending the smoothness framework to autobidding. A crucial component of our framework is the need to balance multiple smoothness parameters, as these parameters differ across agents with different payment sensitivities. We preserve the extendibility of the smoothness framework, which allows us to derive bounds on the POA for fractionally subadditive valuation functions and coarse correlated equilibria. Additionally, our framework allows us to bridge the setting with homogeneous agents and the mixed agent model, as it allows for any composition of agent types. We derive new POA bounds and generalize known results. This framework also allows us to derive bounds for well-supported equilibria of simultaneous FPAs with additive valuation functions and feasible reserve prices, which is especially relevant when the seller can leverage predictions that approximate the agents' valuations.

Chapter 6. This is the sole chapter of Part Three and is based on the following manuscript and originated during a research visit to Sapienza University of Rome in April–May 2024:

- R. Colini-Baldeschi, S. Klumper, T. Kroll, S. Leonardi, G. Schäfer and A. Tsikiridis. Smoothness Meets Autobidding: Tight Price of Anarchy Bounds for Simultaneous First-Price Auctions. ArXiv preprint arXiv:2506.20908, 2025.

Other Publications. The author has also contributed to the following publication, which is beyond the scope of this thesis:

- A. Haret, S. Klumper, J. Maly and G. Schäfer. Committees and equilibria: Multiwinner approval voting through the lens of budgeting games. In *Proceedings of the 25th ACM Conference on Economics and Computation,* pages 51–70, 2024.

In the above-mentioned publication, we apply a game-theoretic perspective to approval-based multiwinner voting. In approval-based multiwinner voting (MWV), the aim is to select a committee of size k from a set of candidates that proportionally represents the opinion of n voters, where each voter either approves a candidate or not. The literature has focused on defining appropriate notions of proportionality and developing voting rules to achieve them. An intuition underlying many notions and voting rules, is that a proportional committee is attained when no group of voters large enough to deserve some seats in the committee would prefer those seats to be occupied by different candidates. We formalize this intuition by introducing the new class of budgeting games. In a budgeting game, players must decide how to allocate a given budget among resources, which results into the activation of a resource when the collectively allocated budget is large enough. Each player aims to maximize the number of activated resources relevant to them. We showed that specific equilibrium notions correspond to proportionality notions for the associated MWV instance. Additionally, we analyzed improving move dynamics in budgeting games, which give rise to new families of voting rules if the dynamics converge.

Chapter 2

Preliminaries

2

In this chapter, we introduce key concepts and general notation. Additional problem-specific details are deferred to the respective chapters. Section 2.2 is relevant for the first two parts of this thesis, and Section 2.3 supports the third part. We note that this chapter does not provide a complete overview of the relevant research areas, and refer the reader to [Nisan et al., 2007] for more details.

2.1 Basic Notation

The following notation will be used throughout this thesis.

- For $n \in \mathbb{N}_{>0}$, we use $[n] = \{1, \ldots, n\}$ to refer to the set of the first n natural numbers (excluding 0).
- Let $N = [n]$. Given an n-dimensional vector $\boldsymbol{x} = (x_1, x_2, \ldots, x_n)$, we use $\boldsymbol{x}_{-i}$ to denote the $(n-1)$-dimensional vector equal to $\boldsymbol{x}$ with the i-th coordinate removed. We use $\boldsymbol{x}_{-S}$ to denote the $(n - |S|)$-dimensional vector equal to $\boldsymbol{x}$ with the i-th coordinate removed for all $i \in S$. Additionally, we use $\boldsymbol{x} = (x_i, \boldsymbol{x}_{-i})$ and $\boldsymbol{x} = (\boldsymbol{x}_S, \boldsymbol{x}_{-S})$ for any $i \in N$ and $S \subseteq N$, respectively.
- Given an n-dimensional vector $\boldsymbol{x}$, we use the bijective function $\sigma : [n] \to [n]$ to denote a desired permutation of the indices. For example, in Chapter 4 we define σ to be a permutation such that $x_{\sigma(1)} \leq x_{\sigma(2)} \leq \ldots \leq x_{\sigma(n)}$.

2.2 Mechanism Design

As introduced in the previous chapter, mechanism design can be thought of as algorithm design with part of the input unknown and distributed over multiple strategic agents. In order to construct a desired solution, this missing input is collected by requesting the strategic agents to declare it. We first introduce some additional notation before we give the formal definition of a mechanism. Suppose the private information is distributed among a set of agents $N = [n]$, and suppose that the set of outcomes (solutions) of the mechanism is A. For the example of single facility location on the real line, the outcome space is naturally defined as $A = \mathbb{R}$. And for the example of a procurement auction of indivisible items, the set of outcomes can be defined as $A = \{0,1\}^n$. In this case, given an outcome $\boldsymbol{x} \in A$, $x_i = 1$ (resp. $x_i = 0$) can be interpreted as agent i's item being sold (resp. not sold). We denote the private information (or type) of each agent $i \in N$ by t_i, and denote the set of all values that the private information of an agent i can attain as T_i, i.e., $t_i \in T_i$. The aggregated private information is denoted by $\boldsymbol{t} = (t_i)_{i \in N}$, and $\boldsymbol{t} \in T = \times_{i \in N} T_i$. We can now introduce the formal definition of a mechanism.

2

Definition 2.2.1. A (direct-revelation[1]) *mechanism* $\mathcal{M}$ consists of an outcome function $f : T \to A$ and a vector of payment functions $\boldsymbol{p} = (p_i)_{i \in N}$ with $p_i : T \to \mathbb{R}$ for each $i \in N$.

We will denote a mechanism by $\mathcal{M}(\boldsymbol{t}) = (f(\boldsymbol{t}), \boldsymbol{p}(\boldsymbol{t}))$. Note that for an auction, the payments are made by the buyers (agents) to the seller (mechanism designer/auctioneer), whereas for a procurement auction the payments are made by the buyer (mechanism designer/auctioneer) to the sellers (agents). In the first part of this thesis, we consider two mechanism design problems without monetary transfers. In this case, we seek mechanisms that only consist of an outcome function, i.e., $\mathcal{M}(\boldsymbol{t}) = f(\boldsymbol{t})$. Alternatively, one can still think of this mechanism as $\mathcal{M}(\boldsymbol{t}) = (f(\boldsymbol{t}), \boldsymbol{p}(\boldsymbol{t}))$, but with $p_i(\boldsymbol{t}) = 0$ for all agents $i \in N$ and all $\boldsymbol{t} \in T$.

2.2.1 Incentive Compatibility Notions

As we have seen in the previous chapter, strategic agents will misreport their private information if this is in their own best interest. Namely, given the declarations $\boldsymbol{t}$ and the output and payments computed by the mechanism, each agent i obtains some utility u_i with $u_i : T \to \mathbb{R}$. And as the agents are strategic, each agent wants to maximize their utility. For the example of a procurement auction with indivisible items, the utility of each agent $i \in N$ is equal to $u_i(\boldsymbol{t}) = p_i(\boldsymbol{t}) - t_i$ if agent i sells their item under the outcome $f(\boldsymbol{t})$, and $u_i(\boldsymbol{t}) = 0$ otherwise. Namely, if an agent sells their item their utility is equal to the payment they receive minus the cost their incur for producing the item. And for GAP in the private graph model, the utility of an agent is equal to the value they have for their assigned task, if any and if compatible. Otherwise, the agent has a utility of 0.

In this thesis, we only study mechanism design problems in the *complete information* setting. In this setting, the agents have access to all the public information, and know the types of utility functions and possible values of the

[1]Agents submit their declarations simultaneously in a single round.

private information (T) of all other agents. The only information that an agent i solely has access to is their own private information t_i. Our goal is to design mechanisms such that it is always in an agent's best interest to declare their true private information. We formally define this in the following definition.

Definition 2.2.2. A mechanism $\mathcal{M}$ is *strategyproof (or truthful)* if for every agent i, every $\boldsymbol{t} \in T$ and every $t'_i \in T_i$, it holds that $u_i(t_i, \boldsymbol{t}_{-i}) \geq u_i(t'_i, \boldsymbol{t}_{-i})$.

In other words, a mechanism is strategyproof if for each agent declaring their true private information is a *dominant strategy*: it guarantees that their utility is maximized, no matter what the other agents do. Note that for the single facility location problem on the real line, given an outcome $\mathcal{M}(\boldsymbol{t}) = y \in \mathbb{R}$, the utility of an agent i is $u_i(\boldsymbol{t}) = |t_i - y|$, i.e, the distance between i's preferred private location and the location of the facility. This is the only problem considered in this thesis for which an agent wants to minimize their utility, i.e., the utility is actually a cost that an agent incurs. In this case, the same incentive compatibility notions can be used by multiplying the utilities of the agents by -1, or the incentive compatibility notions can be defined accordingly by flipping the inequality, i.e., by replacing $\geq$ in Definition 2.2.2 by $\leq$. We will use the latter in Chapter 4.

2

Recall that in the Introduction, we motivated the design of strategyproof mechanisms as (i) it allows us to say something about the quality of the solution and (ii) it is convenient for the agents. One might argue that this is also somewhat satisfied if we design a mechanism such that every agent has a dominant strategy. However, the following theorem states that there is no loss of generality when aiming for direct-revelation mechanisms that are strategyproof.

Theorem 2.2.3: (Revelation Principle). *For every mechanism $\mathcal{M}$ for which every agent has a dominant strategy, there is an equivalent direct-revelation mechanism $\mathcal{M}$ that is strategyproof.*

A strategyproof mechanism might not be immune to all types of manipulations. Namely, if there is an opportunity for the agents to communicate with each other before submitting their declarations, they might collude if they can form a group such that collectively misreporting their private information is in the best interest of every agent in the group. A mechanism is immune to such manipulations if it satisfies the following notion.

Definition 2.2.4. A mechanism $\mathcal{M}$ is *weakly group-strategyproof (WGSP)* if for every set of agents $S \subseteq N$, every $\boldsymbol{t} \in T$ and every $\boldsymbol{t}'_S \in T_S$, there exists an agent $i \in S$ such that $u_i(\boldsymbol{t}'_S, \boldsymbol{t}_{-S}) \leq u_i(\boldsymbol{t}_S, \boldsymbol{t}_{-S})$.

In other words, for any group of agents trying to collude, there is always an agent in the group who will not want to join as they will not benefit. However, this agent might be somewhat "altruistic" and willing to join this collusion if they are not worse off. For such manipulations, we define the following even stronger notion.

Definition 2.2.5. A mechanism $\mathcal{M}$ is *(strongly) group-strategyproof* if for every set of agents $S \subseteq N$, every $\boldsymbol{t} \in T$ and every $\boldsymbol{t}'_S \in T_S$ such that there exists an agent $i \in S$ with $u_i(\boldsymbol{t}'_S, \boldsymbol{t}_{-S}) > u_i(\boldsymbol{t}_S, \boldsymbol{t}_{-S})$, there also exists an agent $j \in S$ such that $u_j(\boldsymbol{t}'_S, \boldsymbol{t}_{-S}) < u_j(\boldsymbol{t}_S, \boldsymbol{t}_{-S})$.

Again, in other words, there is no group of agents that can coordinate their declarations such that there is at least one agent in the group that will benefit, and there is no agent in the group that is worse off.

Randomization. Up to now, we have only considered deterministic mechanisms. One way of expanding the design space of strategyproof mechanisms is to allow for randomization. A *randomized mechanism* is a probability distribution over a set of deterministic mechanisms, and so, the outcome of a randomized mechanism is a probability distribution over outcomes and payments. If a randomized mechanism $\mathcal{M}$ only randomizes over deterministic strategyproof mechanisms, we say that $\mathcal{M}$ is *universally strategyproof.* A reason to allow for randomization is that universally strategyproof mechanisms can be more powerful than deterministic ones (see e.g., [Nisan and Ronen, 1999]). Analogously to the deterministic incentive compatibility notions, stronger notions exists. All randomized mechanisms that we devise in this thesis satisfy the following notion.

Definition 2.2.6. A randomized mechanism $\mathcal{M}$ is *universally WGSP* if it is a probability distribution over a finite set of deterministic WGSP mechanisms.

We note that universal strategyproofness is a stronger notion than strategyproofness in expectation, for which it is sufficient that each agent maximizes their expected utility when declaring their private information. Strategyproofness in expectation can again be more powerful than universal strategyproofness (see e.g., [Dobzinski and Dughmi, 2013]), but at the expense of satisfying a weaker incentive compatibility notion.

2.2.2 Approximation Notions

If we have succeeded in designing a mechanism satisfying one of the desired incentive compatibility notions introduced above, we might wonder how "good" our mechanism actually is. For GAP in the private graph model, never assigning any agent to any task is a strategyproof mechanism. This is obviously not a good mechanism, as the computed assignment never has any value. A well-established benchmark used to evaluate the quality of a mechanism, is the objective value of an optimal solution to the non-strategic version of the problem given declarations $\boldsymbol{t}$. If the goal of the underlying problem is to maximize some objective, we use $\mathrm{SW}^\star(\boldsymbol{t})$ to denote this maximum objective value given declarations $\boldsymbol{t}$. Here, we omit an explicit reference to the public information—which might be needed in order to determine the optimal objective value—for brevity. We then compare this value to the objective value of the outcome computed by our mechanism, denoted by $\mathrm{SW}(\mathcal{M}(\boldsymbol{t}))$. We want to be able to make a statement about this comparison that is true no matter what the declaration $\boldsymbol{t}$ and public information are. In other words, we want to claim some *worst-case guarantee*, as formally defined below.

Definition 2.2.7. A mechanism $\mathcal{M}$ is *ρ-approximate* with $\rho \geq 1$ if for any input $\boldsymbol{t}$ it holds that $\rho \cdot \mathrm{SW}(\mathcal{M}(\boldsymbol{t})) \geq \mathrm{SW}^{\star}(\boldsymbol{t})$.

If the objective of the mechanism design problem at hand is to minimize some social cost function, as for the facility location problem, the definition is defined analogously, i.e., for $\rho \geq 1$ it holds that $\mathrm{SC}(\mathcal{M}(\boldsymbol{t})) \leq \rho \cdot \mathrm{SC}^{\star}(\boldsymbol{t})$.

Learning-Augmented Framework. As introduced in the previous chapter, one way to overcome worst-case guarantees is by augmenting each problem instance with a prediction $\hat{\boldsymbol{t}}$. In the context of mechanism design, this prediction is informative about the private information of the agents, and what $\hat{\boldsymbol{t}}$ specifically is may depend on the problem at hand. We mentioned that in this framework, the goal is to design a strategyproof mechanism that leverages the prediction and achieves improved approximation guarantees when the prediction is perfect, while maintaining close to worst-case approximation guarantee when the prediction is erroneous. As we do not want the approximation guarantee to immediately deteriorate to the worst-case guarantee when a prediction is almost perfect, we define an error measure η that measures the quality of the prediction. The exact definition of η will be introduced in the respective chapter, as it might be problem specific. Additionally, our goal is for the approximation guarantee to smoothly interpolate between the consistency and robustness guarantees as a function of the error measure η. More formally, we consider the following approximation notions in this thesis. Note that we also do not explicitly state the prediction $\hat{\boldsymbol{t}}$ as an input argument for the mechanism, as we assume that the prediction is public information.

Definition 2.2.8. A mechanism $\mathcal{M}$ is:

- *α-consistent* with $\alpha \geq 1$ if for any input $\boldsymbol{t}$ augmented with a perfect prediction $\hat{\boldsymbol{t}}$ it holds that $\alpha \cdot \mathrm{SW}(\mathcal{M}(\boldsymbol{t})) \geq \mathrm{SW}^{\star}(\boldsymbol{t})$,
- *β-robust* with $\beta \geq 1$ if for any input $\boldsymbol{t}$ augmented with an arbitrary prediction $\hat{\boldsymbol{t}}$ it holds that $\beta \cdot \mathrm{SW}(\mathcal{M}(\boldsymbol{t})) \geq \mathrm{SW}^{\star}(\boldsymbol{t})$, and
- *$g(\eta)$-approximate* with $g(\eta) \geq 1$ if for any input $\boldsymbol{t}$ augmented with a prediction $\hat{\boldsymbol{t}}$ with an error of at most η it holds that $g(\eta) \cdot \mathrm{SW}(\mathcal{M}(\boldsymbol{t})) \geq \mathrm{SW}^{\star}(\boldsymbol{t})$.

As for Definition 2.2.7, we define the notions of α-consistency, β-robustness and $g(\eta)$-approximate analogously if the objective for the mechanism design problem at hand is to minimize a social cost function.

Randomization. All approximation notions introduced in this section extend naturally to randomized mechanisms. As a randomized mechanism $\mathcal{M}$ is a probability distribution over outputs and payments, the objective value of a mechanism is replaced with the expected objective value in the respective definitions, e.g., $\mathbb{E}[\mathrm{SW}(\mathcal{M}(\boldsymbol{t}))]$.

2.2.3 Computational Efficiency

Additionally, we require the mechanisms that we seek to be computationally efficient, i.e., to compute the outcome and payments in polynomial time in the input size. All mechanisms derived in this thesis can be implemented in polynomial time in the input size. As this is mostly easy to verify, we omit analyses of the running times. For the mechanisms that we derive in Chapter 5, verifying this might not be trivial due to the computation of the payments and we refer the reader to [Amanatidis et al., 2025] for more details.

2.3 Strategic Games

In this section, we introduce concepts that are relevant in Part Three of this thesis. The chapter in this part studies strategic interactions in a predefined environment, which can be defined as a strategic game.

Definition 2.3.1. A *strategic game* Γ is defined by $\Gamma = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$, where $N = [n]$ is the set of players (or agents), S_i is the strategy set of player $i \in N$, and $u_i : S \to \mathbb{R}$ is the utility function of player $i \in N$ with $S = \times_{i \in N} S_i$.

A strategic game is *finite* if the set of players and the strategy sets are finite. We refer to a vector $\boldsymbol{s} = (s_i)_{i \in N}$ as a *strategy profile.* Again, players want to maximize their utility. As mentioned in the previous section, players can also be associated with a cost function which they want to minimize, in which case $u_i(\boldsymbol{s}) = -c_i(\boldsymbol{s})$ for all $\boldsymbol{s} \in S$. The strategic games that we consider are so-called *one-shot simultaneous games*, as there is one moment in which all players must choose their strategy at the same. Additionally, we only consider strategic games with *full information*, in which players know each other's strategy sets and utility functions.

Equilibrium Notions. Given a strategic game, we are interested in stable outcomes as these might be probable outcomes of the game. In a stable outcome, i.e., strategy profile, no player could have chosen any other strategy in hindsight that would have resulted in a higher utility. Given a finite strategic game $\Gamma = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$, we will define multiple equilibrium notions with ranging degrees of stability. We first consider equilibria in which players choose their strategy independently and deterministically.

Definition 2.3.2. A strategy profile $\boldsymbol{s} \in S$ is a *pure Nash equilibrium (PNE)* if for every player $i \in N$ it holds that $u_i(s_i, \boldsymbol{s}_{-i}) \geq u_i(s_i', \boldsymbol{s}_{-i})$ for every $s_i' \in S_i$.

Example 2.3.3. Consider an auction game among players 1 and 2, i.e., $N = \{1, 2\}$. The players compete over a single item that is allocated according to a second-price auction: the player with the highest bid wins the item and pays as a price the second highest bid, with ties broken in favor of player 1. The strategy s_i of a player i is a bid, which can be any value in $[0.00, 10.00] = S_i$. Player 1 has a value of 4 and player 2 has a value of 3, and both players have quasi-linear utilities leading to:

$$u_1(s_1, s_2) = \begin{cases} 4 - s_2, & \text{if } s_1 \geq s_2, \\ 0, & \text{otherwise,} \end{cases} \quad \text{and} \quad u_2(s_1, s_2) = \begin{cases} 3 - s_1, & \text{if } s_2 > s_1, \\ 0, & \text{otherwise.} \end{cases}$$

Note that the strategy profile $(4, 3)$ is a pure Nash equilibrium. Player 2 has a utility of 0 and can only win the item when bidding above 4, but this would lead to a negative utility. Player 1 wins the item and has a utility of 1, and cannot influence the price while still winning the item.

We also consider randomized strategies. We define Δ_i as the set of all probability distributions over S_i, and $\Delta = \times_{i \in N} \Delta_i$. Let $\sigma \in \Delta$ be a probability distribution over S and let $\boldsymbol{x} \sim \sigma$, i.e., $\boldsymbol{x}$ is a random strategy profile drawn from σ. For ease of notation, we omit the explicit reference to σ and identify $\boldsymbol{x}$ with σ. For the random variable $\boldsymbol{x}_{-i}$, we define:

$$\forall \boldsymbol{s}_{-i} \in S_{-i}: \qquad \mathbb{P}[\boldsymbol{x}_{-i} = \boldsymbol{s}_{-i}] = \sum_{s_i \in S_i} \mathbb{P}[\boldsymbol{x} = \boldsymbol{s}].$$

Definition 2.3.4. Let $\sigma \in \Delta$ be a probability distribution over S and let $\boldsymbol{x} \sim \sigma$. Then:

- σ is a *mixed Nash equilibrium (MNE)* if $\sigma = \prod_{i \in [n]} \sigma_i$ and for every agent $i \in N$ it holds that $\mathbb{E}[u_i(\boldsymbol{x})] \geq \mathbb{E}[u_i(x_i', \boldsymbol{x}_{-i})]$ for all $x_i' \sim \sigma_i'$ with $\sigma_i' \in \Delta_i$,
- σ is a *correlated equilibrium (CE)* if for every agent $i \in N$ it holds that $\mathbb{E}[u_i(s_i, \boldsymbol{x}_{-i}) \mid x_i = s_i] \geq \mathbb{E}[u_i(s_i', \boldsymbol{x}_{-i}) \mid x_i = s_i]$ for all $s_i, s_i', \in S_i$, and,
- σ is a *coarse correlated equilibrium (CCE)* if for every agent $i \in N$ it holds that $\mathbb{E}[u_i(\boldsymbol{x})] \geq \mathbb{E}[u_i(x_i', \boldsymbol{x}_{-i})]$ for all $x_i' \sim \sigma_i'$ with $\sigma_i' \in \Delta_i$.

Note that if $\boldsymbol{s}$ is a pure Nash equilibrium of Γ, then $\boldsymbol{s}$ is also a mixed Nash equilibrium of Γ, and if σ is a mixed Nash equilibrium of Γ, σ is also a correlated equilibrium of Γ, and so forth. Therefore, we have the following hierarchy of equilibrium notions:

$$\mathit{PNE}(\Gamma) \subseteq \mathit{MNE}(\Gamma) \subseteq \mathit{CE}(\Gamma) \subseteq \mathit{CCE}(\Gamma),$$

here, $\mathit{EQ}(\Gamma)$ denotes the set of all equilibria of Γ for $\mathit{EQ} \in \{\mathit{PNE}, \mathit{MNE}, \mathit{CE}, \mathit{CCE}\}$.

Existence of Equilibria. As we are interested in stable outcomes of strategic games, can we be sure that such outcomes always exist? It turns out that we cannot guarantee that every game has a pure Nash equilibrium, and we give an example of such a game below.

Example 2.3.5. The game of Matching Pennies is played among 2 players, call them player 1 and player 2. Each player must choose a strategy that is either heads (H) or tails (T), and the resulting utilities are depicted in Table 2.1. If both players play the same strategy, i.e., $\boldsymbol{s} = (T, T)$ or $\boldsymbol{s} = (H, H)$, then player 1 has a utility of 1 and player 2 has a utility of -1. If the strategies differ, i.e., $\boldsymbol{s} = (T, H)$ or $\boldsymbol{s} = (H, T)$, the roles are reversed and player 1 has a utility of -1 whereas player 2 has a utility of 1. The game has no pure Nash equilibrium as for every strategy profile there is a player with a utility of -1, and this player would have a utility of 1 if they had chosen to play the other strategy.

	Heads (H)	**Tails** (T)
Heads (H)	$(1, -1)$	$(-1, 1)$
Tails (T)	$(-1, 1)$	$(1, -1)$

Table 2.1: Utilities for the game Matching Pennies in Example 2.3.5. The strategies of player 1 (resp. player 2) are depicted vertically (resp. horizontally).

However, the following famous result shows that an equilibrium does exist if we consider randomized strategies.

Theorem 2.3.6: ([Nash, 1951]). *Every finite strategic game has a mixed Nash equilibrium.*

Inefficiency of Equilibria. Finally, if an equilibrium exists, we are interested in quantifying the efficiency loss due to selfish behavior. In order to do so, we need a benchmark to compare the equilibrium with. For a single-item second-price auction game (see Example 2.3.3), we can for example define this benchmark as the social welfare. In this case, the social welfare denotes the (expected) value of the winner of the auction. We define $\boldsymbol{s}^\star$ as a strategy profile that maximizes the social welfare, i.e., a strategy profile for which the player with the highest valuation always wins the item. We can now introduce the price of anarchy, a well-known notion to measure the inefficiency of a class of equilibria. This notion can also be defined for a specific strategic game Γ, but we immediately introduce the notion for a class of strategic games. We note that for the class for which we derive bounds on the price of anarchy (in Chapter 6), the expected social welfare is always greater than 0 in any equilibrium, and thus (2.1) in the definition below is well-defined.

Definition 2.3.7. Given a class of strategic games $\mathcal{G}$ and an equilibrium notion $EQ \in \{PNE, MNE, CE, CCE\}$, we define the *price of anarchy* of EQ with respect to $\mathcal{G}$ as:

$$EQ\text{-}POA(\mathcal{G}) = \sup_{\Gamma \in \mathcal{G}} \sup_{\sigma \in EQ(\Gamma)} \frac{\mathrm{SW}(\boldsymbol{s}^\star)}{\mathbb{E}_{\boldsymbol{x} \sim \sigma}[\mathrm{SW}(\boldsymbol{x})]}. \tag{2.1}$$

Note that it follows from the hierarchy of equilibrium notions that:

$$PNE\text{-}POA(\mathcal{G}) \leq MNE\text{-}POA(\mathcal{G}) \leq CE\text{-}POA(\mathcal{G}) \leq CCE\text{-}POA(\mathcal{G}).$$

Example 2.3.8. Consider the auction game in Example 2.3.3, but suppose that the value of player 2 is very small and equal to ε. The optimal social welfare is equal to 4; the value that player 1 has for the item. However, note that the strategy profile $\boldsymbol{s} = (0, 4)$ is a pure Nash equilibrium. Player 2 wins the item and has a utility of ε, and cannot influence the price while still winning the item. Player 1 has a utility of 0 and can only win the item if $s_1' \geq 4$, but this would not lead to a positive utility. For this class of games $\mathcal{G}$, the $PNE\text{-}POA(\mathcal{G})$ is therefore unbounded as $\frac{\mathrm{SW}(\boldsymbol{s}^\star)}{\mathrm{SW}(\boldsymbol{s})} = \frac{4}{\varepsilon}$ is unbounded as ε approaches 0.

2

Part One

Mechanism Design without Money

Chapter 3

Generalized Assignment Problem in the Private Graph Model

3.1 Introduction

3

In many real-world problems, multiple strategic *agents* hold some private information based on which they make decisions. Consider for example the problem of assigning students to schools, where students might value certain schools more than others. In this case, a *mechanism* can be designed that takes as input the preferences of the students, and computes an assignment in line with a socially desired objective. Oftentimes, it is desirable that the designed mechanism is *strategyproof*, which ensures that it is in the best interest of each agent to reveal their private information truthfully, independently of the other agents. Strategyproofness often imposes strong impossibility results on the worst-case approximation guarantee of the objective, which can be rather disappointing from a practical perspective (see, e.g., [Roughgarden, 2019]). To overcome these limitations, a new line of research initiated by Agrawal et al. [2024] and Xu and Lu [2022] explores how to leverage learning-augmented inputs in the design of mechanisms. While this line of research first emerged in the area of online algorithms (see, e.g., [Lykouris and Vassilvitskii, 2021]), it is particularly appealing in the context of mechanism design when the predictions are informative about the agents' private information.

In this framework of *mechanism design with predictions*, the designer can exploit the predictions to improve the worst-case guarantee of their mechanism. However, the predictions might be inaccurate or even entirely erroneous. As a result, the goal is to design mechanisms that achieve improved guarantees if the prediction is perfect (referred to as *consistency*), while still maintaining a reasonable worst-case guarantee when the prediction is imperfect (referred to as *robustness*). Ideally, the mechanism provides a fine-grained approximation guarantee depending on some measure of the prediction error, which smoothly interpolates between these two extreme cases (referred to as *approximation*).

In this chapter, we study how to leverage predictions in the domain of mechanism design *without money*. As monetary transfers cannot be used to effectively eliminate the incentives of agents to misreport their private information, the problem of designing strategyproof mechanisms has a different flavor (see, e.g., [Schummer and Vohra, 2007, Procaccia and Tennenholtz, 2013]). This domain is especially relevant in certain practical settings in which monetary transfers are infeasible due to practical constraints or not allowed due to ethical and legal issues (see, e.g., [Roughgarden, 2010]).

Generalized Assignment Problem in the Private Graph Model. We focus on the *generalized assignment problem (GAP)*, which encompasses several fundamental special cases, such as matching and multiple knapsack. In this problem, we are given a bipartite graph consisting of a set of strategic agents (or jobs) and a set of resources (or machines). Each agent i has a value v_{ij} and a size s_{ij} for being assigned to a resource j, and each resource has a capacity C_j (in terms of total size) that must not be exceeded. The goal is to compute a feasible assignment of agents to resources such that the overall value is maximized.

The *private graph model* for GAP was introduced by Dughmi and Ghosh [2010]. In this model, the agents' values and sizes are assumed to be public information, but whether or not the value of assigning an agent to a resource can be generated is private information. The latter can naturally be interpreted as *compatibility restrictions* that agents have with respect to the available resources. Note that this model restricts the strategy space of the agents to being able to only misreport their compatibility vector. Despite this restriction, GAP in the private graph model still has several natural applications. For example, consider a department head that needs to assign courses to faculty members. While the faculty's expertise is publicly known, their availability to teach a specific course might be unknown. For additional applications, see, e.g., [Chen et al., 2014].

Output Predictions. We study GAP in the private graph model in a learning-augmented environment, and assume that the optimal assignment with respect to the private compatibilities is given as a prediction. Note that this is information-theoretically weaker than assuming that the actual compatibilities are available as a prediction, as we can always compute an optimal assignment with respect to some given predicted compatibilities (ignoring computational constraints). As such, our prediction belongs to the class of *output predictions* that has recently been studied in various contexts (see, e.g., [Agrawal et al., 2024, Gkatzelis et al., 2025, Christodoulou et al., 2024]). Such output predictions provide aggregate information about the *outcome* (in our case, the optimal assignment), rather than assuming that the private information of the agents (compatibilities) can be predicted. Consequently, under the assumption that predictions are based on historical data, our model aligns better with applications where a complete set of

GAP Variant	Restrictions ($\forall i \in L$, $\forall j \in R$)
Unweighted bipartite matching (UBMP)	$v_{ij} = 1$, $s_{ij} = 1$, $C_j = 1$
Bipartite matching problem (BMP)	$s_{ij} = 1$, $C_j = 1$
Agent value GAP (AVGAP)	$v_{ij} = v_i$
Resource value GAP (RVGAP)	$v_{ij} = v_j$
Value consensus GAP (VCGAP)	$\exists \sigma : \; v_{i\sigma(1)} \geq \ldots, \geq v_{i\sigma(m)}$
Agent size GAP (ASGAP)	$s_{ij} = s_i$
Resource size GAP (RSGAP)	$s_{ij} = s_j$
Restricted multiple knapsack (RMK)	$v_{ij} = v_i$, $s_{ij} = s_i$
Equal RMK (ERMK)	$v_{ij} = s_{ij} = v_i$

Table 3.1: Overview of GAP variants.

input data may not be readily available, due to for instance privacy concerns. We refer the reader to [Christodoulou et al., 2024] for more information on output predictions.

As observed by Dughmi and Ghosh [2010], for the more general model of GAP in which the agents' values are private, no deterministic strategyproof mechanism can achieve a finite approximation guarantee. In fact, we show that this impossibility persists even if the problem is augmented with output predictions (see Section 3.3.3).

3.1.1 Our Contributions

This chapter considers GAP in the private graph model with predictions. We assume that a potentially erroneous prediction of the optimal assignment for the true compatibility graph of an instance is given as part of the input. We derive both deterministic and randomized mechanisms that are (universally) weakly group-strategyproof (WGSP) for different variants of GAP: see Table 3.1 for an overview of GAP variants and see Table 3.2 for an overview of our results. Our mechanisms are parameterized by a *confidence parameter* $\gamma \geq 1$, which determines the trade-off between the respective consistency and robustness guarantees. Choosing a higher confidence value leads to a better consistency but a worse robustness, and vice versa. All our mechanisms provide fine-grained approximation guarantees that smoothly interpolate between the respective consistency and robustness, depending on some natural error parameter η of the prediction. More specifically, $1-\eta$ measures the relative gap between the value of the predicted assignment and an optimal one, where $\eta = 0$ if the prediction is perfect. All respective approximation guarantees can be obtained by multiplying the consistency as in Table 3.2 by a factor of $1-\eta$, where the robustness kicks in whenever

GAP$^+$ variant	**D or R**	**Consistency**	**Robustness**	**Theorem**
BMP, RSGAP	D	$1 + 1/\gamma$	$1 + \gamma$	Thm. 3.4.1
	R	$1 + 1/\gamma$	$\sqrt{2(1+\gamma)}$	Thm. 3.6.1
ERMK	D	$1 + 1/\gamma$	$2 + \gamma$	Thm. 3.5.3
	R	$1 + 1/\gamma$	$\frac{1}{2}(\sqrt{12\gamma + 13} + 1)$	Thm. 3.6.3
VCGAP, ASGAP	R	$1 + 3/\gamma$	$3 + \gamma$	Thm. 3.6.6

Table 3.2: Overview of the consistency and robustness of our (universally) WGSP mechanisms for any confidence parameter $\gamma \geq 1$. The second column denotes if the mechanism is deterministic (D) or randomized (R). Note that AVGAP, RVGAP and RMK are all special cases of VCGAP.

3

it outperforms this guarantee.

We summarize our main results below.

- We prove a lower bound on the best possible trade-off in terms of consistency and robustness achievable by any deterministic strategyproof mechanism for GAP (Section 3.3). More precisely, we show that no deterministic strategyproof mechanism can be $(1 + 1/\gamma)$-consistent and $(1 + \gamma - \epsilon)$-robust for any $\epsilon > 0$. In fact, our lower bound also holds for the special cases of BMP and VCGAP. We also extend our insights to derive a lower bound in terms of consistency and approximation guarantees.

- In light of the lower bound above, our mechanism Boost that we derive for BMP in Section 3.4 achieves the best possible trade-off between consistency and robustness, while satisfying the stronger notion of WGSP. Unlike the deterministic 2-approximate mechanism in [Dughmi and Ghosh, 2010] for the problem without predictions, we crucially do not consider declarations in a fixed order. Instead, our mechanism draws inspiration from the well-known deferred acceptance algorithm by Gale and Shapley [1962]. Here, the agent proposal order is crucial for WGSP and the resource preference order is crucial to improve upon the known guarantee for the problem without predictions. If an edge is in the predicted optimal matching, it potentially has a better ranking in the resource preference order, depending on the confidence parameter and the instance at hand. Our mechanism for BMP extends (with the same approximation guarantee) to many-to-one assignments and RSGAP (Section 3.4.4). In particular, this provides the first deterministic WGSP mechanism that is 2-approximate for RSGAP.

- For GAP, our deterministic mechanism Greedy adds declared edges in a greedy manner (while maintaining feasibility) to an initially empty assignment, according to some order of the declared edges (Section 3.5). The order follows from a specific ranking function, that must be given to Greedy as part of the input. We derive a sufficient condition, called *truth-inducing*, of the ranking function which guarantees that the resulting mechanism is WGSP. For the special case of ERMK, we derive a truth-inducing ranking function that leverages the prediction, resulting in a deterministic WGSP mechanism (Section 3.5.2). The same approach can be used to obtain a 3-approximate WGSP mechanism for ERMK without predictions, for which no polynomial-time deterministic strategyproof mechanism was known.

- For ASGAP and VCGAP, we derive randomized universally WGSP mechanisms in Section 3.6. Our mechanisms randomize over three deterministic mechanisms, consisting of Boost, Greedy and a mechanism that simply follows the prediction. Note that for $\gamma = 1$, our mechanisms retrieve the expected guarantee of the state-of-the-art mechanisms of Chen et al. [2014] for ASGAP, AVGAP and RVGAP. As Greedy is one of the three building blocks in both our mechanisms, it is crucial that for these variants truth-inducing ranking functions exist. Finally, for BMP and ERMK, we also derive randomized universally WGSP mechanisms that achieve the same consistency and outperform the robustness of their respective deterministic counterparts in expectation. In particular, for BMP this provides a separation result showing that randomized mechanisms are more powerful than deterministic ones.

3.1.2 Related Work

Algorithms with predictions represent one perspective within the beyond worst-case paradigm. The primary goal is to overcome existing worst-case guarantees by augmenting each instance of a problem with a prediction, possibly a machine-learned one. Hence, this line of work is also sometimes referred to as learning-augmented algorithms. The conceptual framework that describes the trade-off between α-consistency and β-robustness was introduced by Lykouris and Vassilvitskii [2021] in the context of online algorithms. Since then, online algorithms have remained a major focus (see e.g., [Purohit et al., 2018, Azar et al., 2021, 2022, Banerjee et al., 2022] for some reference works). Thematically relevant to us are the works on online matching (e.g., [Antoniadis et al., 2023b,a, Lavastida et al., 2021a,b, Jin and Ma, 2022, Dinitz et al., 2022]) in non-strategic settings. Other domains that have been studied under the lens of predictions include the re-evaluation of runtime guarantees of algorithms, streaming algorithms, data structures and more. We refer the reader to [Mitzenmacher and Vassilvitskii, 2022] for a survey of some of the earlier works. Moreover, an overview of research articles that appeared on these topics is available at `https://algorithms-with-predictions.github.io`.

Recently, Xu and Lu [2022] and Agrawal et al. [2024] introduced predictions for settings involving strategic agents. In their work, Xu and Lu [2022] showcased four different mechanism design settings with predictions, both with and without monetary transfers. On the other hand, Agrawal et al. [2024] focused solely on strategic facility location. Most subsequent works with strategic considerations have also been in algorithmic mechanism design (see e.g., [Istrate and Bonchis, 2022, Balkanski et al., 2024b, Prasad et al., 2023, Balkanski et al., 2024a, Gkatzelis et al., 2025, Christodoulou et al., 2024, Cohen et al., 2024, Barak et al., 2024]). However, other classical domains have also been revisited in the presence of predictions; see e.g., [Gkatzelis et al., 2022, Chen et al., 2024] on price of anarchy, [Berger et al., 2024] on voting, and [Lu et al., 2024, Caragiannis and Kalantzis, 2024] on auction revenue maximization.

We briefly elaborate on the relation between learning-augmented mechanism design and Bayesian mechanism design. The main difference is the absence of worst-case guarantees in the Bayesian setting, where it is implicitly assumed that one has perfect or relatively good knowledge of the prior distributions when analyzing the expected performance of mechanisms. If this assumption fails, worst-case performance guarantees may not be possible (see also the relevant discussions in [Agrawal et al., 2024] and [Prasad et al., 2023]).

Mechanism design without money has a rich history spanning over fifty years, being deeply rooted in economics and social choice theory. As will be evident in Section 3.4, the works on stable matching (see, e.g., [Gale and Shapley, 1962, Roth, 1982, Hatfield and Milgrom, 2005]) are particularly relevant to us. However, our work aligns more closely with the agenda of *approximate* mechanism design without money set forth by Procaccia and Tennenholtz [2013] and, in particular, the subsequent work by Dughmi and Ghosh [2010]. In their work, Dughmi and Ghosh [2010] introduced the private graph model, which we consider with predictions, and initiated the study of variants of GAP when the agents are strategic (a variant of this model where the resources are strategic instead, was studied by Fadaei and Bichler [2017b]). Dughmi and Ghosh [2010] obtained a deterministic 2-approximate, strategyproof mechanism for BMP and a matching lower bound. Furthermore, they developed randomized strategyproof in expectation mechanisms for special cases of GAP; namely, a 2-approximation for RMK and a 4-approximation for ASGAP and AVGAP. Finally, they proposed a randomized, strategyproof in expectation $O(\log n)$-approximate mechanism for the general case. Subsequently, Chen et al. [2014] improved on these results by devising 4-approximate mechanisms which satisfy universal strategyproofness for all special cases of GAP that are value or size invariant. Additionally, they showed an improved $O(1)$-approximation for GAP.

Beyond the private graph model, for the setting where values are private information but monetary transfers are allowed, Fadaei and Bichler [2017a] devised an $\frac{e}{e-1}$-approximate, strategyproof in expectation mechanism. Finally, from an algorithmic perspective, the best known approximation ratio for GAP is $\frac{e}{e-1} - \epsilon$, for a fixed small $\epsilon > 0$ [Feige and Vondrak, 2006]. On the negative side, Chakrabarty and Goel [2010] have shown that GAP does not admit an approximation better than $^{11}/_{10}$, unless P = NP.

3.2 Preliminaries

An instance of GAP (see Figure 3.1) is defined by a bipartite graph $G = (L \cup R, D)$ consisting of a set $L = [n]$ of $n \geq 1$ agents (or items, jobs) and a set $R = [m]$ of $m \geq 1$ resources (or knapsacks, machines, respectively). Each agent $i \in L$ has a value $v_{ij} > 0$ and a size $s_{ij} > 0$ for being assigned to a resource $j \in R$. Further, each resource $j \in R$ has a capacity $C_j > 0$ that cannot be exceeded. We assume without loss of generality that $s_{ij} \leq C_j$ for every $i \in L$ and $j \in R$. We denote an instance I_{GAP} of GAP by a tuple $I_{\text{GAP}} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C})$ where:

- $G[D] = (L \cup R, D)$ denotes the graph induced by the edge set $D \subseteq L \times R$,
- $\boldsymbol{v} = (v_{ij})_{i \in L, j \in R} \in \mathbb{R}^{n \times m}_{>0}$ denotes the matrix of all agent-resource values,
- $\boldsymbol{s} = (s_{ij})_{i \in L, j \in R} \in \mathbb{R}^{n \times m}_{>0}$ denotes the matrix of all agent-resource sizes, and
- $\boldsymbol{C} = (C_j)_{j \in R} \in \mathbb{R}^{m}_{>0}$ denotes the vector of all resource capacities.

Note that $G[D]$ is not necessarily complete and that D can be any subset of $L \times R$. The bipartite graph $G[D]$ namely encodes compatibilities between agents and resources; we also refer to it as the *compatibility graph*. An agent $i \in L$ is said to be *compatible* with a resource $j \in R$ if $(i, j) \in D$; otherwise i is *incompatible* with j. For example, the compatibility $(i, j) \in D$ might indicate that agent i has access to resource j, or that item i can be assigned to knapsack j, or that job i can be executed on machine j.

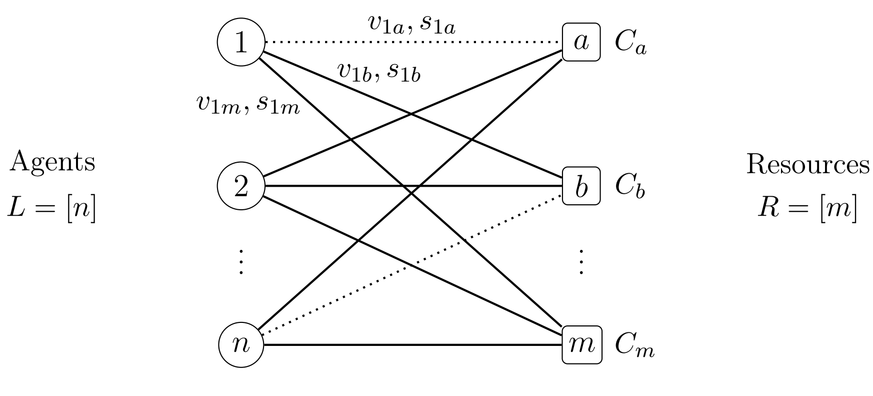

Figure 3.1: Illustration of the components of an instance of GAP with n agents and m resources. The dotted edges $(1, a)$ and (n, b) are incompatible and not contained in the edge set D. Throughout this chapter, we will represent agents and resources as circles and squares, respectively, in figures.

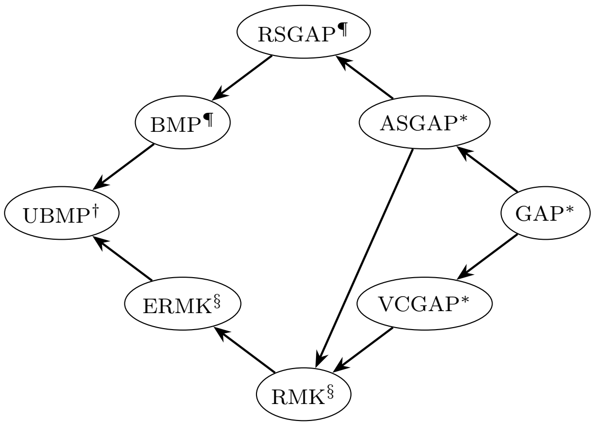

Superscript	**Tractable**	**OPT ⇒ SP**
†	✓	✓
¶	✓	✗
§	✗	✓
$*$	✗	✗

Figure 3.2: Taxonomy of GAP variants where the arrows indicate special cases. The respective superscript (column 1) indicates whether the problem is polynomial-time solvable (column 2) and whether an optimal algorithm gives rise to a strategyproof mechanism that is 1-efficient (column 3).

Given an instance of GAP, the objective is to find an assignment that optimally assigns agents to resources. An *assignment* $M \subseteq L \times R$ is a subset of edges such that each agent $i \in L$ is incident to at most one edge in M. Note that several agents might be assigned to the same resource; we also say that M is a *many-to-one assignment*. If we additionally require that each resource $j \in R$ is incident to at most one edge in M, then M is said to be a *one-to-one assignment* or *matching*. Note that every matching is also an assignment. Given an agent $i \in L$, we use $M(i) = \{j \in R \mid (i,j) \in M\}$ to refer to the resource assigned to i (if any); note that $M(i)$ is a singleton set. We use $M(i) = \emptyset$ if i is unassigned. Similarly, for a resource $j \in R$ we define $M(j) = \{i \in L \mid (i,j) \in M\}$ as the set of agents assigned to j (if any); note that $M(j) = \emptyset$ if no agent is assigned to j. An assignment M is said to be *feasible* for a given compatibility graph $G[D]$ if (1) M is an assignment in the compatibility graph $G[D]$ induced by D, i.e., $M \subseteq D$, and (2) M satisfies all resource capacities constraints, i.e., for each resource $j \in R$, $\sum_{i \in M(j)} s_{ij} \leq C_j$. We define the *value* $v(M)$ of an assignment M as the sum of the values of all edges in M; more formally:

$$v(M) = \sum_{(i,j) \in M} v_{ij}. \tag{3.1}$$

We overload this notation slightly and also write $v(M(j))$ and $v(M(i))$ to refer to the total value of all edges assigned to resource j and agent i, respectively. Further, we define $v_{i\emptyset} = 0$ and $v_{\emptyset j} = 0$ for notational convenience. We use M^*_D to denote a feasible assignment of maximum value in the graph $G[D]$. We also say that M^*_D is an *optimal* assignment with respect to D.

3.2.1 GAP Variants

Below, we give an overview of the different special cases of GAP we consider; see also Figure 3.2 for an illustration of the relationships between these special cases.

- *Bipartite matching problem (BMP).* Here, each agent $i \in L$ has unit size, i.e., $s_{ij} = 1$ for all $j \in R$, and each resource $j \in R$ has unit capacity, i.e., $C_j = 1$. We write $I_{\text{BMP}} = (G[D], \boldsymbol{v})$ to denote an instance of BMP.

- *Value consensus GAP (VCGAP).* Here, the agents have some consensus about the value of the resources in the sense that there exists a permutation σ of the resources such that for each agent $i \in L$ it holds that: $v_{i\sigma(1)} \geq v_{i\sigma(2)} \geq \cdots \geq v_{i\sigma(m)}$. Note that both AVGAP ($v_{ij} = v_i$ for all $j \in R$) and RVGAP ($v_{ij} = v_j$ for all $i \in L$) fall into this case. We write $I_{\text{VCGAP}} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C})$ to denote an instance of VCGAP.

- *Agent size GAP (ASGAP).* Here, each agent $i \in L$ has the same size s_i for all resources, i.e., $s_{ij} = s_i$ for all $j \in R$. We use $I_{\text{ASGAP}} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C})$ to denote an instance of ASGAP.

- *Resource size GAP (RSGAP).* Here, for every resource $j \in R$, all agents have the same size s_j, i.e., $s_{ij} = s_j$ for all $i \in L$. We use $I_{\text{RSGAP}} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C})$ to denote an instance of RSGAP.

- *Restricted multiple knapsack (RMK)*, also known as multiple knapsacks with assignment restrictions (see e.g., [Dawande et al., 2000, Nutov et al., 2006, Aerts et al., 2003]). Here, each agent i has a fixed value $v_i = v_{ij}$ and size $s_i = s_{ij}$ for all $j \in R$. We use $I_{\text{RMK}} = (G[D], (v_i)_{i\in L}, (s_i)_{i\in L}, \boldsymbol{C})$ to denote an instance of RMK. We mostly focus on the special case ERMK of RMK, where $v_i = s_i$ for all i. We use $I_{\text{ERMK}} = (G[D], (v_i = s_i)_{i\in L}, \boldsymbol{C})$ to denote an instance of ERMK.

3.2.2 Private Graph Model

We study GAP in a strategic environment. Namely, we consider the setting in which the agents are strategic and may misreport their compatibilities. This setting is termed the *private graph model* and was introduced by Dughmi and Ghosh [2010]. More specifically, each agent $i \in L$ has a true compatibility set $E_i \subseteq i \times R$ specifying the set of edges that are truly compatible for i. The compatibility set E_i is *private* information, i.e., only known to agent i. In addition, each agent $i \in L$ declares a public compatibility set $D_i \subseteq i \times R$. The interpretation is that i claims to be compatible with resource $j \in R$ if and only if $(i, j) \in D_i$; but these declarations may not be truthful, i.e., it may be that $D_i \neq E_i$. We define $D = \cup_{i\in L} D_i \subseteq L \times R$ as the union of all compatibility sets declared by the agents. Similarly to $G[D]$, we use $G[E]$ to refer to the compatibility graph induced by the true compatibility sets of the agents, i.e., $E = \cup_{i \in L} E_i$. We refer to $G[E]$ as the *private graph model* or *private graph*.

Given an instance I_{GAP} of GAP with private graph $G[E]$ and compatibility declarations D, a deterministic mechanism $\mathcal{M}$ should compute an assignment $M = \mathcal{M}(D)$ that is feasible for I_{GAP}. Then, the *utility* u_i of an agent $i \in L$ is defined as:

$$u_i(D) = \begin{cases} v_{ij}, & \text{if } (i,j) \in \mathcal{M}(D) \cap E_i, \\ 0, & \text{otherwise.} \end{cases} \tag{3.2}$$

Note that the utility of agent i is v_{ij} if (1) i is assigned to resource j, i.e., $(i,j) \in M = \mathcal{M}(D)$, and (2) i is truly compatible with resource j, i.e., $(i,j) \in E_i$. In particular, the utility of i is 0 if i is unassigned in M, or if i is matched to an incompatible resource. We assume that each agent wants to maximize their utility. To this aim, an agent i might misreport their true compatibilities by declaring a compatibility set $D_i \neq E_i$.

3

We therefore seek mechanisms in which agents are never better off misreporting their true compatibilities, i.e., *strategyproof* mechanisms (see Definition 2.2.2). In fact, the mechanisms we devise satisfy an even stronger notion, namely that of *weak group-strategyproofness (WGSP)* (see Definition 2.2.4). Note that if $M = \mathcal{M}(D)$ is the assignment computed by $\mathcal{M}$ for truthfully declared compatibilities, i.e., $D = E$, then its value $v(M)$ (as defined in (3.1)) is equal to the sum of the utilities of the agents. We also devise randomized mechanisms that are *universally WGSP*, see Definition 2.2.6.

3.2.3 GAP with Predictions

In the setting with predictions, we are given an instance $I_{\text{GAP}} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C})$ of GAP, and, in addition, a *predicted assignment* $\hat{M} \subseteq L \times R$. It is important to realize that the predicted assignment $\hat{M}$ is considered to be part of the input. We assume that the predicted assignment respects the capacity constraints, i.e., for all $j \in R$ it holds that $\sum_{(i,j) \in \hat{M}} s_{ij} \leq C_j$. However, $\hat{M}$ can be any assignment in the complete graph $G[L \times R]$, and thus may contain edges which are not in D. We use $I_{\text{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ to refer to an instance of GAP augmented with a predicted assignment $\hat{M}$. We use a similar notation for the various special cases of GAP.

We say that $\hat{M}$ is a *perfect prediction* for I_{GAP^+} if it corresponds to an assignment of maximum value in the graph $G[D]$, i.e., $v(\hat{M} \cap D) = v(M^*_D)$. We define an error parameter that measures the quality of the predicted assignment $\hat{M}$ relative to an optimal assignment M^*_D of $G[D]$. Namely, we define the *prediction error* $\eta(I_{\text{GAP}^+}) \in [0,1]$ of an instance $I_{\text{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ as:

$$\eta(I_{\text{GAP}^+}) = 1 - \frac{v(\hat{M} \cap D)}{v(M^*_D)}. \tag{3.3}$$

Note that with this definition an instance I_{GAP^+} with a perfect prediction has a prediction error of 0. As the value of the predicted assignment $\hat{M}$ deteriorates, the error measure approaches 1. If $\eta(I_{\mathrm{GAP}^+}) = 1$, we must have that $v(\hat{M} \cap D) = 0$ which means that the prediction $\hat{M}$ does not contain any edge that is also in D (recall that the values are assumed to be positive). Given a fixed error parameter $\eta \in [0, 1]$, instances $I_{\mathrm{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ with $\eta(I_{\mathrm{GAP}^+}) \leq \eta$ constitute the class of instances with a prediction error of at most η.

Note that the definition of our error parameter in (3.3) is meaningful as it captures the relative gap between the *values* of the predicted assignment and the optimal one. Alternatively, one could compare structural properties of $\hat{M} \cap D$ and M^*_D. However, this seems less suitable in our context: For example, under an error notion that is not value-based, a predicted assignment may only miss one edge of an optimal assignment, i.e., $|M^*_D \setminus (\hat{M} \cap D)| = 1$, but still be of relatively low value if this missing edge is valuable. Further, a predicted assignment might contain none of the edges of the optimal assignment, i.e., $(\hat{M} \cap D) \cap M^*_D = \emptyset$, but still be very useful when its value is close to optimal; in fact, $(\hat{M} \cap D)$ might even be an optimal matching that is disjoint from M^*_D (because the optimal assignment might not be unique). Finally, note that accounting for the value of edges in $\hat{M} \setminus D$ in a prediction error notion is not informative, as our goal is to compute a feasible assignment M (i.e., $M \subseteq D$). All these cases are captured by the definition of our prediction error as in (3.3).

We use $(I_{\mathrm{GAP}^+}, G[E])$ to refer to an instance I_{GAP^+} of GAP with predictions in the private graph model $G[E]$. We note that all input data of $I_{\mathrm{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ is public information accessible by the mechanism, while the private graph $G[E]$ is private information. For the sake of conciseness, we often omit input parameters which remain fixed; in fact, most of the time it will be sufficient to refer explicitly to the compatibility declarations D only. We will consider the approximation notions as introduced in Definition 2.2.8, which extend naturally to randomized mechanisms.

3.2.4 Stable Matching

We introduce the notions and results from stable matching theory needed in this chapter. Consider a complete bipartite graph $G = (L \cup R, E)$ consisting of a set of agents L, a set of resources R and a set of edges $E = L \times R$. Each agent $i \in L$ has a strict total preference order $\succ_i$ over $E_i \cup \{\emptyset\}$, where $E_i = \{i\} \times R$. Given two distinct edges $e, e' \in E_i$, agent i prefers e over e' if $e \succ_i e'$. The position of $\emptyset$ in the order indicates whether an edge is acceptable or not: $e \in E_i$ is *acceptable* if $e \succ_i \emptyset$; otherwise e is *unacceptable*. An agent would rather be unmatched than be matched through an unacceptable edge. Similarly, each resource $j \in R$ has a strict total preference order $\succ_j$ over $E_j \cup \{\emptyset\}$, where $E_j = L \times \{j\}$. All notions introduced above naturally extend to a resource j with preference order $\succ_j$. We refer to $(G[L \times R], (\succ_i)_{i \in L}, (\succ_j)_{j \in R})$ as a *standard preference system* (i.e., if G is complete and all preference orders are total and strict).

Here, a matching $M \subseteq E$ is a subset of *acceptable* edges such that no two distinct edges in M share a common resource or agent. We say that an edge $e = (i, j) \in E$ *blocks* M if (1) i prefers to be matched through e instead of $(i, M(i))$, i.e., $e \succ_i (i, M(i))$, and (2) j prefers to be matched through e instead of $(M(j), j)$, i.e., $e \succ_j (M(j), j)$. A matching M is *stable* if it is not blocked by any edge in E.

In their seminal work, Gale and Shapley [1962] proposed an algorithm, also known as the *Gale-Shapley algorithm* or *deferred acceptance algorithm*, that computes a stable matching for any given standard preference system. This result sparked the development of stable matching theory. One of the consequences that will be useful for us is the following: if the agents are strategic and can misreport their preference orders arbitrarily, then the agent-proposing deferred acceptance algorithm guarantees WGSP. This result is stated below and is attributed to Roth [1982] (SP) and, independently, to Dubins and Freedman [1981] (WGSP). Later, Gale and Sotomayor [1985] gave a greatly simplified proof of WGSP.

3

Theorem 3.2.1. *Let $(G[L \times R], (\succ_i)_{i\in L}, (\succ_j)_{j\in R})$ be a standard preference system. Then, the agent-proposing deferred acceptance algorithm is weakly group-strategyproof.*[1]

Note that for this result to hold, it is crucial that only the agents can misreport their preferences. In particular, the preferences of the resources are assumed to be fixed (i.e., public).

3.2.5 Lexicographic Extensions and Sort Operators

The notions defined here will be useful throughout this chapter. Let $X = \{x_1, \dots, x_n\}$ be a set of $n \geq 1$ elements and assume that each element $x_i \in X$ is associated with $k \geq 1$ numerical values $z_1(x_i), \dots, z_k(x_i)$. We define $\succeq'$ as the partial order over X that we obtain by comparing the elements in X lexicographically with respect to $(z_1, \dots, z_k)$, i.e., $\forall i, j \in [n]$:

$$x_i \succeq' x_j \quad \Leftrightarrow \quad \big(z_1(x_i), \dots, z_k(x_i)\big) \succeq^{\text{lex}} \big(z_1(x_j), \dots, z_k(x_j)\big). \tag{3.4}$$

We also say that $\succeq'$ is the *extended lexicographic order* of X with respect to $(z_1, \dots, z_k)$. We write $\succ'$ instead of $\succeq'$ if the order is strict. We also introduce an operator ORDER that orders the elements in X by lexicographic decreasing order of their values $(z_1, \dots, z_k)$. More formally, let $\mathcal{S}_n$ denote the set of all permutations of $[n]$. Given X and $(z_1, \dots, z_k)$ as above, we define $\text{ORDER}(X, (z_1, \dots, z_k)) = (\pi_1, \dots, \pi_n) \in \mathcal{S}_n$ such that for all $i, j \in [n]$ with $i < j$,

[1]See [Dubins and Freedman, 1981, Roth, 1982, Gale and Sotomayor, 1985].

we have that $\big(z_1(x_{\pi_i}), \ldots, z_k(x_{\pi_i})\big) \succeq^{\text{lex}} \big(z_1(x_{\pi_j}), \ldots, z_k(x_{\pi_j})\big)$. Further, we define an operator SORT that sorts the elements in X according to this order, namely:

$$\text{SORT}(X, (z_1, \ldots, z_k)) = \langle x_{\pi_1}, \ldots, x_{\pi_n} \rangle. \tag{3.5}$$

3.3 Impossibility Results and the Baseline Mechanism

In Section 3.3.1, we prove a lower bound on the best possible trade-off in terms of consistency and robustness guarantees achievable by any deterministic strategyproof mechanism for GAP$^+$. We also derive a lower bound in terms of the error parameter η. Then, in Section 3.3.2, we introduce a trivial mechanism, called TRUST, that serves as a baseline mechanism in subsequent sections. Finally, we conclude this section with Section 3.3.3, by pointing out an impossibility result for the neighboring model of private values.

3.3.1 Impossibility Results

All the lower bounds that we prove in this subsection hold for both BMP$^+$ and VCGAP$^+$ in the private graph model. Clearly, these lower bounds extend to all variants of GAP$^+$ that contain BMP$^+$ or VCGAP$^+$ as a special case (see Figure 3.2).

Theorem 3.3.1. *Let $\gamma \geq 1$. Then, no deterministic strategyproof mechanism for BMP$^+$ can be $(1 + 1/\gamma)$-consistent and $(1 + \gamma - \epsilon)$-robust for any $\epsilon > 0$.*

Proof: The proof is by contradiction. Let $\gamma \geq 1$ be fixed arbitrarily and suppose that there exists a deterministic strategyproof mechanism $\mathcal{M}$ that is $(1 + 1/\gamma)$-consistent and $(1 + \gamma - \epsilon)$-robust for some $\epsilon \in (0, \gamma]$ (note that for $\epsilon > \gamma$, we immediately obtain a contraction). Fix $\bar{\epsilon}$ to be a constant in $(0, \frac{\epsilon}{3+2\gamma-2\epsilon})$.

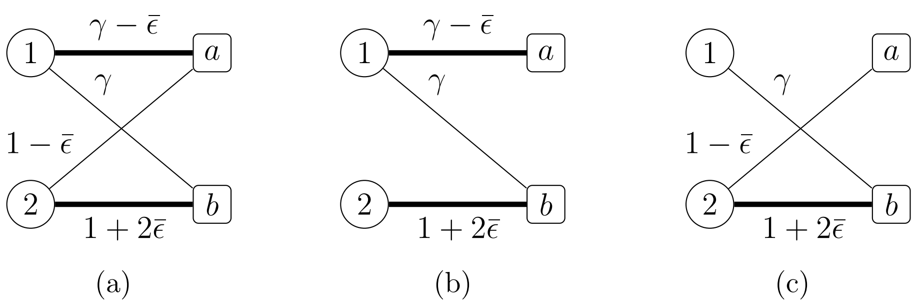

Figure 3.3: Instances used in the proof of Theorem 3.3.1.

3

3

Consider the instance $I_{\text{BMP}^+} = (G[D], \boldsymbol{v}, \hat{M})$ with private graph $G[E]$ depicted in Figure 3.3a, where the compatibility declarations are truthful, i.e., $D = E$. Note that the values $\boldsymbol{v}$ of this instance depend on γ and $\bar{\epsilon}$. Here, the predicted matching is $\hat{M} = \{(1,a),(2,b)\}$ (indicated in bold). Recall that we use $\mathcal{M}(D)$ to denote the matching computed by $\mathcal{M}$ for compatibility declarations D. The optimal matching in $G[D]$ is $M_D^* = \hat{M}$. Because $\mathcal{M}$ is $(1 + 1/\gamma)$-consistent and $\bar{\epsilon} > 0$, we have:

$$\Big(1 + \frac{1}{\gamma}\Big) v(\mathcal{M}(D)) \geq v(M_D^*) = 1 + \gamma + \bar{\epsilon} > 1 + \gamma.$$

Dividing both sides by $(1+\gamma)/\gamma$ leads to $v(\mathcal{M}(D)) > \gamma$. Note that there are only two matchings in $G[D]$ that *always* have a value strictly greater than γ, namely $\{(1,b),(2,a)\}$ and $\{(1,a),(2,b)\}$. Additionally, for the matching $\{(2,b)\}$, this is true for $\gamma < 1+2\bar{\epsilon}$. In particular, the above implies that the matching output by $\mathcal{M}$ on D must be either $\mathcal{M}(D) = M_1 = \{(1,b),(2,a)\}$ or $\mathcal{M}(D) = M_2 \ni (2,b)$. We consider these two cases.

<u>Case 1:</u> $\mathcal{M}(D) = M_1$. Note that the utility of agent 2 with respect to D is $u_2(\mathcal{M}(D)) = u_2(M_1) = v_{2a} = 1 - \bar{\epsilon}$. Consider the compatibility declarations $D' = D \setminus \{(2,a)\}$ that we obtain from D if agent 2 deviates by hiding their edge $(2,a)$. The respective instance is depicted in Figure 3.3b. Suppose we run $\mathcal{M}$ on D'. Since $\mathcal{M}$ is strategyproof, edge $(2,b)$ cannot be contained in $\mathcal{M}(D')$. (To see this, note that otherwise the utility of agent 2 with respect to D' would be $v_{2b} = 1 + 2\bar{\epsilon} > 1 - \bar{\epsilon} = u_2(\mathcal{M}(D))$, contradicting that $\mathcal{M}$ is strategyproof.) But then, we must have $v(\mathcal{M}(D')) \leq \max\{v_{1a}, v_{1b}\} = \gamma$. On the other hand, because $M_{D'}^* = \hat{M}$ and $\mathcal{M}$ is $(1 + 1/\gamma)$-consistent, it must again hold that $v(\mathcal{M}(D')) > \gamma$, leading to a contradiction.

<u>Case 2:</u> $\mathcal{M}(D) = M_2 \ni (2,b)$. Consider the compatibility declarations $D' = D \setminus \{(1,a)\}$ that we obtain from D if agent 1 deviates by hiding their edge $(1,a)$. The respective instance is depicted in Figure 3.3c. Suppose we run $\mathcal{M}$ on D'. Analogously to the argument given above, since $\mathcal{M}$ is strategyproof, edge $(1,b)$ cannot be contained in $\mathcal{M}(D')$. (To see this, note that otherwise the utility of agent 1 with respect to D' would be $v_{1b} = \gamma > \gamma - \bar{\epsilon} \geq u_1(\mathcal{M}(D))$, contradicting that $\mathcal{M}$ is strategyproof.) But then, we must have $v(\mathcal{M}(D')) \leq \max\{v_{2a}, v_{2b}\} = 1 + 2\bar{\epsilon}$. On the other hand, because $v(M_{D'}^*) > v(\hat{M} \cap D')$ and $\mathcal{M}$ is $(1 + \gamma - \epsilon)$-robust, it must hold that $(1 + \gamma - \epsilon) v(\mathcal{M}(D')) \geq v(M_{D'}^*) = 1 + \gamma - \bar{\epsilon}$. Dividing both sides by $(1 + \gamma - \epsilon)$, we obtain:

$$v(\mathcal{M}(D')) \geq \frac{1 + \gamma - \bar{\epsilon}}{1 + \gamma - \epsilon} > 1 + 2\bar{\epsilon},$$

where the last inequality can be verified to hold if $\bar{\epsilon} < \frac{\epsilon}{3+2\gamma-2\epsilon}$ (which is a restriction we impose on the choice of $\bar{\epsilon}$). We thus obtain a contradiction. □

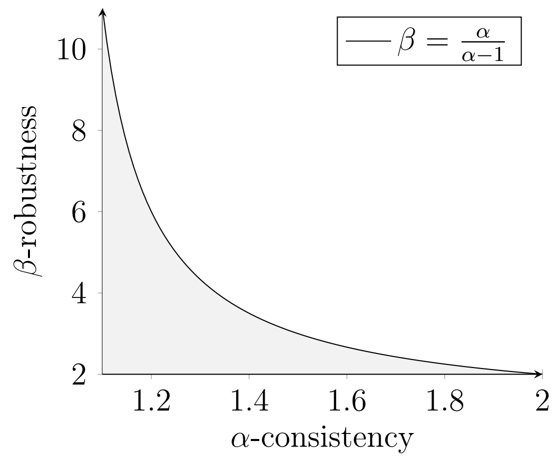

Figure 3.4: Impossibility trade-off in terms of α-consistency and β-robustness: no deterministic strategyproof mechanism for BMP$^+$ or VCGAP$^+$ can achieve a combination of α and β in the gray area.

An illustration of the impossibility trade-off between consistency and robustness proven in Theorem 3.3.1 is given in Figure 3.4. Ideally, we would want a mechanism to be 1-consistent. However, note that as $\gamma \to \infty$, the consistency guarantee converges to 1, but the robustness will be unbounded in this case. So it follows from Theorem 3.3.1 that no deterministic strategyproof mechanism can achieve 1-consistency and bounded robustness. Ideally, we would also want a mechanism to be 2-robust, as this is the best-possible approximation guarantee for the setting without predictions. But for any fixed $\epsilon > 0$ in Theorem 3.3.1, we need $\gamma = 1 + \epsilon$ to achieve 2-robustness. However, this leads to an impossibility of $(1 + \frac{1}{1+\epsilon})$-consistency for any $\epsilon > 0$. In particular, there is no deterministic strategyproof mechanism that is 2-robust and has a consistency strictly smaller than 2. And so, in order to utilize the predictions and find deterministic strategyproof mechanisms that achieve a consistency strictly smaller than 2, we need to settle for a robustness strictly greater than 2.

As mentioned above, a result of Theorem 3.3.1 is that any deterministic strategyproof mechanism that is 1-consistent will have unbounded robustness. Recall that robustness is defined with respect to an arbitrary prediction quality. And so, Theorem 3.3.1 gives no insight on what approximation guarantee a deterministic strategyproof 1-consistent mechanism can achieve for certain prediction errors. This is what we establish in the next theorem for our error parameter η as defined in (3.3).

Theorem 3.3.2. *No deterministic strategyproof mechanism for BMP$^+$ can be* 1-*consistent and* $\frac{1}{1-\eta+\epsilon}$-*approximate for any* $\eta \in (0, 1)$ *and some* $\epsilon \in (0, \eta]$.

Proof: The proof is by contradiction. Let $\eta \in (0, 1)$ be fixed arbitrarily and suppose that there exists a deterministic strategyproof mechanism $\mathcal{M}$ that is

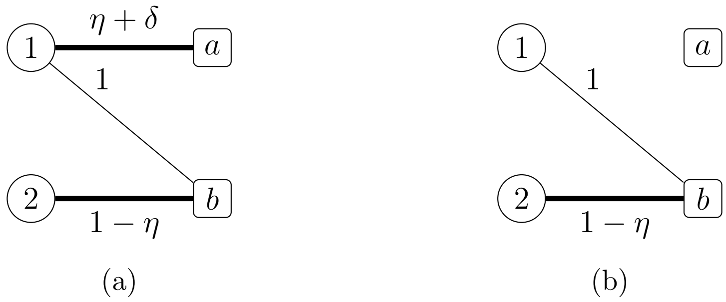

Figure 3.5: Instances used in the proof of Theorem 3.3.2.

1-consistent and $\frac{1}{1-\eta+\epsilon}$-approximate for η and some $\epsilon \in (0, \eta]$.

Fix $\delta \in (0, 1-\eta)$ and consider the instance $I_{\mathrm{BMP}^+} = (G[D], \boldsymbol{v}, \hat{M})$ with private graph $G[E]$ depicted in Figure 3.5a, where the compatibility declarations are truthful. The predicted matching is $\hat{M} = \{(1, a), (2, b)\}$. The optimal matching in $G[D]$ is $M^*_D = \hat{M} \cap D$, and so $\mathcal{M}$ must be 1-approximate for this instance, i.e., $\mathcal{M}(D) = \{(1, a), (2, b)\}$.

Consider the compatibility declarations $D' = D \setminus \{(1, a)\}$ that we obtain from D if agent 1 deviates by hiding their edge $(1, a)$. The respective instance is depicted in Figure 3.5b. Suppose we run $\mathcal{M}$ on D'. Since $\mathcal{M}$ is strategyproof, edge $(1, b)$ cannot be contained in $\mathcal{M}(D')$. Therefore, we must have $v(\mathcal{M}(D')) \leq v_{2b} = 1 - \eta$, which implies that:

$$\frac{1}{1-\eta+\epsilon}(1-\eta) \geq \frac{1}{1-\eta+\epsilon} v(\mathcal{M}(D')) \geq v(M^*_{D'}) = 1.$$

The second inequality follows from the fact that the error for this instance is η, and so, by assumption, $\mathcal{M}$ must be $(1/(1-\eta+\epsilon))$-approximate. However, this leads to a contradiction as $\epsilon > 0$. □

As we seek mechanisms that achieve bounded robustness, we combine the ideas of Theorems 3.3.1 and 3.3.2 in Theorem 3.3.3, and establish a lower bound in terms of our error parameter η for any deterministic strategyproof mechanism that is $(1 + 1/\gamma)$-consistent.

Theorem 3.3.3. *Let $\gamma \geq 1$. No deterministic strategyproof mechanism for BMP$^+$ can be $(1 + 1/\gamma)$-consistent and $\frac{1}{1-\eta+\epsilon}$-approximate for any $\eta \in (0, \gamma/1+\gamma)$ with $\epsilon \in \big(0, \min\{\frac{2}{1+\gamma}, \frac{2}{2+\gamma}\big(\gamma - \eta(1+\gamma)\big)\}\big)$.*

Proof: The proof is by contradiction. Let $\gamma \geq 1$ be fixed arbitrarily and suppose that there exists a deterministic strategyproof mechanism $\mathcal{M}$ that is $(1 + 1/\gamma)$-consistent and $\frac{1}{1-\eta+\epsilon}$-approximate for some $\eta \in (0, \gamma/1+\gamma)$. Further, let $\epsilon \in (0, \min\{\frac{2}{1+\gamma}, \frac{2}{2+\gamma}(\gamma - \eta(1+\gamma))\})$.

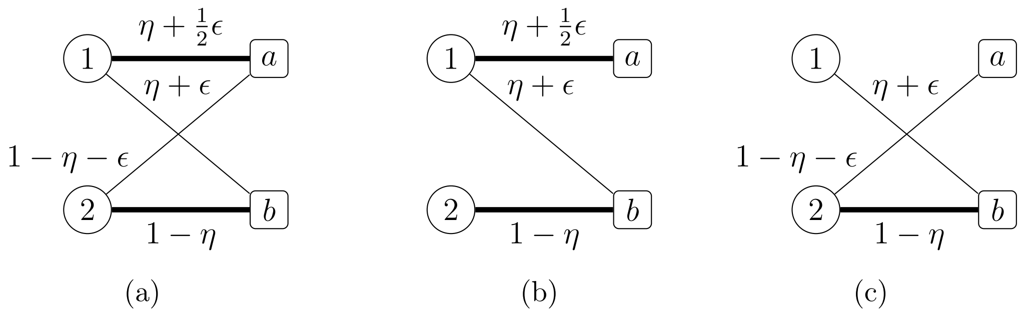

Figure 3.6: Instances used in the proof of Theorem 3.3.3.

Consider the instance $I_{\text{BMP}^+} = (G[D], \boldsymbol{v}, \hat{M})$ with private graph $G[E]$ depicted in Figure 3.6a and assume that compatibility declarations are truthful. Here, the predicted matching is $\hat{M} = \{(1, a), (2, b)\}$.

First, consider the instance depicted in Figure 3.6b that we obtain if agent 2 hides their edge $(2, a)$. Let $D' = D \setminus \{(2, a)\}$ be the corresponding declarations. Note that the optimal matching in $G[D']$ is $M^*_{D'} = \hat{M}$ and has value $v(M^*_{D'}) = 1 + \frac{1}{2}\epsilon$, as $\epsilon < 2/(1+\gamma)$ and $\eta < \gamma/(1+\gamma)$. Thus, for $\mathcal{M}$ to be $(1 + {}^1\!/\!_\gamma)$-consistent, we must have $v(\mathcal{M}(D')) \geq (\frac{\gamma}{1+\gamma})(1 + \frac{1}{2}\epsilon)$. Therefore, by choice of ϵ, it must be that $(2, b) \in \mathcal{M}(D')$.

Next, consider the truthful declarations $D = E$ as depicted in Figure 3.6a. As $\mathcal{M}$ is strategyproof, it must be that $(2, b) \in \mathcal{M}(D)$. Otherwise, agent 2 could increase their utility by hiding edge $(2, a)$, leading to the case as in Figure 3.6b.

Finally, consider the instance depicted in Figure 3.6c that we obtain if agent 1 hides their edge $(1, a)$. Let $D' = D \setminus \{(1, a)\}$ be the corresponding declarations. As $\mathcal{M}$ is strategyproof, it must be that $(1, b) \notin \mathcal{M}(D')$, because $(1, b) \notin \mathcal{M}(D)$ as $(2, b) \in \mathcal{M}(D)$. This implies that $v(\mathcal{M}(D')) \leq \max(v_{2a}, v_{2b}) = 1 - \eta$. Note that the optimal matching in $G[D']$ is $M^*_{D'} = \{(1, b), (2, a)\}$ and has value $v(M^*_{D'}) = 1$. Note that in this instance $\eta = 1 - v(\hat{M} \cap D')/v(M^*_{D'})$ and so $\mathcal{M}$ must be $1/(1 - \eta + \epsilon)$-approximate. We get:

$$\frac{1}{1 - \eta + \epsilon}(1 - \eta) \geq \frac{1}{1 - \eta + \epsilon} v(\mathcal{M}(D')) \geq v(M^*_{D'}) = 1,$$

leading to a contradiction as $\epsilon > 0$. □

Remark 3.3.4. Theorems 3.3.1, 3.3.2 and 3.3.3 also hold for VCGAP$^+$. Namely, the instances used in the proofs are all instances of VCGAP$^+$ with $\sigma(1) = b$, $\sigma(2) = a$ and unit sizes and capacities.

3.3.2 Baseline Mechanism

We now introduce a naïve mechanism for GAP$^+$ in the private graph model, which simply adheres to the prediction. Given an instance $I_{\text{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$, the mechanism returns the assignment $\hat{M} \cap D$. We call this mechanism Trust (see Mechanism 1).

MECHANISM 1: Trust(I_{GAP^+})

Input: An instance $I_{\text{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$.
Output: A feasible assignment for I_{GAP^+}.
return $\hat{M} \cap D$

It is trivial to see that Trust is 1-consistent. It is also not hard to prove that Trust is WGSP and achieves the best possible approximation guarantee of Theorem 3.3.2.

Theorem 3.3.5. *Consider the class of instances of GAP$^+$ in the private graph model with a prediction error of at most $\eta \in [0, 1)$. Then, Trust is weakly group-strategyproof and achieves an approximation guarantee of $1/(1-\eta)$, which is best possible.*

Proof: Let $(I_{\text{GAP}^+}, G[E])$ be an instance of GAP$^+$ in the private graph model. Note that each agent $i \in L$ will be included with edge (i, j) in the assignment M computed by Trust(I_{GAP^+}) if and only if $(i, \hat{M}(i)) \in D_i$ and $\hat{M}(i) = j$. Therefore, for each agent i with $(i, \hat{M}(i)) \in E_i$, it is a dominant strategy to include $(i, \hat{M}(i))$ in their set D_i of compatibility declarations. All other agents i with $(i, \hat{M}(i)) \notin E_i$ receive zero utility, independently of their compatibility declarations. This proves that Trust is WGSP.

Also, by the definition of the error parameter $\eta(I_{\text{GAP}^+})$, the approximation guarantee of Trust follows trivially:

$$v(\hat{M} \cap D) = (1 - \eta(I_{\text{GAP}^+}))v(M_D^*) \geq (1 - \eta)v(M_D^*),$$

matching the best possible approximation guarantee of Theorem 3.3.2. □

We conclude from Theorem 3.3.5 that Trust realizes our strongest notion of incentive compatibility, i.e., WGSP, and even achieves the best possible consistency and approximation guarantees. But the point is, that Trust completely fails to achieve any bounded robustness guarantee. This demonstrates that the actual challenge in deriving strategyproof mechanisms for GAP$^+$ in the private graph model is to achieve the best possible trade-off in terms of consistency/approximation *and* robustness guarantees; without the latter, the whole problem becomes trivial (as Trust is the best possible mechanism). Despite this deficiency, and perhaps surprisingly, we will use this non-robust mechanism TRUST as an important building block in our randomized mechanisms described in Section 3.6.

We now present a simple observation that we already used in the proof of Theorem 3.3.5. Because we will reuse it several times in the upcoming sections, we summarize it in the following lemma and refer to it as the Lifting Lemma.

Lemma 3.3.6: (Lifting Lemma). *Consider the class of instances of GAP$^+$ in the private graph model with a prediction error of at most $\eta \in [0, 1)$. If for every instance I_{GAP^+}, $\mathcal{M}(I_{GAP^+})$ returns an assignment M such that $\alpha \cdot \mathbb{E}[v(M)] \geq v(\hat{M} \cap D)$ for some $\alpha \geq 1$, then $\mathcal{M}$ is $\alpha/(1-\eta)$-approximate (in expectation).*

Proof: By definition, we have that $v(\hat{M} \cap D) = (1 - \eta(I_{\text{GAP}^+}))v(M^*_D)$. Also, only instances of I_{GAP^+} are considered with $\eta(I_{\text{GAP}^+})) \leq \eta$, and so:

$$\alpha \cdot \mathbb{E}[v(M)] \geq v(\hat{M} \cap D) = (1 - \eta(I_{\text{GAP}^+}))v(M^*_D) \geq (1 - \eta)v(M^*_D).$$

□

3.3.3 Beyond the Private Graph Model

We conclude this section with a discussion about GAP with private valuations. In this problem, the size s_{ij} of an agent $i \in L$ for a resource $j \in R$ and the capacity C_j of a resource j is still public information. However, the value v_{ij} of an agent i for a resource j is private information that the agent must declare to a mechanism. For this problem, it is known that deterministic strategyproof mechanisms are unable to provide bounded approximation guarantees (see [Dughmi and Ghosh, 2010]). Even if this problem, GAP with private valuations, is augmented with a prediction $\hat{M}$ of the optimal assignment, this limitation cannot be overcome. This is illustrated by the following example.

Example 3.3.7. Let $\mathcal{M}$ be a deterministic strategyproof mechanism that achieves both a bounded consistency and robustness, i.e., $\mathcal{M}$ is α-consistent and β-robust with $1 \leq \alpha < \infty$ and $1 \leq \beta < \infty$. Consider an instance with two agents, agent 1 and agent 2, and one resource a. Let the private valuations be

Figure 3.7: Instances used in Example 3.3.7 of GAP with private valuations augmented with a predicted assignment.

$v_{1a} = 1$ and $v_{2a} = {}^{1}/_{\alpha} - \epsilon$ with $0 < \epsilon < \frac{1}{\alpha}$, and consider truthful declarations. This instance is depicted in Figure 3.7a. For a perfect prediction $\hat{M} = \{(1, a)\}$, $\mathcal{M}$ must return $\{(1, a)\}$ to achieve α-consistency. Now consider the same instance in Figure 3.7b, with the only difference the declared value for $(2, a)$ being $\beta + \epsilon$. As now the prediction $\hat{M} = \{(1, a)\}$ is imperfect, $\mathcal{M}$ must return $\{(2, a)\}$ to achieve β-robustness. However, this contradicts strategyproofness, as agent 2 will unilaterally deviate from the instance depicted in Figure 3.7a and declare a value of $\beta + \epsilon$ in order to improve their utility. So the best a deterministic strategyproof mechanism can do is always return the prediction, leading to 1-consistency and unbounded robustness.

Recently, Christodoulou et al. [2024] considered the house allocation problem enhanced with an assignment prediction. In this problem, the aim is to assign each agent $i \in L$ a house $j \in R$ with $|L| = |R|$ and maximize social welfare, i.e., maximize the total value of the assignment. They consider *unit-sum* valuations (for each agent $i \in L$, $\sum_{j \in R} v_{ij} = 1$) and *unit-range* valuations (for each agent $i \in L$, $v_{ij} \in [0, 1]$ for all $j \in R$ and $\exists k, h \in R$ with $v_{ik} = 1$ and $v_{ih} = 0$). For these valuation functions, they show that bounded approximation guarantees can be achieved[2].

3.4 Bipartite Matching Problem

We introduce our mechanism for BMP$^+$ in the private graph model. Our mechanism, called Boost, is inspired by the *deferred acceptance algorithm* by Gale and Shapley [1962]. Boost is parameterized by some $\gamma \geq 1$, which we term the *confidence parameter* and which can be set by the user. It is important to realize that all properties proved in this subsection hold for any choice of $\gamma \geq 1$. Boost also serves as an important component of our randomized mechanisms for other GAP variants, which we will consider in later subsections.

3.4.1 Boost Mechanism

A description of Boost can be found in Mechanism 2. As input, Boost receives an instance $I_{\text{BMP}^+} = (G[D], \boldsymbol{v}, \hat{M})$ of BMP$^+$ and a confidence parameter $\gamma \geq 1$. During its entire execution, Boost maintains a tentative matching $M \subseteq D$, a set of active agents $A \subseteq L$ and a set of remaining proposals P_i for each agent i. Initially, the matching is empty, i.e., $M = \emptyset$, all agents are active, i.e., $A = L$, and for each agent i their set of proposals is equal to their declaration, i.e., $P_i = D_i$. An agent $i \in L$ is *active* if it is not tentatively matched to any resource, i.e.,

[2] Without restrictions on valuations, our impossibility result can be easily extended to problems where $|L| = |R|$. To see this, add a resource b to Figure 3.7 with $v_{1b} < \epsilon$ and $v_{2b} < {}^{\epsilon}/_{\beta}$, and let the prediction be $\hat{M} = \{(1, a), (2, b)\}$.

MECHANISM 2: BOOST(I_{BMP^+}, γ)

Input: An instance $I_{\text{BMP}^+} = (G[D], \boldsymbol{v}, \hat{M})$, confidence parameter $\gamma \geq 1$.
Output: A feasible matching for I_{BMP^+}.

```
1  set M = ∅, A = L and P_i = D_i for all i ∈ L
2  while A ≠ ∅ do
3      choose i ∈ A and let j ∈ arg max_{j:(i,j)∈P_i} v_ij
4      if (i, j) ≻_j (M(j), j) then
5          A = A \ {i}
6          if M(j) ≠ ∅ then
7              M = M \ {(M(j), j)}
8              if P_{M(j)} ≠ ∅ then A = A ∪ M(j)
9          end
10         M = M ∪ {(i, j)}
11     end
12     P_i = P_i \ {(i, j)}
13     if P_i = ∅ then  A = A \ {i}
14 end
15 return M
```


$M(i) = \emptyset$, and if i has a remaining proposal to make, i.e., $P_i \neq \emptyset$. In each iteration of the while loop, BOOST chooses an arbitrary active agent $i \in A$ who then makes an offer to an adjacent resource $j \in R$ by following a specific proposal order:

Agent Proposal Order: Each agent $i \in L$ maintains an order on their set of incident edges D_i by sorting them according to non-increasing values v_{ij}. We assume that ties are resolved according to a fixed tie-breaking rule τ_i. Note that the choice of the edge (i, j) of maximum value v_{ij} in Line 3 is uniquely determined by this order.

The main idea behind BOOST is that when an agent i proposes to a resource j, the value v_{ij} is *boosted* only if the edge (i, j) is part of the predicted matching $\hat{M}$. More formally, given an agent i and a declared edge $(i, j) \in D_i$, we define the *offer* $\theta_{ij}(\gamma, \hat{M})$ for resource j as:

$$\theta_{ij}(\gamma, \hat{M}) := \begin{cases} v_{ij}, & \text{if } (i,j) \notin \hat{M}, \\ \gamma v_{ij}, & \text{if } (i,j) \in \hat{M}. \end{cases} \tag{3.6}$$

For brevity, we use θ_{ij} and omit explicit reference to γ and $\hat{M}$ when clear from context. Intuitively, this way BOOST increases the chance that an agent proposing through a predicted edge is accepted (see below) by increasing the offered value by a factor of $\gamma \geq 1$.

Suppose that a resource j receives an offer θ_{ij} from an agent i. Then j *accepts* i's offer if θ_{ij} is the largest offer that j has received so far; otherwise j *rejects* i. We define $\theta_{\emptyset j} = 0$ to indicate that the highest offer that j received is zero if j is still unmatched, i.e., $M(j) = \emptyset$. To this aim, each resource $j \in R$ maintains a fixed preference order over their set of incident edges:

Resource Preference Order: Each resource $j \in R$ maintains an order on their set of incident edges $D_j = \{(i,j) \in D\}$ by sorting them according to non-increasing offer values θ_{ij}. We assume that ties are resolved according to a fixed tie-breaking rule τ_j. Note that the comparison in Line 4 is done with respect to this order.

If an agent i is accepted by a resource j, then i becomes *tentatively matched* to j, i.e., (i,j) is added to M, and i is no longer active, i.e., i is removed from A. Also, if there is some agent k that was tentatively matched to j before, then k is rejected by j, i.e., (k,j) is removed from M, and k becomes active again if k has remaining proposals to make, i.e., k is added to A if $P_k \neq \emptyset$. Whenever an agent proposes and gets rejected, it moves on to the next proposal (if any) according to their proposal order. Note that an agent proposes at most once to each adjacent resource.

3

The mechanism terminates when no active agents are left, i.e., $A = \emptyset$. The current matching becomes definite and is returned by the mechanism. Note that we do not specify how an agent i is chosen from the set of active agents A in Line 3. In fact, any choice will work here. For example, a natural choice is to always choose an active agent $i \in A$ whose next offer is largest.

Intuitively, the confidence parameter $\gamma \geq 1$ specifies to which extent BOOST follows the prediction. On the one extreme, for $\gamma = 1$ BOOST ignores the prediction, which is the best choice in terms of achieving optimal robustness (at the expense of no improved consistency). As γ increases, BOOST follows the prediction more and more. On the other extreme, for $\gamma \to \infty$ BOOST becomes the mechanism TRUST (as introduced in Section 3.3) and simply returns the predicted matching. Naturally, this is the best choice in terms of achieving optimal consistency (at the expense of unbounded robustness).

A crucial difference between BOOST and the existing mechanism for BMP by Dughmi and Ghosh [2010] is that BOOST maintains a tentative matching until it terminates, whereas the known 2-approximate mechanism greedily and permanently matches declared edges according to non-increasing values, while maintaining feasibility. This mechanism coincides with BOOST if $\gamma = 1$, i.e., BOOST ignores the prediction and the next proposal is determined by the highest valued remaining edge among the active agents (assuming that ties are broken in the same way). It is not hard to see that for a mechanism to be strategyproof, it must consider declared edges of an agent i in decreasing order of values. However, to achieve a consistency better than 2, it is crucial that preference is given to predicted edges. These are conflicting objectives if one would want to use a fixed

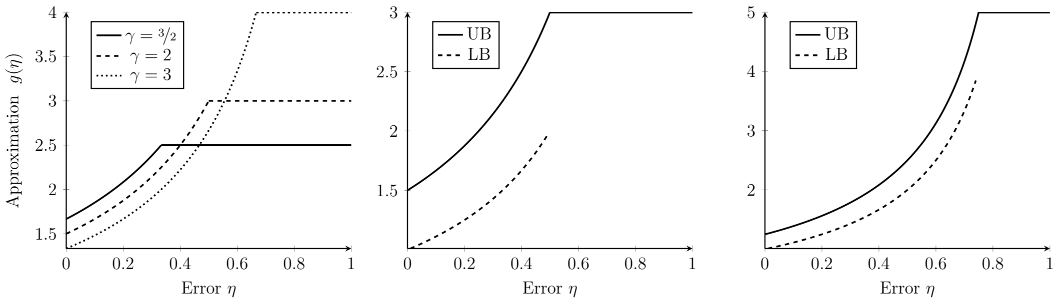

(a) (b) (c)

Figure 3.8: Approximation guarantee $g(\eta)$ as a function of η. (a) For $\gamma \in \{3/2, 2, 3\}$. (b) Upper vs. lower bound for $\gamma = 2$. (c) Upper vs. lower bound for $\gamma = 4$.

global order on declared edges.[3] It therefore seems that tentative matches are needed in order to comply with these conflicting objectives, which is what the mechanism Boost does.

The following is the main result of this section:

Theorem 3.4.1. *Consider the class of instances of BMP$^+$ in the private graph model with a prediction error of at most $\eta \in [0, 1]$. For every confidence parameter $\gamma \geq 1$, Boost is weakly group-strategyproof and achieves an approximation guarantee of:*

$$g(\eta, \gamma) = \begin{cases} \frac{1+\gamma}{\gamma(1-\eta)}, & \text{if } \eta \leq 1 - 1/\gamma, \\ 1+\gamma, & \text{otherwise.} \end{cases} \tag{3.7}$$

In particular, Boost is $(1 + 1/\gamma)$-consistent and $(1 + \gamma)$-robust, which is best possible.

Note that Boost is not only strategyproof, but it satisfies the stronger incentive compatibility notion of WGSP, which we prove in Section 3.4.2. Also, in light of the lower bound given in Theorem 3.3.1, Boost achieves the best possible trade-off in terms of consistency and robustness guarantees. Note that the robustness guarantee kicks in when $\eta \geq 1 - 1/\gamma$. For the range $\eta \in [0, 1 - 1/\gamma]$, as η increases, the approximation interpolates between the consistency and robustness guarantee (see Figure 3.8a). For $\eta \in (0, 1 - 1/\gamma)$, the upper bound for Boost as stated above is off by a factor of $1 + 1/\gamma$ from the lower bound proven in Theorem 3.3.3. As γ increases, this gap between the upper and lower bound decreases (see Figure 3.8b and Figure 3.8c).

[3] Note that the greedy mechanism in [Dughmi and Ghosh, 2010] is not strategyproof when declared edges are considered according to non-increasing γ-boosted values, e.g., consider truthful reports with $v_{1a} = 1 + \epsilon$, $v_{1b} = 1$ and $v_{2b} = 1/2$, prediction $\hat{M} = \{(1, b)\}$ and $\gamma > 1 + \epsilon$.

3.4.2 Weak Group-Strategyproofness

In this section, we show that BOOST is WGSP. The idea behind the proof is to reduce the setting to a one-to-one stable matching problem with standard preferences and argue that BOOST is an instantiation of the (agent-proposing) deferred acceptance algorithm by Gale and Shapley [1962]. Our mechanism inherits the incentive compatibility properties of the deferred acceptance algorithm by Theorem 3.2.1.

In the proof below, it will be crucial that both the agents and the resources use fixed tie-breaking rules $(\tau_i)_{i\in L}$ and $(\tau_j)_{j\in R}$, respectively, as assumed above.

Proof of Theorem 3.4.1 (Weak group-strategyproofness): Let $\gamma \geq 1$ be fixed arbitrarily. Suppose we are given an instance $I_{\text{BMP}^+} = (G[D], \boldsymbol{v}, \hat{M})$ of BMP^+ with compatibility declarations D and a private graph $G[E]$. Based on this, we construct a standard preference system $(G[L \times R], (\succ_i)_{i\in L}, (\succ_j)_{j\in R})$ such that:

3

1. the preference order $\succ_i$ of each agent $i \in L$ mimics i's proposal order over D_i, and

2. the preference order $\succ_j$ of each resource $j \in R$ encodes j's preference order over D_j.

In order to be able to apply Theorem 3.2.1, it is crucial that only the agents can manipulate their preference orders (corresponding to compatibility declarations). The preference orders of the resources must be independent of the declarations and remain fixed.

We start with the definition of the preference orders of the agents.

Agent Preference Order: We define the preference order $\succ_i$ of each agent i as follows. Let D_i be the compatibility set of agent i. Let $\succ_i'$ be the extended lexicographic order of D_i with respect to (v_{ij}, τ_i) as defined in (3.4) in Section 3.2.5, where τ_i refers to the tie-breaking rule used for agent i. (Note that this ensures that $\succ_i'$ is a strict order.) Then $\succ_i$ is defined over the set of all edges in $\{i\} \times R$ as follows:

$$\succ_i: \quad \langle \ \succ_i' \ , \ \ \emptyset \ , \ \ (\{i\} \times R) \setminus D_i \ \ \rangle. \tag{3.8}$$

That is, $\succ_i$ first orders the edges in D_i as in $\succ_i'$, succeeded by $\emptyset$, succeeded by any order of the edges in $(\{i\} \times R) \setminus D_i$.

By construction, e is acceptable ($e \succ_i \emptyset$) if and only if e is declared ($e \in D_i$). Also, all declared edges in D_i are ordered by non-increasing values v_{ij}, which is crucial to align the preferences of the agents with their utilities as defined in (3.2). The order on the set of unacceptable edges in $(\{i\} \times R) \setminus D_i$ is irrelevant and any order will do.

We continue with the definition of the preference orders of the resources.

Resource Preference Order: We define the preference order $\succ_j$ of each resource j as follows. Let $\succ'_j$ be the extended lexicographic order of $L \times \{j\}$ with respect to (θ_{ij}, τ_j) as defined in (3.4) in Section 3.2.5, where τ_j refers to the tie-breaking rule used for resource j. (Note that this ensures that $\succ'_j$ is a strict order.) Then $\succ_j$ is simply defined as $\succ'_j$ over the set of all edges in $L \times \{j\}$.

Note that by construction, this order only depends on $\boldsymbol{v}$, $\hat{M}$ and γ. In particular, $\succ_j$ is independent of the compatibility declarations as required.

Based on the above definitions, running Boost on $(I_{\mathrm{BMP}^+}, \gamma)$ as input is equivalent to running the agent-proposing deferred acceptance algorithm on the standard preference system $(G[L \times R], (\succ_i)_{i \in L}, (\succ_j)_{j \in R})$. Thus, it follows from Theorem 3.2.1 that Boost is weakly group-strategyproof. □

3.4.3 Consistency, Robustness and Approximation

In this section, we prove the consistency, robustness and approximation guarantee of Boost. The following lemma will turn out to be useful.

Lemma 3.4.2. *Let $\gamma \geq 1$. Let $I_{BMP^+} = (G[D], \boldsymbol{v}, \hat{M})$ be an instance of BMP$^+$ and let M be the matching returned by Boost(I_{BMP^+}, γ). Then:*

$$v(M_D^*) \leq 2v(M) + (\gamma - 1)v(M \cap \hat{M}).$$

Proof: Let $\gamma \geq 1$ and let M_D^* be an optimal matching with respect to D. We prove that the value of each edge in M_D^* can be covered by the value of an edge in M output by Boost $(I_{\mathrm{BMP}^+}, \gamma)$. More precisely, we define a mapping $g : M_D^* \to M$ together with some scalars $(\alpha_e)_{e \in M_D^*}$ such that for each edge $e \in M_D^*$ it holds that $v_e \leq \alpha_e \cdot v_{g(e)}$ with $\alpha_e \geq 1$. We also say that e is α_e*-covered* by edge $g(e) \in M$.

Let $e = (i,j) \in M_D^*$. If $e \in M$, we define $g(e) = e$ and $\alpha_e = 1$. Suppose $e = (i,j) \notin M$. We distinguish the following cases.

1. $\exists k \in M(i)$ with $v_{ik} \geq v_{ij}$. We define $g(e) = (i,k)$ and $\alpha_e = 1$.
2. $\exists k \in M(i)$ with $v_{ik} < v_{ij}$. In this case, agent i first proposes to resource j and only later to resource k. In particular, j must have rejected the offer of i immediately or subsequently. Thus, there must be an agent ℓ with $(\ell, j) \in M$ whose offer is preferred over that of i. We define $g(e) = (\ell, j)$ in this case. The definition of α_e depends on whether (i,j) or (ℓ,j) is part of the predicted matching $\hat{M}$.
   (a) $(i,j) \notin \hat{M}$ and $(\ell,j) \notin \hat{M}$: We have $v_{\ell j} \geq v_{ij}$ and define $\alpha_e = 1$.
   (b) $(i,j) \notin \hat{M}$ and $(\ell,j) \in \hat{M}$: We have $\gamma v_{\ell j} \geq v_{ij}$ and define $\alpha_e = \gamma$.
   (c) $(i,j) \in \hat{M}$ and $(\ell,j) \notin \hat{M}$: We have $v_{\ell j} \geq \gamma v_{ij}$ and define $\alpha_e = 1$.

3. $\nexists k \in M(i)$. Note that agent i proposed to resource j but was rejected at some stage. Similarly to the previous case, this implies that there exists some agent ℓ with $(\ell, j) \in M$ whose offer is preferred over that of i. We can follow the same line of arguments as in the previous case and define $g(e) = (\ell, j)$ and α_e in the same way.

Note that the mapping g defined above maps each edge $e \in M_D^*$ either to itself, i.e., $g(e) = e \in M$, or to an edge $g(e) = f \in M$ that is adjacent to e. Also, because M_D^* is a matching, there are at most two edges in M_D^* which are adjacent to an edge $f \in M$. And so, each edge $f \in M$ covers at most two edges in M_D^*. Moreover, if edge $f = (\ell, j) = g(e) \in M$ γ-covers an edge $e = (i, j) \in M_D^*$ (Cases (2b) and (3b) above), f and e must share a common resource (endpoint) j. In particular, the other edge in M_D^* that is mapped to f (if any) has endpoint ℓ and must be 1-covered by f (Case 1). Using the above observations we get:

$$\begin{aligned} v(M_D^*) = \sum_{e \in M_D^*} v_e &\leq \sum_{e \in M_D^*} \alpha_e v_{g(e)} \\ &\leq \sum_{f \in M \setminus \hat{M}} 2v_f + \sum_{f \in M \cap \hat{M}} (1+\gamma) v_f \\ &= 2v(M) + (\gamma - 1) v(\hat{M} \cap M), \end{aligned}$$

which completes the proof. □

Now, the proof that Boost is $(1+\gamma)$-robust follows easily from Lemma 3.4.2.

Proof of Theorem 3.4.1 (Robustness): Let $\gamma \geq 1$. Let $I_{\mathrm{BMP}^+} = (G[D], \boldsymbol{v}, \hat{M})$ be an instance of BMP$^+$ and let M be the matching returned by Boost$(I_{\mathrm{BMP}^+}, \gamma)$. Further, let M_D^* be an optimal matching with respect to D. By Lemma 3.4.2:

$$v(M_D^*) \leq 2v(M) + (\gamma - 1) v(M \cap \hat{M}) \leq 2v(M) + (\gamma - 1) v(M) \leq (1+\gamma) v(M),$$

which proves that Boost is $(1+\gamma)$-robust. □

The next lemma shows that the matching computed by Boost is a $(1 + 1/\gamma)$-approximation of the predicted matching $\hat{M}$ intersected with the declared edges D. It will be useful when proving the approximation guarantee of Boost. The proof is similar to that of Lemma 3.4.2.

Lemma 3.4.3. *Let $\gamma \geq 1$. Let $I_{BMP^+} = (G[D], \boldsymbol{v}, \hat{M})$ be an instance of BMP$^+$ and let M be the matching returned by Boost(I_{BMP^+}, γ). Then:*

$$v(\hat{M} \cap D) \leq (1 + 1/\gamma) v(M).$$

Proof: Let $\gamma \geq 1$. We now prove that the value of each edge in $\hat{M} \cap D$ can be covered by the value of an edge in the matching M output by Boost $(I_{\mathrm{BMP}^+}, \gamma)$.

We define a mapping $g : \hat{M} \cap D \to M$ together with some scalars $(\alpha_e)_{e \in \hat{M} \cap D}$ such that for each edge $e \in \hat{M} \cap D$ it holds that $\alpha_e \cdot v_e \leq v_{g(e)}$ with $\alpha_e \geq 1$. We also say that e is $(1/\alpha_e)$-*covered* by edge $g(e) \in M$.

Let $e = (i, j) \in \hat{M} \cap D$. If $e \in M$, we define $g(e) = e$ and $\alpha_e = 1$. Suppose $e = (i, j) \notin M$. We distinguish the following cases.

1. $\exists k \in M(i)$ with $v_{ik} \geq v_{ij}$. We define $g(e) = (i, k)$ and $\alpha_e = 1$.
2. $\exists k \in M(i)$ with $v_{ik} < v_{ij}$. Note that agent i first proposes to resource j and only later to resource k. In particular, j must have rejected the offer of i immediately or subsequently. Thus, there must be an agent ℓ with $(\ell, j) \in M$ whose offer is preferred over that of i. Recall that $(i, j) \in \hat{M} \cap D$ and thus i made a boosted offer of γv_{ij}. On the other hand, $(\ell, j) \notin \hat{M} \cap D$ and thus ℓ offers $v_{\ell j}$. Therefore, $v_{\ell j} \geq \gamma v_{ij}$. We define $g(e) = (\ell, j)$ and $\alpha_e = \gamma$.
3. $\nexists k \in M(i)$. Note that agent i proposed to resource j at some stage but was rejected (immediately or subsequently). Similarly to the previous case, this implies that there exists some agent ℓ with $(\ell, j) \in M$ and $v_{\ell j} \geq \gamma v_{ij}$. We define $g(e) = (\ell, j)$ and $\alpha_e = \gamma$.

3

Note that the mapping g defined above maps each edge $e \in \hat{M} \cap D$ either to itself or to an edge in M that is adjacent to e. Therefore, each edge in M covers at most two edges in $\hat{M} \cap D$, as $\hat{M} \cap D$ is a matching. Moreover, if edge $f = (\ell, j) = g(e) \in M$ $(1/\gamma)$-covers an edge $e = (i, j) \in \hat{M} \cap D$ (Cases (2) and (3) above), f and e must share a common resource j. In particular, the other edge in $\hat{M} \cap D$ that is mapped to f (if any) has endpoint ℓ and must be 1-covered by f (Case 1). Using the above observations, we can now prove the claim:

$$v(\hat{M} \cap D) = \sum_{e \in \hat{M} \cap D} v_e \leq \sum_{e \in \hat{M} \cap D} \frac{v_{g(e)}}{\alpha_e} \leq \Big(1 + \frac{1}{\gamma}\Big) \sum_{f \in M} v_f = \Big(1 + \frac{1}{\gamma}\Big) v(M).$$

□

We can now complete the proof of Theorem 3.4.1.

Proof of Theorem 3.4.1 (Approximation): Let $\gamma \geq 1$ be fixed arbitrarily. Consider an instance $I_{\mathrm{BMP}^+} = (G[D], \boldsymbol{v}, \hat{M})$ of BMP^+ with prediction error $\eta(I_{\mathrm{BMP}^+}) \leq \eta$. Let M be the matching returned by $\textsc{Boost}(I_{\mathrm{BMP}^+}, \gamma)$. Note that by Lemma 3.4.3, it holds that $(1 + {}^1\!/\!_\gamma) v(M) \geq v(\hat{M} \cap D)$. Now, by using the Lifting Lemma (Lemma 3.3.6), we conclude that $\textsc{Boost}$ is $(1 + {}^1\!/\!_\gamma)/(1 - \eta)$-approximate. Further, the robustness guarantee of $1 + \gamma$ holds independently of the prediction error η. The claimed bound on the approximation guarantee $g(\eta, \gamma)$ as in (3.7) now follows by combining these two bounds. □

3

3.4.4 Extensions of Boost

Boost is rather versatile, in the sense that it can be adapted to handle more general GAP variants while retaining its WGSP property. We summarize a few extensions below.

(E1) Boost can also be run with a many-to-one assignment as input prediction. We exploit this in Section 3.6. The only change is that the offer function in (3.6) is defined with respect to a predicted many-to-one assignment $\hat{M}$. The proof of WGSP in Theorem 3.4.1 continues to hold without change.

(E2) Boost can also handle settings where each agent has unit size, all resources have integer capacities, and the prediction is a many-to-one matchings. Also here, the offer function in (3.6) is defined with respect to a predicted many-to-one matching $\hat{M}$. Further, each resource j now accepts the at most C_j highest offers among the set of proposing agents and rejects the remaining ones. The resulting adaptation of Boost remains WGSP. An easy way to see this is by realizing that this adapted mechanism mimics Boost on the instance obtained from the reduction described next.

(E3) Boost can also be used to handle instances of RSGAP$^+$ by a simple reduction to BMP$^+$. Recall that for an instance $I_{\text{RSGAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ of RSGAP$^+$ it holds that all agents have the same size s_j with respect to a resource $j \in R$, i.e., $s_{ij} = s_j$ for all $i \in L$. It is not hard to see that we can reduce I_{RSGAP^+} to an equivalent instance I_{BMP^+} of BMP$^+$: For each resource $j \in R$, we introduce $m_j = \lfloor C_j/s_j \rfloor$ copies $j_1, \dots, j_{m_j}$. Each copy j_ℓ with $\ell \in [m_j]$ inherits the set of edges D_j incident to j and the value matrix $\boldsymbol{v}'$ is defined accordingly. Similarly, for each resource $j \in R$ with $k_j \leq m_j$ predicted edges $(i_1, j), (i_2, j), \dots, (i_{k_j}, j) \in \hat{M}(j)$, edges $(i_1, j_1), (i_2, j_2), \dots, (i_{k_j}, j_{k_j})$ are added to the predicted matching $\hat{M}'$. Now, each many-to-one assignment in I_{RSGAP^+} corresponds to a matching in I^+_{BMP}. It is not hard to prove that the two instances are equivalent.

The latter observation leads to the following corollary.

Corollary 3.4.4. *Consider the class of instances of RSGAP$^+$ in the private graph model with a prediction error of at most η. For any confidence parameter $\gamma \geq 1$, Boost is weakly group-strategyproof and achieves an approximation guarantee of:*

$$g(\eta, \gamma) = \begin{cases} \frac{1+\gamma}{\gamma(1-\eta)}, & \text{if } \eta \leq 1 - 1/\gamma, \\ 1 + \gamma, & \text{otherwise.} \end{cases}$$

Note that, by choosing $\gamma = 1$, Corollary 3.4.4 implies that Boost is a WGSP and 2-approximate mechanism for RSGAP in the private graph model. To the

MECHANISM 3: GREEDY($I_{\text{GAP}^+}, \boldsymbol{z}$)

Input: An instance $I_{\text{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$, ranking function $\boldsymbol{z} : L \times R \to \mathbb{R}^k$ for some $k \in \mathbb{N}$.

Output: A feasible assignment for I_{GAP^+}.

set $M = \emptyset$ and $\mathcal{L} = \text{SORT}(D, (z(e))_{e \in D})$
while $\mathcal{L} \neq \langle\rangle$ **do**
  let (i, j) be the first edge of $\mathcal{L}$ and remove it
  if $\sum_{t \in M(j) \cup \{i\}} s_{tj} \leq C_j$ **then**
    $M = M \cup \{(i, j)\}$ remove all edges of agent i from $\mathcal{L}$
  end
end
return M

best of our knowledge, the previously best-known mechanism for this problem was the randomized, universally strategyproof, 4-approximate mechanism by Chen et al. [2014].

3.5 Deterministic Greedy Mechanisms for GAP

We start this section by introducing a greedy mechanism for GAP$^+$ that is used as a building block for several mechanisms that will follow. After that, we provide a first application of this mechanism and derive a deterministic mechanism for ERMK$^+$.

3.5.1 A Template of Greedy Mechanisms

On a high level, the template of our mechanism GREEDY works as follows: The mechanism first orders all declared edges according to some specific ranking, which is given as part of the input. According to this order, the mechanism then greedily adds as many edges as possible to an initially empty assignment, while maintaining feasibility. A description of GREEDY can be found in Mechanism 3.

Specifically, GREEDY receives as input an instance $I_{\text{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ of GAP$^+$ and a *ranking function* $\boldsymbol{z} : L \times R \to \mathbb{R}^k$ for some $k \in \mathbb{N}$. It then uses the SORT operator as defined in (3.5) to sort the set of declared edges D in lexicographic decreasing order according to their values $(z_1, \ldots, z_k)$. As a result, the list $\mathcal{L} = \langle e_{\pi_1}, \ldots, e_{\pi_{|D|}} \rangle$ output by SORT satisfies $\boldsymbol{z}(e_{\pi_i}) \succeq^{\text{lex}} \boldsymbol{z}(e_{\pi_j})$ for all $i < j$. GREEDY then processes the edges in this order by always removing the first element (i, j) from $\mathcal{L}$, and greedily assigns agent i to resource j whenever this maintains the feasibility of the constructed assignment M. If i can be assigned to j, the assignment M is updated accordingly, i.e., $M = M \cup \{(i, j)\}$, and all edges of i are removed from the list $\mathcal{L}$. GREEDY terminates if there are no more edges in $\mathcal{L}$, i.e., all agents have either been assigned to a resources or all of their declared edges have been considered.

It is important to realize that GREEDY coupled with an arbitrary ranking function $\boldsymbol{z}$ may not result in a strategyproof mechanism for GAP$^+$ in general. For a ranking function $\boldsymbol{z}$ coupled with GREEDY, Definition 3.5.1 captures a sufficient condition for weak group-strategyproofness, as we show in Theorem 3.5.2.

Definition 3.5.1. Consider some class of instances of GAP$^+$ in the private graph model. We say that a ranking function $\boldsymbol{z}$ is *truth-inducing* if for every instance I_{GAP^+} of this class it holds that:

1. The extended lexicographic order of $L \times R$ with respect to $\boldsymbol{z}$ is strict and total.

2. For every agent $i \in L$ and every $e = (i,j)$, $e' = (i,j') \in \{i\} \times R$ with $\boldsymbol{z}(e) \succeq^{\text{lex}} \boldsymbol{z}(e')$, it holds that $v_e \geq v_{e'}$.

3

Theorem 3.5.2. *Consider some class of instances of GAP$^+$ in the private graph model and let $\boldsymbol{z}$ be a ranking function that is truth-inducing with respect to this class. Then, GREEDY coupled with $\boldsymbol{z}$ is a weakly group-strategyproof mechanism.*

Proof: First, note that as $\boldsymbol{z}$ is truth-inducing, there is a strict and total extended lexicographic order of $L \times R$. As GREEDY always considers declared edges according to this order, for two edges e and f both in D and D', either e is always considered before f or vice verse. Secondly, note that by construction, GREEDY never unassigns an edge during its entire execution. And lastly, consider an instance $I_{\text{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ with truth-inducing ranking function $\boldsymbol{z}$, edge $e \in D$ and an instance $I'_{\text{GAP}^+} = (G[D \backslash \{e\}], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$. Let $M = \text{GREEDY}(I_{\text{GAP}^+}, \boldsymbol{z})$ and $M' = \text{GREEDY}(I'_{\text{GAP}^+}, \boldsymbol{z})$. Note that by construction, if $e \notin M$, then $M = M'$. In other words, removing a single unassigned edge from the instance has no influence on the execution of GREEDY.

Now, consider an instance $I_{\text{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ with private graph $G[E]$. Let $\boldsymbol{z}$ be a ranking function that is truth-inducing for I_{GAP^+}. Consider an agent $i \in L$ with $D_i = E_i$. Let D'_i be a deviation of agent i with I'_{GAP^+} the corresponding instance of this unilateral deviation of i. Let $M = \text{GREEDY}(I_{\text{GAP}^+}, \boldsymbol{z})$ and let $M' = \text{GREEDY}(I'_{\text{GAP}^+}, \boldsymbol{z})$.

As GREEDY never unassigns an edge, agent i cannot benefit from reporting edges in D'_i that are not in E_i. If such an edge $(i,j) \in D'_i \backslash E_i$ is in M', this edge will never be unassigned leading to a utility of zero for agent i. And if these edges are unassigned in M', they have no influence of the execution of GREEDY as $\boldsymbol{z}$ is truth-inducing for I_{GAP^+}. And so, the utility of agent i will not strictly increase in this case. Therefore, consider deviations such that $D'_i \subset E_i$. Again, as $\boldsymbol{z}$ is truth-inducing for I_{GAP^+}, hiding an edge $e \in E_i$ that is not in M has no influence on the the execution of Greedy. So consider the case in which $\exists e = (i,j) \in M$ and $e \notin D'_i$. However, note that for edges $f = (i,k) \in D'_i$ with $z(f) \succeq^{\text{lex}} z(e)$, nothing has changed in the execution of Greedy when these edges are considered as also $f \in E_i$. And so, these edges will also not be assigned in M'. As $\boldsymbol{z}$ is truth-inducing, these are all the edges $f \in D'_i$ with $v_f \geq v_e$. Therefore, the utility of agent i will also not strictly increase.

For weak group-strategyproofness, consider a subset $S \subseteq L$ of agents such that $\forall i \in S$ it holds that $D_i = E_i$. Let D'_S be a group deviation of S with corresponding instance I'_{GAP^+}. Let $M = \text{Greedy}(I_{\text{GAP}^+}, \boldsymbol{z})$ and let $M' = \text{Greedy}(I'_{\text{GAP}^+}, \boldsymbol{z})$. By the same reasoning as above, the utility of an agent $i \in S$ cannot strictly increase when reporting edges in D'_i that are not in E_i. Therefore, consider group deviations such that $D'_i \subset E_i$ for all $i \in S$. Again, if for all $i \in S$ it holds that all $e \in E_i \backslash D'_i$ are not in M, this has no influence on the the execution of Greedy. So consider the case in which $\exists i \in S$, $\exists e = (i,j) \in M$ and $e \notin D'_i$. Of all such agents, let i be the first agent that got an edge $e = (i,j) \in M$ assigned by Greedy. Again, note that for edges $f = (i,k) \in D'_i$ with $z(f) \succeq^{\text{lex}} z(e)$, nothing has changed in the execution of Greedy when these edges are considered as also $f \in E_i$. And so, these edges will also not be assigned in M'. As $\boldsymbol{z}$ is truth-inducing, these are all the edges $f \in D'_i$ with $v_f \geq v_e$. Therefore, the utility of agent i will not strictly increase, and no such i will join such a group deviation, proving weak group-strategyproofness. □ 3

Finally, we stress that, given an instance $I_{\text{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$, Greedy is not necessarily dependent on the predicted assignment $\hat{M}$; it can handle a non-augmented instance of GAP as well. However, the flexibility is in place for the accompanying ranking function $\boldsymbol{z}$ to use $\hat{M}$ in a beneficial manner. We present an implementation of this concept in the following section.

3.5.2 Equal Restricted Multiple Knapsack

In this section, we use the greedy template to derive a deterministic WGSP mechanism for ERMK$^+$. Recall that in this GAP variant every agent i has a single value v_i and a single size s_i, and that $v_i = s_i$. As ERMK is strongly NP-hard as shown by Dawande et al. [2000] and we want to derive a polynomial-time mechanism, we aim to find a truth-inducing ranking function $\boldsymbol{z}$ such that Greedy achieves a constant approximation guarantee. Note that we need to find a ranking function with a domain of $L \times R$, but during an execution of Greedy only the declared edges will need to be ordered. Our ranking function $\boldsymbol{z}$ combines our boosted offer notion $\theta_{ij}(\gamma, \hat{M})$ with a specific tie-breaking rule to favor edges in the predicted assignment $\hat{M}$. This allows us to derive improved approximation guarantees if the prediction error is small, while at the same time retaining bounded robustness if the prediction is erroneous.

Let $\gamma \geq 1$ be fixed arbitrarily. We define the ranking function $\boldsymbol{z} : L \times R \to \mathbb{R}^4$

as follows: Let the boosted offer $\theta_{ij}(\gamma, \hat{M})$ be defined as in (3.6). Also, let $\mathbb{1}_{(i,j)\in\hat{M}}$ be the indicator function which is 1 if and only if $(i,j) \in \hat{M}$. Then, for each $(i,j) \in L \times R$ we define:

$$\boldsymbol{z}((i,j)) := \Big(\theta_{ij}(\gamma, \hat{M}),\ \mathbb{1}_{(i,j)\in\hat{M}},\ -i,\ -j\Big). \tag{3.9}$$

The intuition behind our ranking function is to rank the edges by their γ-boosted value. Recall that for ERMK we have for each agent $i \in L$, $v_i = v_{ij}$ for all resources $j \in R$. In particular, for $\gamma > 1$, the first-order criterion $\theta_{ij}(\gamma, \hat{M})$ already ensures that if there is a predicted edge of agent i, it is ordered before all the other (non-predicted) edges. Crucially, the second-order criterion ensures that this property also holds for $\gamma = 1$. More generally, whenever $\theta_e(\gamma, \hat{M}) = \theta_{e'}(\gamma, \hat{M})$ we make sure that priority is given to the edge in $\hat{M}$ (if any). If any ties remain, they are broken in increasing index of first i and then j.

We use Greedy-by-Theta to refer to the mechanism that we derive by coupling Greedy with the ranking function $\boldsymbol{z}$ as defined in (3.9).

3

Theorem 3.5.3. *Consider the class of instances of ERMK$^+$ in the private graph model with a prediction error of at most η. For every confidence parameter $\gamma \geq 1$, Greedy-by-Theta is weakly group-strategyproof and has an approximation guarantee of:*

$$g(\hat{\eta}, \gamma) = \begin{cases} \frac{1+\gamma}{\gamma(1-\eta)}, & \text{if } \eta \leq 1 - \frac{1+\gamma}{\gamma(2+\gamma)}, \\ 2+\gamma, & \text{otherwise.} \end{cases} \tag{3.10}$$

In particular, Greedy-by-Theta is $(1 + 1/\gamma)$-consistent and $(2 + \gamma)$-robust.

Notably, even for $\gamma = 1$ the preference we give to the predictions in case of ties leads to an improved approximation guarantee if the prediction error η is small, i.e., for $\eta < 1/3$. Also note that, by choosing $\gamma = 1$, Theorem 3.5.3 implies that Greedy-by-Theta is a WGSP and 3-approximate mechanism for ERMK in the private graph model. To the best of our knowledge, no deterministic strategyproof mechanism was known prior to our work.

The remainder of this section is devoted to the proof of Theorem 3.5.3, and starts by showing WGSP.

Proof of Theorem 3.5.3 (Weak group-strategyproofness): We show that the ranking function $\boldsymbol{z}$ is truth-inducing (Definition 3.5.1) for the class of ERMK$^+$ instances. Let $\gamma \geq 1$ and let $I_{\text{ERMK}^+} = (G[D], (v_i = s_i)_{i\in L}, \boldsymbol{C}, \hat{M})$ be any instance of ERMK$^+$. Note that the extended lexicographic order with respect to $\boldsymbol{z}$ as defined in (3.9) is strict and total. Also, for each agent $i \in L$ and all (i,j), $(i,j') \in \{i\} \times R$, we have that $\theta_{ij}(\gamma, \hat{M}) \geq \theta_{ij'}(\gamma, \hat{M})$ trivially implies $v_{ij} \geq v_{ij'}$ because $v_{ij} = v_{ij'} = v_i$. Since the ranking function $\boldsymbol{z}$ is truth-inducing, weak group-strategyproofness follows from Theorem 3.5.2. □

In order to prove that GREEDY-BY-THETA is $(2+\gamma)$-robust, we prove two auxiliary lemmas. The first lemma shows that GREEDY-BY-THETA computes an assignment M that provides a 'utilization guarantee' for each resource which has agents assigned in the optimal assignment that are unassigned in M.[4]

Lemma 3.5.4. *Let $\gamma \geq 1$. Let $I_{ERMK^+} = (G[D], (v_i = s_i)_{i \in L}, \boldsymbol{C}, \hat{M})$ be an instance of ERMK$^+$ and let $M = $ GREEDY-BY-THETA(I_{ERMK^+}, γ) be the computed assignment. For every resource $j \in R$ for which there exists some agent $i \in M^*_D(j)$ with $M(i) = \emptyset$, it holds that:*

$$C_j \leq 2v(M(j)) + (\gamma - 1)v(M(j) \cap \hat{M}(j)).$$

Proof: Let $j \in R$ and let $i \in M^*_D(j)$ with $M(i) = \emptyset$. Since $i \in M^*_D(j)$ and $M(i) = \emptyset$, in the iteration in which edge (i,j) was considered the if-statement in Line 4 failed and (i,j) was not added to M. Let $T \subset L$ denote the set of agents that were assigned to j just before this iteration. Note that $T \neq \emptyset$, since $i \in M^*_D(j)$ and so $s_i = v_i \leq C_j$. Thus, we must have:

3

$$C_j < v(T) + v_i \leq v(M(j)) + v_i, \tag{3.11}$$

where the second inequality holds because $T \subseteq M(j)$, as GREEDY only adds edges to M.

Fix an arbitrary agent $k \in T$. Because GREEDY-BY-THETA considered edge (k,j) before edge (i,j), by the first-order ranking of $\boldsymbol{z}$ in (3.9) we have that $\theta_{ij} \leq \theta_{kj}$. We prove the following useful upper bound on θ_{kj}:

$$\theta_{kj} \leq v(M(j)) + (\gamma - 1)v(M(j) \cap \hat{M}(j)). \tag{3.12}$$

There are two cases to consider.

<u>Case 1:</u> $k \in \hat{M}(j)$. Here, $\theta_{kj} = \gamma v_k \leq \gamma \cdot v(M(j) \cap \hat{M}(j)) \leq v(M(j)) + (\gamma - 1) \cdot v(M(j) \cap \hat{M}(j))$. The first equality follows by the definition of θ_{kj} in (3.6) and the first inequality is true since $k \in M(j) \cap \hat{M}(j)$.

<u>Case 2:</u> $k \notin \hat{M}(j)$. Similarly, $\theta_{kj} = v_k \leq v(M(j)) \leq v(M(j)) + (\gamma - 1) \cdot v(M(j) \cap \hat{M}(j))$, where the last inequality is true since $\gamma \geq 1$.

We can now expand (3.11) as follows:

$$\begin{aligned} C_j &< v(M(j)) + v_i \leq v(M(j)) + \theta_{ij} \leq v(M(j)) + \theta_{kj} \\ &\leq 2v(M(j)) + (\gamma - 1)v(M(j) \cap \hat{M}(j)), \end{aligned}$$

[4]This is a generalization of a key idea used in [Dawande et al., 2000] to our setting *with* predictions. They show that for each unassigned item (agent) in an assignment computed by their greedy algorithm, each knapsack (resource) in which this item fits is at least half-full. For $\gamma = 1$, our mechanism retains this property.

where the second inequality holds by (3.6) and as $\gamma \geq 1$. The third inequality holds because $\theta_{ij} \leq \theta_{kj}$ as argued above, and the last inequality follows by applying (3.12). This concludes the proof. $\square$

The second auxiliary lemma will not only be useful in the robustness proof of Greedy-by-Theta, but it will also be useful in Section 3.6 when we devise a randomized mechanism.

Lemma 3.5.5. *Let $\gamma \geq 1$. Let $I_{ERMK^+} = (G[D], (v_i = s_i)_{i \in L}, \boldsymbol{C}, \hat{M})$ be an instance of ERMK$^+$ and let $M =$ Greedy-by-Theta(I_{ERMK^+}, γ) be the computed assignment. Then:*

$$v(M_D^*) \leq 3v(M) + (\gamma - 1)v(M \cap \hat{M}).$$

Proof: We define S as the set of agents that are unassigned and assigned in M and M_D^*, respectively, i.e., $S := \{i \in L \mid M(i) = \emptyset \wedge M_D^*(i) \neq \emptyset\}$. We define T as the set of resources for which not all, but some agents assigned in the optimal assignment are contained in S, i.e., $T := \{j \in R \mid S \cap M_D^*(j) \neq \emptyset\}$. We have:

$$\begin{aligned} \sum_{i \in S} v_i = \sum_{j \in T} v(S \cap M_D^*(j)) \leq \sum_{j \in T} C_j &\leq \sum_{j \in T} \Big(2v(M(j)) + (\gamma - 1)v(M(j) \cap \hat{M}(j)\Big) \\ &= 2\sum_{j \in T} v(M(j)) + (\gamma - 1)\sum_{j \in T} v(M(j) \cap \hat{M}(j)) \\ &\leq 2v(M) + (\gamma - 1)v(M \cap M). \end{aligned} \tag{3.13}$$

Here, the first equality follows from the definitions of S and T. The first inequality is due to the feasibility of the assignment M_D^*. The second inequality follows by applying Lemma 3.5.4 for each resource $j \in T$. (Note that the precondition of Lemma 3.5.4 is satisfied for each resource $j \in T$: under the assignment M_D^* there exists at least one agent $i \in M_D^*(j)$ with $M(i) = \emptyset$.) Finally, the last inequality holds since $T \subseteq R$ and $\gamma \geq 1$.

Furthermore, the definition of $S \subseteq L$ implies that:

$$v(M_D^*) = \sum_{\substack{i \in L \setminus S: \\ M_D^*(i) \neq \emptyset}} v_i + \sum_{i \in S} v_i \leq \sum_{\substack{i \in L: \\ M(i) \neq \emptyset}} v_i + \sum_{i \in S} v_i = v(M) + \sum_{i \in S} v_i. \tag{3.14}$$

By summing (3.13) and (3.14), the lemma follows. $\square$

Now, the proof that Greedy-by-Theta is $(2 + \gamma)$-robust follows easily.

Proof of Theorem 3.5.3 (Robustness): Let $\gamma \geq 1$ be fixed arbitrarily and let $I_{\text{ERMK}^+} = (G[D], (v_i = s_i)_{i \in L}, \boldsymbol{C}, \hat{M})$ be an instance of ERMK$^+$. Further, let M

be the assignment computed by Greedy-by-Theta($I_{\text{ERMK}^+}, \gamma$) and let M_D^* be an optimal assignment. Then:

$$v(M_D^*) \leq 3v(M) + (\gamma - 1)v(M \cap \hat{M}) \leq 3v(M) + (\gamma - 1)v(M) = (2 + \gamma)v(M),$$

where the first inequality follows by Lemma 3.5.5. This concludes the proof. □

In the remainder of this section, we establish the approximation guarantee of Greedy-by-Theta as stated in Theorem 3.5.3. We first show that whenever Greedy-by-Theta makes a 'mistake', meaning that for an edge $(i, j) \in \hat{M}$ it chooses not to assign agent i to their predicted resource j, then resource j is already $\gamma/\gamma+1$-utilized.

Lemma 3.5.6. *Let $\gamma \geq 1$. Let $I_{ERMK^+} = (G[D], (v_i = s_i)_{i \in L}, \boldsymbol{C}, \hat{M})$ be an instance of ERMK$^+$ and let $M =$ Greedy-by-Theta(I_{ERMK^+}, γ) be the computed assignment. Then, every resource $j \in R$ with $((\hat{M}(j) \cap D) \setminus M(j)) \neq \emptyset$ satisfies:*

$$C_j \leq (1 + 1/\gamma)v(M(j)).$$

3

Proof: Fix an arbitrary agent $i \in (\hat{M}(j) \cap D) \setminus M(j)$ and denote by $T \subset L$ the set of agents assigned to resource j before edge (i, j) was considered. Fix an arbitrary agent $k \in T \setminus (\hat{M}(j) \cap D)$.[5] We have:

$$C_j < v(T) + v_i \leq v(M(j)) + v_i < v(M(j)) + \frac{v_k}{\gamma} \leq \left(1 + \frac{1}{\gamma}\right) v(M(j)).$$

The first inequality holds since $i \notin M(j)$ and $i \in \hat{M}(j)$, so edge (i, j) was considered by Greedy-by-Theta but the if statement in Line 4 failed. The second inequality follows because $T \subseteq M(j)$ as Greedy only adds edges to M. The third inequality follows because Greedy-by-Theta considered edge (k, j) before (i, j), $(i, j) \in \hat{M}$ and by the ranking $\boldsymbol{z}$ (as defined in (3.9)). Finally, the last inequality is true since $k \in M(j)$, concluding the proof. □

The following lemma will further facilitate our proof. It implies that Greedy-by-Theta will never reject a predicted edge for a resource $j \in R$, unless it has already selected an agent not in $\hat{M}(j)$ in a previous iteration.

Lemma 3.5.7. *Let $\gamma \geq 1$. Let $I_{ERMK^+} = (G[D], (v_i = s_i)_{i \in L}, \boldsymbol{C}, \hat{M})$ be an instance of ERMK$^+$ and let $M =$ Greedy-by-Theta(I_{ERMK^+}, γ) be the computed assignment. Then, there is no resource $j \in R$ such that $M(j) \subset \hat{M}(j) \cap D$.*

[5] Note that $T \setminus (\hat{M}(j) \cap D) \neq \emptyset$. Otherwise, by feasibility of $\hat{M} \cap D$, we would have that $v(T) + v_i \leq C_j$ and (i, j) would have been added to M, contradicting that $i \notin M(j)$.

Proof: For contradiction, suppose that there exists a resource $j \in R$ with $M(j) \subset \hat{M}(j) \cap D$. Fix an arbitrary agent $i \in (\hat{M}(j) \cap D) \setminus M(j)$. Note that edge (i, j) is the first edge of agent i that is considered by GREEDY-BY-THETA due to its ranking function $\boldsymbol{z}$ (see (3.9)). Namely, $\gamma v_i = \theta_{ij}(\gamma, \hat{M}) \geq \theta_{i\ell}(\gamma, \hat{M}) = v_i$ for all edges $(i, \ell) \in D_i \setminus \{(i, j)\}$ as $i \in \hat{M}(j)$ and $\hat{M}$ is a many-to-one assignment. Furthermore, in case of ties (which can occur when $\gamma = 1$), edge (i, j) still precedes all other edges by the second-order criterion of $\boldsymbol{z}$.

Denote by $T \subset L$ the set of agents that were already assigned to resource j before edge (i, j) was considered. We have:

$$C_j < v(T) + v_i \leq v(M(j)) + v_i \leq v((\hat{M}(j) \cap D) \setminus \{i\}) + v_i = v(\hat{M}(j) \cap D).$$

Again, note that the first inequality holds since (i, j) is the first edge of agent i considered and the if statement in Line 4 failed for edge (i, j), and the second inequality follows as $T \subseteq M(j)$ (GREEDY only adds edges to M). The third inequality follows by assumption, as $M(j) \subseteq (\hat{M}(j) \cap D) \setminus \{i\}$.

3

As for instances of ERMK$^+$ it holds that $v_i = s_i \; \forall i \in L$, the above inequality implies that $\hat{M} \cap D$ is an infeasible assignment, which is a contradiction. □

We now have all the ingredients to prove that the assignment computed by GREEDY-BY-THETA is a $(1 + {}^1\!/\!_\gamma)$-approximation with respect to the predicted assignment. This is the final building block to complete the proof of the approximation guarantee.

Lemma 3.5.8. *Let $\gamma \geq 1$. Let $I_{ERMK^+} = (G[D], (v_i = s_i)_{i \in L}, \boldsymbol{C}, \hat{M})$ be an instance of ERMK$^+$ and let $M =$ GREEDY-BY-THETA(I_{ERMK^+}, γ) be the computed assignment. Then:*

$$(1 + {}^1\!/\!_\gamma) v(M) \geq v(\hat{M} \cap D).$$

Proof: Define $S_1 := \{j \in R \mid M(j) \supseteq \hat{M}(j) \cap D\}$ and $S_2 := \{j \in R \mid (\hat{M}(j) \cap D) \setminus M(j) \neq \emptyset\}$. Note that S_1 and S_2 form a partition of R, i.e., $R = S_1 \cup S_2$ and $S_1 \cap S_2 = \emptyset$, because there is no resource $j \in R$ with $M(j) \subset \hat{M}(j) \cap D$ by Lemma 3.5.7.

By the definition of S_1, we obtain:

$$\sum_{j \in S_1} v(\hat{M}(j) \cap D) \leq \sum_{j \in S_1} v(M(j)) < \left(1 + \frac{1}{\gamma}\right) \sum_{j \in S_1} v(M(j)). \tag{3.15}$$

Further, we have:

$$\sum_{j \in S_2} v(\hat{M}(j) \cap D) \leq \sum_{j \in S_2} C_j \leq \left(1 + \frac{1}{\gamma}\right) \sum_{j \in S_2} v(M(j)), \tag{3.16}$$

where the first inequality follows because $v_i = s_i$ for all agents $i \in L$ and by the feasibility of $\hat{M} \cap D$. The second inequality follows by applying Lemma 3.5.6 for each resource $j \in S_2$. By summing (3.15) and (3.16), the proof follows. □

We can now complete the proof of Theorem 3.5.3.

Proof of Theorem 3.5.3 (Approximation): Let $\gamma \geq 1$ be fixed arbitrarily. Consider an instance $I_{\text{ERMK}^+} = (G[D], (v_i = s_i)_{i \in L}, \boldsymbol{C}, \hat{M})$ of ERMK$^+$ with prediction error $\eta(I_{\text{ERMK}^+}) \leq \eta$ and let $M = \text{Greedy-by-Theta}(I_{\text{ERMK}^+}, \gamma)$. Note that $(1 + 1/\gamma)v(M) \geq v(\hat{M} \cap D)$ by Lemma 3.5.8, and using the Lifting Lemma (Lemma 3.3.6) we conclude that Greedy-by-Theta achieves an approximation guarantee of $(1 + 1/\gamma)/(1 - \eta)$. Further, the robustness guarantee of $2 + \gamma$ holds independently of the prediction error η. The claimed bound on the approximation guarantee $g(\eta, \gamma)$ in (3.10) now follows by combining these two bounds. □

3.6 Randomized Mechanisms for GAP Variants

3

In this section, we devise randomized mechanisms for multiple variants of GAP with predictions. All randomized mechanisms in this section attain the stronger property of universal WGSP (rather than just universal strategyproofness).

A common thread of all mechanisms in this section is that they may return the outcome of Trust with some probability. As discussed in Section 3.3.2, Trust alone lacks a robustness guarantee. However, we can obtain randomized mechanisms with (finite) expected robustness guarantees by randomizing between Trust and other mechanisms. First, we present our methodology in Section 3.6.1 by applying it to BMP$^+$. Then, we show in Section 3.6.2 that the same idea can be applied to ERMK$^+$. Finally, we derive our randomized mechanisms for the more general variants of GAP$^+$ in Section 3.6.3.

3.6.1 Separation Result for Matching via Randomization

We introduce our mechanism Boost-or-Trust for BMP$^+$ (see Mechanism 4). As the name of our mechanism suggests, it randomizes over the two deterministic mechanisms Boost and Trust. The factor used to boost the values of the predicted edges in Boost is now a function of the confidence parameter γ and is equal to $\delta(\gamma) = \sqrt{2(1+\gamma)} - 1$. Note that $\delta(\gamma) \geq 1$ for all $\gamma \geq 1$. Boost and Trust are executed with probabilities $p = 2/(1 + \delta(\gamma))$ and $1 - p$, respectively. Note that $p \in (0, 1]$ for all $\gamma \geq 1$. If there is confidence that the quality of the prediction $\hat{M}$ is high, Trust can be used to obtain good approximation guarantees for small errors. Boost-or-Trust is able to leverage this and achieve an improved robustness guarantee by sometimes taking a risk and executing Trust, while retaining the approximation guarantee of Theorem 3.4.1 by balancing this risk and sometimes executing BOOST.[6]

MECHANISM 4: BOOST-OR-TRUST(I_{BMP^+}, γ)

Input: An instance $I_{\text{BMP}^+} = (G[D], \boldsymbol{v}, \hat{M})$, confidence parameter $\gamma \geq 1$.
Output: A feasible matching for I_{BMP^+}.

set $\delta(\gamma) = \sqrt{2(1+\gamma)} - 1$ and $p = 2/(1+\delta(\gamma))$
$M_1 = \text{BOOST}(I_{\text{BMP}^+}, \delta(\gamma))$
$M_2 = \text{TRUST}(I_{\text{BMP}^+}, D)$
return M_1 with probability p and M_2 with probability $1-p$

Theorem 3.6.1. *Consider the class of instances of BMP$^+$ in the private graph model with a prediction error of at most η. For every confidence parameter $\gamma \geq 1$, BOOST-OR-TRUST is universally weakly group-strategyproof and has an expected approximation guarantee of:*

$$g(\eta, \gamma) = \begin{cases} \frac{1+\gamma}{\gamma(1-\eta)}, & \text{if } \eta \leq 1 - \frac{\sqrt{2(1+\gamma)}}{2\gamma}, \\ \sqrt{2(1+\gamma)}, & \text{otherwise.} \end{cases} \tag{3.17}$$

In particular, BOOST-OR-TRUST is $(1 + 1/\gamma)$-consistent and $\sqrt{2(1+\gamma)}$-robust (both in expectation).

Remark 3.6.2. As previously argued (Corollary 3.4.4), BOOST can be adapted to handle instances of RSGAP$^+$ with the same approximation guarantees as for BMP$^+$ and remain WGSP. Hence, Theorem 3.6.1 also holds for RSGAP$^+$.

Recall that we concluded from Theorems 3.3.1 and 3.4.1 that BOOST achieves the optimal trade-off between consistency and robustness among all deterministic strategyproof mechanisms. However, as shown in Theorem 3.6.1, the trade-off of the (expected) consistency and robustness achieved by BOOST-OR-TRUST are strictly better than that of any deterministic mechanism (see Figure 3.9). This implies a separation between randomized and deterministic WGSP mechanisms for BMP$^+$ in the private graph model.

Proof of Theorem 3.6.1: Let $\gamma \geq 1$ be fixed arbitrarily. Consider an instance $I_{\text{BMP}^+} = (G[D], \boldsymbol{v}, \hat{M})$ of BMP$^+$ in the private graph model with prediction error $\eta(I_{\text{BMP}^+}) \leq \eta$. Observe that, both matchings that are potentially returned by BOOST-OR-TRUST are the outcome of deterministic WGSP mechanisms, leading

[6]Randomizing between TRUST with probability $p = (\gamma-1)/(1+\gamma)$ and the greedy mechanism in [Dughmi and Ghosh, 2010] with probability $1-p$ is a universally WGSP mechanism that is $(1 + 1/\gamma)$-consistent and $(1+\gamma)$-robust (in expectation). Note that BOOST-OR-TRUST achieves a superior trade-off for $\gamma > 1$.

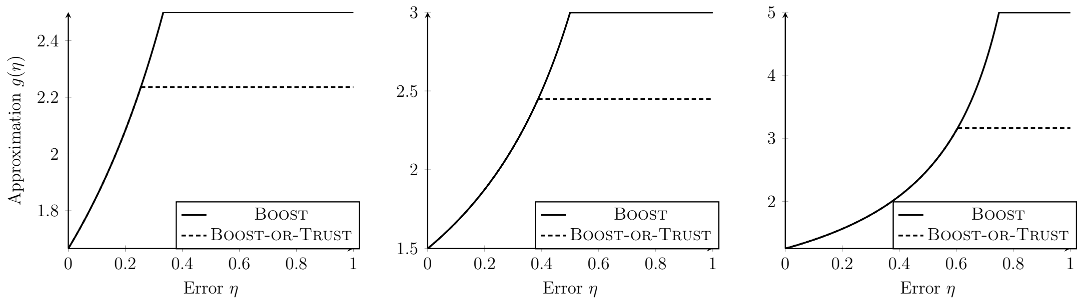

(a) $\gamma = 3/2$ (b) $\gamma = 2$ (c) $\gamma = 4$

Figure 3.9: Approximation guarantees of Boost and Boost-or-Trust (in expectation) as a function of η for various values of γ. The approximation guarantees coincide for small errors, and when the robustness guarantee kicks in the relative improvement of Boost-or-Trust increases as γ increases.

to universal weak group-strategyproofness of Boost-or-Trust. In particular, Boost and Trust are WGSP by Theorems 3.4.1 and 3.3.5, respectively.

We continue by showing that Boost-or-Trust is $\sqrt{2(\gamma+1)}$-robust. Let $M_1 = \textsc{Boost}(I_{\text{BMP}^+}, \delta(\gamma))$ and $M_2 = \textsc{Trust}(I_{\text{BMP}^+})$. We have:

$$\begin{aligned} v(M_D^*) &\leq 2v(M_1) + (\delta(\gamma)-1)v(M_1 \cap \hat{M}) \leq 2v(M_1) + (\delta(\gamma)-1)v(M_2) \\ &= (1+\delta(\gamma))\left(\frac{2}{1+\delta(\gamma)}v(M_1) + \frac{\delta(\gamma)-1}{1+\delta(\gamma)}v(M_2)\right) \\ &= (1+\delta(\gamma))\Big(pv(M_1) + (1-p)v(M_2)\Big) \\ &= (1+\delta(\gamma))\mathbb{E}\left[v\left(\textsc{Boost-or-Trust}\left(I_{\text{BMP}^+}, \gamma\right)\right)\right]. \end{aligned}$$

The first and second inequality follow from Lemma 3.4.2 and the fact that $(M_1 \cap \hat{M}) \subseteq (\hat{M} \cap D) = M_2$, respectively. The penultimate equality is true since, by Line 1 of Boost-or-Trust, we have that $p = 2/(1+\delta(\gamma))$. Finally, since $1+\delta(\gamma) = \sqrt{2(1+\gamma)}$, the robustness guarantee follows.

We now show that Boost-or-Trust achieves an approximation guarantee of $g(\eta, \gamma)$ as in (3.17). We have:

$$\begin{aligned} \mathbb{E}[v(\textsc{Boost-or-Trust}(I_{\text{BMP}^+}, \gamma))] &= pv(M_1) + (1-p)v(\hat{M} \cap D) \\ &\geq \frac{1+\delta(\gamma)-p}{1+\delta(\gamma)}v(\hat{M} \cap D), \end{aligned}$$

where the equality holds since $M_2 = \hat{M} \cap D$ and the inequality follows by applying Lemma 3.4.3 for M_1. After setting $p = 2/(1+\delta(\gamma))$ and $\delta(\gamma) = \sqrt{2(1+\gamma)} - 1$ in the above, as Line 1 of Boost-or-Trust prescribes, it is a matter of simple calculus to obtain $(1+1/\gamma)\mathbb{E}[v(\textsc{Boost-or-Trust}(I_{\text{BMP}^+}, \gamma))] \geq v(\hat{M} \cap D)$. Thus, using the Lifting Lemma (Lemma 3.3.6), we conclude that Boost-or-Trust is $(1 + 1/\gamma)/(1 - \eta)$-approximate in expectation. Furthermore, the expected robustness guarantee of $\sqrt{2(1+\gamma)}$ holds independently of the prediction error η. The claimed bound on the expected approximation guarantee $g(\eta, \gamma)$ now follows by combining these two bounds. □

3.6.2 Randomized Mechanism for ERMK

Using the technique from Section 3.6.1, we give a randomized universally WGSP mechanism, called Greedy-or-Trust (see Mechanism 5), for ERMK$^+$ which randomizes over Greedy-by-Theta and Trust. The factor used to boost the values of the predicted edges in the ranking function $\boldsymbol{z}$ depends on the confidence parameter γ and is defined as $\delta(\gamma) = (\sqrt{12\gamma + 13} - 3)/2$. Note that $\delta(\gamma) \geq 1$ for all $\gamma \geq 1$. The two deterministic mechanisms are executed with probabilities $p = 3/(2 + \delta(\gamma))$ and $1 - p$. Note that $p \in (0, 1]$ for all $\gamma \geq 1$.

MECHANISM 5: Greedy-or-Trust($I_{\text{ERMK}^+}, \gamma$)

Input: An instance $I_{\text{ERMK}^+} = (G[D], (v_i = s_i)_{i \in L}, \boldsymbol{C})$, confidence parameter $\gamma \geq 1$.

Output: A feasible assignment for I_{ERMK^+}.

set $\delta(\gamma) = (\sqrt{12\gamma + 13} - 3)/2$ and $p = 3/(2 + \delta(\gamma))$
$M_1 =$ Greedy-by-Theta($I_{\text{ERMK}^+}, \delta(\gamma)$)
$M_2 =$ Trust(I_{ERMK^+})
return M_1 with probability p and M_2 with probability $1 - p$

Theorem 3.6.3. *Consider the class of instances of ERMK$^+$ in the private graph model with a prediction error of at most η. For every confidence parameter $\gamma \geq 1$, Greedy-or-Trust is universally weakly group-strategyproof and has an expected approximation guarantee of:*

$$g(\eta, \gamma) = \begin{cases} \frac{1+\gamma}{\gamma(1-\eta)}, & \text{if } \eta \leq 1 - \frac{2(1+\gamma)}{\gamma(\sqrt{12\gamma+13}+1)} \\ \frac{\sqrt{12\gamma+13}+1}{2}, & \text{otherwise.} \end{cases} \tag{3.18}$$

In particular, Greedy-or-Trust is $(1 + 1/\gamma)$-consistent and $\frac{1}{2}(\sqrt{12\gamma + 13} + 1)$-robust (both in expectation).

Proof: Consider an instance $I_{\text{ERMK}^+} = (G[D], (v_i = s_i)_{i \in L}, \boldsymbol{C})$ of ERMK$^+$ in the private graph model with prediction error $\eta(I_{\text{BMP}^+}) \leq \eta$ and an arbitrary $\gamma \geq 1$. Greedy-or-Trust is universally WGSP since both deterministic mechanisms in its support, i.e., Greedy-by-Theta and Trust, are WGSP by Theorems 3.5.3 and 3.3.5, respectively.

Let $M_1 = \textsc{Greedy-by-Theta}(I_{\text{ERMK}^+}, \delta(\gamma))$ and $M_2 = \textsc{Trust}(I_{\text{ERMK}^+})$. We continue by showing that Greedy-or-Trust is $\frac{1}{2}(\sqrt{12\gamma+13}+1)$-robust:

$$\begin{aligned} v(M^*_D) &\leq 3v(M_1) + (\delta(\gamma)-1)v(M_1 \cap \hat{M}) \leq 3v(M_1) + (\delta(\gamma)-1)v(M_2) \\ &= (2+\delta(\gamma))\left(\frac{3}{2+\delta(\gamma)}v(M_1) + \frac{\delta(\gamma)-1}{2+\delta(\gamma)}v(M_2)\right) \\ &= (2+\delta(\gamma))\Big(pv(M_1) + (1-p)v(M_2)\Big) \\ &= (2+\delta(\gamma))\mathbb{E}\left[\textsc{Greedy-or-Trust}(I_{\text{ERMK}^+}, \gamma)\right]. \end{aligned}$$

The first and second inequality follow from Lemma 3.5.5 and the fact that $(M_1 \cap \hat{M}) \subseteq (\hat{M} \cap D) = M_2$, respectively. The penultimate equality is true since $p = 3/(2+\delta(\gamma))$, as Greedy-or-Trust prescribes. The expected robustness guarantee then follows as $\delta(\gamma)+2 = \frac{1}{2}(\sqrt{12\gamma+13}+1)$.

To show the claimed approximation $g(\eta,\gamma)$ as in (3.18), note that:

$$\begin{aligned} \mathbb{E}[\textsc{Greedy-or-Trust}(I_{\text{ERMK}^+}, \delta(\gamma))] &= pv(M_1) + (1-p)v(\hat{M} \cap D) \\ &\geq v(\hat{M} \cap D)\left(\frac{p\delta(\gamma)}{1+\delta(\gamma)} + 1 - p\right) \\ &= v(\hat{M} \cap D)\left(1 - \frac{3}{(1+\delta(\gamma))(2+\delta(\gamma))}\right), \end{aligned}$$

where the equality holds since $M_2 = \hat{M} \cap D$ and the inequality follows by applying Lemma 3.5.8 for M_1. The last equality is true since $p = 3/(2+\delta(\gamma))$, as Greedy-or-Trust prescribes. Observe that, for $\delta(\gamma) = \frac{1}{2}(\sqrt{12\gamma+13}-3)$ it holds that:

$$\frac{(1+\delta(\gamma))(2+\delta(\gamma))}{3} = \frac{(\sqrt{12\gamma+13}-1)(\sqrt{12\gamma+13}+1)}{12} = \frac{12\gamma+12}{12} = \gamma+1.$$

Therefore, $(1+{}^1\!/\!_\gamma)\mathbb{E}[\textsc{Greedy-or-Trust}(I_{\text{ERMK}^+}, \delta(\gamma))] \geq v(\hat{M} \cap D)$ follows by rearranging terms. Using the Lifting Lemma (Lemma 3.3.6), we conclude that Greedy-or-Trust is $(1+{}^1\!/\!_\gamma)/(1-\eta)$-approximate in expectation. And as the expected robustness guarantee of $\frac{1}{2}(\sqrt{12\gamma+13}+1)$ holds independently of the prediction error η, the claimed bound $g(\eta,\gamma)$ follows. □

3.6.3 More General Variants of GAP

In this section we devise randomized universally WGSP mechanisms for two variants of GAP$^+$, namely VCGAP$^+$ and ASGAP$^+$. Recall that for VCGAP the resources can be ordered from most to least valuable such that every agent agrees with this ordering, and that AVGAP, RVGAP and RMK are special cases of

VCGAP. Also recall that for ASGAP every agent i has a single size s_i. While Dughmi and Ghosh [2010] and Chen et al. [2014] studied multiple GAP variants in the private graph model, to the best of our knowledge, the VCGAP has not been considered in the literature and no deterministic strategyproof $O(1)$-approximate mechanism for the private graph model is known. Also for ASGAP in the private graph model no deterministic strategyproof $O(1)$-approximate mechanism is known.

The mechanisms that we derive for both GAP variants will randomize over Boost, Greedy and Trust. We therefore start with introducing the last remaining building blocks by applying the greedy template to both variants.

Two WGSP Greedy Mechanisms. In Section 3.5.2 we demonstrated how Greedy combined with an appropriate ranking function can serve as a WGSP mechanism for the special case of ERMK$^+$. Here, we present two different instantiations of ranking functions, one for VCGAP and one for ASGAP, which can be coupled with Greedy to obtain a WGSP mechanism for their respective GAP variant.

We present our ranking function for VCGAP first, and define the function $\boldsymbol{z}_{\text{VCGAP}} : L \times R \mapsto \mathbb{R}^3$ as follows: Let $\sigma = (\sigma(1), \dots, \sigma(m))$ be the ordinal consensus permutation of resources in R for the considered instance. Then, for every pair $(i, j) \in L \times R$ we define:

$$\boldsymbol{z}_{\text{VCGAP}}((i,j)) \coloneqq \Big(-\sigma(j), {}^{v_{ij}}/_{s_{ij}}, -i \Big). \tag{3.19}$$

This function instructs Greedy to rank the edges so that the edges to the most 'preferred' resources, according to the ordinal consensus σ, are considered first. Then, among the edges to a specific resource, the second-order criterion instructs Greedy to give precedence to the edge with the highest value per size ratio. If any ties remain, they are broken in increasing index of i.

For ASGAP, we define the function $\boldsymbol{z}_{\text{ASGAP}} : L \times R \mapsto \mathbb{R}^3$ as follows: For every pair $(i, j) \in L \times R$ we define:

$$\boldsymbol{z}_{\text{ASGAP}}((i,j)) \coloneqq \Big({}^{v_{ij}}/_{s_{ij}}, -i, -j \Big). \tag{3.20}$$

This ranking function instructs Greedy to prioritize edges with the highest value per size ratio. Then, in case of ties, they are broken in increasing index of first i and then j.

Throughout this section, when invoking Greedy for VCGAP$^+$ and ASGAP$^+$ instances, we will refer to the pairing of Greedy with the corresponding ranking functions $\boldsymbol{z}_{\text{VCGAP}}$ and $\boldsymbol{z}_{\text{ASGAP}}$ respectively. We first show that each of these two instantiations of Greedy is a WGSP mechanism for its respective GAP variant.

Lemma 3.6.4. *Mechanism* Greedy *coupled with ranking function* $\boldsymbol{z}_{VCGAP}$ *(or* $\boldsymbol{z}_{ASGAP}$*) is a weakly group-strategyproof mechanism for instances of VCGAP*$^+$ *(or ASGAP*$^+$*, respectively).*

Proof: We show that both ranking functions are truth-inducing (Definition 3.5.1).

First, consider the ranking function $\boldsymbol{z}_{\text{VCGAP}}$ and the class of VCGAP$^+$ instances. Consider an instance $I_{\text{VCGAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ and let σ be the permutation of resources that indicates the consensus of their value among the agents. By definition, the extended lexicographic order with respect to $\boldsymbol{z}_{\text{VCGAP}}$ as defined in (3.19) is strict and total. Furthermore, by the definition of σ, it is true that for every agent $i \in L$ and all $e = (i,j), e' = (i,j') \in \{i\} \times R$, $\sigma(j) \leq \sigma(j')$, i.e., $\boldsymbol{z}(e) \succeq^{\text{lex}} \boldsymbol{z}(e')$, implies that $v_{ij} \geq v_{ij'}$. Therefore, the ranking function $\boldsymbol{z}_{\text{VCGAP}}$ is truth-inducing and weak group-strategyproofness of this greedy mechanism follows from Theorem 3.5.2.

Now, consider the ranking function $\boldsymbol{z}_{\text{ASGAP}}$ and the class of ASGAP$^+$ instances. Consider an instance $I_{\text{ASGAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$. By definition, the extended lexicographic order with respect to $\boldsymbol{z}_{\text{VCGAP}}$ as defined in (3.20) is strict and total. Furthermore, for every agent $i \in L$ and all $e = (i,j), e' = (i,j') \in \{i\} \times R$, since $s_{ij} = s_{ij'}$, $\frac{v_{ij}}{s_{ij}} \geq \frac{v_{ij'}}{s_{ij'}}$, i.e., $\boldsymbol{z}(e) \succeq^{\text{lex}} \boldsymbol{z}(e')$, directly implies that $v_{ij} \geq v_{ij'}$. Therefore, the ranking function $\boldsymbol{z}_{\text{ASGAP}}$ is truth-inducing and weak group-strategyproofness of this greedy mechanism follows from Theorem 3.5.2. □

3

Note that both $\boldsymbol{z}_{\text{VCGAP}}$ and $\boldsymbol{z}_{\text{ASGAP}}$ do not use the predicted assignment $\hat{M}$ in any way. Furthermore, the greedy mechanisms described above do not guarantee worst-case approximation guarantees when run as stand-alone mechanisms. However, as stated in the following observation, both ranking functions do ensure that agents are processed by Greedy in an efficient way.

Observation 3.6.5. For an instance $I_{\text{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ of VCGAP$^+$ (or ASGAP$^+$), the corresponding ranking function $\boldsymbol{z}$ as defined in (3.19) (or (3.20), respectively) satisfies the following property: for every resource $j \in R$, and every $(i,j), (k,j) \in D_j$ with $\boldsymbol{z}((i,j)) \succeq^{\text{lex}} \boldsymbol{z}((k,j))$, it holds that $\frac{v_{ij}}{s_{ij}} \geq \frac{v_{kj}}{s_{kj}}$.

Two Universally WGSP Mechanisms. We present our randomized mechanisms for VCGAP$^+$ and ASGAP$^+$. Intuitively, for these GAP variants, edges can be of crucial value to the constructed assignment in two ways. Namely, edges can have a large value or a large value over cost ratio (efficiency). To incorporate this, our mechanism randomizes, among others, over Greedy coupled with the ranking functions introduced above, which processes edges in order of efficiency (as argued in Observation 3.6.5). Our mechanism also randomizes over Boost, which processes edges in order of their values. While neither of these mechanisms achieves a bounded approximation guarantee by itself, their probabilistic combination does in expectation. In fact, this is the key idea that Chen et al. [2014]

MECHANISM 6: BOOST-OR-GREEDY-OR-TRUST $(I_{\mathrm{GAP}^+}, \gamma, \boldsymbol{z})$

Input: An instance $I_{\mathrm{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$, confidence parameter $\gamma \geq 1$, ranking function $\boldsymbol{z} : L \times R \mapsto \mathbb{R}^k$ for some $k \in \mathbb{N}$.

Output: A feasible assignment for I_{GAP^+}.

1 construct an instance $I_{\mathrm{BMP}^+} = (G[D], \boldsymbol{v}, \hat{M})$ and set $p = 2/(3 + \gamma)$
2 $M_1 = \text{BOOST}(I_{\mathrm{BMP}^+}, \gamma)$
3 $M_2 = \text{GREEDY}(I_{\mathrm{GAP}^+}, \boldsymbol{z})$
4 $M_3 = \text{TRUST}(I_{\mathrm{GAP}^+})$
5 **return** M_1 or M_2 both with probability p and M_3 with probability $1 - 2p$

used to devise the current state-of-the-art universally strategyproof mechanism for ASGAP and special cases of VCGAP. The deferred acceptance mechanism that Chen et al. [2014] use can be cast into our GREEDY template as it simply follows a specific ranking of the declared edges such that agents are never unmatched (see the proof of Theorem 1 in [Chen et al., 2014]). Inspired by this idea, we instead randomize over our mechanisms BOOST, GREEDY and TRUST.

We refer to the resulting mechanism as BOOST-OR-GREEDY-OR-TRUST (see Mechanism 6). Note that the assignments M_1 and M_2 computed by BOOST and GREEDY, respectively, are returned with positive probability $p = 2/(3 + \gamma)$ for any confidence parameter $\gamma \geq 1$. On the other hand, the assignment computed by TRUST is returned with probability $1 - 2p$, which is only positive if $\gamma > 1$, i.e., when there is some confidence in the prediction. A subtle point of BOOST-OR-GREEDY-OR-TRUST is that the predicted assignment $\hat{M}$ of the constructed instance I_{BMP^+} in Line 1 is a many-to-one assignment. However, as argued in (E1) in Section 3.4.4, BOOST can handle such instances as input as well. Notably, M_1 being a one-to-one assignment output by BOOST suffices to prove bounded approximation guarantees for BOOST-OR-GREEDY-OR-TRUST.

Theorem 3.6.6. *Consider the class of instances of ASGAP$^+$ (or VCGAP$^+$, respectively) in the private graph model with a prediction error of at most η. For every confidence parameter $\gamma \geq 1$, BOOST-OR-GREEDY-OR-TRUST is universally weakly group-strategyproof and has an expected approximation guarantee of:*

$$g(\eta, \gamma) = \begin{cases} \frac{3+\gamma}{\gamma(1-\eta)}, & \text{if } \eta \leq 1 - \frac{1}{\gamma}, \\ 3 + \gamma, & \text{otherwise.} \end{cases} \tag{3.21}$$

In particular, BOOST-OR-GREEDY-OR-TRUST is $(1 + 3/\gamma)$-consistent and $(3 + \gamma)$-robust (both in expectation).

The remainder of this section is dedicated to the proof of Theorem 3.6.6. Note that universal WGSP of BOOST-OR-GREEDY-OR-TRUST follows trivially.

Proof of Theorem 3.6.6 (Universal weak group-strategyproofness): Fix $\gamma \geq 1$ arbitrarily. Observe that, all assignments that are potentially returned by BOOST-OR-GREEDY-OR-TRUST are the outcome of deterministic WGSP mechanisms, leading to universal weak group-strategyproofness of BOOST-OR-GREEDY-OR-TRUST. In particular, GREEDY and TRUST are WGSP by Lemma 3.6.4 and Theorem 3.3.5, respectively, and in Section 3.4.4 we showed that BOOST is also WGSP when given a many-to-one predicted assignment as input (E1). □

Consistency, Robustness and Approximation. Next, we turn our attention to the expected approximation guarantee of Theorem 3.6.6. To that end, we introduce some auxiliary notation. For an arbitrary $\gamma \geq 1$ and instance $I_{\text{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$, let M_1 and M_2 be the computed assignment of BOOST and GREEDY, respectively, as prescribed by BOOST-OR-GREEDY-OR-TRUST. Furthermore, let M be any feasible assignment for I_{GAP^+}. We denote by $S(M) \subseteq L$ the set of agents who have been assigned to a resource of smaller value under both M_1 and M_2 compared to their value under M. Formally:

3

$$S(M) = \left\{ i \in L \;\middle|\; v_{iM_1(i)} < v_{iM(i)} \wedge v_{iM_2(i)} < v_{iM(i)} \right\}.$$

Note that $S(M)$ depends on both M_1 and M_2, but we omit these arguments for ease of notation. In order to prove the expected approximation guarantee of Theorem 3.6.6, we first state a simple observation implied by a previous work. Then, we establish two technical lemmas, Lemma 3.6.8 and Lemma 3.6.9, and with these lemmas in place, the expected approximation guarantee will follow easily.

Observation 3.6.7: (Implied by Theorem 1 in [Chen et al., 2014]). Let $\gamma \geq 1$. Let $I_{\text{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ be an instance of VCGAP$^+$ (or ASGAP$^+$, respectively) and let M_1 and M_2 be the assignments computed by BOOST and GREEDY as prescribed by BOOST-OR-GREEDY-OR-TRUST. Furthermore, let M be a feasible assignment for I_{GAP^+}. Then:

$$v(M_1) + v(M_2) \geq \sum_{j \in R} v(M(j) \setminus S(M)).$$

We continue with the two technical lemmas. The first lemma exploits the fact that GREEDY considers edges in order of efficiency.

Lemma 3.6.8. *Let $\gamma \geq 1$. Let $I_{GAP^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ be an instance of VCGAP$^+$ (or ASGAP$^+$, respectively) and let M_1 and M_2 be the assignments computed by BOOST and GREEDY as prescribed by BOOST-OR-GREEDY-OR-TRUST.*

Furthermore, let M *be any feasible assignment for* I_{GAP^+} *and let* $j \in R$ *be a resource with* $M(j) \cap S(M) \neq \emptyset$. *Finally, let* $\ell \in \arg\max_{i \in M(j) \cap S(M)} s_{ij}$. *It holds that:*

$$v\big(M(j) \cap S(M)\big) \leq v(M_2(j)) + v_{\ell j}.$$

Proof: Consider any agent $k \in M(j) \cap S(M)$. By definition of $S(M)$, it holds that $v_{kj} > v_{kM_2(k)}$. Observe that, for both ranking functions $\boldsymbol{z}_{\text{VCGAP}}$ and $\boldsymbol{z}_{\text{VCGAP}}$, this implies that $(k, j) \succeq^{\text{lex}} (k, M_2(k))$. Therefore, agent k was first considered by Greedy to be matched to resource j, but was rejected and was instead matched to resource $M_2(k)$.

Let $T(k) \subseteq M_2(j)$ be the set of agents that were assigned to resource j by Greedy *before* edge (k, j) was considered. Note that $T(k)$ is not empty, since $s_{kj} \leq C_j$ by assumption. We have:

$$\sum_{i \in T(k)} s_{ij} + s_{kj} > C_j \geq \sum_{i \in M(j) \cap S(M)} s_{ij}, \tag{3.22}$$

where the first inequality follows from the fact that Greedy considered edge (k, j) but the if statement in Line 4 failed. The second inequality follows from the feasibility of M.

Further, by Observation 3.6.5, for each agent $i \in T(k)$ it holds that $v_{kj} s_{ij} \leq v_{ij} s_{kj}$. By summing over all agents in $T(k)$ and rearranging terms, we obtain that:

$$v_{kj} \leq s_{kj} \frac{\sum_{i \in T(k)} v_{ij}}{\sum_{i \in T(k)} s_{ij}}. \tag{3.23}$$

Recall that $\ell = \arg\max_{i \in M(j) \cap S(M)} s_{ij}$. We can conclude the following:

$$\begin{aligned} v(M(j) \cap S(M)) - v_{\ell j} &= \sum_{k \in (M(j) \cap S(M)) \setminus \{\ell\}} v_{kj} \leq \sum_{k \in (M(j) \cap S(M)) \setminus \{\ell\}} s_{kj} \frac{\sum_{i \in T(k)} v_{ij}}{\sum_{i \in T(k)} s_{ij}} \\ &\leq v(M_2(j)) \cdot \sum_{k \in (M(j) \cap S(M)) \setminus \{\ell\}} \frac{s_{kj}}{\sum_{i \in T(k)} s_{ij}} \\ &\leq v(M_2(j)) \cdot \sum_{k \in (M(j) \cap S(M)) \setminus \{\ell\}} \frac{s_{kj}}{\sum_{i \in (M(j) \cap S(M))} s_{ij} - s_{kj}} \\ &\leq v(M_2(j)) \cdot \sum_{k \in (M(j) \cap S(M)) \setminus \{\ell\}} \frac{s_{kj}}{\sum_{i \in (M(j) \cap S(M))} s_{ij} - s_{\ell j}} \\ &= v(M_2(j)). \end{aligned}$$

The first inequality follows from (3.23) and the second inequality from the fact that $T(k) \subseteq M_2(j)$ for each agent $k \in (M(j) \cap S(M)) \setminus \{\ell\}$. Then, the next inequality follows by applying (3.22) for each agent k in the summation. Finally, the last inequality follows from the definition of agent ℓ. The lemma follows. □

The following second technical lemma is the last ingredient for the proof of the approximation guarantee of Theorem 3.6.6.

Lemma 3.6.9. *Let $\gamma \geq 1$. Let $I_{GAP^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ be an instance of VCGAP$^+$ (or ASGAP$^+$, respectively) and let M_1 and M_2 be the assignments computed by* Boost *and* Greedy *as prescribed by* Boost-or-Greedy-or-Trust. *Furthermore, let M be any feasible assignment for I_{GAP^+}. For each $j \in R$ with $M(j) \cap S(M) \neq \emptyset$, fix an agent $\ell(j) \in \arg\max_{i \in M(j) \cap S(M)} s_{ij}$. Then:*

$$v(M) \leq v(M_1) + 2v(M_2) + \sum_{\substack{j \in R: \\ M(j) \cap S(M) \neq \emptyset}} v_{\ell(j)j}. \tag{3.24}$$

Proof: We start by rewriting $v(M)$ and then apply Observation 3.6.7:

$$\begin{aligned} v(M) &= \sum_{j \in R} v(M(j)) = \sum_{j \in R} v(M(j) \setminus S(M)) + \sum_{j \in R} v(M(j) \cap S(M)) \\ &\leq v(M_1) + v(M_2) + \sum_{j \in R} v(M(j) \cap S(M)) \\ &\leq v(M_1) + 2v(M_2) + \sum_{\substack{j \in R: \\ M(j) \cap S(M) \neq \emptyset}} v_{\ell(j)j}. \end{aligned}$$

The last inequality follows by applying Lemma 3.6.8 for each $j \in R$ with $M(j) \cap S(M) \neq \emptyset$, concluding the proof. □

Proof of Theorem 3.6.6 (Robustness): Let $\gamma \geq 1$. Let $I_{\text{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ be an instance of VCGAP$^+$ (or ASGAP$^+$, respectively). Let M_1, M_2 and M_3 be the assignments computed by Boost, Greedy and Trust as prescribed by Boost-or-Greedy-or-Trust. Let M be the assignment computed by Boost-or-Greedy-or-Trust.

Consider an edge (i, j) so that $i \in \arg\max_{k \in M_D^*(j) \cap S(M_D^*)} s_{kj}$. By the construction of Boost and the definition of $S(M_D^*)$, it must be that agent i proposed to resource $M_D^*(i) = j$. However, resource j (eventually) rejected the offer θ_{ij} of agent i and instead opted for the offer of agent $M_1(j)$. Therefore, by definition of θ as in (3.6), it holds that:

$$v_{ij} \leq \theta_{ij} \leq \theta_{M_1(j)j} = v(M_1(j)) + \mathbb{1}_{M_1(j) \in \hat{M}(j)} \cdot (\gamma - 1)v(M_1(j)). \tag{3.25}$$

We continue by applying Lemma 3.6.9 to M_D^* and each $j \in R$ with $M_D^*(j) \cap S(M_D^*) \neq \emptyset$ and agent $\ell(j) \in \arg\max_{i \in M_D^*(j) \cap S(M_D^*)} s_{ij}$. Thus, we can expand

(3.24) as follows:

$$\begin{aligned}
v(M_D^*) &\leq v(M_1) + 2v(M_2) + \sum_{\substack{j\in R:\\ M_D^*(j)\cap S(M_D^*)\neq\emptyset}} v_{\ell(j)j} \\
&\leq v(M_1) + 2v(M_2) + \sum_{\substack{j\in R:\\ M_D^*(j)\cap S(M_D^*)\neq\emptyset}} \Big(v(M_1(j)) + \mathbb{1}_{M_1(j)\in\hat{M}(j)} \cdot (\gamma-1)v(M_1(j))\Big) \\
&\leq v(M_1) + 2v(M_2) + \sum_{j\in R} \Big(v(M_1(j)) + \mathbb{1}_{M_1(j)\in\hat{M}(j)} \cdot (\gamma-1)v(M_1(j))\Big) \\
&= 2v(M_1) + 2v(M_2) + (\gamma-1)\sum_{j\in R} \mathbb{1}_{M_1(j)\in\hat{M}(j)} v(M_1(j)) \\
&= 2v(M_1) + 2v(M_2) + (\gamma-1)v(M_1\cap\hat{M}) \\
&\leq 2v(M_1) + 2v(M_2) + (\gamma-1)v(M_3) \\
&= (3+\gamma)\Big(pv(M_1) + pv(M_2) + (1-2p)v(M_3)\Big) = (3+\gamma)\mathbb{E}[M].
\end{aligned}$$

Here, the second inequality follows from (3.25) and the third inequality follows as $\gamma \geq 1$. The last inequality follows from the fact that $M_1 \cap \hat{M} \subseteq D \cap \hat{M} = M_3$. Finally, as Boost-or-Greedy-or-Trust sets $p = 2/(3+\gamma)$, the penultimate equality and the claimed expected robustness guarantee follow. □

We can now complete the proof of Theorem 3.6.6.

Proof of Theorem 3.6.6 (Approximation): Let $\gamma \geq 1$ be fixed arbitrarily. Consider an instance $I_{\mathrm{GAP}^+} = (G[D], \boldsymbol{v}, \boldsymbol{s}, \boldsymbol{C}, \hat{M})$ of VCGAP$^+$ (or ASGAP$^+$, respectively) with prediction error $\eta(I_{\mathrm{GAP}^+}) \leq \eta$. Let M_1 and M_2 be the assignments computed by Boost and Greedy as prescribed by Boost-or-Greedy-or-Trust and let M be the assignment returned by Boost-or-Greedy-or-Trust. For notational convenience, we use $\hat{M}_D \coloneqq \hat{M} \cap D$ and, for each $j \in R$, $\hat{M}_D(j) \coloneqq \hat{M}(j) \cap D_j$.

Consider an edge (i,j) so that $i \in \arg\max_{k\in\hat{M}_D(j)\cap S(\hat{M}_D)} s_{kj}$. By the construction of Boost and the definition of $S(\hat{M}_D)$, it must be that agent i proposed to resource j for which i is predicted to be assigned. However, resource j (eventually) rejected the offer θ_{ij} of agent i and instead opted for the offer of agent $M_1(j)$. Therefore, by definition of θ as in (3.6) and since $(i,j) \in \hat{M}$, it holds that:

$$v_{ij} = \frac{\theta_{ij}}{\gamma} \leq \frac{\theta_{M_1(j)j}}{\gamma} \leq \frac{\gamma v(M_1(j))}{\gamma} = v(M_1(j)). \tag{3.26}$$

We now apply Lemma 3.6.9 for $\hat{M}_D$ and each $j \in R$ with $\hat{M}_D(j) \cap S(\hat{M}_D) \neq \emptyset$

3

and agent $\ell(j) \in \arg\max_{\hat{M}_D(j) \cap S(\hat{M}_D)} s_{ij}$. We expand (3.24) as follows:

$$v(\hat{M}_D) \leq v(M_1) + 2v(M_2) + \sum_{\substack{j \in R: \\ \hat{M}_D(j) \cap S(\hat{M}_D) \neq \emptyset}} v_{\ell(j)j} \leq v(M_1) + 2v(M_2) + \sum_{\substack{j \in R: \\ \hat{M}_D(j) \cap S(\hat{M}_D) \neq \emptyset}} v(M_1(j))$$
$$\leq v(M_1) + 2v(M_2) + \sum_{j \in R} v(M_1(j)) = 2v(M_1) + 2v(M_2).$$

The second inequality follows by applying (3.26) for each resource $j \in R$ with $\hat{M}_D(j) \cap S(\hat{M}_D) \neq \emptyset$ and agent $\ell(j)$ as specified above. Finally, we get:

$$\mathbb{E}[M] = pv(M_1) + pv(M_2) + (1 - 2p)v(\hat{M}_D) \geq v(\hat{M}_D)\left(\frac{p}{2} + 1 - 2p\right) = v(\hat{M}_D)\frac{\gamma}{3 + \gamma},$$

where the first equality follows from the fact that $\hat{M}_D = \textsc{Trust}(I_{\text{GAP}^+})$ and the first inequality follows from the expression above as $p > 0$ for every $\gamma \geq 1$. The last equality is due to the fact that $p = 2/(3 + \gamma)$ as defined in Line 1 of Boost-or-Greedy-or-Trust. By rearranging terms we obtain that $(3 + \gamma)/\gamma \cdot \mathbb{E}[M] \geq v(\hat{M}_D)$. Now, by using the Lifting Lemma (Lemma 3.3.6), we conclude that Boost-or-Greedy-or-Trust achieves an expected approximation guarantee of $(3 + \gamma)/(\gamma(1 - \eta))$. Furthermore, the expected robustness guarantee of $3 + \gamma$ holds independently of the prediction error η. The claimed bound on the expected approximation guarantee $g(\eta, \gamma)$ in (3.21) now follows by combining these two bounds. □

3

3.7 Conclusion

This chapter contributes to the emerging line of research on approximate mechanism design with predictions for problems without monetary transfers. More specifically, it studies the generalized assignment problem in the private graph model, which was introduced by Dughmi and Ghosh [2010], in which each problem instance is augmented with a prediction of the optimal assignment. All mechanisms derived in this chapter satisfy weak group-strategyproofness and achieve approximation guarantees that smoothly interpolate between the respective consistency and robustness guarantees, depending on a natural error parameter.

We believe that this chapter offers a comprehensive treatment of leveraging predictions for different variants of the problem. A crucial element throughout the chapter is our deterministic mechanism Boost, initially designed for the bipartite matching problem for which it achieves the optimal consistency-robustness trade-off. For the randomized mechanisms that we derive for more general variants of the problem, Boost proves to be a crucial component that leverages the prediction in a non-trivial way. The other component that our randomized mechanisms use to trivially leverage the prediction is our mechanism Trust, that simply adheres to the prediction. One could randomize over Trust and any known strategyproof mechanism for the problem without predictions, but, as we show for the bipartite matching problem, leveraging the prediction in a non-trivial way can give rise to a superior consistency-robustness trade-off. Additional to the aforementioned mechanisms Boost and Trust, our randomized mechanisms for the more general variants also have a greedy mechanism as a component that does not leverage the prediction. It would be interesting to see if further improvements in the consistency-robustness trade-offs are possible if this greedy, or another, mechanism also leverages the prediction. In fact, improving the expected approximation guarantees for the more general variants are also intriguing open problems in the setting without predictions, and developments here probably also give rise to improvements in the setting with predictions. Another open problem that remains is the approximation guarantee of Boost for the bipartite matching problem, which might still be improved for small values of the prediction error. Finally, an interesting new direction would be to study problems with more complex preferences of the agents and resources in a learning-augmented environment.

Chapter 4

Facility Location with Outliers

4.1 Introduction

Suppose you are working in the public sector and your task is to decide on the location of a new public facility. You could determine the location without asking for the opinion of the citizens that will end up using the facility, but this might result in a very unpopular decision. Ideally though, you would like to choose a location that is aligned with the preferences of the citizens—everyone wants the location to be as close as possible to their preferred location. Just as the last time when you were faced with this problem, you will ask the citizens to report their preferences in a way that incentivizes them to declare their *true* preferences. However, given all the preferred locations of last time in Figure 4.1, this method of decision-making led to the location of the facility as marked in red. Although citizens truthfully reported their preferred locations, this decision resulted in a long commute to the facility for a majority of the citizens. Therefore, you have decided that this time you will leave out this type of 'outliers' in order to make a better decision for the majority of the citizens. After all, not considering some outliers should simplify the problem, shouldn't it?

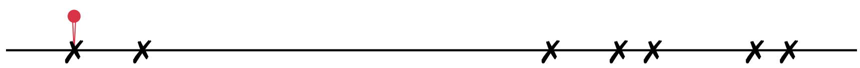

Figure 4.1: Preferred locations of the citizens for the new public facility.

Problems like the one described above fall into the category of *truthful facility location problems*, which have been extensively studied since the work of Moulin [1980]. Furthermore, the seminal paper of Procaccia and Tennenholtz [2013] established facility location problems as *the* paradigm for *mechanism design without money*. Since then, a wide variety of models, settings, and mechanisms have been studied; see, for example, the survey of Chan et al. [2021] for an overview of different models. In this work, we revisit the two foundational models of truthful facility location, the one of Moulin [1980] and the one of Procaccia and Tennenholtz [2013], and study them under the presence of outliers (see [Charikar et al., 2001]). Although outliers have been a popular consideration in algorithm design in general, to the best of our knowledge, they have not been considered in the context of mechanism design. In this paper, we take the first steps towards understanding the impact of outliers on a fundamental mechanism design problem.

4.1.1 Our Contributions

Our main contributions are as follows.

1. We introduce the notion of outliers in mechanism design problems. We envision that outliers can be meaningfully incorporated into *many* mechanism design problems, which opens up a whole new domain of mechanism design with outliers. Studying the impact of outliers is not only theoretically appealing but also practically relevant, especially in applications involving agents with extreme or atypical preferences.

2. We use single facility location on the real line as a first, natural test case for our setting of mechanism design with outliers. We derive tight bounds for deterministic strategyproof (SP) mechanisms for the two most-studied objectives, i.e., utilitarian and egalitarian social cost. We provide a complete picture of the impact of outliers and our results reveal some counter-intuitive phenomena.

3. We further enrich our model by incorporating output predictions, contributing to the recent emerging line of research on learning-augmented mechanism design. We derive a mechanism with an optimal consistency-robustness trade-off for the utilitarian objective and, unlike the problem without outliers, prove an impossibility result for the egalitarian objective.

Next, we elaborate on these three points in more detail.

New Domain: Mechanism Design with Outliers. In a mechanism design problem with outliers, the social cost objective does not aim to optimize over the set of *all* agents, but it only accounts for a certain number of them. More specifically, the input of the problem includes an extra integer parameter z that denotes the number of agents that are excluded from the social cost objective. Put differently, if there are n agents in total, the objective only accounts for the social cost of the "best" set of $n - z$ agents, i.e., those that contribute the least to the social cost objective; formal definitions will be given below. As noted above, this perspective of disregarding z outliers can also be applied to other objective functions.

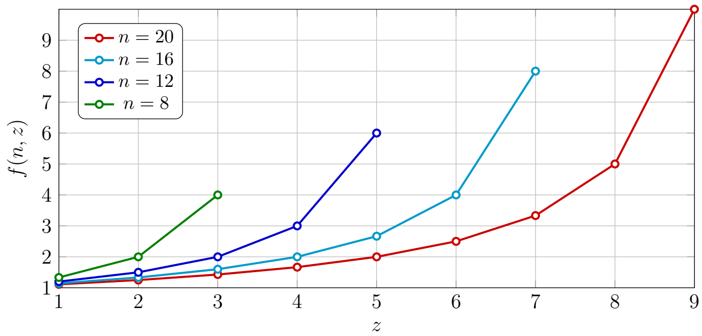

Figure 4.2: Illustration of the best possible approximation guarantee $f(n,z) = \frac{n}{n-2z}$ for the utilitarian objective for different values of an even number of agents n. For $z > \lfloor \frac{n-1}{2} \rfloor$, the problem is inapproximable.

Facility Location with Outliers. In facility location terms, the problem with outliers reduces to the following setting. We have n agents, where each one of them has a private preferred location on the real line. Our goal is to choose a single location on the line for a facility such that a social cost objective, depending on the distances between the facility and the locations of the agents, is minimized, while ensuring that the agents declare their true locations. We consider the two most prominently studied social cost objectives, namely the *utilitarian social cost* (i.e., total cost of the agents) and the *egalitarian social cost* (i.e., maximum cost of the agents). However, in our setting with outliers, the contributions of the z agents with the largest costs are disregarded from the social cost objective; it is not hard to see that these agents will be the ones farthest away from the facility.

Although it may seem intuitively plausible that outliers should only improve the efficiency of a mechanism—since we are free to disregard some of the agents—our results reveal that the opposite is true. For example, the achievable approximation guarantee for utilitarian social cost deteriorates as the number of outliers increases (see Figure 4.2). The crux of the issue is that a strategyproof mechanism *cannot* simply ignore the z outliers chosen by the optimal solution. A partial overview of the bounds we derive in this chapter is given in Table 4.1.

Our first set of technical results fully characterizes the landscape of deterministic SP mechanisms with outliers for both the utilitarian and egalitarian objective. We begin by showing that if we can discard $z \geq n/2$ of the agents, then there is no SP mechanism that achieves a bounded approximation guarantee. This impossibility result holds for both objectives. We then turn our attention to the egalitarian objective and prove that the mechanism choosing the $(z+1)$-*th order statistic*, i.e., the $(z+1)$-th smallest location reported by the agents, as the location of the facility is SP and 2-approximate. We show that this is best possible for every $z < n/2$ by providing a matching lower bound. After that, we focus on the utilitarian objective. To this end, we consider the SP mechanism which always chooses the location of the median agent. Analyzing the approximation

4

Objective	Upper Bound		Lower Bound	
Utilitarian	$\frac{n}{n-2z} \mid \frac{n-1}{n-2z+1}$	(Thm. 4.4.2)	$\frac{n}{n-2z} \mid \frac{n-1}{n-2z+1}$	(Thm. 4.4.4)
Egalitarian	2	(Thm. 4.3.1)	2	(Thm. 4.3.2)

Table 4.1: Partial overview of our bounds. For the utilitarian objective the bounds depend on the parity of n; we use (even n | odd n) to denote these two cases.

is technically more involved, and we derive an upper bound as a function of the number of agents n and outliers z (see Table 4.1). Finally, we also show that this is best possible.

Integrating Predictions. We extend our model by incorporating predictions. In this setting, we assume that the mechanism has access to some (possibly erroneous) prediction of the optimal location of the facility (i.e., with respect to the social cost objective adapted to the outlier setting). This kind of prediction, termed *output advice* in [Christodoulou et al., 2024], has recently become one of the standard benchmarks in mechanism design with predictions. We note that from an information-theoretic perspective, this augmentation is minimal, as the predicted optimal location represents an aggregate of the private locations of $n-z$ agents.

4

Our second set of technical results examines how mechanisms can leverage these predictions in the setting of facility location with outliers to achieve improved approximation guarantees. More precisely, we are interested in designing mechanisms that achieve the best possible trade-off between *consistency* and *robustness*. Surprisingly, we derive strong impossibility results for the egalitarian objective: we prove that no SP mechanism can achieve a bounded robustness and a consistency better than 2 (Theorem 4.5.1). In light of our 2-approximate SP mechanism with outliers mentioned above, this means that this type of prediction does not help at all. However, for the utilitarian objective and for $z \leq n/3$ outliers, we obtain positive results. We derive a SP mechanism, called In-Range, that is 1-consistent and achieves the best possible robustness guarantee (Theorem 4.5.4). Additionally, we show that the approximation of In-Range smoothly interpolates between these two extremes depending on a natural error parameter quantifying the accuracy of the prediction. If the number of outliers is $n/3 < z < n/2$, we show that no SP mechanism can achieve 1-consistency and a bounded robustness (Theorem 4.5.2). We also derive positive results for this range of n and z.

4.1.2 Related Work

Facility location. A big variety of different models have been proposed that studied the number of facilities whose location needs to be determined [Procaccia and Tennenholtz, 2013, Lu et al., 2010, Fotakis and Tzamos, 2014], agents with different types of preferences for the facilities (optional [Chen et al., 2020, Kanellopoulos et al., 2023b, Li et al., 2020, Serafino and Ventre, 2016], fractional [Fong et al., 2018], or hybrid [Feigenbaum and Sethuraman, 2015]) and obnoxious facilities [Cheng et al., 2013]. Furthermore, there exist models with other limitations or features: the facilities can only be built at specific fixed locations [Feldman et al., 2016, Gai et al., 2024, Kanellopoulos et al., 2023a, Xu et al., 2021]; there are limited resources that allow only some of the available facilities to be built [Deligkas et al., 2023]; there is limited information available during the decision process [Chan et al., 2023, Filos-Ratsikas et al., 2024]. One specific model that looks similar to ours is the *capacitated* facility location problem [Aziz et al., 2020b,a, Auricchio et al., 2024a,b]. In this model, there is one facility (or more), and each facility is associated with a capacity, i.e., an upper bound on the number of agents it can serve. There, part of the mechanism design problem is to *choose* which agents will be served by a facility. In our case, although we have a "capacity" on the number of agents we consider in the social cost objective, we do not exclude any agent from being served by the facility.

Mechanism Design with Predictions. We refer to Chapter 3 for a more comprehensive overview of related work on learning-augmented algorithms. This line of work aims to overcome worst-case guarantees by leveraging predictions. As previously mentioned, Xu and Lu [2022] and Agrawal et al. [2024] simultaneously proposed the study of strategic environments augmented with predictions about private information or aggregations thereof. In fact, strategic facility location can be viewed as the paradigmatic problem for this setting as it is studied in both papers, with Agrawal et al. [2024] deriving deterministic, SP mechanisms that achieve the best-possible consistency-robustness trade-offs in two dimensions for both the utilitarian and egalitarian objectives. Follow-up work on strategic facility location in this environment includes [Christodoulou et al., 2024], in which the universal error notion that we adopt is proposed, and [Barak et al., 2024, Balkanski et al., 2024a, Chen et al., 2024].

Clustering with Outliers. The problem of optimization with outliers was introduced by Charikar et al. [2001], who studied several well-known clustering problems —k-median, k-means, facility location— in high dimensions with outliers and provided polynomial-time algorithms that achieve constant approximation guarantees. Building upon this, there was a long line of research that studied the abovementioned problems improving the approximation guarantees using a variety of techniques, ranging from local search to iterative rounding and parameterized algorithms [Almanza et al., 2022, Chen, 2008, Friggstad et al., 2019, Dabas and Gupta, 2021, Krishnaswamy et al., 2018, Zhang et al., 2021, Agrawal et al., 2023, Maity et al., 2024, Feng et al., 2019, Goyal et al., 2020, Gupta et al., 2017].

4.2 Preliminaries

We start by introducing the non-strategic problem of single facility location on the real line with outliers, which follows the model in [Charikar et al., 2001].

Facility Location with Outliers. In the standard single facility location problem on the real line, we must determine the location $y \in \mathbb{R}$ of a facility that serves a set $N = [n]$ of n agents. Each agent $i \in N$ has a location $x_i \in \mathbb{R}$ and if y is chosen as the location of the facility, agent i incurs a cost that is equal to the distance between i's location and the location of the facility, i.e., $|y - x_i|$. Two well-studied objectives are minimizing the *utilitarian* social cost, which is the total cost of all agents, i.e., $\sum_{i \in N} |y - x_i|$, and minimizing the *egalitarian* social cost, which is the maximum cost of an agent, i.e., $\max_{i \in N} |y - x_i|$. However, scenarios exist for which it is desirable to only consider a subset of the agents in the objective, and we refer to this problem as the problem with *outliers*. In this case, for any $n \geq 2$, we are additionally given an integer parameter[1] $z \in \{1, 2, \dots, n-1\}$ representing the number of outliers. The goal is still to find a location y that minimizes the social cost objective, but only for a set of $n - z$ agents which can be chosen freely. Let the profile $\boldsymbol{x} = (x_i)_{i \in N} \in \mathbb{R}^n$ denote the vector of locations of the n agents. If the location y is chosen, then the set of outliers can be computed by solving the following optimization problems for the utilitarian and egalitarian objective, respectively:

4

$$\mathrm{SC}(y, \boldsymbol{x}, z) := \min_{\substack{S \subset N: \\ |S| = n - z}} \sum_{i \in S} |y - x_i| \quad \text{and} \quad \mathrm{MC}(y, \boldsymbol{x}, z) := \min_{\substack{S \subset N: \\ |S| = n - z}} \max_{i \in S} |y - x_i|.$$

For both objectives, we use $S^\star(\boldsymbol{x}, z)$ to denote an optimal set of non-outliers and $y^\star(\boldsymbol{x}, z)$ to denote an optimal location, i.e., for the utilitarian objective $y^\star(\boldsymbol{x}, z) \in \arg\min_{y \in \mathbb{R}} \mathrm{SC}(y, \boldsymbol{x}, z)$.[2] For brevity, we sometimes use $y^\star$ and $S^\star$ and omit explicit reference to $\boldsymbol{x}$ and z when clear from context. Moreover, we use $\mathrm{SC}^\star(\boldsymbol{x}, z) = \min_{y \in \mathbb{R}} \mathrm{SC}(y, \boldsymbol{x}, z)$ and $\mathrm{MC}^\star(\boldsymbol{x}, z) = \min_{y \in \mathbb{R}} \mathrm{MC}(y, \boldsymbol{x}, z)$ to denote the optimal value of the objectives, given a profile $\boldsymbol{x}$ and z outliers. Given a profile $\boldsymbol{x}$, we will use σ to refer to the indices in increasing order of value, i.e., $x_{\sigma(1)} \leq x_{\sigma(2)} \leq \dots \leq x_{\sigma(n)}$.

Strategic Agents. In the strategic version of the problem, the preferred location $p_i \in \mathbb{R}$ of an agent $i \in N$ is *private information* and only known to agent i. Therefore, each agent $i \in N$ declares a location $x_i \in \mathbb{R}$. Given a profile of declared locations $\boldsymbol{x} \in \mathbb{R}^n$ and the number of outliers $z \in \{1, 2, \dots, n-1\}$,

[1] Note that if $z = n$, both objectives are equal to 0, and if $z = 0$, we retrieve the problem without outliers.

[2] Note that an optimal solution is not necessarily unique; if this is the case, we will clarify which optimal solution we refer to.

we seek a *mechanism* $\mathcal{M}$ that chooses the location $y \in \mathbb{R}$ of the facility, i.e., $\mathcal{M} : (\mathbb{R}^n, \mathbb{N}_{>0}) \to \mathbb{R}$. If y is chosen as the location of the facility, agent i incurs a cost of $|y - p_i|$. As the agents are strategic, they will misreport their private preferred location if this reduces their incurred cost. We therefore seek *strategyproof* mechanisms in which no agent can reduce their incurred cost by misreporting, see also Definition 2.2.2 with cost functions. Note that we do not require a mechanism to determine which agents are the outliers and all n agents can still make use of the facility. Namely, given a profile $\boldsymbol{x}$, the number of outliers z and the location y chosen by the mechanism, the set of non-outliers minimizing the social cost objective can easily be computed. Moulin [1980] established the following characterization result for deterministic strategyproof mechanisms which will be useful in our setting as well.

Theorem 4.2.1: ([Moulin, 1980]). *A deterministic mechanism $\mathcal{M}$ for the facility location problem on the line is strategyproof if and only if there exist $n+1$ real numbers $\alpha_1, \dots, \alpha_{n+1} \in \mathbb{R} \cup \{-\infty, +\infty\}$, called* phantom points*, such that:*

$$\forall \boldsymbol{x} = (x_1, \dots, x_n) \in \mathbb{R}^n : \qquad \mathcal{M}(\boldsymbol{x}) = \textit{MEDIAN}(x_1, \dots, x_n, \alpha_1, \dots, \alpha_{n+1}). \quad (4.1)$$

Note that the above class of mechanisms in particular also contains the class of *k-order statistic mechanisms* for any $k \in [n]$: by defining $n + 1 - k$ many phantom points to be $-\infty$ and the others to be $+\infty$, $\text{MEDIAN}(x_1, \dots, x_n, \alpha_1, \dots, \alpha_{n+1})$ always corresponds to the k-th smallest element of $\boldsymbol{x}$. Some proofs will use this observation. Furthermore, the original characterization result by Moulin [1980] is proven for agents with single-peaked preferences. Our setting of Euclidean distances on the real line is a special case of this. However, in terms of characterization result, it is known that Theorem 4.2.1 remains valid for this case as well (see [Border and Jordan, 1983] and the discussion in [Peters et al., 1993]).

The following corollary also follows from Theorem 4.2.1 and will be useful when proving impossibility results for both objectives.

Corollary 4.2.2. *Consider a profile $\boldsymbol{x}$ and z outliers and let $\mathcal{M}$ be a strategyproof mechanism with $\mathcal{M}(\boldsymbol{x}, z) = y$. Then, for any $i \in N$ with $x_i < y$ or $y < x_i$, it holds that $\mathcal{M}((x_i', \boldsymbol{x}_{-i}), z) = y$ for all $x_i' \in [x_i, y]$ and $x_i' \in [y, x_i]$, respectively.*

Given the utilitarian objective, we say that a mechanism $\mathcal{M}$ is ρ-approximate, with $\rho \geq 1$, if for any input $(\boldsymbol{x}, z)$ it holds that $\text{SC}(\mathcal{M}(\boldsymbol{x}, z), \boldsymbol{x}, z) \leq \rho \cdot \text{SC}^\star(\boldsymbol{x}, z)$. The definition for the egalitarian objective is defined analogously. For notational convenience, we use $\text{SC}(\mathcal{M}(\boldsymbol{x}, z))$ and $\text{MC}(\mathcal{M}(\boldsymbol{x}, z))$ instead of $\text{SC}(\mathcal{M}(\boldsymbol{x}, z), \boldsymbol{x}, z)$ and $\text{MC}(\mathcal{M}(\boldsymbol{x}, z), \boldsymbol{x}, z)$, respectively, in the remainder of this chapter.

Impossibility Result. Note that if $z = n - 1$, both objectives minimize the cost of a single agent, which can trivially be solved optimally by any mechanism deterministically choosing a k-th order statistic. However, if the number of outliers is at least half the number of agents, we cannot hope to achieve any bounded approximation guarantee, as we show in the lemma below.

Theorem 4.2.3. *Consider the utilitarian or egalitarian social cost objective and let $n \geq 4$ and $z \in \{\lceil \frac{n}{2} \rceil, \lceil \frac{n}{2} \rceil + 1, \ldots, n-2\}$. Then, there is no deterministic strategyproof mechanism that achieves a bounded approximation guarantee.*

Proof: The main idea behind this negative result is that for a profile with an optimal solution with a social cost of zero, any mechanism that achieves a bounded approximation guarantee must also choose a location with a social cost of zero. We leverage this in our proof by sequentially constructing profiles, each time only changing the location of one agent, and due to strategyproofness, the mechanism *cannot* move the facility.

Let $n \geq 4$ and $z \geq \lceil \frac{n}{2} \rceil$ and towards a contradiction, assume that there exists a strategyproof mechanism $\mathcal{M}$ that achieves a bounded approximation guarantee. Consider the profiles depicted in Figure 4.3, in which the numbers below the locations (dots) indicate the number of agents with that location.

First, consider the profile in Figure 4.3a. In this case, there are two possible optimal locations $y^\star$ with a social cost of 0, namely the leftmost and rightmost location. To see this, note that there is at most $n - 2\lfloor \frac{n}{2} \rfloor \leq 1$ agent located in the middle, so this agent together with the leftmost (or the rightmost) cluster of agents can be disregarded in the objective as $z \geq \lceil \frac{n}{2} \rceil$. Therefore, $\mathcal{M}$ will have to choose one of the optimal locations to achieve a bounded approximation. Assume w.l.o.g. that $\mathcal{M}$ chooses the rightmost location as depicted in Figure 4.3a.

Secondly, consider the profile in Figure 4.3b with $\delta = 0$. In this case, the leftmost location is the only optimal location $y^\star$ with a social cost of 0, as this is the only location equal to at least the location of $n - z$ agents. Therefore, $\mathcal{M}$ will have to choose this location in order to achieve a bounded approximation guarantee, as depicted in Figure 4.3b.

Now consider the profile in Figure 4.3b with $\delta = 1$, where one agent moved their location from the rightmost cluster to the second rightmost cluster. Note that in this case, the location of the second rightmost cluster is also an optimal location. However, as $\mathcal{M}$ is strategyproof, the location y must remain unchanged by Corollary 4.2.2. The same argument holds when considering $\delta = 1, 2, \ldots, z + \lfloor \frac{n}{2} \rfloor - n$ consecutively.

Finally, consider the last agent with a location in the rightmost cluster in Figure 4.3b with $\delta = z + \lfloor \frac{n}{2} \rfloor - n$, who moved their location to the second rightmost cluster, resulting in the profile depicted in Figure 4.3a. But for the profile depicted in Figure 4.3a, $\mathcal{M}$ chooses the location of the rightmost cluster as the location of the facility. This contradicts that $\mathcal{M}$ is strategyproof (Corollary 4.2.2), as the location of the facility should remain unchanged after this deviation. This concludes the proof. □

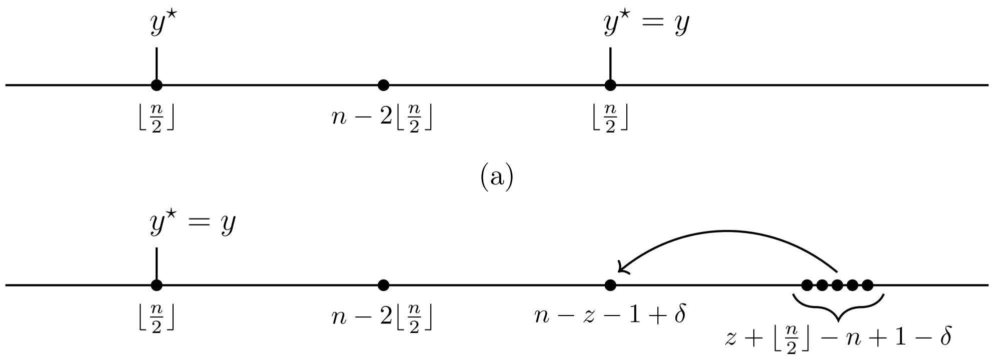

(b) $\delta \in \{0, 1, \dots, z + \lfloor \frac{n}{2} \rfloor - n\}$

Figure 4.3: Profiles used in the proof of Theorem 4.2.3. The numbers below the locations (dots) indicate the number of agents with this location.

In light of this impossibility result and optimality for $z = n - 1$, we only consider profiles with $n \geq 3$ agents and z outliers such that $1 \leq z \leq \lfloor \frac{n-1}{2} \rfloor$ in the remainder of this chapter.

Augmenting Mechanisms with Predictions. In the setting with predictions, besides $\boldsymbol{x}$ and z, a mechanism is additionally given a predicted location $\hat{y}$ of the optimal facility as input. We use $(\boldsymbol{x}, z, \hat{y})$ to refer to input augmented with a prediction. Given input $(\boldsymbol{x}, z, \hat{y})$, we say that $\hat{y}$ is a *perfect prediction* if $\hat{y}$ corresponds to a location of an optimal facility, i.e., for the utilitarian objective this holds if $\mathrm{SC}(\hat{y}, \boldsymbol{x}, z) = \mathrm{SC}^\star(\boldsymbol{x}, z)$. Depending on the quality of the prediction, we consider the approximation notions as in Definition 2.2.8 for cost functions.

To define an approximation notion beyond the two extremes of a perfect prediction and any arbitrary prediction, we define an error that measures the quality of the prediction (see, e.g., [Christodoulou et al., 2024, Colini-Baldeschi et al., 2024, Gkatzelis et al., 2025]), as follows:[3]

$$\eta(\boldsymbol{x}, z, \hat{y}) = \frac{\mathrm{SC}(\hat{y}, \boldsymbol{x}, z)}{\mathrm{SC}^\star(\boldsymbol{x}, z)}. \tag{4.2}$$

Given this definition, input $(\boldsymbol{x}, z, \hat{y})$ with a perfect prediction has a prediction error of 1. As the quality of the prediction deteriorates, the error measure increases, possibly to ∞. Our goal is to construct a mechanism that achieves an approximation guarantee that smoothly interpolates between the consistency and robustness guarantees as a function of the error measure.

[3] If $\mathrm{SC}^\star(\boldsymbol{x}, z) = 0$, we define $\eta(\boldsymbol{x}, z, \hat{y}) = 1$ if $\mathrm{SC}(\hat{y}, \boldsymbol{x}, z) = 0$ and $\eta(\boldsymbol{x}, z, \hat{y}) = \infty$ otherwise.

4.3 Egalitarian Objective with Outliers

In this section, we focus on the egalitarian objective and prove a tight bound of 2 for deterministic strategyproof mechanisms.

An optimal solution to the non-strategic problem for a profile $\boldsymbol{x}$ and z outliers disregards the locations x_i in the objective that are among the smallest and largest values. This could be z of the smallest values, i.e., all locations x_i with $i = \sigma(k)$ and $k < z + 1$, or z of the largest values, i.e., all locations x_i with $i = \sigma(k)$ and $k > n - z$, or any combination in between. Therefore, in order for a mechanism $\mathcal{M}(\boldsymbol{x}, z) = y$ to achieve a bounded approximation guarantee, $\mathcal{M}$ must choose a location y such that $x_{\sigma(z+1)} \leq y \leq x_{\sigma(n-z)}$. For example, if $\mathcal{M}$ would output $y < x_{\sigma(z+1)}$, it could be that $x_{\sigma(z)} < x_{\sigma(z+1)} = x_{\sigma(z+2)} = \ldots = x_{\sigma(n)} = y^\star$, leading to an egalitarian social cost of 0 by disregarding the z leftmost locations. But as $z \leq \lfloor \frac{n-1}{2} \rfloor$, the location y chosen by $\mathcal{M}$ has a positive egalitarian social cost as there are $n - z$ locations with a positive distance to y, i.e., $x_{\sigma(k)} - y > 0$ for $k \in \{z+1, z+2, \ldots, n\}$, leading to an unbounded approximation guarantee.

Therefore, consider the mechanism Left-z that chooses the $(z+1)$-th order statistic as the location of the facility:[4]

4

$$\text{Left-z}(\boldsymbol{x}, z) = x_{\sigma(z+1)}.$$

Theorem 4.3.1. *Let $z \leq \lfloor \frac{n-1}{2} \rfloor$. Then, Left-z is strategyproof and 2-approximate for the egalitarian objective.*

Proof: Strategyproofness of Left-z follows from Theorem 4.2.1. For the approximation guarantee, consider an arbitrary profile $(\boldsymbol{x}, z)$ and let $y = \text{Left-z}(\boldsymbol{x}, z)$. Fix an optimal solution $y^\star := y^\star(\boldsymbol{x}, z)$ and let $S^\star := S^\star(\boldsymbol{x}, z)$ be the optimal set of non-outliers with respect to $y^\star$. We first prove that $\max_{i \in S^\star} |y - x_i| \leq 2 \cdot \text{MC}^\star(\boldsymbol{x}, z)$. We consider two cases, depending on the relative values of y and $y^\star$.

Case 1: $y^\star \leq y$. Note that $\text{MC}^\star(\boldsymbol{x}, z) \geq y - y^\star$, as there are at least z agents with a location of at least $y = x_{\sigma(z+1)}$, i.e., $y - y^\star \leq x_i - y^\star$ for $i \in \{\sigma(z+2), \sigma(z+3), \ldots, \sigma(n)\}$. For an illustrative example, see Figure 4.4. For locations $x_i \leq y^\star$ with $i \in S^\star$, it holds that $y - x_i = y^\star - x_i + y - y^\star \leq 2 \cdot \text{MC}^\star(\boldsymbol{x}, z)$. For locations $x_i \geq y$ with $i \in S^\star$, it holds that $x_i - y \leq x_i - y^\star \leq \text{MC}^\star(\boldsymbol{x}, z)$. Finally, for locations $y^\star < x_i < y$ with $i \in S^\star$, it holds that $y - x_i \leq y - y^\star \leq \text{MC}^\star(\boldsymbol{x}, z)$.

Case 2: $y^\star > y$. This case is symmetrical. Again note that $\text{MC}^\star(\boldsymbol{x}, z) \geq y^\star - y$, as there are at least z agents with a location of at most $y = x_{\sigma(z+1)}$, i.e., $y^\star - y \leq y^\star - x_i$ for $i \in \{\sigma(1), \sigma(2), \ldots, \sigma(z)\}$. For locations $x_i \geq y^\star$ with $i \in S^\star$, it holds that $x_i - y = x_i - y^\star + y^\star - y \leq 2 \cdot \text{MC}^\star(\boldsymbol{x}, z)$. For locations $x_i \leq y$ with $i \in S^\star$, it holds that $y - x_i \leq y^\star - x_i \leq \text{MC}^\star(\boldsymbol{x}, z)$. Finally, for locations $y < x_i < y^\star$ with $i \in S^\star$, it holds that $x_i - y \leq y^\star - y \leq \text{MC}^\star(\boldsymbol{x}, z)$.

[4] In fact, any choice of the k-th order statistic with $z + 1 \leq k \leq n - z$ would work.

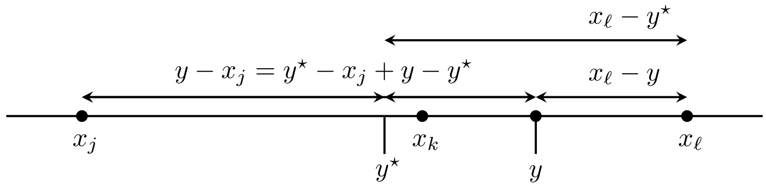

Figure 4.4: Example of Case 1 in the proof of Theorem 4.3.1.

We can now complete the proof by observing that:

$$\mathrm{MC}(y, \boldsymbol{x}, z) = \min_{\substack{S \subset N: \\ |S| = n-z}} \max_{i \in S} |y - x_i| \leq \max_{i \in S^\star} |y - x_i| \leq 2 \cdot \mathrm{MC}^\star(\boldsymbol{x}, z),$$

where the first inequality follows by the fact that $|S^\star| = n - z$, i.e., $S^\star$ is a feasible solution. □

We conclude this section by proving that, for the egalitarian objective, Left-z achieves the best-possible approximation guarantee among all deterministic strategyproof mechanisms.

Theorem 4.3.2. *Let $z \leq \lfloor \frac{n-1}{2} \rfloor$. Then, there is no deterministic strategyproof mechanism that is better than 2-approximate for the egalitarian objective.*

Proof: The proof is inspired by the lower bound of Procaccia and Tennenholtz [2013]. Let $z \leq \lfloor \frac{n-1}{2} \rfloor$ and $\varepsilon > 0$. Towards a contradiction, assume that there exists a deterministic strategyproof mechanism $\mathcal{M}$ that is $(2 - \varepsilon)$-approximate. We will again consider a sequence of profiles, which are depicted in Figure 4.5.

The profile in Figure 4.5a is such that there are z agents with a location of 0 (leftmost cluster), $n - 2z$ agents with a location of $\frac{1}{2}$, and z agents with a location of 1 (rightmost cluster). Note that $n - 2z \geq 1$ as $z \leq \lfloor \frac{n-1}{2} \rfloor$. Also note that all agents from either the leftmost or rightmost cluster can be disregarded in the objective. Therefore, there are two optimal locations $y^\star$, namely at $\frac{1}{4}$ and $\frac{3}{4}$, with an egalitarian social cost of $\frac{1}{4}$. As $\mathcal{M}$ is $(2 - \varepsilon)$-approximate, it must be that y is located in $[\frac{\varepsilon}{4}, \frac{1}{2} - \frac{\varepsilon}{4}]$ or in $[\frac{1}{2} + \frac{\varepsilon}{4}, 1 - \frac{\varepsilon}{4}]$, depicted by the curly brackets in Figure 4.5a. Assume w.l.o.g. that $\mathcal{M}$ chooses a location y in $[\frac{\varepsilon}{4}, \frac{1}{2} - \frac{\varepsilon}{4}]$.

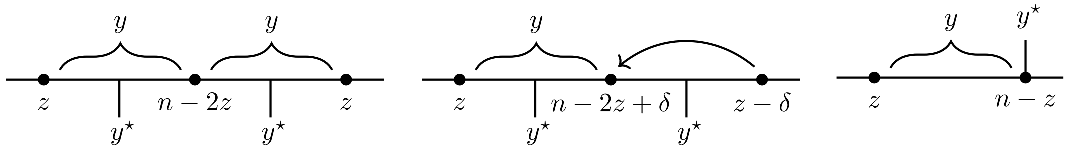

(a) (b) $\delta \in \{1, 2, \ldots, z - 1\}$ (c)

Figure 4.5: Profiles used in the proof of Theorem 4.3.2.

Now, consider the profile in Figure 4.5b with $\delta = 1$: the location of one agent from the rightmost cluster moved from 1 to $\frac{1}{2}$. Note that as $\mathcal{M}$ is strategyproof, y remains unchanged by Corollary 4.2.2. The same reasoning holds when considering $\delta = 2, 3, \ldots, z-1$ consecutively.

Finally, consider the profile in Figure 4.5c in which all the locations of the z agents from the right cluster moved to $\frac{1}{2}$. Note that y remains unchanged as $\mathcal{M}$ is strategyproof (Corolarry 4.2.2). However, as the left cluster now contains the locations of z agents, all agents from this cluster can be disregarded in the objective. Therefore, the only optimal location $y^\star$ is the location of the right cluster at $\frac{1}{2}$ with an egalitarian social cost of 0, as depicted in Figure 4.5c. This contradicts that $\mathcal{M}$ is $(2-\varepsilon)$-approximate as $y < \frac{1}{2}$, so y has a positive social cost, concluding the proof. □

4.4 Utilitarian Objective with Outliers

An optimal solution to the non-strategic problem for the utilitarian objective has a similar structure as the one described for the egalitarian objective: if there are z outliers, an optimal solution disregards z_ℓ and $z_r = z - z_\ell$ of the leftmost (smallest) and rightmost (largest) locations of $\boldsymbol{x}$, respectively, with $z_\ell \in \{0, 1, \ldots, z\}$. We can therefore identify a set of locations that contains at least one optimal solution for any number of z outliers with $1 \leq z \leq n-2$, which will turn out to be useful when deriving the approximation guarantees of our mechanisms.

Lemma 4.4.1. *Consider an arbitrary profile $\boldsymbol{x}$ and z outliers with $1 \leq z \leq n-2$. Consider the set of locations:*

$$O := \left\{ x_{\sigma(\lceil \frac{n-z+1}{2} \rceil)}, x_{\sigma(\lceil \frac{n-z}{2} \rceil + 1)}, \ldots, x_{\sigma(\lceil \frac{n-z}{2} \rceil) + z} \right\}. \tag{4.3}$$

Then, there exists a location $y^\star := y^\star(\boldsymbol{x}, z) \in O$ which minimizes the utilitarian objective, and $|O| = z$ if $n - z$ is even and $|O| = z + 1$ otherwise. Furthermore, assuming the optimal set of non-outliers $S^\star := S^\star(\boldsymbol{x}, z)$ is with respect to $y^\star$, the following statements are true.

1. *For every $i \in N \setminus S^\star$ with $x_i \leq y^\star$ and every j with $x_j < x_i$, it holds that $j \in N \setminus S^\star$.*
2. *For every $i \in N \setminus S^\star$ with $x_i \geq y^\star$ and every j with $x_j > x_i$, it holds that $j \in N \setminus S^\star$.*

Proof: For ease of notation, suppose that $x_1 \leq x_2 \leq \ldots \leq x_n$ and for brevity, let $y^\star := y^\star(\boldsymbol{x}, z)$ and let $S^\star := S^\star(\boldsymbol{x}, z)$ be the optimal set of non-outliers with respect to $y^\star$. We first show that each optimal solution disregards some of the leftmost and rightmost locations of $\boldsymbol{x}$ in the objective.

Consider an agent $i \in N \setminus S^\star$ and assume w.l.o.g. that $x_i \leq y^\star$. Towards a contradiction, suppose that there exists an agent $j \in S^\star$ with $x_j < x_i$. This contradicts optimality of $(y^\star, S^\star)$, as $(y^\star, S')$ with $S' = S^\star \cup \{i\} \setminus \{j\}$ has a lower utilitarian social cost. Namely, for $k \in S^\star \cap S'$ the contribution $|x_k - y^\star|$ to the social cost remains unchanged. But for $j \in S^\star$ and $i \in S'$, it holds that $|x_j - y^\star| > |x_i - y^\star|$. Therefore, given an optimal solution $(y^\star, S^\star)$, there exists a $z_\ell \in \{0, 1, \ldots, z\}$ such that $i \notin S^\star$ for $i \leq z_\ell$ and $i > n - (z - z_\ell)$. Then, the optimal location $y^\star$ is between the $\lceil \frac{n-z}{2} \rceil$-th and $\lceil \frac{n-z+1}{2} \rceil$-th order statistic of the locations corresponding to the $n - z$ agents in $S^\star$ [Procaccia and Tennenholtz, 2013].

Specifically, for $n - z$ odd and $z_\ell = 0$ this leads to $y^\star = x_{\lceil \frac{n-z}{2} \rceil} = x_{\lceil \frac{n-z+1}{2} \rceil}$ and as z_ℓ increases by 1, the order statistic of the location equal to $y^\star$ increases by 1 until $z_\ell = z$ and $y^\star = x_{\lceil \frac{n-z}{2} \rceil + z}$. And so, $|O| = z+1$ in this case. For $n-z$ even and $z_\ell = 0$, this leads to $x_{\lceil \frac{n-z}{2} \rceil} \leq y^\star \leq x_{\lceil \frac{n-z+1}{2} \rceil}$ and as z_ℓ increases by 1, so do both order statistics of this interval until $z_\ell = z$ and $x_{\lceil \frac{n-z}{2} \rceil + z} \leq y^\star \leq x_{\lceil \frac{n-z+1}{2} \rceil + z}$. Note that by definition, for $z_\ell = 0$ only $x_{\lceil \frac{n-z+1}{2} \rceil} \in O$ and for $z_\ell = z$ only $x_{\lceil \frac{n-z}{2} \rceil + z} \in O$, leading to $|O| = z$ in this case. □

As for the problem without outliers, let us consider the mechanism LEFT-MEDIAN that chooses the left median as the location of the facility:

$$\text{LEFT-MEDIAN}(\boldsymbol{x}) = x_{\sigma(\lceil \frac{n}{2} \rceil)}.$$

Note that if $z = 1$ and n is odd, it follows from Lemma 4.4.1 that there is always an optimal location $y^\star$ in O with $O = \left\{x_{\sigma(\lceil \frac{n-1+1}{2} \rceil)}, x_{\sigma(\lceil \frac{n-1}{2} \rceil + 1)}\right\} = \left\{x_{\sigma(\lceil \frac{n}{2} \rceil)}\right\}$. Therefore, for this specific case, LEFT-MEDIAN is 1-approximate, i.e., optimal. For all other combinations of n and z LEFT-MEDIAN is no longer 1-approximate and as it turns out, for a fixed number of agents n the approximation guarantee increases to $z + 1$ if n is even, and z if n is odd, as z increases to $\lfloor \frac{n-1}{2} \rfloor$.

Theorem 4.4.2. *Let $z \in \{1, 2, \ldots, \lfloor \frac{n-1}{2} \rfloor\}$. Then, LEFT-MEDIAN is strategyproof and has an approximation guarantee for the utilitarian objective of:*

$$f(n, z) = \begin{cases} \frac{n-1}{n-2z+1}, & \text{if } n \text{ odd,} \\ \frac{n}{n-2z}, & \text{otherwise.} \end{cases} \tag{4.4}$$

We will use the following fact in order to prove the approximation guarantee of Theorem 4.4.2, as it will turn out to be useful to simplify and upper bound the approximation in our analysis.

Fact 4.4.3. Let $x \geq y > 0$ and $z \in [0, y)$. Then $\frac{x}{y} \leq \frac{x-z}{y-z}$.

Proof of Theorem 4.4.2: Strategyproofness follows from Theorem 4.2.1. For the approximation guarantee, consider a profile $(\boldsymbol{x}, z)$ and let $y = \text{LEFT-MEDIAN}(\boldsymbol{x})$.

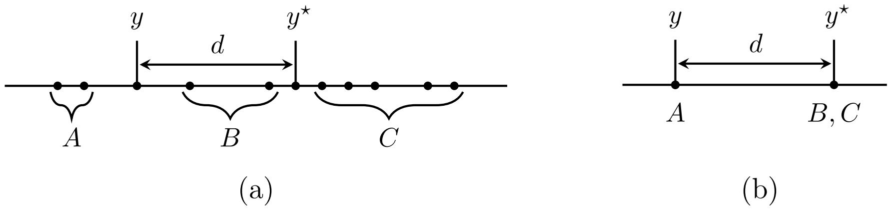

Figure 4.6: Example showing how the approximation guarantee is upper bounded in the proof of Theorem 4.4.2, where (b) illustrates how the locations of agents $i \in S^\star$ have moved w.r.t. (a).

As for any combination of n and z it holds that $f(n,z) \geq 1$, the approximation guarantee follows if there exists an optimal solution with location y. Therefore, assume that there is no optimal solution with location y and assume w.l.o.g. that $y < y^\star := y^\star(\boldsymbol{x}, z)$ and let $d = |y^\star - y|$. Let $S^\star := S^\star(\boldsymbol{x}, z)$ be the set of non-outliers for $y^\star$.

We will upper bound the approximation guarantee by evaluating the utilitarian objective of y with respect to $S^\star$. Note that:

$$\mathrm{SC}(y, \boldsymbol{x}, z) = \min_{\substack{S \subseteq N:\\ |S| = n-z}} \sum_{i \in S} |y - x_i| \leq \sum_{i \in S^\star} |y - x_i|,$$

since $|S^\star| = n - z$. We have:

$$\begin{aligned} \frac{\mathrm{SC}(y, \boldsymbol{x}, z)}{\mathrm{SC}^\star(\boldsymbol{x}, z)} &\leq \frac{\sum_{i \in S^\star} |y - x_i|}{\sum_{i \in S^\star} |y^\star - x_i|} \\ &= \frac{\sum_{\substack{i \in S^\star:\\ x_i < y}} |y - x_i| + \sum_{\substack{i \in S^\star:\\ y < x_i \leq y^\star}} |y - x_i| + \sum_{\substack{i \in S^\star:\\ x_i > y^\star}} |y - x_i|}{\sum_{\substack{i \in S^\star:\\ x_i \leq y}} |y^\star - x_i| + \sum_{\substack{i \in S^\star:\\ y < x_i < y^\star}} |y^\star - x_i| + \sum_{\substack{i \in S^\star:\\ x_i > y^\star}} |y^\star - x_i|} \\ &\leq \frac{\sum_{\substack{i \in S^\star:\\ y < x_i}} |y^\star - y|}{\sum_{\substack{i \in S^\star:\\ x_i \leq y}} |y^\star - y|} = \frac{\sum_{\substack{i \in S^\star:\\ y < x_i}} d}{\sum_{\substack{i \in S^\star:\\ x_i \leq y}} d} = \frac{|\{i \in S^\star : y < x_i\}|}{|\{i \in S^\star : x_i \leq y\}|}. \end{aligned} \tag{4.5}$$

Here, the first equality follows by partitioning the agents $i \in S^\star$ depending on their location x_i with respect to y and $y^\star$. The second inequality follows by Fact 4.4.3 by (a) moving agents $i \in S^\star$ with $x_i < y$ (A in Figure 4.6a) to y (A in Figure 4.6b), (b) moving agents $i \in S^\star$ with $y < x_i \leq y^\star$ (B in Figure 4.6a) to $y^\star$ (B in Figure 4.6b), and (c) moving agents $i \in S^\star$ with $x_i > y^\star$ (C in Figure 4.6a) to $y^\star$ (C in Figure 4.6b). The last two equalities follow by replacing $|y^\star - y|$ with d and multiplying both the numerator and denominator by $1/d$.

In order to further upper bound the ratio in (4.5), we want to consider the maximum number of agents $i \in S^\star$ with $y < x_i$. Using Lemma 4.4.1, this number is maximized if $y^\star = x_{\sigma(\lceil \frac{n-z}{2} \rceil + z)}$ and $y = x_{\sigma(\lceil \frac{n}{2} \rceil)} < x_{\sigma(\lceil \frac{n}{2} \rceil + 1)}$; note that as $z \leq \lfloor \frac{n-1}{2} \rfloor$, it holds that $\sigma(\lceil \frac{n}{2} \rceil) \in S^\star$ by Lemma 4.4.1. We define π as the number of agents $i \in S^\star$ with $\sigma(\lceil \frac{n}{2} \rceil) < \sigma(i) < \sigma(\lceil \frac{n-z}{2} \rceil + z)$. The exact value of π depends on the parity of z and n. Let $i^\star = \sigma(\lceil \frac{n-z}{2} \rceil + z)$ and consider the following cases.

Case 1: n is odd and z is even. In this case, $\pi = \frac{z}{2} - 1 = \frac{z-2}{2}$. Furthermore, there are $(n-z-1)/2$ agents $i \in S^\star$ with $\sigma(i) < i^\star$ and $(n-z-1)/2$ agents $i \in S^\star$ with $\sigma(i) > i^\star$. Combining this with (4.5) leads to an upper bound of:

$$\frac{\frac{n-z-1}{2} + 1 + \frac{z-2}{2}}{\frac{n-z-1}{2} - \frac{z-2}{2}} = \frac{n-1}{n-2z+1}.$$

Case 2: n is odd and z is odd. Assume that $z \geq 3$, as for $z = 1$ and n odd LEFT-MEDIAN is 1-approximate. We have that $\pi = \frac{z-1}{2} - 1 = \frac{z-3}{2}$. Furthermore, there are $((n-z)/2) - 1$ agents $i \in S^\star$ with $\sigma(i) < i^\star$ and $(n-z)/2$ agents $i \in S^\star$ with $\sigma(i) > i^\star$. Combining this with (4.5) leads to an upper bound of:

$$\frac{\frac{n-z}{2} + 1 + \frac{z-3}{2}}{\frac{n-z}{2} - 1 - \frac{z-3}{2}} = \frac{n-1}{n-2z+1}.$$

Case 3: n is even and z is even. In this case, $\pi = \frac{z}{2} - 1 = \frac{z-2}{2}$. Furthermore, there are $((n-z)/2) - 1$ agents $i \in S^\star$ with $\sigma(i) < i^\star$ and $(n-z)/2$ agents $i \in S^\star$ with $\sigma(i) > i^\star$. Combining this with (4.5) leads to an upper bound of:

$$\frac{\frac{n-z}{2} + 1 + \frac{z-2}{2}}{\frac{n-z}{2} - 1 - \frac{z-2}{2}} = \frac{n}{n-2z}.$$

Case 4: n is even and z is odd. In this case, $\pi = \frac{z+1}{2} - 1 = \frac{z-1}{2}$. Furthermore, there are $(n-z-1)/2$ agents $i \in S^\star$ with $\sigma(i) < i^\star$ and $(n-z-1)/2$ agents $i \in S^\star$ with $\sigma(i) > i^\star$. Combining this with (4.5) leads to an upper bound of:

$$\frac{\frac{n-z-1}{2} + 1 + \frac{z-1}{2}}{\frac{n-z-1}{2} - \frac{z-1}{2}} = \frac{n}{n-2z}.$$

This concludes the proof.

□

We conclude this section by deriving a lower bound for deterministic strategyproof mechanisms, which matches the upper bound in Theorem 4.4.2 for all values of n and z. In the proof, we utilize the characterization of Moulin [1980].

Theorem 4.4.4. *Let $z \in \{1, 2, \dots, \lfloor \frac{n-1}{2} \rfloor\}$. Then, there is no deterministic strategyproof mechanism with an approximation guarantee for the utilitarian objective better than:*

$$f(n,z) = \begin{cases} \frac{n-1}{n-2z+1}, & \text{if } n \text{ odd,} \\ \frac{n}{n-2z}, & \text{otherwise.} \end{cases} \tag{4.6}$$

4

(a) Profile $\boldsymbol{x}^1$ (b) Profile $\boldsymbol{x}^2$ (c) Profile $\boldsymbol{x}^3$ (d) Profile $\boldsymbol{x}^4$

Figure 4.7: Profiles used in the proof of Theorem 4.4.4.

Proof: Recall that Theorem 4.2.1 states that a deterministic mechanism is strategyproof if and only if there exist $n+1$ real numbers $\alpha_1, \ldots, \alpha_{n+1} \in \mathbb{R} \cup \{-\infty, +\infty\}$ such that for every $(x_1, \ldots, x_n) \in \mathbb{R}^n$, the mechanism returns the median of $(x_1, \ldots, x_n, \alpha_1, \ldots, \alpha_{n+1})$. We will utilize this, together with the four profiles depicted in Figure 4.7, in order to prove that any deterministic strategyproof mechanism cannot perform better than the claimed bound. Consider a deterministic mechanism $\mathcal{M}$ that picks $n+1$ real numbers $\alpha_1, \ldots, \alpha_{n+1} \in \mathbb{R} \cup \{-\infty, +\infty\}$; in what follows we assume w.l.o.g. that $\alpha_i \leq \alpha_{i+1}$ for $i \in [n]$. Note that the same α_i's must be used for every profile $(x_1, \ldots, x_n) \in \mathbb{R}^n$. Consider the following two cases.

Case 1: Assume that all α_i's are equal and either $-\infty$ or $+\infty$. Consider a profile with $x_1 = x_2 = \ldots = x_n = 0$, and note that the optimal social cost is zero. In addition, observe that the median of the $n+1$-many α-points and the n-many x-points is an α-point. Thus, we get that $\mathcal{M}$ chooses a location with positive social cost, leading to an unbounded approximation guarantee.

Case 2: Assume that there exist a $k \in [n]$ such that $\alpha_{n-k+1} < \alpha_{n-k+2}$. Next we will "embed" the four profiles from Figure 4.7 in the interval $(\alpha_{n-k+1}, \alpha_{n-k+2})$ and we will show that $\mathcal{M}$ will achieve the claimed bound in one of the profiles. This means that for all four profiles, we will assume that $x_i \in (\alpha_{n-k+1}, \alpha_{n-k+2})$ for every $i \in [n]$. Observe that in this case, the median of the α-points and the x-points will be k-th order statistic of $(x_1, x_2, \ldots, x_n)$, i.e., it will be the k-th smallest x-point. In addition, observe the following.

- If $k \leq z$, $\mathcal{M}$ achieves an unbounded approximation guarantee for the profile $\boldsymbol{x}^1$ depicted in Figure 4.7a. This is because the optimal solution locates the facility where $n-z$ agents are located in Figure 4.7a, thus achieving a utilitarian social cost of 0. On the other hand, $\mathcal{M}$ locates the facility where z agents are located in Figure 4.7a, leading to a positive social cost as $z \leq \lfloor \frac{n-1}{2} \rfloor$.

- Symmetrically, if $k \geq n-z+1$, $\mathcal{M}$ achieves an unbounded approximation guarantee for the profile $\boldsymbol{x}^4$ depicted in Figure 4.7d. This is because the optimal solution locates the facility where $n-z$ agents are located in Figure 4.7d, thus achieving a utilitarian social cost of 0. On the other hand, $\mathcal{M}$ locates the facility where z agents are located in Figure 4.7d, leading to a positive social cost.

It remains to consider the cases where $z+1 \leq k \leq n-z$. Assume that $k \leq n/2$ and consider the profile $\boldsymbol{x}^2$ in Figure 4.7b; if $k > n/2$, then symmetric arguments can be applied to the profile $\boldsymbol{x}^3$ in Figure 4.7c. Let d denote the distance between the location where $\lfloor \frac{n}{2} \rfloor$ agents are located and the location where $\lceil \frac{n-2z}{2} \rceil$ agents are located. Let the distance between the location where $\lceil \frac{n-2z}{2} \rceil$ agents are located and the location where z agents are located be larger than d. Observe the following two facts.

- In this case, the optimal solution locates the facility at the location where $\lfloor \frac{n}{2} \rfloor$ agents are located, leading to a utilitarian social cost of $\lceil \frac{n-2z}{2} \rceil \cdot d$.
- The k-th order statistic is where $\lceil \frac{n-2z}{2} \rceil$ agents are located. Hence the social cost of $\mathcal{M}$ will be $\lfloor \frac{n}{2} \rfloor \cdot d$.

Combining the two facts above, we get that the approximation guarantee of $\mathcal{M}$ is $\lfloor \frac{n}{2} \rfloor / \lceil \frac{n-2z}{2} \rceil$, which is $\frac{n}{n-2z}$ if n is even and $\frac{n-1}{n-2z+1}$ if n is odd. This concludes the proof. □

4.5 Augmenting with Output Predictions

4

In this section, we augment our base model with predictions and provide several tight bounds for deterministic strategyproof mechanisms under the consistency-robustness dimension.

4.5.1 Impossibility Result for the Egalitarian Objective

Augmenting facility location problems with a prediction of the optimal location is common in the literature of learning-augmented mechanism design. For the single facility location problem without outliers on the real line, Agrawal et al. [2024] obtained a mechanism that is 1-consistent and 2-robust, achieving the best of both worlds. However, when the problem with outliers is augmented with a prediction of the optimal location, achieving a consistency guarantee that is better than the optimal worst-case guarantee of 2 (Theorem 4.3.1) inevitably leads to an unbounded robustness guarantee, as shown in Theorem 4.5.1 below.

Theorem 4.5.1. *Let $z \leq \lfloor \frac{n-1}{2} \rfloor$. Then, there is no deterministic strategyproof mechanism augmented with a prediction $\hat{y}$ that is better than 2-consistent and achieves a finite robustness for the egalitarian objective.*

Proof: Let $z \leq \lfloor \frac{n-1}{2} \rfloor$. Towards a contradiction, assume that there exists a deterministic strategyproof mechanism $\mathcal{M}$ augmented with $\hat{y}$ that is $(2-\varepsilon)$-consistent, with $\varepsilon > 0$, and achieves a finite robustness. Consider the profiles depicted in Figure 4.8.

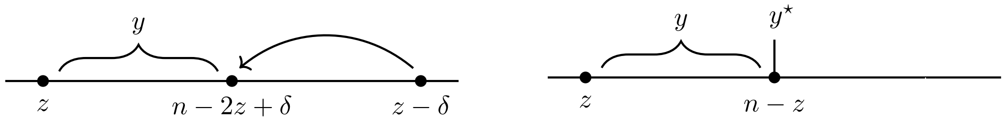

(a) $\delta \in \{0, 1, 2, \ldots, z-1\}$ (b)

Figure 4.8: Profiles used in the proof of Theorem 4.5.1.

The profile in Figure 4.8a with $\delta = 0$ has 3 clusters of locations: there are z agents with a location of 0, $n - 2z \geq 1$ agents with a location of $\frac{1}{2}$ and z agents with a location of 1. Note that all agents from either the leftmost cluster at 0 or the rightmost cluster at 1 can be disregarded in the objective function. Therefore, there are two optimal locations $y^\star$ at $\frac{1}{4}$ and $\frac{3}{4}$ with an egalitarian social cost of $\frac{1}{4}$. Consider the case of a perfect prediction, i.e., $\hat{y} = \frac{1}{4}$ or $\hat{y} = \frac{3}{4}$. In either case, it must be that $\mathcal{M}$ places the facility y in $[\frac{\varepsilon}{4}, \frac{1}{2} - \frac{\varepsilon}{4}]$ or in $[\frac{1}{2} + \frac{\varepsilon}{4}, 1 - \frac{\varepsilon}{4}]$, as $\mathcal{M}$ is $(2-\varepsilon)$-consistent. Assume w.l.o.g. that $y \in [\frac{\varepsilon}{4}, \frac{1}{2} - \frac{\varepsilon}{4}]$, depicted by the curly bracket in Figure 4.8a.

Now, consider the profile in Figure 4.8a with $\delta = 1$: one agent with a location of 1 moved to $\frac{1}{2}$. Note that as $\mathcal{M}$ is strategyproof, the location y remains unchanged by Corollary 4.2.2. The same reasoning holds for $\delta = 2, 3, \ldots, z-1$ consecutively.

Finally, consider the profile in Figure 4.8b in which all z agents with a location of 1 have moved to $\frac{1}{2}$. As the left cluster at 0 now contains z agents, all agents from this cluster can be disregarded, leading to only one optimal location $y^\star$ at $\frac{1}{2}$ with an egalitarian objective of 0. However, as $\mathcal{M}$ is strategyproof the location $y < \frac{1}{2}$ still remains unchanged by Corollary 4.2.2, so y has a positive cost. This contradicts robustness, concluding the proof. □

Additionally, augmenting the problem with other natural predictions leads to the same negative result. One example is a prediction $\hat{z}_\ell$ that indicates how many of the leftmost (smallest) x_i's are outliers (which implies how many of the rightmost (largest) x_i's are outliers, i.e., $\hat{z}_r = z - \hat{z}_\ell$). In this case, the two perfect predictions considered in the instance depicted in Figure 4.8a with $\delta = 0$ of Theorem 4.5.1 are $\hat{z}_\ell = 0$ and $\hat{z}_\ell = z$. Another example is a prediction indicating which z agents are outliers. In this case, the two perfect predictions considered in the instance depicted in Figure 4.8a with $\delta = 0$ are that either all z agents with a location at 0 or all z agents with a location at 1 are outliers. An interesting follow-up question is to determine what type of prediction, if any, would be useful under the egalitarian objective, or whether there is an underlying property of the problem that rules out such a positive result.

4.5.2 Positive Results for the Utilitarian Objective

Without predictions, we have shown in Theorem 4.2.3 that no deterministic strategyproof mechanism exists if $z \geq n/2$. When augmenting the problem with a prediction $\hat{y}$ of the optimal location, this inapproximability persists for the robustness guarantee. Before presenting our positive results, we first elaborate on this. In fact, if the number of outliers is *greater than one-third* of the total number of agents, no deterministic strategyproof mechanism augmented with a prediction $\hat{y}$ can be 1-consistent and achieve a finite robustness. This follows from the more general theorem below, which considers a relaxed consistency guarantee of $\frac{z}{n-2z} > 1$ for $z > 1$ and $n \in \{2z+1, 2z+2, \ldots, 3z-1\}$.

Theorem 4.5.2. *Let $z > 1$ and $n \in \{2z+1, 2z+2, \ldots, 3z-1\}$. Then, there is no deterministic strategyproof mechanism augmented with a prediction $\hat{y}$ that is better than $\frac{z}{n-2z}$-consistent and achieves a finite robustness for the utilitarian objective.*

Proof: Towards a contradiction, assume that there exists a deterministic strategyproof mechanism $\mathcal{M}$ that is better than $\frac{z}{n-2z}$-consistent and achieves a finite robustness. Consider the profiles depicted in Figure 4.9.

The profile in Figure 4.9a with $\delta = 0$ has 3 clusters of locations: there are z agents with a location of 0, $n - 2z \geq 1$ agents with a location of $z \cdot d > 0$ and z agents with a location of $(z+1)d$. As $n < 3z$, there is one optimal location $y^\star$ at $(z+1)d$ with a utilitarian objective of $(n-2z)d$. Consider a perfect prediction $\hat{y}$. As $\mathcal{M}$ is better than $\frac{z}{n-2z}$-consistent, it must be that $\mathcal{M}$ places the facility at y such that the social cost is smaller than $z \cdot d$. To achieve this, $\mathcal{M}$ must place the facility at y such that $y > z \cdot d$.

Now, consider the profile in Figure 4.9a with $\delta = 1$ in which one agent with a location of 0 moved to $z \cdot d$. As $\mathcal{M}$ is strategyproof, it must be that y remains unchanged by Corollary 4.2.2. The same reasoning holds when considering $\delta = 2, 3, \ldots, z-1$ consecutively.

Finally, consider the profile in Figure 4.9b in which all z agents with a location of 0 moved to $z \cdot d$. In this case, there is one optimal location $y^\star$ at $z \cdot d$ with a utilitarian objective of 0. So in order for $\mathcal{M}$ to achieve a bounded robustness, note that $\hat{y} \neq y^\star$, it must place the facility at $y^\star$. However, this contradicts

(a) $\delta \in \{0, 1, 2, \ldots, z-1\}$ (b)

Figure 4.9: Profiles used in the proof of Theorem 4.5.2.

strategyproofness of $\mathcal{M}$ by Corollary 4.2.2 as $y > z \cdot d$, concluding the proof. $\square$

Note that for $z > 1$ and $n \in \{z+2, z+3, \ldots, 2z\}$, it follows from Theorem 4.2.3 that there is no deterministic strategyproof mechanism augmented with a prediction $\hat{y}$ that is 1-consistent and achieves a finite robustness. Finally, if the number of outliers is at most one-third of the total number of agents, we show that no deterministic strategyproof mechanism augmented with a prediction $\hat{y}$ can be 1-consistent and achieve a specific finite robustness depending on the number of agents and outliers.

Theorem 4.5.3. *Let $n \geq 3z$. Then, there is no deterministic strategyproof mechanism augmented with a prediction $\hat{y}$ that is 1-consistent and better than $f(n,z)$-robust for the utilitarian objective with:*

$$f(n,z) = \begin{cases} \frac{n+z-1}{n-3z+1}, & \text{if } n-z \text{ is odd,} \\ \frac{n+z-2}{n-3z+2}, & \text{otherwise.} \end{cases} \tag{4.7}$$

Proof: Towards a contradiction, assume that there exists a deterministic strategyproof mechanism $\mathcal{M}$ augmented with a prediction $\hat{y}$ that is 1-consistent and better than $f(n,z)$-robust as in (4.7). Consider the profiles in Figure 4.10.

First, assume that $n-z$ is odd and consider the profile in Figure 4.10a with $\delta = 0$, $d_1 > d_2 + d_3$ and $0 < d_2 < d_3$. There are $\frac{n-z-1}{2} = z + \theta$, with $\theta \geq 0$, agents with a location of 0, z agents with a location of d_1, $1+\theta$ agents with a location of $d_1 + d_2$, and z agents with a location of $d_1 + d_2 + d_3$. It is easy to see that the optimal location of the facility is at $d_1 + d_2$, as the optimal location is either at d_1 or at $d_1 + d_2$ by Lemma 4.4.1. Consider a perfect prediction $\hat{y}$ located at $d_1 + d_2$. As $\mathcal{M}$ is 1-consistent, it must be that $\mathcal{M}$ places the facility y at $\hat{y}$, as depicted in Figure 4.10a.

Now, consider the profile in Figure 4.10a with $\delta = 1$ in which one agent with a location of 0 moved to d_1. As $\mathcal{M}$ is strategyproof, it must be that y remains unchanged by Corollary 4.2.2. The same reasoning holds when considering $\delta = 2, 3, \ldots, z + \theta - 1$ consecutively.

Finally, consider the profile in Figure 4.10b in which all $z + \theta$ agents with a location of 0 moved to d_1. In this case, the optimal location $y^\star$ is located at d_1 with a utilitarian objective of $(1+\theta)d_2$. Again as $\mathcal{M}$ is strategyproof, it

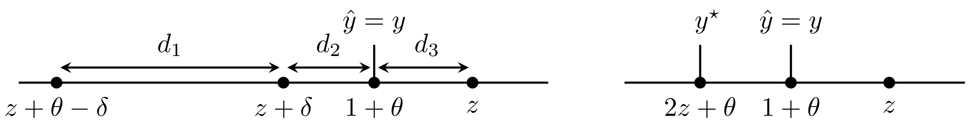

(a) $\delta \in \{0, 1, 2, \ldots, z+\theta-1\}$ (b)

Figure 4.10: Profiles with $n-z$ odd used in the proof of Theorem 4.5.3.

must be that y remains unchanged by Corollary 4.2.2. The social cost of y is $(2z+\theta)d_2$, as $d_3 > d_2$. And as $y^\star \neq \hat{y}$, this leads to a robustness guarantee of $\frac{(2z+\theta)d_2}{(1+\theta)d_2} = \frac{2z+\theta}{1+\theta} = \frac{n+z-1}{n-3z+1}$, as $\theta = \frac{n-z-1}{2} - z$. This contradicts that $\mathcal{M}$ is better than $f(n,z)$-robust as in (4.7).

If $n-z$ is even,[5] consider the same profile in Figure 4.10a with $\delta = 0$, but with $z-1$ agents with a location of d_1 and with $\theta = \frac{n-z}{2} - z \geq 0$. The optimal location is still located at $d_1 + d_2$ and given a perfect prediction $\hat{y} = d_1 + d_2$, $\mathcal{M}$ must still place the facility y at $\hat{y}$ in order to be 1-consistent. Again, moving all $z+\theta$ agents with a location of 0 consecutively to d_1, this leads to the profile depicted in Figure 4.10b with $2z+\theta-1$ agents with a location of d_1. And as $\mathcal{M}$ is strategyproof, $\mathcal{M}$ must again place the facility y at d_1+d_2 by Corollary 4.2.2. In this case, the optimal location $y^\star$ is located at d_1 with a utilitarian objective of $(1+\theta)d_2$, and the social cost of y is $(2z+\theta-1)d_2$, as $d_3 > d_2$. As $y^\star \neq \hat{y}$, this leads to a robustness guarantee of $\frac{(2z+\theta-1)d_2}{(1+\theta)d_2} = \frac{2z+\theta-1}{1+\theta} = \frac{n+z-2}{n-3z+2}$, resulting in a contradiction. This concludes the proof. □

Mechanism In-Range

There is a subtlety why a deterministic strategyproof mechanism cannot achieve 1-consistency and a finite robustness for $n < 3z$. Consider a profile $\boldsymbol{x}$ with 5 locations and $z = 2$ outliers. By Lemma 4.4.1, the optimal location is equal to $x_{\sigma(2)}$, $x_{\sigma(3)}$ or $x_{\sigma(4)}$. However, if there is only one optimal location $y^\star = x_{\sigma(4)}$, then $\sigma(2) \notin S^\star(\boldsymbol{x}, z)$, making it potentially impossible to relate the optimal utilitarian social cost to the social cost when placing the facility at $x_{\sigma(2)}$ (when $x_{\sigma(2)}$ is much smaller than $x_{\sigma(3)}$). Whereas if the number of agents increases by 1 to $n = 6 = 3z$, there exists an optimal location that is equal to either $x_{\sigma(3)}$ or $x_{\sigma(4)}$. Whichever of these two locations is optimal, the other location will be regarded in the optimal utilitarian social cost, i.e., $\sigma(3), \sigma(4) \in S^\star(\boldsymbol{x}, z)$, and therefore the optimal utilitarian social cost can always be related to the social cost when placing the facility at any of these two locations.

So given input $(\boldsymbol{x}, z, \hat{y})$, a mechanism $\mathcal{M}$ is ensured to have a bounded robustness guarantee if $\mathcal{M}$ never chooses a location for the facility that is smaller than $x_{\sigma(z+1)}$ or larger than $x_{\sigma(n-z)}$. Our mechanism In-Range, introduced below, satisfies this property by choosing the prediction $\hat{y}$ as the location of the facility if (i) $x_{\sigma(z+1)} \leq \hat{y} \leq x_{\sigma(n-z)}$, (ii) if $\hat{y}$ is larger than or equal to the smallest value in the set O (Lemma 4.4.1), and (iii) if $\hat{y}$ is smaller than or equal to the largest value in the set O. Note that this interval is constructed in the first line of In-Range and that this interval is well defined. Additionally, note that for $n \geq 3z$, (i) is subsumed by (ii) and (iii).

[5] Note that for $z = 1$, $f(n,z) = 1$ which is a trivial bound on the robustness guarantee.

4

MECHANISM 7 IN-RANGE$(\boldsymbol{x}, z, \hat{y})$

```
let ℓ = max{⌈(n−z+1)/2⌉, z+1} and r = min{⌈(n−z)/2⌉ + z, n−z}
if x_σ(ℓ) ≤ ŷ ≤ x_σ(r) then
    return ŷ
end
if ŷ < x_σ(ℓ) then
    return x_σ(ℓ)
end
else
    return x_σ(r)
end
```

The main result of this section is the following theorem. Recall that the error measure $\eta(\boldsymbol{x}, z, \hat{y})$ is defined as the ratio of the minimum utilitarian social cost for $\hat{y}$ and the optimal utilitarian social cost.

Theorem 4.5.4. *Let $n \geq 3z$. Then, for any input $(\boldsymbol{x}, z, \hat{y})$ with $\eta(\boldsymbol{x}, z, \hat{y}) \leq \eta$, IN-RANGE is strategyproof and achieves an approximation guarantee for the utilitarian objective of:*

$$f(n, z, \eta) = \begin{cases} \min\{\eta, \frac{n+z-1}{n-3z+1}\}, & \textit{if } n-z \textit{ is odd,} \\ \min\{\eta, \frac{n+z-2}{n-3z+2}\}, & \textit{otherwise.} \end{cases} \tag{4.8}$$

In particular, IN-RANGE is 1-consistent and $f(n, z)$-robust as in (4.7)*, which is best possible.*

We introduce some axillary notation in order to prove the approximation guarantee of IN-RANGE. Given input $(\boldsymbol{x}, z, \hat{y})$, let IN-RANGE$(\boldsymbol{x}, z, \hat{y}) = y$. Define $i(y^\star)$ such that $x_{\sigma(i(y^\star))} = y^\star$ (note that such a $y^\star$ always exists). We define $i(y) = i(y^\star)$ if there exists a $y^\star_\ell$ and $y^\star_r$ such that $y \in [y^\star_\ell, y^\star_r]$. If there exists a location in $\boldsymbol{x}$ at which the facility y is placed, i.e., $\exists k \in [n]$ such that $y = x_{\sigma(k)}$, we define $i(y)$ as the maximum (resp. minimum) index such that $x_{\sigma(i(y))} = y$ if $y^\star < y$ (resp. $y^\star > y$). Otherwise, if $y^\star < y$, we define $i(y)$ such that $x_{\sigma(i(y))} > y$ and $x_{\sigma(i(y)-1)} < y$. Symmetrically, if $y^\star > y$, we define $i(y)$ such that $x_{\sigma(i(y))} < y$ and $x_{\sigma(i(y)+1)} > y$. Finally, we define δ as the difference in indices:

$$\delta = |i(y^\star) - i(y)|. \tag{4.9}$$

Note that $\delta \in \{0, 1, \ldots, z\}$ if $n - z$ is odd and $\delta \in \{0, 1, \ldots, z-1\}$ if $n - z$ is even by construction of IN-RANGE and Lemma 4.4.1. If $\delta = 0$, the optimal utilitarian social cost for y and $y^\star$ coincide and as δ grows, the number of agents with a location between y and $y^\star$ potentially increases and potentially leads to a worse approximation guarantee.

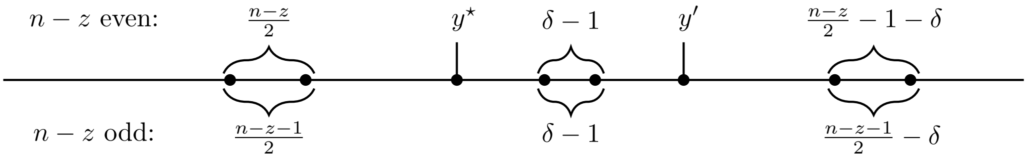

Figure 4.11: Part of the profile used in the proof of Theorem 4.5.4.

Proof of Theorem 4.5.4: We first show that IN-RANGE is strategyproof. Let $(\boldsymbol{x}, z, \hat{y})$ be arbitrary and let $y = \text{IN-RANGE}(\boldsymbol{x}, z, \hat{y})$. Consider an agent $i \in N$ with $x_i = p_i$ and a unilateral deviation of i to x'_i. First, consider the case that $y = \hat{y} \in [x_{\sigma(\ell)}, x_{\sigma(r)}]$. The deviation of agent i can only potentially move y to the right (resp. left) if $x_i \leq x_{\sigma(\ell)}$ (resp. $x_i \geq x_{\sigma(\ell)}$), but this would increase the cost of agent i. Secondly, consider the case that $y = x_{\sigma(\ell)} > \hat{y}$. Also in this case, the deviation of agent i can only potentially move y to the right (resp. left) if $x_i \leq x_{\sigma(\ell)}$ (resp. $x_i \geq x_{\sigma(\ell)}$), but this would increase the cost of agent i. Similar reasoning holds for $y = x_{\sigma(r)} < \hat{y}$.

We now show that IN-RANGE is $f(n, z, \eta)$-approximate as in (4.8). For $n \geq 3z$ it holds that $\ell = \lceil \frac{n-z+1}{2} \rceil$ and $r = \lceil \frac{n-z}{2} \rceil + z$, so $\ell \leq r$. It therefore follows by Lemma 4.4.1 that IN-RANGE is 1-consistent. Note that if $n-z$ is even and the perfect prediction $\hat{y} \notin [x_{\sigma(\ell)}, x_{\sigma(r)}]$, IN-RANGE still returns an optimal solution by construction by returning $x_{\sigma(\ell)}$ or $x_{\sigma(r)}$. Now consider an imperfect prediction $\hat{y}$ with $\hat{y} \in [x_{\sigma(\ell)}, x_{\sigma(r)}]$. In this case, IN-RANGE returns $\hat{y}$ and by definition of $\eta(\boldsymbol{x}, z, \hat{y})$ and as $\eta(\boldsymbol{x}, z, \hat{y}) \leq \eta$, it follows that:

$$\text{SC}(\hat{y}, \boldsymbol{x}, z) = \eta(\boldsymbol{x}, z, \hat{y}) \cdot \text{SC}^\star(\boldsymbol{x}, z) \leq \eta \cdot \text{SC}^\star(\boldsymbol{x}, z) = f(n, z, \eta) \cdot \text{SC}^\star(\boldsymbol{x}, z).$$

The last equality holds as the robustness upper bounds the approximation guarantee, as shown below.

Consider any imperfect prediction $\hat{y}$. First, consider the case that $n - z$ is odd and additionally, assume w.l.o.g. that $y^\star < y$ (symmetric reasoning holds for $y^\star > y$). We upper bound the utilitarian social cost of y in two ways. First, as in the proof of Theorem 4.4.2, we evaluate the social cost of y w.r.t. the set of non-outliers $S^\star$ of an optimal location $y^\star$. Secondly, we upper bound the social cost of y by considering the social cost of $y' = i(y)$ w.r.t. the set $S^\star$. Note that this upper bounds the social cost as, in the case of $y^\star < y$, it holds that $y \leq y'$ and there are more agents $i \in S^\star$ with $x_i \leq y$ than there are with $x_i > y$. Together, this leads to first inequality in (4.10) below.

As $n - z$ is odd, there are $\frac{n-z-1}{2}$ agents $i \in S^\star$ with $\sigma(i) < i(y^\star)$ and $\frac{n-z-1}{2}$ agents $i \in S^\star$ with $\sigma(i) > i(y^\star)$, as depicted in Figure 4.11. Furthermore, let $\delta = |i(y^\star) - i(y)| = |i(y^\star) - i(y')|$ as in (4.9). Then, by Lemma 4.4.1, there are $\delta - 1 \geq 0$ agents $i \in S^\star$ with $i(y^\star) < \sigma(i) < i(y')$, as $\hat{y}$ is not a perfect prediction, which is also depicted in Figure 4.11. Again as in the proof of Theorem 4.4.2, we upper bound the approximation guarantee by moving all agents $i \in S^\star$ with

$x_i < y'$ to $y^\star$, and all agents $i \in S^\star$ with $x_i > y'$ to y', leading to:

$$\frac{\mathrm{SC}(y, \boldsymbol{x}, z)}{\mathrm{SC}^\star(\boldsymbol{x}, z)} \leq \frac{\sum_{i \in S^\star} |y' - x_i|}{\mathrm{SC}^\star(\boldsymbol{x}, z)} \leq \frac{\sum_{\substack{i \in S^\star: \\ x_i < y'}} |y^\star - y'|}{\sum_{\substack{i \in S^\star: \\ x_i \geq y'}} |y^\star - y'|} = \frac{|\{i \in S^\star : x_i < y'\}|}{|\{i \in S^\star : x_i \geq y'\}|}$$
$$\leq \frac{\frac{n-z-1}{2} + \delta}{\frac{n-z-1}{2} + 1 - \delta} \leq \frac{n + z - 1}{n - 3z + 1}. \tag{4.10}$$

Here, the last inequality follows as $\delta \leq z$ if $n - z$ is odd. Additionally, note that if the imperfect prediction $\hat{y} \notin [x_{\sigma(\ell)}, x_{\sigma(r)}]$, we assumed that $y^\star < y$ it must be that $\hat{y} > x_{\sigma(r)} = y$. In this case $\eta \geq \frac{\mathrm{SC}(y, \boldsymbol{x}, z)}{\mathrm{SC}^\star(\boldsymbol{x}, z)}$, as both y and $\hat{y}$ disregard the z leftmost agents in order to minimize the social cost, and y is the median of the corresponding set of non-outliers. Similar reasoning holds when $n - z$ even, i.e., when there are $\frac{n-z}{2}$ locations left of $y^\star$ and $\frac{n-z}{2} - 1 - \delta$ locations right of y' in Figure 4.11, leading to:

$$\frac{\mathrm{SC}(y, \boldsymbol{x}, z)}{\mathrm{SC}^\star(\boldsymbol{x}, z)} \leq \frac{\frac{n-z}{2} + \delta}{\frac{n-z}{2} - \delta} \leq \frac{n + z - 2}{n - 3z + 2},$$

where the last inequality follows as $\delta \leq z - 1$ if $n - z$ is even. Note that optimality of the consistency-robustness trade-off follows by Theorem 4.5.3, concluding the proof. □

Fine-Grained Analysis for $n < 3z$

We consider the approximation guarantee of In-Range separately for $n \in \{2z + 1, 2z + 2, \ldots, 3z - 1\}$, as it cannot be stated as $f(n, z, \eta)$ of Theorem 4.5.4. We elaborate on this with the following illustrative example. Consider the profile $(\boldsymbol{x}, z = 3, \hat{y})$ depicted in Figure 4.12. Note that the optimal utilitarian social cost is 1 and that the minimum utilitarian social cost of $\hat{y}$ is 3.2, leading to $\eta(\boldsymbol{x}, z, \hat{y}) = 3.2$. However, the minimum utilitarian social cost of y chosen by In-Range is 4. This is because for $n < 3z$, it is no longer true that the set of non-outliers minimizing the utilitarian social cost objective for y and $\hat{y}$ is the same. In this example, $\hat{y}$ is contained in the set O of potential optimal locations

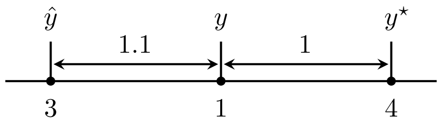

Figure 4.12: Profile with eight agents and three outliers, i.e., $n < 3z$, for which In-Range$(\boldsymbol{x}, z, \hat{y}) = y$ has a worse guarantee than $\eta(\boldsymbol{x}, z, \hat{y})$.

defined by Lemma 4.4.1, and returning the prediction if it is bounded by the smallest and largest values of O would lead to an approximation guarantee of η. But as explained earlier, this mechanism would fail to achieve any bounded robustness guarantee.

In order to derive a fine-grained approximation guarantee of In-Range for these values of n and z, and not just a robustness guarantee, we will use δ as defined in (4.9) more prominently. Intuitively, a mechanisms potentially achieves a good approximation guarantee if it chooses a location close to an optimal location $y^\star$. If the location y chosen by the mechanism is not optimal, then $\delta - 1$ indicates the number of agents with a location between y and $y^\star$, and so δ can be interpreted as an error measure of the mechanism. We therefore introduce some additional axillary notation to define this error measure for perfect and imperfect predictions.

For $n \in \{2z{+}1, 2z{+}2, \dots, 3z{-}1\}$ it holds that $\ell = \max\{\lceil \frac{n-z+1}{2} \rceil, z{+}1\} = z{+}1$ and $r = \min\{\lceil \frac{n-z}{2} \rceil + z, n-z\} = n - z$, so $\ell = z + 1 \leq n - z = r$. Therefore, In-Range is only 1-consistent is some cases and in all other cases, the consistency guarantee depends on how much the thresholds ℓ and r differ from the thresholds of Lemma 4.4.1. We define this by determining the difference δ^c in the k-th order statistics used for the thresholds:

$$\delta^c = z + 1 - \left\lceil \frac{n-z+1}{2} \right\rceil = \left\lceil \frac{n-z}{2} \right\rceil + z - (n-z).$$

4

Note that the difference of the lower and upper thresholds is equal. On the other hand, the robustness guarantee improves compared to the guarantee in (4.7) for $n \geq 3z$. This is due to the change in thresholds which leads to a smaller upper bound of δ. Namely, in this case $\delta \leq \delta^r$ with:

$$\delta^r = \left\lceil \frac{n-z}{2} \right\rceil + z - (z+1) = n - z - \left\lceil \frac{n-z+1}{2} \right\rceil.$$

Note that $\delta \leq \delta^c{+}\delta^r$, and that the consistency guarantee is always better than the robustness guarantee as $\delta^c \leq \delta^r$. We now present the approximation guarantee of In-Range for $n < 3z$.

Theorem 4.5.5. *Let $n \in \{2z+1, 2z+2, \dots, 3z-1\}$. Then, for any input $(\boldsymbol{x}, z, \hat{y})$, In-Range is strategyproof and achieves an approximation guarantee for the utilitarian objective of:*

$$f(n, z, \delta) = \begin{cases} \frac{\frac{n-z-1}{2}+\delta}{\frac{n-z-1}{2}+1-\delta}, & \text{if } n - z \text{ is odd,} \\ \frac{\frac{n-z}{2}+\delta}{\frac{n-z}{2}-\delta}, & \text{otherwise.} \end{cases} \tag{4.11}$$

In particular, In-Range is $f(n, z, \delta^c)$-consistent and $f(n, z, \delta^r)$-robust.

The proof of Theorem 4.5.5 follows analogously to the proof of Theorem 4.5.4. Note that if the input $(\boldsymbol{x}, z, \hat{y})$ of IN-RANGE is such that $\eta(\boldsymbol{x}, z, \hat{y}) \leq \eta$, and $\hat{y} \in [x_{\sigma(\ell)}, x_{\sigma(r)}]$ or $\exists y_\ell^\star, y_r^\star$ such that $\hat{y} \in [y_\ell^\star, y_r^\star]$ and $[y_\ell^\star, y_r^\star] \cap [x_{\sigma(\ell)}, x_{\sigma(r)}] \neq \emptyset$, IN-RANGE achieves an approximation of η for $n \in \{2z+1, 2z+2, \ldots, 3z-1\}$.

Confidence Parameter

If $n \geq 3z$, then IN-RANGE is 1-consistent and $f(n, z)$-robust as in (4.7). And so, the robustness guarantee is worse than the worst-case guarantee of Theorem 4.4.2 without predictions. If there is some uncertainty about the quality of the prediction, one might want to settle for a worse consistency and a better robustness guarantee. For $n < 3z$, we have already seen that IN-RANGE can achieve a trade-off between consistency and robustness by altering the thresholds that determine if the prediction is chosen as the location of the facility. This leads to a natural confidence parameter of the prediction. We introduce the confidence parameter γ, which influences the thresholds defined in Line 1 of IN-RANGE. Namely, we add γ to ℓ and subtract γ from r. If $z > 1$ and depending on the parity of n, γ attains a feasible value if $\gamma \in \{0, 1, \ldots, \gamma_{\max}\}$, with $\gamma_{\max}$ defined as:

4

$$\gamma_{\max} = \begin{cases} \frac{z}{2} - 1, & \text{if } n \text{ is even and } z \text{ is even,} \\ \lfloor \frac{z}{2} \rfloor, & \text{otherwise.} \end{cases}$$

Here, $\gamma = 0$ can be interpreted as full confidence in the prediction, i.e., the thresholds used in Line 1 of IN-RANGE are not adjusted. As γ increases, the left (ℓ) and right (r) thresholds become larger and smaller, respectively, modeling less confidence in the prediction. Less confidence, i.e., a larger value of γ, will lead to a worse consistency but an improved robustness. If $\gamma = \gamma_{\max}$, this translates to no confidence in the prediction and IN-RANGE returns the median and achieves the approximation guarantee as in Theorem 4.4.2.

4.6 Conclusion

In this chapter, we introduced the study of mechanism design with outliers and considered single facility location on the real line as a test case for this perspective. We considered the utilitarian and egalitarian social cost, which are the two most-studied objectives, and provided a complete picture of the impact of outliers by deriving deterministic strategyproof mechanisms achieving tight bounds. One counter-intuitive phenomenon of our results is that under the utilitarian objective, the approximation guarantee deteriorates as the number of outliers increases. Other counter-intuitive phenomena occur when we enrich our model by incorporating output predictions. Unlike the problem without outliers, we prove an impossibility result under the egalitarian objective. Furthermore, while we derive a deterministic strategyproof mechanism with an optimal consistency-robustness trade-off under the utilitarian objective when $n \geq 3z$, we prove that for $n < 3z$ no such mechanism can achieve 1-consistency and finite robustness. We would like to note that all of our mechanisms actually satisfy the stronger requirement of group-strategyproofness. For the setting without predictions, this is a direct consequence of the characterization result of Moulin [1980], as all our mechanisms are k-th order statistics. For the setting with predictions, this can be proven from first principles.

An interesting research question that directly follows from our work concerns the generality of our strong negative result for $n \leq 2z$. While we believe that this result can be extended to other social cost objectives and to mechanisms that are strategyproof in expectation, it would be interesting to examine whether this lower bound even persists for a class of mechanism design problems. Furthermore, a natural direction for future work is to apply this new perspective of outliers to other mechanism design problems. Do other problems exhibit similar phenomena, such as the degradation of approximation guarantees as the number of outliers increases, or impossibility results when incorporating predictions that are contrary to the settings without outliers. Or are there mechanism design problems in which outliers can be leveraged to improve approximation guarantees? Gaining insights in either direction would contribute to our understanding of this new perspective.

Part Two

Budget-Feasible Mechanism Design

Chapter 5

Knapsack Procurement Auctions with Partial Allocations

5.1 Introduction

5

This chapter considers knapsack procurement auctions in which the auctioneer has a strict budget constraint on the total payment. In this setting, we are given a set of agents $N = [n]$, where each agent $i \in N$ has a *private* costs c_i for the service that they offer, and the auctioneer has a publicly known valuation v_i for this service. As the auctioneer makes a payment for a procured service, the agents may strategically misreport their costs in order to obtain higher payments. The auctioneer must determine the *allocation* and *payments*— that is, which services to procure and at what price—while ensuring that the total payment does not exceed the budget B; such payments are said to be *budget feasible*. Under these considerations, a natural goal for the auctioneer is to come up with a truthful and budget-feasible mechanism that computes the allocation and payments, and maximizes the procured value.

This model was originally proposed by Singer [2010] for *indivisible* agents, i.e., the allocation is a binary decision for each agent (procured or not). Given that even the non-strategic version of such budget-constrained problems tend to be NP-hard, the main focus is on deriving mechanisms which optimize the approximation ratio of the procured value with respect to the auctioneer's optimal potential value, i.e., the optimal value of the non-strategic problem. Since the work of Singer [2010], a line of research emerged devoted to obtaining improved results for the original model, as well as results for extensions of the original model. These extensions include, among others, additional feasibility constraints, richer objectives, more general valuation functions and additional assumptions, enhancing our understanding for the indivisible setting.

In this chapter, we study two settings that have received much less attention in the literature. In both setting, the allocation is not binary decision of procuring agents entirely or not at all, but the auctioneer has more flexibility and is allowed to partially procure the services offered by each agent. We additionally assume that the total value obtained by the auctioneer is equal to the sum of the procured values of each individual agent.

Agents with Multiple Levels of Service. We first consider the setting in which each agent offers a service that consists of multiple levels, where more levels can be interpreted as a higher quality of the service. We refer to this model as the *k-level* model, where k denotes the maximum number of levels. Hence, the auctioneer can choose not to hire an agent, or hire the first x levels for some integer x, or to hire the agent entirely, i.e., all k levels. Furthermore, the valuation function associated with each agent is concave, meaning that the marginal value of each level of service is non-increasing. This setting was first introduced by Chan and Chen [2014], in which each agent offered multiple copies of the same good with each additional copy having a smaller marginal value. In their work, it was assumed that the cost of a single level is arbitrary, meaning that it is plausible that the auctioneer can only afford to hire a single level of service of a single agent. Chan and Chen [2014] proposed randomized, truthful, and budget-feasible mechanisms for this setting, with approximation guarantees that depend on the total number of levels/copies. The crucial difference with our setting is that we assume that the auctioneer's budget is large enough to afford any single individual agent entirely, which is in line with the indivisible setting where the auctioneer can afford to hire any single agent.

5

Divisible Agents. A closely related setting is the one in which agents are offering a divisible service, e.g., offering their time. We refer to this model as the *divisible* model. In this case, it is reasonable to assume that the auctioneer can hire each agent for any fraction of the service that they are offering. Again, the valuation function associated with each agent is assumed to be concave, meaning that the marginal gain is non-increasing in the fraction of the procured service. Anari et al. [2018] were the first to study the divisible setting under the so-called *large market* assumption, which in the context of budget-feasible mechanism design roughly means that the cost of the entire service of each agent is much smaller than the budget of the auctioneer. While this assumption is justified in certain settings, it doesn't cover other relevant settings in which the agents' costs may differ vastly or be close to the budget. Additionally, Anari et al. [2018] noticed that in the divisible setting, no truthful and budget-feasible mechanism with a finite approximation guarantee exists without any restriction on the costs. Recently, we revisited this problem without the large market assumption in [Klumper and Schäfer, 2022], but under the much milder assumption that the auctioneer can afford to hire any individual agent entirely. Note that this is standard in the literature for the indivisible setting, but for the divisible setting this assumption does restrict the bidding space of the agents. Under this cost restriction, an instance of the divisible model with linear valuation functions results in the fractional relaxation of the problem introduced in [Singer, 2010].

The two aforementioned settings of procurement auctions have a number of practical applications in various domains. As previously mentioned, the divisible setting would, for example, be useful to model the time availability of a worker in the context of crowdsourcing. The k-level setting can, for example, 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. As mentioned above, we study the above settings under the assumption that the auctioneer can afford each individual agent entirely. For the divisible model, this assumption is necessary in order to achieve any non-trivial approximation guarantee (see [Anari et al., 2018]), and for the k-level model, this assumption circumvents a strong lower bound of Chan and Chen [2014] which is linear to the maximum number of levels offered by any agent. In both settings, our assumptions are much weaker than the large market assumptions often made in the literature (see, e.g., [Anari et al., 2018, Jalaly and Tardos, 2018]). Finally, we believe concave valuation functions are both interesting and relevant to applications as such functions can be used to model diminishing returns of the auctioneer, i.e., the marginal increase in value decreases as the allocation of an agent increases. This motivates the main question that we address in this chapter: can we derive truthful and budget-feasible mechanisms with attractive approximation guarantees if the valuation functions are concave and when partial allocations are allowed?

5.1.1 Our Contributions

This chapter considers truthful and budget-feasible mechanisms for procurement auctions when partial allocations are allowed. Depending on the type of valuation functions, our mechanisms achieve a constant factor approximation: see Table 5.1 for an overview of our results. We elaborate on our main results below.

- We first consider the k-level model in Section 5.3, and introduce our mechanism Greedy-and-Prune (see Mechanism 9) that is specifically tailored to linear valuation functions and is inspired by our mechanism in [Klumper and Schäfer, 2022]. In fact, Greedy-and-Prune is the discretized version and as k grows, the approximation guarantee of Greedy-and-Prune improves (decreases) and converges to $1 + \varphi \approx 2.62$, which was the previously best-known result for the divisible model. Greedy-and-Prune starts by greedily adding single levels of service until adding an additional level exceeds an α fraction of the value of an optimal fractional solution. Then, an upper bound on the declared costs of agents with a positive allocation is imposed, and if exceeded, this leads to the agent being discarded completely.

Model	**Upper Bound**		**Lower Bound**
	Linear	**Concave**	**(Linear)**
k-level ($1 < k < \infty$)	$(1+\beta)/\beta$ (Thm. 5.3.1)	$2+\sqrt{3}$ (Thm. 5.3.6)	$1+\sqrt{2}$ ([Chen et al., 2011])
Divisible ($k = \infty$)	2 (Thm. 5.4.3)	$1+\varphi+\varepsilon$ (Thm. 5.4.15)	$e/(e-1)$ ([Anari et al., 2018])

Table 5.1: Overview of the guarantees of our mechanisms, depending on the model and type of valuation functions. Note that for the k-level model with linear valuations, the optimal value of β depends on k (for details, see Table 5.2), and note that a lower bound for linear valuations also holds for concave valuations.

- Next, we present our mechanism Sort-&-Reject (see Mechanism 10) for concave and non-decreasing valuation functions in Section 5.3.2, achieving a first constant factor approximation of $2+\sqrt{3} \approx 3.73$. The main idea behind our mechanism is to start from an optimal fractional solution, and greedily discard single levels of service one by one until a carefully chosen stopping condition is met.

- We then consider the setting of divisible agents in Section 5.4. For linear valuation functions, we improve upon the state-of-the-art by deriving a 2-approximate mechanism called Prune-and-Assign (see Mechanism 12), establishing that the best-possible approximation is in the interval $[e/(e-1), 2]$, where the lower bound is due to Anari et al. [2018]. Additionally, we establish a separation between the model with divisible agents and its indivisible counterpart, for which a lower bound of $1+\sqrt{2}$ is known due to Chen et al. [2011]. Our mechanism is inspired by the randomized 2-approximate mechanism in [Gravin et al., 2020] for the indivisible setting.

- For the setting with divisible agents and concave and non-decreasing valuation functions, we exploit the correspondence with the setting of k-levels of service: the latter converges to the former as $k \to \infty$. We use a similar mechanism as Greedy-and-Prune, but tailored to concave valuations, as a subroutine in our mechanism Chunk-and-Solve (see Mechanism 14), after applying discretization to the instance with divisible agents. The approximation guarantee of Chunk-and-Solve depends on the discretization step, i.e., on k. When the discretization step is equal to the number of agents, this leads to an approximation guarantee of $2(1+\varphi) \approx 5.24$, as we prove in Section 5.4.2. However, when k is large enough, Chunk-and-Solve can achieve an approximation guarantee of $1+\varphi+\varepsilon \approx 2.62$. We introduced Chunk-and-Solve in [Amanatidis et al., 2025], in which we achieved a first constant approximation of approximately 7.46, but significantly improve upon this with a subtle change in the allocation rule and a more involved analysis of the approximation guarantee. The proof of the $(1 + \varphi + \varepsilon)$-approximation is non-trivial and given in Section 5.4.3. It is mainly based on upper bounding the value of an optimal solution of a fractional concave knapsack problem by the value of an optimal solution to its discretized counterpart, i.e., the k-bounded knapsack problem. Here, $s \geq 3$ is an integer and $k = n \cdot s$, with n the number of agents in the instance at hand.

5.1.2 Related Work

The design of truthful budget-feasible mechanisms for indivisible agents 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 $1 + \sqrt{2} \approx 2.41$ for deterministic mechanisms and a lower bound of 2 for randomized mechanisms. After almost a decade, Gravin et al. [2020] gave a deterministic 3-approximate mechanism, which is currently the best known guarantee, along with a lower bound of 3 when the guarantee is with respect to the optimal non-strategic fractional solution. Regarding randomized mechanisms, Gravin et al. [2020] settled the question by providing a randomized 2-approximate mechanism, matching the lower bound in [Chen et al., 2011]. Finally, the question has also been settled under the large market assumption by Anari et al. [2018], who extended their $\frac{e}{e-1} \approx 1.58$ mechanism for the setting with divisible agents to the indivisible setting. Klumper and Schäfer [2022] studied the divisible setting without the large market assumption, but under the assumption that the private cost of each agent is bounded by the budget and gave, among other results, a deterministic $\frac{3+\sqrt{5}}{2}$-approximate mechanism for linear valuation functions.

5

For indivisible agents, the problem has also been studied under richer valuation functions. This line of work was also initiated by Singer [2010], who gave a randomized 112 approximation for submodular valuations. Once again, this result was improved by Chen et al. [2011] to a 7.91 approximation, and the same authors devised a deterministic mechanism with a 8.34 approximation. The approximation for randomized mechanisms was later improved to 5 by Jalaly and Tardos [2018]. More recently, Balkanski et al. [2022] presented a novel method for designing budget-feasible mechanisms, namely in the form of descending clock auctions, and derived a 4.75-approximate mechanism for monotone submodular valuations. Very recently, these guarantees were further improved by Han et al. [2023], who devised a randomized and a deterministic mechanism with approximation guarantees of 4.3 and 4.45, respectively, under the clock auction framework.

Beyond monotone submodular valuation functions, it becomes significantly harder to obtain truthful mechanisms with small constant approximation guarantees. Namely, for non-monotone submodular valuations, the first randomized mechanism that runs in polynomial time was due to Amanatidis et al. [2019] with an approximation guarantee of 505. This guarantee was improved to 64 by Balkanski et al. [2022] who provided the first deterministic mechanism for the problem and Huang et al. [2023] who gave a further improvement of $(3+\sqrt{5})^2$ for randomized mechanisms. The state-of-art randomized mechanism for non-monotone submodular valuation functions is due to Han et al. [2023], achieving an approximation ratio of 12. Richer valuation functions, such as XOS valuation functions (see [Bei et al., 2017, Amanatidis et al., 2017, Neogi et al., 2024]) and subadditive valuation functions (see [Dobzinski et al., 2011, Bei et al., 2017, Balkanski et al., 2022, Neogi et al., 2024]), have been extensively studied. Very recently, Neogi et al. [2024] gave the first explicit construction of a mechanism achieving a constant approximation for subadditive valuation functions, whereas previously only an existential result was known due to Bei et al. [2017].

Settings with additional combinatorial constrains have also been studied (see [Amanatidis et al., 2016, Leonardi et al., 2017, Amanatidis et al., 2019, Huang et al., 2023]). The setting in which the auctioneer wants to get a set of heterogeneous tasks done and each task requires that the hired agent has a certain skill, has been studied as well (see [Goel et al., 2014, Jalaly and Tardos, 2018]). Recently, Li et al. [2022] studied facility location problems under the lens of budget-feasibility, in which facilities have private opening costs. Finally, the problem has been studied in a beyond worst-case analysis setting by Rubinstein and Zhao [2023].

For further details, we refer the interested reader to the recent survey of Liu et al. [2024].

5.2 Preliminaries

We first define the standard budget-feasible mechanism design model which constitutes the basis of the more general models considered in this paper. The k-level model is introduced in Section 5.2.2 and the divisible model in Section 5.2.3.

5.2.1 Basic Model

We consider a procurement auction consisting of a set of agents $N = [n]$ and an auctioneer (buyer) who has an available budget $B \in \mathbb{R}_{>0}$. Each agent $i \in N$ offers a service (or good) and has a private cost $c_i \in \mathbb{R}_{>0}$, representing their true cost for providing this service in full. The auctioneer derives some value $v_i \in \mathbb{R}_{>0}$ from the full service of agent i, which is assumed to be public information, and wants to purchase services so that the procured value is maximized while not exceeding their budget. An instance I of this problem will be denoted by a tuple $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B)$. Whenever part of the input is clear from context, we omit its explicit reference for conciseness (e.g., often we refer to an instance simply by its corresponding cost vector $\boldsymbol{c}$).

A deterministic mechanism $\mathcal{M}$ in this setting consists of an allocation rule $\boldsymbol{x} : \mathbb{R}^n_{\geq 0} \to \mathbb{R}^n_{\geq 0}$ and a payment rule $\boldsymbol{p} : \mathbb{R}^n_{\geq 0} \to \mathbb{R}^n_{\geq 0}$. To begin with, the auctioneer collects a profile $\boldsymbol{b} = (b_i)_{i \in N} \in \mathbb{R}^n_{\geq 0}$ of declared costs from the agents. Here, b_i denotes the cost declared by agent $i \in N$, which may differ from their true cost c_i. Given the declarations, the auctioneer determines an allocation (hiring scheme) $\boldsymbol{x}(\boldsymbol{b}) = (x_1(\boldsymbol{b}), \dots, x_n(\boldsymbol{b}))$, where $x_i(\boldsymbol{b}) \in \mathbb{R}_{\geq 0}$ is the allocation decision for agent i, i.e., to what extent agent i is hired. Generally, we distinguish between the *divisible* and *indivisible* agent setting by means of the corresponding allocation rule. In the divisible setting, each agent i can be allocated fractionally, i.e., $x_i(\boldsymbol{b}) \in \mathbb{R}_{\geq 0}$, and in the indivisible setting, each agent i can only be allocated integrally, i.e., $x_i(\boldsymbol{b}) \in \mathbb{N}_{\geq 0}$. Given a (possibly fractional) allocation $\boldsymbol{x}$, we define $W(\boldsymbol{x}) = \{i \in N \mid x_i > 0\}$ as the set of agents who are positively allocated under $\boldsymbol{x}$. The auctioneer also determines a vector of payments $\boldsymbol{p}(\boldsymbol{b}) = (p_1(\boldsymbol{b}), \dots, p_n(\boldsymbol{b}))$, where $p_i(\boldsymbol{b})$ is the payment agent i will receive for the extend to which they are allocated/hired.

We assume that agents have quasi-linear utilities, i.e., for a deterministic mechanism $\mathcal{M}(\boldsymbol{b}) = (\boldsymbol{x}(\boldsymbol{b}), \boldsymbol{p}(\boldsymbol{b}))$, the utility of agent $i \in N$ given the profile $\boldsymbol{b}$ is $u_i(\boldsymbol{b}) = p_i(\boldsymbol{b}) - c_i \cdot x_i(\boldsymbol{b})$. The goal of each agent is to maximize their utility, and as the agents are strategic, this may lead to agents misreporting their true cost in order to obtain a higher utility. We seek mechanisms that are truthful (see Definition 2.2.2) and satisfy the following two properties.

- *Individual rationality:* For any declared profile $\boldsymbol{b}$, each agent $i \in N$ has a non-negative utility, i.e., $u_i(\boldsymbol{b}) \geq 0$.

- *Budget-feasibility:* For any declared profile $\boldsymbol{b}$, the sum of all payments made by the auctioneer does not exceed their budget, i.e., $\sum_{i \in N} p_i(\boldsymbol{b}) \leq B$.

Given an allocation $\boldsymbol{x}$, the total value that the auctioneer obtains is denoted by $v(\boldsymbol{x})$. The exact form of this function depends on the studied model, which will be defined in the subsections below.

All models considered in this chapter are single-parameter settings and so the characterization of Myerson [1981] applies.[1] As we seek truthful mechanisms, it is therefore sufficient to focus on the class of mechanisms with *monotone non-increasing* (called *monotone* for short) allocation rules. An allocation rule is monotone if the allocation of an agent increases as their declaration decreases, more formally: if for every agent $i \in N$, every profile $\boldsymbol{b}$, and all $b'_i \leq b_i$, it holds that $x_i(b'_i, \boldsymbol{b}_{-i}) \geq x_i(\boldsymbol{b})$. We will use this together with Theorem 5.2.1 below in order to design truthful and individually rational mechanisms.

[1] We refer the reader to [Apt and Heering, 2022] for a rigorous treatment of the uniqueness property of Myerson's characterization result.

5

Theorem 5.2.1: ([Archer and Tardos, 2001, Myerson, 1981]). *A monotone non-increasing allocation rule $\boldsymbol{x}(\boldsymbol{b})$ admits a payment rule $p(\boldsymbol{b})$ such that $(\boldsymbol{x}(\boldsymbol{b}), p(\boldsymbol{b}))$ is truthful and individually rational if and only if for each agent $i \in N$ and any profile $\boldsymbol{b}_{-i}$, it holds that $\int_0^\infty x_i(z, \boldsymbol{b}_{-i})dz < \infty$. In this case, we can take the payment rule to be:*

$$p_i(\boldsymbol{b}) = b_i x_i(\boldsymbol{b}) + \int_{b_i}^{\infty} x_i(z, \boldsymbol{b}_{-i})dz\,. \tag{5.1}$$

We will exclusively derive monotone allocation rules that are implemented with the payment rule as defined in (5.1). Therefore, in the remainder of this chapter, we adopt the convention of referring to the true cost profile $\boldsymbol{c}$ of the agents as input rather than distinguishing it from the declared cost profile $\boldsymbol{b}$. Finally, whenever tie-breaking is needed in any of our mechanisms, we assume that ties are resolved according to a fixed tie-breaking rule.

5.2.2 Multiple Levels of Service

A natural extension of the standard model introduced above is the model with multiple levels of service (see also [Chan and Chen, 2014]). We will refer to this model as the *k-level* model: Suppose each agent $i \in N$ sells $k \geq 1$ levels of service and has an associated valuation function $v_i : \{0, \ldots, k\} \to \mathbb{R}_{\geq 0}$ which is public information.[2] Here, $v_i(j)$ denotes the auctioneer's value for the first j levels of service of agent i. In this setting, the auctioneer can buy any j levels of service of an agent with $j \in \{0, 1, \ldots, k\}$. Therefore, the allocation rule is constrained to $\boldsymbol{x} : \mathbb{R}^n_{\geq 0} \to \{0, \ldots, k\}^n$.

5

Valuation Functions. Without loss of generality, we assume that each valuation function v_i is normalized such that $v_i(0) = 0$. We study the general class of *concave* and non-decreasing valuation functions, i.e., for each agent $i \in N$, it holds that $v_i(j) \geq v_i(j-1)$ and $v_i(j)-v_i(j-1) \geq v_i(j+1)-v_i(j)$ for all $j = 1, \ldots, k-1$. We also define the *j-th marginal valuation* of agent i as $m_i(j) := v_i(j) - v_i(j-1)$, for $j \in \{1, \ldots, k\}$, with the convention that $m_i(0) = 0$. So for each agent i the marginal valuations are non-increasing. Given a profile $\boldsymbol{c}$, the total value that the auctioneer derives from an allocation $\boldsymbol{x}(\boldsymbol{c})$ is defined by the separable concave function $v(\boldsymbol{x}(\boldsymbol{c})) = \sum_{i \in N} v_i(x_i(\boldsymbol{c}))$.

Cost Restrictions: In the k-level model, the true cost and the declared cost of each agent $i \in N$ are with respect to a *single* level of service of agent i. Furthermore, the total incurred cost of each agent i is linear, i.e., the true cost of agent i for j levels of service is $j \cdot c_i$, aligning with the utility of agent i as defined

[2] Our results very easily extend to the setting where there is a different (public) k_i associated with each agent i. We use a common k for the sake of presentation.

in Section 5.2.1. An instance I of the k-level model will be denoted by the tuple $I = (N, \boldsymbol{c}, (v_i)_{i \in N}, B, k)$. In this chapter, we only consider instances I in which the auctioneer can afford all levels of service of each single agent, i.e., for each agent $i \in N$ it holds that $k \cdot c_i \leq B$. Note that under this assumption, we can constrain the space of the (declared) cost profiles to $[0, \frac{B}{k}]^n$, i.e., we assume that any agent who violates the respective condition is discarded up front from the input by for example running a pre-processing step that removes such agents. In contrast, in the setting of Chan and Chen [2014] the auctioneer is only guaranteed to be able to afford one level of service of each single agent, i.e., for each agent $i \in N$ it holds that $c_i \leq B$. We derive an almost tight result for this setting in [Amanatidis et al., 2025], the details of which will not be covered in this chapter.

Benchmark. The performance of a mechanism is measured by comparing $v(\boldsymbol{x}(\boldsymbol{c}))$ with the underlying (non-strategic) combinatorial optimization problem, which is commonly referred to as the *k-bounded knapsack problem* (see, e.g., [Martello and Toth, 1990] for a classification of knapsack problems):

$$\texttt{opt}^k(\boldsymbol{c}) := \max \sum_{i=1}^{n} v_i(x_i), \quad \text{s.t.} \quad \sum_{i=1}^{n} c_i x_i \leq B, \quad x_i \in \{0, \ldots, k\} \quad \forall i \in N. \tag{5.2}$$

The k-bounded knapsack problem is NP-hard in general, since for $k = 1$ it reduces to the well-known *0-1 knapsack problem.* We consider the relaxation of the above problem as a proxy for $\texttt{opt}^k(\boldsymbol{c})$. Namely, we consider the *fractional k-bounded knapsack problem* that follows from the k-bounded knapsack problem defined in (5.2) by relaxing the integrality constraint:

5

$$\begin{aligned} \texttt{opt}_{\mathrm{F}}^k(\boldsymbol{c}) := \max \; & \sum_{i=1}^{n} v_i(\lfloor x_i \rfloor) + m_i(\lceil x_i \rceil)(x_i - \lfloor x_i \rfloor) \\ \text{s.t.} \; & \sum_{i=1}^{n} c_i x_i \leq B, \quad x_i \in [0, k] \quad \forall i \in N. \end{aligned} \tag{5.3}$$

Naturally, it holds that $\texttt{opt}_{\mathrm{F}}^k(\boldsymbol{c}) \geq \texttt{opt}^k(\boldsymbol{c})$.

It is not hard to see that due to the fact that the objective is separable concave, $\texttt{opt}_{\mathrm{F}}^k(\boldsymbol{c})$ inherits the well-known properties of its one-dimensional analogue. This includes the fact that an optimal solution can be computed by an adaptation of the standard greedy algorithm that sorts the elements in non-increasing order of efficiency (marginal value-per-cost) and selects as many elements as possible [Hochbaum, 1995]. To the best of our knowledge, separable concave functions are the widest class of objectives for which the fractional k-bounded knapsack problem admits optimal solutions with this structure which we crucially utilize. For completeness, we state the algorithm as Algorithm 8, along with an easy fact that we will use.

Algorithm 8: Fractional k-Bounded Knapsack Problem

Input: An instance $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B, k)$.
Output: An optimal fractional solution $\boldsymbol{x}^*$.

initialize an empty list $\mathcal{L}$ of $k \cdot n$ elements
for $i \in N$ **do**
  for $j \in \{1, \ldots, k\}$ **do**
    add $\frac{m_i(j)}{c_i}$ to $\mathcal{L}$
  end
end
sort $\mathcal{L}$ in decreasing order and let $\boldsymbol{x^*} = \boldsymbol{0}$ and $j = 1$
while $\sum_{i\in N} c_i x_i^* < B$ **do**
  let $\ell \in [n]$ be the agent corresponding to the j-th entry in $\mathcal{L}$
  set $x_\ell^* = x_\ell^* + \min\left(\frac{B - \sum_{i\in N} c_i x_i^*}{c_\ell}, 1\right)$ and $j = j + 1$
end
return $\boldsymbol{x^*}$

Fact 5.2.2. Given an instance $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B, k)$ of the fractional k-bounded knapsack problem, Algorithm 8 computes an optimal solution $\boldsymbol{x}^*$ that has at most one coordinate with a non-integral value.

5.2.3 Divisible Agents

Another natural extension of the standard model introduced above is the model with divisible agents (see also [Anari et al., 2018]). We will refer to this model as the *divisible* model: Suppose each agent $i \in N$ sells an arbitrary fraction of their full service and has an associated valuation function $v_i : [0, 1] \to \mathbb{R}_{\geq 0}$ which is public information. Here, $v_i(x)$ represents how valuable a fraction $x \in [0, 1]$ of the full service of agent i is to the auctioneer. In this setting the auctioneer can buy any fraction x_i of the full service of an agent i with $x_i \in [0, 1]$. Therefore, the allocation rule is constrained to $\boldsymbol{x} : \mathbb{R}^n_{\geq 0} \to [0, 1]^n$.

Valuation Functions. Also here, we assume without loss of generality that each valuation function v_i is normalized such that $v_i(0) = 0$. We study the general class of concave and non-decreasing valuation functions, i.e., for each agent $i \in N$ and $0 \leq x \leq z \leq 1$ it holds that $v_i(x) \leq v_i(z)$ and for $y = (1 - \alpha)x + \alpha z$ with $\alpha \in [0, 1]$ it holds that $v_i(y) = v_i((1 - \alpha)x + \alpha z) \geq (1 - \alpha)v_i(x) + \alpha v_i(z)$. Given a profile $\boldsymbol{c}$, the total value that the auctioneer derives from an allocation $\boldsymbol{x}(\boldsymbol{c})$ is still additive over the different agents, i.e., we define $v(\boldsymbol{x}(\boldsymbol{c})) = \sum_{i\in N} v_i(x_i(\boldsymbol{c}))$. Following from the assumptions on the valuation functions, the following fact shows that the value-over-cost ratios of smaller allocations have a better efficiency.

Fact 5.2.3. Let $f : [0, 1] \to \mathbb{R}_{\geq 0}$ be a concave and non-decreasing function with $f(0) = 0$ and let $0 \leq x < y < z \leq 1$. Then:

$$\frac{f(y) - f(x)}{y - x} \geq \frac{f(z) - f(x)}{z - x} \geq \frac{f(z) - f(y)}{z - y}.$$

Cost Restrictions. In the divisible model, the true cost and the declared cost of each agent $i \in N$ are with respect to the full service of agent i. In this chapter, we only consider instances I in which the auctioneer can afford the full service of each single agent, i.e., for each agent $i \in N$ it holds that $c_i \leq B$. Note that under this assumption, we can also constrain the space of the (declared) cost profiles as in Section 5.2.2, but in this case to $[0, B]^n$. Also note that, as mentioned in [Anari et al., 2018], there is no mechanism that satisfies the three properties introduced in Section 5.2.1 and achieve a bounded approximation guarantee of α, if there is no assumption on the fraction that the auctioneer is able to afford of each single agent.

Benchmark. As above, the performance of a mechanism is measured by comparing $v(\boldsymbol{x}(\boldsymbol{c}))$ with the underlying (non-strategic) combinatorial optimization problem, which we refer to as the *fractional concave knapsack problem*:

$$\texttt{opt}_{\mathrm{F}}(\boldsymbol{c}) := \max \sum_{i=1}^{n} v_i(x_i) \ \text{ s.t. } \ \sum_{i=1}^{n} c_i x_i \leq B, \quad x_i \in [0, 1] \ \ \forall i \in N. \qquad (5.4)$$

5

5.3 Mechanisms for Multiple Levels of Service

We first consider the k-level model and derive two greedy mechanisms. In Section 5.3.1, we introduce our mechanism Greedy-and-Prune that is specifically tailored to linear valuation functions and is inspired by our work in [Klumper and Schäfer, 2022]. In fact, Greedy-and-Prune is the discretized version, and as the levels of service k grows, the approximation guarantee of Greedy-and-Prune improves (decreases) and converges to $1 + \varphi \approx 2.62$, which was the previously best-known result for the divisible model. In Section 5.3.2, we consider concave valuation functions and introduce our mechanism Sort-&-Reject. The approximation guarantee of Sort-&-Reject remains fixed as the levels of service k grows. This mechanism will also be used in our Chunk-and-Solve mechanism for the divisible model (see Section 5.4.2). Both our mechanisms fall into the family of mechanisms which truncate a greedy optimal solution, as is often the case in the literature of budget-feasible mechanism design.

5.3.1 Linear Valuations

In this section, we derive a mechanism, called GREEDY-AND-PRUNE, for instances $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B, k)$ of the k-level model with linear valuations. In this case, it holds that the marginal valuations of an agent $i \in N$ are equal, i.e., $m_i(j) = v_i(j) - v_i(j-1) = v_i(j+1) - v_i(j) = m_i(j+1)$ for all $j = 1, \ldots, k-1$. Recall that in the k-level model, the true and the declared cost of each agent $i \in N$ are with respect to a single level of service and that we only consider instances I with $c_i \cdot k \leq B$ for each agent $i \in N$.

As mentioned in Section 5.2.1, we will need to design a mechanism with a monotone allocation rule and compute the payments according to (5.1) in order for our mechanism to be truthful and individually rational (Theorem 5.2.1). When computing the payment according to (5.1) in case of a truthful declaration of an agent i, the first term in (5.1) is equal to the true cost incurred by agent i (area of the left dashed rectangle in Figure 5.1a) and the second term can be interpreted as the amount agent i is overpaid in order to ensure truthfulness (which can be upper bounded by the area of the right dashed rectangle in Figure 5.1a). We want this second term to be as small as possible, so that more budget is left to spend on other agents in order to acquire more value and achieve a better approximation guarantee.

5

Our mechanism realizes this by first computing the allocation $\boldsymbol{x}$, and subsequently imposing a *threat* for each agent with a positive allocation. More specifically, if the declared cost of an agent $i \in N$ with $x_i > 0$ exceeds this threat denoted by τ_i, the allocation x_i of agent i is set to 0 (see also Figure 5.1b). Thus, for each agent i for which our mechanism initially computes a positive allocation, our goal is to define the threat τ_i in such a way that it is close to the true cost, while remaining at least as large as the true cost. This way, we try to limit the amount that an agent i is overpaid, while ensuring that agent i is not deselected. Note that the threat τ_i imposes an upper bound on the *threshold bid* of agent i, which is the largest cost agent i can declare such that agent i wins ($x_i > 0$).

As we will use the threats to upper bound the payments (as illustrated in Figure 5.1), the threat τ_i of an agent $i \in N$ must not increase when agent i declares a higher cost. Note that bounding the payments with τ_i can still be done if τ_i decreases when agent i declares a higher cost. In our mechanism, we use the following threat τ_i for an agent i, which is independent of the declared cost of i:

$$\tau_i(\boldsymbol{c}_{-i}) = \frac{m_i(1)}{\alpha(1+\beta)\,\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i})} B, \tag{5.5}$$

where $\alpha \in (0, 1]$ and $\beta > 0$ are parameters of GREEDY-AND-PRUNE. Note that $\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i})$ is the value of an optimal solution to the (non-strategic) fractional k-bounded knapsack problem (as defined in (5.3)) for the set of agents $N \setminus \{i\}$.

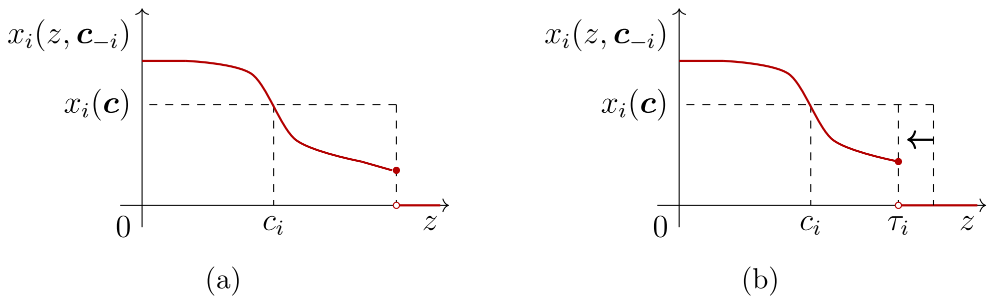

Figure 5.1: The areas of the dashed rectangles in (a) provide an upper bound for the payment of agent i given declaration c_i. Our mechanism tries to decrease this upper bound by imposing a threat τ_i so that the area of the dashed right rectangle decreases, as depicted in (b).

Additionally, we define for each agent $i \in N$:

$$\rho_i(\boldsymbol{c}_{-i}) = \frac{v_i(k)}{\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i})}, \tag{5.6}$$

which represents some measure of how valuable agent i individually is compared to the set of agents $N \setminus \{i\}$. Note that an agent i has no influence on their respective measure ρ_i.

We can now introduce Greedy-and-Prune (see Mechanism 9), which will either select one very valuable agent, or select agents in a greedy manner according to marginal value-per-cost ratios (efficiencies). We define the most valuable agent i^* as the agent with the highest ρ_i as defined in (5.6). If this value ρ_{i^*} is high enough compared to some predefined threshold β, we allocate all k levels of agent i^*, i.e., $x_{i^*} = k$, and no levels of any other agent $i \neq i^*$, i.e., $x_i = 0$. Otherwise, Greedy-and-Prune computes the allocation greedily. This is done by first renaming the agents in order of non-increasing marginal value-per-cost ratios. Then, agents are allocated to the extend of k levels in this order, until allocating k levels to agent ℓ results in the value of the tentative allocation exceeding another predefined threshold. Agent ℓ is then allocated to the extend of $j \in \{0, 1, \ldots, k-1\}$ levels, such that the predefined threshold is not exceeded for $x_\ell = j$ but would be exceeded for $x_\ell = j + 1$. The intuition behind this greedy scheme is that the allocated levels of these agents are the most valuable (efficient) ones: recall that the valuation functions are linear, so the marginal valuations of each agent are constant. Additionally, these are also the levels of service of the agents that Algorithm 8 allocates first in order to compute an optimal solution to the (non-strategic) fractional k-bounded knapsack problem. Greedy-and-Prune then imposes the threat τ_i to each agent $i \in N$ with positive allocation, i.e., $x_i > 0$, and finally computes the payments according to (5.1) in Theorem 5.2.1.

The main result of this section is the following theorem:

MECHANISM 9: Greedy-and-Prune(I)

Input: An instance $I = (N, \boldsymbol{c}, (v_i)_{i \in N}, B, k)$.
Output: An allocation vector $\boldsymbol{x}$ and a payment vector $\boldsymbol{p}$.

let $i^* \in \arg\max_{i \in N} \rho_i(\boldsymbol{c}_{-i})$
if $\rho_{i^*}(\boldsymbol{c}_{-i}) \geq \beta$ **then**
  set $x_{i^*} = k$ and $x_i = 0$ for $i \in N \setminus \{i^*\}$
end
else
  rename agents in N such that $\frac{m_1(1)}{c_1} \geq \frac{m_2(1)}{c_2} \geq \cdots \geq \frac{m_n(1)}{c_n}$
  compute ℓ and j such that $\sum_{i=1}^{\ell-1} v_i(k) < \alpha \, \mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) \leq \sum_{i=1}^{\ell} v_i(k)$ and $\sum_{i=1}^{\ell-1} v_i(k) + v_\ell(j) < \alpha \, \mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) \leq \sum_{i=1}^{\ell-1} v_i(k) + v_\ell(j+1)$
  set $x_i = k$ for $1 \leq i \leq \ell - 1$, $x_\ell = j$ and $x_i = 0$ for $\ell + 1 \leq i \leq n$
  for $i \leq \ell$ **do**
    if $c_i > \tau_i$ **then**
      set $x_i = 0$
    end
  end
end
for $i \in N$ **do**
  compute payment p_i according to (5.1)
end
return $(\boldsymbol{x}, \boldsymbol{p})$

Theorem 5.3.1. *Consider the class of instances of the k-level model with linear valuation functions. Then, for $\alpha \in (0,1]$ and $\beta > 0$ such that $\alpha = \frac{\beta}{1+\beta} + \frac{\beta}{k}$ and $\alpha(1+\beta) \leq 1 - \alpha$, Greedy-and-Prune is a truthful, individually rational and budget-feasible mechanism that is $(1+\beta)/\beta$-approximate.*

We start by showing that the allocation rule of Greedy-and-Prune is monotone.

Lemma 5.3.2. *The allocation rule of Greedy-and-Prune is monotone for $\alpha \in (0,1)$ and $\beta > 0$.*

Proof: Let $\boldsymbol{c}$ be a cost profile. Note that the two threshold values used in Greedy-and-Prune also depend on the input parameters α and β, which are fixed and do not depend on the declarations $\boldsymbol{c}$. As an agent has a positive allocation by either being allocated as i^* or by being allocated in the greedy order, we consider the following two cases.

Case 1: $\rho_{i^*}(\boldsymbol{c}_{-i^*}) \geq \beta$. In this case, only agent i^* has a positive allocation ($x_{i^*}(\boldsymbol{c}) = k$) and i^* is therefore the only agent for which we have to examine a unilateral deviation. Suppose that agent i^* unilaterally deviates and decreases their declared cost to $c'_{i^*} < c_{i^*}$, and let $\boldsymbol{c}' = (c'_{i^*}\boldsymbol{c}_{-i^*})$. Note that such a deviation has no influence on which agent is chosen as i^* and does not alter the if condition in Line 2. To see this, notice that the value $\texttt{opt}^k_{\mathrm{F}}(\boldsymbol{c}'_{-i^*}) = \texttt{opt}^k_{\mathrm{F}}(\boldsymbol{c}_{-i^*})$ does not depend on the declaration of i^*, whereas $\texttt{opt}^k_{\mathrm{F}}(\boldsymbol{c}'_{-i^*}) \geq \texttt{opt}^k_{\mathrm{F}}(\boldsymbol{c}_{-i^*})$, for any $i \neq i^*$. Therefore, $\rho_{i^*}(\boldsymbol{c}'_{-i^*}) = \rho_{i^*}(\boldsymbol{c}_{-i^*})$, whereas $\rho_i(\boldsymbol{c}'_{-i}) \leq \rho_i(\boldsymbol{c}_{-i})$, for any $i \neq i^*$, and agent i^* remains the same by consistent tie-breaking. And so, agent i^* will remain the sole agent with a positive allocation and $x_{i^*}(\boldsymbol{c}') = x_{i^*}(\boldsymbol{c}) = k$ (where in a slight abuse of notation we introduce arguments to the components of $\boldsymbol{x}$ to distinguish between the two runs).

Case 2: $\rho_{i^*}(\boldsymbol{c}_{-i^*}) < \beta$. In this case, the initial allocation is computed in the else statement of Greedy-and-Prune, and multiple agents may have a positive allocation. (Note that if no agent has a positive allocation, we do not need to examine any unilateral deviation.) Fix an agent $i \in N$ with $x_i(\boldsymbol{c}) > 0$, which implies that $c_i \leq \tau_i$. Suppose agent i unilaterally deviates to $c'_i < c_i$ and let $\boldsymbol{c}' = (c'_i, \boldsymbol{c}_{-i})$. First of all, again note that for every $j \in N \setminus \{i\}$, it holds that $\rho_j(\boldsymbol{c}'_{-j}) \leq \rho_j(\boldsymbol{c}_{-j})$, whereas $\rho_i(\boldsymbol{c}'_{-i}) = \rho_i(\boldsymbol{c}_{-i})$. Therefore, under $\boldsymbol{c}'$, the if statement in Line 2 also evaluates to false and Greedy-and-Prune executes the else statement. Moreover, agent i may get a smaller index assigned in Line 6 as $c'_i \leq c_i$. Therefore, agent i may have a higher initial allocation, as $\texttt{opt}^k_{\mathrm{F}}(\boldsymbol{c}') \geq \texttt{opt}^k_{\mathrm{F}}(\boldsymbol{c})$. Finally, note that the threat τ_i is independent of the declaration of agent i, and so $c'_i < c_i \leq \tau_i$. Therefore, $x_i(\boldsymbol{c}') \geq x_i(\boldsymbol{c})$, proving monotonicity of the allocation rule of Greedy-and-Prune. □

5

Lemma 5.3.3. *For $\alpha \in (0,1)$ and $\beta > 0$, Greedy-and-Prune is budget feasible.*

Proof: Let $\boldsymbol{c}$ be a cost profile. We consider the following two cases.

Case 1: $\rho_{i^*}(\boldsymbol{c}_{-i^*}) \geq \beta$. In this case, only agent i^* has a positive allocation and $x_{i^*}(\boldsymbol{c}) = k$. Therefore:

$$\sum_{i \in N} p_i(\boldsymbol{c}) = p_{i^*}(\boldsymbol{c}) \leq k \cdot \frac{B}{k} = B,$$

as the allocation rule of Greedy-and-Prune is monotone (Lemma 5.3.2), so we can bound the payment of agent i^* by $x_{i^*}(\boldsymbol{c}) = k$ times the threshold bid of i^*. Note that we can upper bound the threshold bid by $\frac{B}{k}$, as we run a pre-processing step that removes agents with a declaration greater than $\frac{B}{k}$.

Case 2: $\rho_{i^*}(\boldsymbol{c}_{-i^*}) < \beta$. In this case, the allocation is computed in the else statement of Greedy-and-Prune. We again bound the payment of an agent $i \in N$ by $x_i(\boldsymbol{c})$ times the threshold bid of i. Note that in this case, we can use τ_i as an upper bound of the threshold bid of any agent $i \in N$, as τ_i is independent of the declaration c_i. Rename the agents as in Line 6 of Greedy-and-Prune, then:

$$\sum_{i \in N} p_i(\boldsymbol{c}) \leq \sum_{i=1}^{\ell} x_i(\boldsymbol{c}) \cdot \tau_i = \sum_{i=1}^{\ell} x_i(\boldsymbol{c}) \frac{m_i(1)B}{\alpha(1+\beta)\,\texttt{opt}_{\mathrm{F}}^{k}(\mathbf{c}_{-i})}$$
$$\leq \sum_{i=1}^{\ell} x_i(\boldsymbol{c}) m_i(1) \cdot \frac{B}{\alpha\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})} < B,$$

where the first inequality holds as $x_i(\boldsymbol{c}) = 0$ for $i \geq \ell + 1$. The second inequality follows as for each agent $i \in N$ it holds that $(1+\beta)\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i}) > \texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i}) + v_i(k) \geq \texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$, as the if condition in Line 2 evaluated to false. Finally, the last inequality follows as $\sum_{i=1}^{\ell} x_i(\boldsymbol{c}) m_i(1) = \sum_{i=1}^{\ell-1} v_i(k) + v_\ell(j) < \alpha\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$ by construction, proving that Greedy-and-Prune is budget feasible. □

In order to prove the approximation guarantee of Greedy-and-Prune, we will make use of the following lemma.

Lemma 5.3.4. *Given an instance $I = (N, \boldsymbol{c}, (v_i)_{i \in N}, B, k)$ of the k-bounded knapsack problem with linear valuations and $\alpha \in (0,1)$. Let $n = |N|$ and assume that the agents are ordered such that $\frac{m_1(1)}{c_1} \geq \frac{m_2(1)}{c_2} \geq \cdots \geq \frac{m_n(1)}{c_n}$. Let $\ell \in [n]$ such that $\sum_{i=1}^{\ell-1} v_i(k) < \alpha\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) \leq \sum_{i=1}^{\ell} v_i(k)$. Then, for $i \in [\ell]$, it holds that:*

$$\frac{c_i}{m_i(1)}(1-\alpha)\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) < B.$$

Proof: Let $\boldsymbol{x}^*$ be the optimal solution computed by Algorithm 8 when ties are broken in order of increasing indices of the agents. Note that we can split $\boldsymbol{x}^*$ into $\boldsymbol{x} = (x_1^*, \ldots, x_{\ell-1}^*, x_\ell, 0, \ldots, 0)$ and $\boldsymbol{y} = (0, \ldots, 0, x_\ell^* - x_\ell, x_{\ell+1}^*, \ldots, x_n^*)$ with $x_i = k = x_i^*$ for $i < \ell$ and $x_\ell \in \{0, 1, \ldots, x_\ell^* - 1\}$ such that:

(1) $\boldsymbol{x}^* = \boldsymbol{x} + \boldsymbol{y}$,

(2) $v(\boldsymbol{x}) < \alpha\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$, and

(3) $v(\mathbf{y}) > (1-\alpha)\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$.

By feasibility of $\boldsymbol{x}^*$, and thus of $\boldsymbol{y}$, it follows that for $i \in [\ell]$:

$$B \geq \sum_{i=\ell}^{n} c_i y_i = \sum_{i=\ell}^{n} \frac{c_i}{m_i(1)} m_i(1) y_i \geq \frac{c_\ell}{m_\ell(1)} \sum_{i=\ell}^{n} m_i(1) y_i$$
$$= \frac{c_\ell}{m_\ell(1)} v(\boldsymbol{y}) > \frac{c_\ell}{m_\ell(1)}(1-\alpha)\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) \geq \frac{c_i}{m_i(1)}(1-\alpha)\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}),$$

where the second inequality follows by the ordering of the agents and the third inequality follows from (3). Note that the last inequality also follows by the ordering of the agents, concluding the proof. □

With Lemma 5.3.4 in place, the proof of the approximation guarantee of Greedy-and-Prune follows easily.

Lemma 5.3.5. *Given an instance $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B, k)$ with linear valuations. Let $\alpha \in (0,1)$ and $\beta > 0$ such that $\alpha = \frac{\beta}{1+\beta} + \frac{\beta}{k}$ and $\alpha(1+\beta) \leq 1-\alpha$. Then, Greedy-and-Prune is $(1+\beta)/\beta$-approximate.*

Proof: Let $\boldsymbol{x} = \boldsymbol{x}(\boldsymbol{c})$ be the allocation computed by Greedy-and-Prune. We consider the following two cases.

Case 1: $\rho_{i^*}(\boldsymbol{c}_{-i^*}) \geq \beta$. In this case only agent i^* has a positive allocation of $x_{i^*} = k$. Furthermore, it holds that:

$$v_{i^*}(k) \geq \beta\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i^*}) \geq \beta\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}) - \beta v_{i^*}(k),$$

and rearranging terms leads to $v(\boldsymbol{x}) = v_{i^*}(k) \geq \beta/(1+\beta)\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c})$.

Case 2: $\rho_{i^*}(\boldsymbol{c}_{-i^*}) < \beta$. In this case, the allocation is computed in the else statement of Greedy-and-Prune. In order for Greedy-and-Prune to achieve the best possible approximation guarantee, we want no agent $i \leq \ell$, according to the ordering of the agents as in Line 6, to be deselected due to the imposed threat, i.e. we want that $c_i \leq \tau_i$. As $\alpha(1+\beta) \leq 1-\alpha$ and $\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i}) \leq \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c})$, it holds that:

$$\alpha(1+\beta)\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i}) \leq (1-\alpha)\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}),$$

which for $i \leq \ell$ leads to:

$$\tau_i(\boldsymbol{c}_{-i}) = \frac{m_i(1)}{\alpha(1+\beta)\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i})} B \geq \frac{m_i(1)}{(1-\alpha)\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c})} B > c_i.$$

Here, the last inequality follows from Lemma 5.3.4 and the ordering of the agents. Therefore, by construction of $\boldsymbol{x}$:

$$\begin{aligned} v(\boldsymbol{x}) &\geq \alpha\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}) - m_\ell(1) \geq \alpha\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}) - \frac{\beta}{k}\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-\ell}) \\ &\geq \left(\alpha - \frac{\beta}{k}\right) \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}) = \frac{\beta}{1+\beta}\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}), \end{aligned}$$

where the second inequality follows by definition of $\rho_\ell(\boldsymbol{c}_{-i})$ as the valuations are linear. The third inequality follows as $\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-\ell}) \leq \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c})$, and the last inequality follows by definition of α. As $\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}) \geq \mathtt{opt}^k(\boldsymbol{c})$, this concludes the proof. □

Together, Lemmas 5.3.2, 5.3.3 and 5.3.5 prove Theorem 5.3.1.

5

k	$=$	2	3	7	8	50	100	200
β	$\approx$	0.343	0.398	0.494	0.506	0.596	0.607	0.612
$\frac{1+\beta}{\beta}$	$\approx$	3.916	3.515	3.026	2.977	2.679	2.649	2.633

Table 5.2: Optimal parameters β and approximation guarantees $\frac{1+\beta}{\beta}$ for certain levels of service k for Greedy-and-Prune such that Theorem 5.3.1 holds.

Given an instance I of the problem with k levels of service, one can compute the parameters α and β such that Theorem 5.3.1 holds and achieves the best possible approximation guarantee. This boils down to solving $\beta(\beta+k+2) = {}^{k}/_{\beta+1}$ for $\beta > 0$, for which solutions for specific values of k can be found in Table 5.2. Note that for $k = 3$, the approximation guarantee is smaller than $2+\sqrt{3} \approx 3.73$, which is the approximation guarantee we derive in the next section for concave valuation functions, which include linear valuation functions, for each level of service k. However, only for $k \geq 8$ levels of service does the approximation guarantee drop below 3, which is the approximation guarantee of the best-known mechanism for the model of indivisible agents ($k = 1$) by Gravin et al. [2020]. In the analysis of the approximation guarantee, Gravin et al. [2020] compare the value of the allocation computed by their mechanism to the optimal value of the (non-strategic) fractional knapsack problem. As the optimal value of the fractional knapsack problem is greater than or equal to the optimal value of the k-bounded knapsack problem, we can obtain an approximation guarantee of 3 by running the mechanism of Gravin et al. [2020] for instances I with $k \leq 7$ levels of service. To see this, note that any feasible solution for an instance I with indivisible agents, i.e., $k = 1$, is also a feasible solution of the same instance I but with k levels of service. Finally, note that as $k \to \infty$, the optimal approximation guarantee of Greedy-and-Prune converges to $1 + \varphi \approx 2.62$, which was the previously best-known result for the divisible model by Klumper and Schäfer [2022].

5.3.2 Concave Valuations

In this section, we derive a mechanism, called Sort-&-Reject, for instances of the k-level model with concave and non-decreasing valuation functions. We first introduce our mechanism Sort-&-Reject and prove that it is monotone, after which we prove its approximation guarantee. The most challenging part is to prove budget-feasibility of the mechanism, to which the final subsection is dedicated.

Sort-&-Reject: A Truthful Mechanism

The main idea underlying our mechanism Sort-&-Reject (see Mechanism 10) is in similar spirit as Greedy-and-Prune: If there is an agent i^* whose maximum value $v_{i^*}(k)$ is high enough (in a certain sense), then we simply allocate all k levels of this agent and allocate no levels of any other agent. We define the most valuable agent i^* in a similar way (see Line 1 of Sort-&-Reject), and again compare the value of all k levels of i^* to some predefined threshold.

Otherwise, we compute an allocation using the following greedy procedure: We first compute an optimal allocation $\boldsymbol{x}^*(\boldsymbol{c})$ to the corresponding (non-strategic) fractional k-bounded knapsack problem, and use the integral part of this solution as an initial allocation. The intuition is that this allocation is close to the optimal fractional solution because $\boldsymbol{x}^*(\boldsymbol{c})$ has at most one fractional component (Fact 5.2.2). We then repeatedly discard the worst level of service (in terms of marginal value-per-cost) of an agent from this allocation, until the total value of our allocation would drop below some other predefined threshold, i.e., an α-fraction of the optimal solution $\boldsymbol{x}^*(\boldsymbol{c})$. Finally, the payments are computed according to (5.1) in Theorem 5.2.1.

We need some more notation for the formal description of our mechanism: Given an allocation $\boldsymbol{x}$, we denote by $\ell(\boldsymbol{x})$ the agent whose $x_{\ell(\boldsymbol{x})}$-th level of service is the least valuable in $\boldsymbol{x}$, in terms of their 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 $x_{\ell(\boldsymbol{x})}$-th ratio of agent $\ell(\boldsymbol{x})$. When $\boldsymbol{x}$ is clear from the context, we refer to this agent simply as ℓ.

The main result of this section is the following theorem:

5

Theorem 5.3.6. *Consider the class of instances of the k-level model with concave and non-decreasing valuation functions. Then, for $\alpha = (2+\sqrt{3})^{-1}$, the mechanism Sort-&-Reject is truthful, individually rational, budget-feasible and $(2+\sqrt{3})$-approximate.*

We now prove several lemmas to establish the properties stated in Theorem 5.3.6. We first show that Sort-&-Reject is monotone and establishes the approximation guarantee mentioned in Theorem 5.3.6 for certain ranges of the input parameter α.

The following property directly follows by construction of the mechanism.

Fact 5.3.7. Given an instance $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B, k)$ of the k-level model. Let $\boldsymbol{x}$ be computed by Sort-&-Reject(I) and let $\boldsymbol{x}^*$ be an optimal solution to the fractional k-bounded knapsack problem of I computed by Algorithm 8, using similar tie-breaking. Then, $x_i \leq x_i^*$ for every $i \in N$.

We now show that the allocation rule of Sort-&-Reject is monotone.

MECHANISM 10: SORT-&-REJECT(I)

Input: An instance $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B, k)$.
Output: An allocation vector $\boldsymbol{x}$ and a payment vector $\boldsymbol{p}$.

1 let $i^* \in \arg\max_{i\in N} v_i(k)/\, \texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i})$
2 **if** $v_{i^*}(k) \geq \frac{\alpha}{1-\alpha} \cdot \texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i^*})$ **then**
3   set $x_{i^*} = 1$ and $x_i = 0$ for $i \in N \setminus \{i^*\}$
4 **end**
5 **else**
6   compute an optimal fractional allocation $\boldsymbol{x}^*(\boldsymbol{c})$ using Algorithm 8
7   initialize $\boldsymbol{x} = (\lfloor x_1^*(\boldsymbol{c})\rfloor, \dots, \lfloor x_n^*(\boldsymbol{c})\rfloor)$
8   **for** $i \in W(\boldsymbol{x})$ **do**
9   **for** $j \in \{1, \dots, x_i\}$ **do**
10   add the marginal value-per-cost ratio $m_i(j)/c_i$ to a list $\mathcal{L}$
11   **end**
12   **end**
13   sort $\mathcal{L}$ in non-increasing order and let ℓ be the agent corresponding to the last marginal value-per-cost in $\mathcal{L}$
14   **while** $v(\boldsymbol{x}) - m_\ell(x_\ell) \geq \alpha\, \texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$ **do**
15   set $x_\ell = x_\ell - 1$
16   remove the last element from $\mathcal{L}$ and update ℓ
17   **end**
18 **end**
19 **for** $i \in N$ **do**
20   compute payment p_i according to (5.1)
21 **end**
22 **return** $(\boldsymbol{x}, \boldsymbol{p})$

Lemma 5.3.8. *For $\alpha \in (0, 1)$, the allocation rule of* SORT-&-REJECT *is monotone.*

Proof: Let $\boldsymbol{c}$ be a cost profile. Note that the input parameter α does not depend on $\boldsymbol{c}$. We distinguish the following two cases.

Case 1: $v_{i^*}(k) \geq \frac{\alpha}{1-\alpha} \cdot \texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i^*})$. In this case, only agent i^* has a positive allocation ($x_{i^*}(\boldsymbol{c}) = k$). Suppose that i^* unilaterally deviates and decreases their cost to $c'_{i^*} < c_{i^*}$. Such a deviation has no influence on which agent is chosen in Line 1 as i^* and does not alter the if condition in Line 2. To see that, notice that the value $\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i^*})$ does not depend on the bid of i^*, whereas $\texttt{opt}_{\mathrm{F}}^{k}\left(\left(c'_{i^*}, \boldsymbol{c}_{-i^*}\right)_{-i}\right) \geq \texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i})$, for any $i \neq i^*$. Therefore, for any such deviation, i^* will remain the sole winner and will be hired for k levels of service, i.e., $x_{i^*}(c'_{i^*}, \boldsymbol{c}_{-i^*}) = x_{i^*}(\boldsymbol{c}) = k$. No other agent was winning in this case, hence there is no need to examine deviations by other agents.

Case 2: $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \cdot \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i^*})$. Here, the allocation rule of SORT-&-REJECT can allocate to multiple agents, which are referred to as the set of winning agents $W(\boldsymbol{x}(\boldsymbol{c}))$. Note that because of our cost restriction $c_i \leq B/k$ for all $i \in N$, $W(\boldsymbol{x}(\boldsymbol{c}))$ starts as a non-empty set and our while condition in Line 14 guarantees that it remains so. Fix an agent $i \in W(\boldsymbol{x}(\boldsymbol{c}))$ and suppose i unilaterally deviates and declares a cost of $c_i' < c_i$. First of all, note that for every $j \in N \setminus \{i\}$ it holds that:

$$\frac{v_j(k)}{\mathtt{opt}_{\mathrm{F}}^k(c_i', \boldsymbol{c}_{-\{j,i\}})} \leq \frac{v_j(k)}{\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-j})} \leq \frac{v_{i^*}(k)}{\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i^*})},$$

whereas for i we have that:

$$\frac{v_i(k)}{\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i})} \leq \frac{v_{i^*}(k)}{\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i^*})}.$$

Therefore, even if some agent $\hat{\imath} \neq i^*$ is chosen in Line 1 under the profile $(c_i', \boldsymbol{c}_{-i})$, SORT-&-REJECT will always execute the else case as the if condition in Line 2 remains false. Moreover, by the way $\boldsymbol{x}^*(c_i', \boldsymbol{c}_{-i})$ is constructed, the corresponding initial integral solution $\boldsymbol{x}$ in Line 7 will have an i-th coordinate at least as large as before. Now consider what happens to the while condition of the mechanism. On the one hand, it can only be that $\mathtt{opt}_{\mathrm{F}}^k(c_i', \boldsymbol{c}_{-i}) \geq \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c})$. On the other hand, the marginal value-per-cost ratios of agent i under the profile $(c_i', \boldsymbol{c}_{-i})$ can only have a better position in the ordering constructed by SORT-&-REJECT. Therefore, agent i will be hired to at least the same extent or more, i.e., for the final allocations we have $x_i(c_i', \boldsymbol{c}_i) \geq x_i(\boldsymbol{c})$, proving monotonicity. □

5

Since the payments are computed according to (5.1), we conclude that the mechanism SORT-&-REJECT is truthful and individually rational. We continue by showing that SORT-&-REJECT achieves the claimed approximation guarantee.

Lemma 5.3.9. *Given an instance $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B, k)$ of the k-level model. Let $\alpha \in (0, (3-\sqrt{5})/2]$. Then, SORT-&-REJECT is $(1/\alpha)$-approximate, i.e., for the allocation $\boldsymbol{x}$ computed by SORT-&-REJECT(I), it holds that $v(\boldsymbol{x}) \geq \alpha\, \mathtt{opt}^k(\boldsymbol{c})$.*

Proof: Given a cost profile $\boldsymbol{c}$, let $\boldsymbol{x}(\boldsymbol{c})$ be the allocation computed by the mechanism SORT-&-REJECT. We will prove the claimed guarantee against the optimal value of the fractional k-bounded knapsack problem of the instance, i.e., we will show that $v(\boldsymbol{x}(\boldsymbol{c})) \geq \alpha\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c})$. This establishes our guarantee since $\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}) \geq \mathtt{opt}^k(\boldsymbol{c})$. Again we distinguish the following two cases.

Case 1: $v_{i^*}(k) \geq \frac{\alpha}{1-\alpha} \cdot \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i^*})$. We directly have that:

$$v_{i^*}(k) \geq \frac{\alpha}{1-\alpha} \cdot \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i^*}) \geq \frac{\alpha}{1-\alpha}\big(\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}) - v_{i^*}(k)\big),$$

where the second inequality follows from the fact that $\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i}) + v_i(k) \geq \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c})$ for all $i \in N$. Rearranging terms leads to $v_{i^*}(k) \geq \alpha\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c})$ which concludes the case since $v_{i^*}(k) = v(\boldsymbol{x}(\boldsymbol{c}))$.

Case 2: $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \cdot \mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i^*})$. In this case, whenever the while loop in Line 14 of SORT-&-REJECT runs at least once, we immediately obtain that $v(\boldsymbol{x}(\boldsymbol{c})) \geq \alpha\, \mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$ for any $\alpha < 1$. Therefore, consider an instance for which the while condition evaluates to false the first time it is checked. We argue that, for $\alpha \leq \frac{3-\sqrt{5}}{2}$, we still obtain that $v(\boldsymbol{x}(\boldsymbol{c})) \geq \alpha\, \mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$ and the lemma follows. Towards a contradiction, suppose that this is not the case. Let f be the last agent of which a level of service is (fractionally) included in $\boldsymbol{x}^*(\boldsymbol{c})$ by Algorithm 8. Using Fact 5.2.2, this implies that:

$$\begin{aligned}\alpha\, \mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) > v(\boldsymbol{x}(\boldsymbol{c})) &= \mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) - m_f(\lceil x_f^* \rceil)\, \frac{B - \sum_{i \in N} \lfloor x_i^* \rfloor c_i}{c_f} \\ &\geq \mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) - m_f(\lceil x_f^* \rceil)\,,\end{aligned}$$

which, by rearranging terms, yields:

$$m_f(\lceil x_f^* \rceil) > (1-\alpha)\, \mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})\,. \tag{5.7}$$

At the same time, we have:

$$\frac{m_f(\lceil x_f^* \rceil)}{\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})} \leq \frac{v_f(k)}{\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-f})} \leq \frac{v_{i^*}(k)}{\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i^*})} < \frac{\alpha}{1-\alpha}\,, \tag{5.8}$$

where the first inequality follows from straightforward bounds on both the numerator and the denominator The second inequality holds by the choice of i^* and the third inequality follows by assumption (Case 2).

5

By combining inequalities (5.7) and (5.8), we obtain that α must be such that $\frac{\alpha}{1-\alpha} > 1 - \alpha$. However, this is not satisfied for $\alpha \in \left(0, \frac{3-\sqrt{5}}{2}\right]$, resulting in a contradiction. □

Budget-Feasibility of SORT-&-REJECT

It remains to be proven that SORT-&-REJECT is budget-feasible. We introduce some auxiliary notation: Consider a cost profile $\boldsymbol{c}$ and an agent $i \in W(\boldsymbol{x}(\boldsymbol{c}))$. Let $j \in \{1, \ldots, x_i(\boldsymbol{c})\}$ be any level of service among the ones allocated to agent i. Intuitively, we refer to the *critical payment* $p_{ij}(\boldsymbol{c}_{-i})$ of the j-th level of service of agent i as the largest cost that i can declare such that their j-th level of service is still allocated by SORT-&-REJECT (see Figure 5.2 for an illustrative example). More formally, we define $Q_{ij}(\boldsymbol{c}_{-i})$ as the set of all points q satisfying $\lim_{z \to q^-} x_i(z, \boldsymbol{c}_{-i}) \geq j$ and $\lim_{z \to q^+} x_i(z, \boldsymbol{c}_{-i}) \leq j$ and let $p_{ij}(\boldsymbol{c}_{-i}) = \sup(Q_{ij}(\boldsymbol{c}_{-i}))$. Such a number q must always exist. To see the latter, note that $x_i(c_i, \boldsymbol{c}_{-i}) \geq j$, which implies that $c_i \leq q$, and that $x_i(z, \boldsymbol{c}_{-i}) = 0 < j$ for all $z > \frac{B}{k}$ (we assume that any agent declaring a cost greater than $\frac{B}{k}$ is discarded up front in Section 5.2.2), which implies that $q \leq \frac{B}{k}$. Given that $Q_{ij}(\boldsymbol{c}_{-i})$ is nonempty and bounded

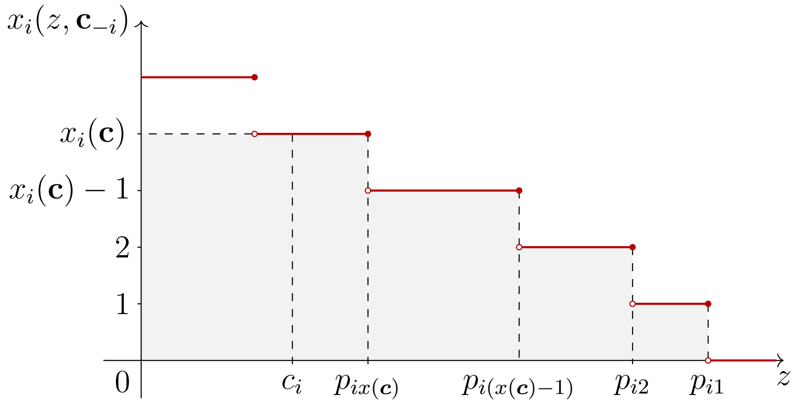

Figure 5.2: Illustrative example of the critical payments of an agent i for a cost profile $\boldsymbol{c}$.

from above, $p_{ij}(\boldsymbol{c}_{-i})$ always exists. It is easy to see that the payment of an agent i can be written as the sum over these critical payments of the levels of service allocated to i.

Lemma 5.3.10. *Let $\mathbf{c}$ be a cost profile and assume $i \in W(\boldsymbol{x}(\boldsymbol{c}))$. It holds that:*

$$p_i(\boldsymbol{c}) = \sum_{j=1}^{x_i(\boldsymbol{c})} p_{ij}(\boldsymbol{c}_{-i}).$$

Proof: Starting with (5.1) from Theorem 5.2.1, we obtain that:

$$\begin{aligned} p_i(\boldsymbol{c}) &= c_i x_i(\boldsymbol{c}) + \int_{c_i}^{\infty} x_i(z, \boldsymbol{c}_{-i}) dz \\ &= c_i x_i(\boldsymbol{c}) + \sum_{j=1}^{x_i(\boldsymbol{c})-1} \int_{p_{i(j+1)}(\boldsymbol{c}_{-i})}^{p_{ij}(\boldsymbol{c}_{-i})} j dz + \int_{c_i}^{p_{ix_i(\boldsymbol{c})}(\boldsymbol{c}_{-i})} x_i(\boldsymbol{c}) dz = \sum_{j=1}^{x_i(\boldsymbol{c})} p_{ij}(\boldsymbol{c}_{-i}) \,. \end{aligned}$$

□

Lemma 5.3.11 is the main technical tool needed to establish budget-feasibility for the else part of the Sort-&-Reject mechanism.

Lemma 5.3.11. *Let $\alpha \in (0,1)$ and let $\mathbf{c}$ be a cost profile of an instance I with $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i^*})$. Let ℓ be the index of the agent with smallest marginal value-per-cost ratio in $\mathcal{L}$ when Sort-&-Reject terminates. Then:*

$$\sum_{i=1}^{n} p_i(\boldsymbol{c}) < \frac{B}{1-\alpha} \left(\frac{m_\ell(x_\ell(\boldsymbol{c}))}{\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-\ell})} + \frac{\alpha}{1-\alpha} \right) .$$

5

Using Lemma 5.3.11, we can determine a range for values of α for which Sort-&-Reject is budget-feasible.

Lemma 5.3.12. *Let* $\alpha \in \left(0, \frac{1}{2+\sqrt{3}}\right]$. *Then, Sort-&-Reject is budget-feasible.*

Proof: Let $\boldsymbol{c}$ be a cost profile. We again consider the following two cases.
Case 1: $v_{i^*}(k) \geq \frac{\alpha}{1-\alpha}\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i^*})$. In this case only agent i^* has a positive allocation and is allocated k levels of service. By Lemma 5.3.10, we obtain:

$$\sum_{i=1}^{n} p_i(\boldsymbol{c}) = p_{i^*}(\boldsymbol{c}) = \sum_{j=1}^{k} p_{i^*j}(\boldsymbol{c}_{-i^*}) = k \cdot p_{i^*k}(\boldsymbol{c}_{-i^*}) \leq k\,\frac{B}{k} = B\,.$$

The third equality follows from the definition of the critical payments and from the fact that the allocation of i^* will only be smaller than k when another agent becomes i^* or when $c_i' > \frac{B}{k}$. The inequality follows from the fact that $p_{i^*k}(\boldsymbol{c}_{-i^*}) \leq \frac{B}{k}$ as argued above.
Case 2: $v_{i^*}(k) < \frac{\alpha}{1-\alpha}\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i^*})$. In this case, we invoke Lemma 5.3.11 leading to:

$$\begin{aligned}\sum_{i=1}^{n} p_i(\boldsymbol{c}) &< \frac{B}{1-\alpha}\left(\frac{m_\ell(x_\ell(\boldsymbol{c}))}{\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-\ell})} + \frac{\alpha}{1-\alpha}\right) \leq \frac{B}{1-\alpha}\left(\frac{v_\ell(k)}{\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-\ell})} + \frac{\alpha}{1-\alpha}\right)\\ &\leq \frac{B}{1-\alpha}\left(\frac{v_{i^*}(k)}{\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i^*})} + \frac{\alpha}{1-\alpha}\right) < \frac{B}{1-\alpha}\left(\frac{\alpha}{1-\alpha} + \frac{\alpha}{1-\alpha}\right)\\ &= B\frac{2\alpha}{(1-\alpha)^2}\,.\end{aligned}$$

The second inequality follows by observing that $m_\ell(x_\ell(\boldsymbol{c})) \leq v_\ell(k)$ since $v_\ell(\cdot)$ is non-decreasing. The next two inequalities are due to the definition of i^* and by assumption (Case 2). To obtain budget-feasibility, we must ensure that $\frac{2\alpha}{(1-\alpha)^2} \leq 1$ or, equivalently, that $\alpha \leq \frac{1}{2+\sqrt{3}}$. This concludes the proof. □

Theorem 5.3.6 now follows by Lemma 5.3.12, together with Lemmas 5.3.8 and 5.3.9. The rest of this section is devoted to proving Lemma 5.3.11. We start by presenting three auxiliary statements, which will prove to be useful in our analysis. The purpose of these statements is to characterize and give upper bounds on the individual payments of winning agents, whenever the else part of Sort-&-Reject is executed. We begin with Lemma 5.3.13, in which we derive an upper bound on the costs of winning agents.

Lemma 5.3.13. *Let* $\alpha \in (0,1)$ *and let* $\boldsymbol{c}$ *be a cost profile of an instance* I *with* $v_{i^*}(k) < \frac{\alpha}{1-\alpha}\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i^*})$. *Let* ℓ *be the index of the agent with smallest marginal value-per-cost ratio in* $\mathcal{L}$ *when Sort-&-Reject terminates. Then:*

$$c_\ell < \frac{B}{1-\alpha} \cdot \frac{m_\ell(x_\ell(\boldsymbol{c}))}{\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})}\,.$$

5

Proof: Let $\boldsymbol{x}$ be the allocation vector computed by Sort-&-Reject for the cost profile $\boldsymbol{c}$ and let $\boldsymbol{x}^* := \boldsymbol{x}^*(\boldsymbol{c})$ for brevity. Observe that, since by assumption $v_{i^*}(k) < \frac{\alpha}{1-\alpha}\,\mathtt{opt}^k_{\mathrm{F}}(\boldsymbol{c}_{-i^*})$, it must be that $v(\boldsymbol{x}) - m_\ell(x_\ell) < \alpha\,\mathtt{opt}^k_{\mathrm{F}}(\boldsymbol{c})$, which implies that:

$$\mathtt{opt}^k_{\mathrm{F}}(\boldsymbol{c}) - v(\boldsymbol{x}) + m_\ell(x_\ell) > (1-\alpha)\,\mathtt{opt}^k_{\mathrm{F}}(\boldsymbol{c})\,. \tag{5.9}$$

By Fact 5.2.2, there exists at most one agent in $W(\boldsymbol{x}^*)$ with a non-integer allocation. We denote that agent by f (if no such agent exists, let f be an arbitrary agent in $W(\boldsymbol{x}^*)$). Note that it is possible that $f = \ell$. We have:

$$\begin{aligned}
B &\geq \sum_{i\in W(\boldsymbol{x}^*)} c_i x_i^* \geq c_\ell(x_\ell^* - x_\ell + 1) + \sum_{i\in W(\boldsymbol{x}^*)\setminus\{\ell\}} c_i(x_i^* - x_i) = \sum_{j=x_\ell}^{\lfloor x_\ell^*\rfloor} \frac{c_\ell}{m_\ell(j)} m_\ell(j)\\
&+ \sum_{i\in W(\boldsymbol{x}^*)\setminus\{\ell\}} \sum_{j=x_i+1}^{\lfloor x_i^*\rfloor} \frac{c_i}{m_i(j)} m_i(j) + \frac{c_f}{m_f(\lceil x_f^*\rceil)}(x_f^* - \lfloor x_f^*\rfloor) m_f(\lceil x_f^*\rceil)\\
&\geq \frac{c_\ell}{m_\ell(x_\ell)}\left(\sum_{j=x_\ell}^{\lfloor x_\ell^*\rfloor} m_\ell(j) + \sum_{i\in W(\boldsymbol{x}^*)\setminus\{\ell\}} \sum_{j=x_i+1}^{\lfloor x_i^*\rfloor} m_i(j) \quad + \quad (x_f^* - \lfloor x_f^*\rfloor) m_f(\lceil x_f^*\rceil)\right)\\
&= \frac{c_\ell}{m_\ell(x_\ell)}\left(m_\ell(x_\ell) + \sum_{i\in W(\boldsymbol{x}^*)} \sum_{j=x_i+1}^{\lfloor x_i^*\rfloor} m_i(j) \quad + \quad (x_f^* - \lfloor x_f^*\rfloor) m_f(\lceil x_f^*\rceil)\right)\\
&= \frac{c_\ell}{m_\ell(x_\ell)}\left(\mathtt{opt}^k_{\mathrm{F}}(\boldsymbol{c}) - v(\boldsymbol{x}) + m_\ell(x_\ell)\right) > \frac{c_\ell}{m_\ell(x_\ell)}(1-\alpha)\,\mathtt{opt}^k_{\mathrm{F}}(\boldsymbol{c})\,.
\end{aligned}$$

The first inequality follows by the feasibility of $\boldsymbol{x}^*$, whereas the second inequality follows by the fact that $x_\ell \geq 1$ and, thus, every cost now has a smaller or equal coefficient than before. The next inequality is due to the marginal value-per-cost ordering that Algorithm 8 uses to built $\boldsymbol{x}^*$ and that Sort-&-Reject uses (in the reverse order) to obtain $\boldsymbol{x}$. Finally, the last inequality is due to (5.9). The lemma follows by rearranging terms. □

We now proceed to obtain an upper bound on the payments each winning agent receives for each level of service.

Lemma 5.3.14. *Let $\alpha \in (0,1)$ and let $\boldsymbol{c}$ be a cost profile of an instance I with $v_{i^*}(k) < \frac{\alpha}{1-\alpha}\,\mathtt{opt}^k_{\mathrm{F}}(\boldsymbol{c}_{-i^*})$. Moreover, let $\boldsymbol{x}(\boldsymbol{c})$ be the allocation computed by Sort-&-Reject and let $i \in W(\boldsymbol{x}(\boldsymbol{c}))$. For $1 \leq j \leq x_i(\boldsymbol{c})$, it holds that:*

$$p_{ij}(\boldsymbol{c}_{-i}) \leq \frac{B}{1-\alpha} \cdot \frac{m_i(j)}{\mathtt{opt}^k_{\mathrm{F}}(\boldsymbol{c}_{-i})}\,.$$

Proof: Let $p \in Q_{ij}(\boldsymbol{c}_{-i})$, i.e., p is a cost declaration that guarantees agent i at least j levels of service, and let λ be the index of the agent with the smallest marginal

5

value-per-cost ratio in the allocation computed by SORT-&-REJECT with input $(p, \boldsymbol{c}_{-i})$ (as opposed to ℓ when the input is $\boldsymbol{c}$). By the definition of p and λ, we have $m_i(j)/p \geq m_\lambda(x_\lambda(p, \boldsymbol{c}_{-i}))/c_\lambda$. By rearranging and applying Lemma 5.3.13 for the profile $(p, \boldsymbol{c}_{-i})$ and λ, we get:

$$p \leq \frac{c_\lambda m_i(j)}{m_\lambda(x_\lambda(p, \boldsymbol{c}_{-i}))} < \frac{B}{1-\alpha} \cdot \frac{m_i(j)}{\mathtt{opt}_{\mathrm{F}}^k(p, \boldsymbol{c}_{-i})} \leq \frac{B}{1-\alpha} \cdot \frac{m_i(j)}{\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i})},$$

where the last inequality follows by the monotonicity of $\mathtt{opt}_{\mathrm{F}}^k(\cdot)$ with respect to subinstances. This implies that:

$$p_{ij}(\boldsymbol{c}_{-i}) = \sup(Q_{ij}(\boldsymbol{c}_{-i})) \leq \frac{B}{1-\alpha} \cdot \frac{m_i(j)}{\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i})},$$

as claimed. □

The final component needed for the proof of Lemma 5.3.11 is a lower bound on the value of an optimal solution to the fractional k-bounded knapsack problem when one agent is excluded.

5

Lemma 5.3.15. *Let $\alpha \in (0,1)$ and let $\mathbf{c}$ be a cost profile of an instance I with $v_{i^*}(k) < \frac{\alpha}{1-\alpha}\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i^*})$. Moreover, let $\boldsymbol{x}(\boldsymbol{c})$ be computed by SORT-&-REJECT and let ℓ be the index of the agent with smallest marginal value-per-cost ratio in $\mathcal{L}$ when SORT-&-REJECT terminates. For every agent $i \in N$ it holds that:*

$$\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i}) > \frac{1-\alpha}{\alpha}\left(v(\boldsymbol{x}(\boldsymbol{c})) - m_\ell(x_\ell)\right).$$

Proof: By the stopping condition of the while loop in Line 14, we have:

$$\begin{aligned} v(\boldsymbol{x}(\boldsymbol{c})) - m_\ell(x_\ell) &< \alpha\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}) \leq \alpha \left(\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i}) + v_i(k)\right) \\ &\leq \alpha\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i}) \left(1 + \frac{v_{i^*}(k)}{\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i^*})}\right) \\ &< \alpha\, \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i}) \left(1 + \frac{\alpha}{1-\alpha}\right) = \frac{\alpha}{1-\alpha}\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i}). \end{aligned}$$

The third inequality follows by the definition of i^*, whereas the last inequality follows directly by the assumption of the lemma. □

We finally present the proof of Lemma 5.3.11.

Proof of Lemma 5.3.11: We can upper bound the total payments as follows:

$$
\begin{aligned}
&\sum_{i=1}^{n} p_i(\boldsymbol{c}) = \sum_{i \in W(\boldsymbol{x}(\boldsymbol{c}))} p_i(\boldsymbol{c}) \;=\; \sum_{i \in W(\boldsymbol{x}(\boldsymbol{c}))} \sum_{j=1}^{x_i(\boldsymbol{c})} p_{ij}(\boldsymbol{c}_{-i}) \\
&\leq \frac{B}{1-\alpha} \sum_{i \in W(\boldsymbol{x}(\boldsymbol{c}))} \sum_{j=1}^{x_i(\boldsymbol{c})} \frac{m_i(j)}{\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i})} = \frac{B}{1-\alpha} \sum_{i \in W(\boldsymbol{x}(\boldsymbol{c}))} \frac{v_i(x_i(\boldsymbol{c}))}{\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i})} \\
&= \frac{B}{1-\alpha} \left(\sum_{i \in W(\boldsymbol{x}(\boldsymbol{c})) \setminus \{\ell\}} \frac{v_i(x_i(\boldsymbol{c}))}{\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i})} \;+\; \frac{v_\ell(x_\ell(\boldsymbol{c})) - m_\ell(x_\ell(\boldsymbol{c}))}{\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-\ell})} \;+\; \frac{m_\ell(x_\ell(\boldsymbol{c}))}{\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-\ell})} \right) \\
&< \frac{B}{1-\alpha} \left(\frac{\alpha}{1-\alpha} \cdot \frac{\sum_{i \in W(\boldsymbol{x}(\boldsymbol{c}))} v_i(x_i(\boldsymbol{c})) - m_\ell(x_\ell(\boldsymbol{c}))}{v(x(\boldsymbol{c})) - m_\ell(x_\ell(\boldsymbol{c}))} + \frac{m_\ell(x_\ell(\boldsymbol{c}))}{\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-\ell})} \right) \\
&= \frac{B}{1-\alpha} \left(\frac{\alpha}{1-\alpha} + \frac{m_\ell(x_\ell(\boldsymbol{c}))}{\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-\ell})} \right) .
\end{aligned}
$$

The first equality reflects the fact that $p_{ij}(\boldsymbol{c}_{-i}) = 0$ for $i \notin W(\boldsymbol{x}(\boldsymbol{c})$, whereas the second equality is due to Lemma 5.3.10. The first inequality follows by applying Lemma 5.3.14 for every agent $i \in W(\boldsymbol{x}(\boldsymbol{c}))$ and every $j \in \{1 \ldots, x_i(\boldsymbol{c})\}$. Finally, the second inequality follows by applying Lemma 5.3.15 to the denominators of all the terms except from $m_\ell(x_\ell(\boldsymbol{c}))/\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-\ell})$. □

5

5.4 Mechanisms for Divisible Agents

We now turn our attention to the divisible agent model: recall that in this model the (declared) cost represents the cost for the full service of an agent. In Section 5.4.1, we introduce our mechanism Prune-and-Assign that is specifically tailored to linear valuation functions and is inspired by Gravin et al. [2020]. Prune-and-Assign improves the best-known approximation of $1+\varphi \approx 2.62$ by Klumper and Schäfer [2022] to 2. Then, in Section 5.4.2, we consider concave valuation functions and improve upon the first constant factor approximation that we derived in [Amanatidis et al., 2025]. As in [Amanatidis et al., 2025], our mechanism Chunk-and-Solve first reduces an instance of the divisible model to an instance of the k-level model through discretization. If the discretization step is equal to the number of agents n, i.e., $k = n$, then Chunk-and-Solve is $2(1+\varphi)$-approximate. With a more involved analysis, we finally show in Section 5.4.3 that Chunk-and-Solve is $(1+\varphi+\varepsilon)$-approximate if a large enough discretization step is chosen.

MECHANISM 11: PRUNING(I) ([Gravin et al., 2020])

Input: An instance $I = (N, \boldsymbol{c}, (v_i)_{i \in N}, B)$ with $\frac{v_1}{c_1} \geq \cdots \geq \frac{v_n}{c_n}$.
Output: A provisional allocation $\boldsymbol{x}$ and a value-over-cost rate r.

set $r = \max\{v_i \mid i \in N\}/B$
for $i \in N$ **do**
  set $\bar{x}_i = 1$ if $\frac{v_i}{c_i} \geq r$ and $\bar{x}_i = 0$ otherwise
end
set $\ell = \arg\max\{i \mid \bar{x}_i = 1\}$
while $rB < \sum_{i=1}^{\ell} v_i - \max\{v_i \mid 1 \leq i \leq \ell\}$ **do**
  continuously increase rate r
  if $\frac{v_\ell}{c_\ell} \leq r$ **then**
    set $\bar{x}_\ell = 0$ and $\ell = \ell - 1$
  end
end
return $(\bar{\boldsymbol{x}}, r)$

5.4.1 Linear Valuations

In this section, we present a budget-feasible mechanism that is 2-approximate for the divisible model with linear valuations, thereby establishing that the optimal approximation guarantee for this problem lies in the interval $[e/(e-1), 2]$. To avoid unnecessarily heavy notation, let $v_i := v_i(1)$ for all $i \in N$, and thus, $v_i(x) = x \cdot v_i$ for $x \in [0, 1]$.

Phase 1: Pruning Mechanism

The PRUNING mechanism of Gravin et al. [2020], which constitutes as a crucial building block for both their deterministic 3-approximate mechanism and their randomized 2-approximate mechanism for indivisible agents, and it will turn out to be a useful starting point for the divisible setting as well.

Given a profile $\boldsymbol{c}$, PRUNING (see Mechanism 11) computes an allocation $\bar{\boldsymbol{x}}(\boldsymbol{c})$, which we refer to as the *provisional allocation*, and a positive quantity $r(\boldsymbol{c})$, which we refer to as the *rate*. We assume that the input of PRUNING, i.e., an instance I, already has the agents relabeled by their decreasing value-per-cost ratio, i.e., $\frac{v_1}{c_1} \geq \frac{v_2}{c_2} \geq \cdots \geq \frac{v_n}{c_n}$. It has been showed that PRUNING is monotone, and that an even stronger property holds: each agent i that is a winner in the provisional allocation, i.e., $\bar{x}_i = 1$, cannot alter the outcome of PRUNING unilaterally while remaining a winner in the provisional allocation.

Lemma 5.4.1: (implied by Lemma 3.1 in [Gravin et al., 2020]). *Let $\boldsymbol{c}$ be a cost profile. Consider an agent $i \in N$ with $\bar{x}_i(\boldsymbol{c}) = 1$. Then, for all c'_i such that $\bar{x}_i(c'_i, \boldsymbol{c}_{-i}) = 1$, it holds that $\bar{\boldsymbol{x}}(c'_i, \boldsymbol{c}_{-i}) = \bar{\boldsymbol{x}}(\boldsymbol{c})$ and $r(c'_i, \boldsymbol{c}_{-i}) = r(\boldsymbol{c})$.*

Given this property, Pruning can be used as a first filtering step to discard inefficient agents, followed by a subsequent allocation scheme which takes $(\bar{\boldsymbol{x}}(\boldsymbol{c}), r(\boldsymbol{c}))$ as input. If the subsequent allocation scheme is monotone, then the sequential composition of Pruning with this allocation scheme is monotone as well. This *composability property* is proven in Lemma 3.1 in [Gravin et al., 2020].

Let $(\bar{\boldsymbol{x}}, r)$ be the output of Pruning for a cost profile $\boldsymbol{c}$. Given $\bar{\boldsymbol{x}}$, we define S as the set of agents that are provisionally allocated, i^* as the agent in S of highest value (where ties are broken lexicographically), and T as the set of remaining agents. Formally:

$$S = \{i \in N \mid \bar{x}_i = 1\}, \quad i^* \in \arg\max\{v_i \mid i \in S\}, \quad \text{and} \quad T = S \setminus \{i^*\}\,. \tag{5.10}$$

Note that the definitions of S, i^* and T depend on $\bar{\boldsymbol{x}}$ (and thus the cost profile $\boldsymbol{c}$). For notational convenience, we do not state this reference explicitly if it is clear from the context. The following properties will be useful in our analysis.

Lemma 5.4.2: (Lemma 3.2 in [Gravin et al., 2020]). *Given a profile $\boldsymbol{c}$, let $(\bar{\boldsymbol{x}}, r)$ be the output of Pruning. Let $S = T \cup \{i^*\}$ be defined as in* (5.10) *with respect to $\bar{\boldsymbol{x}}$. We have:*

(i) $c_i \le \frac{v_i}{r} \le B$ *for all* $i \in S$,

(ii) $v(T) \le rB < v(S)$, *and*

(iii) $\mathtt{opt}_{\mathrm{F}}(\boldsymbol{c}) \le v(S) + r \cdot (B - c(S))$.

Phase 2: The Allocation Scheme

Our mechanism combines Pruning with the allocation scheme defined in (5.12) below. We refer to the resulting mechanism as *Prune-and-Assign* (see Mechanism 12). In order to define the allocation scheme, we need the following constants:

$$q = \frac{1}{2}\frac{v(S) - rB}{\min\{v_{i^*}, v(T)\}}, \qquad q_{i^*} = \begin{cases} \frac{1}{2} - q, & \text{if } v_{i^*} \le v(T), \\ \frac{1}{2}, & \text{otherwise,} \end{cases}$$

and $q_i = 1 - q_{i^*} - q$ for all $i \in T$. Note that for all agents $i \in T$, the constant q_i is the same. It is not hard to prove that $q \in \left[0, \frac{1}{2}\right]$ (see Lemma 5.1 in [Gravin et al., 2020]). The constants above are chosen so that:

$$\frac{rB}{2} = q_{i^*} v_{i^*} + (1 - q_{i^*} - q) v(T)\,. \tag{5.11}$$

Given r and S, we define the (fractional) allocation function $x_i(\boldsymbol{c})$ for $i \in S$ in Line 3 of Prune-and-Assign, so that it only depends on agent i's declared cost c_i. We will slightly abuse notation and write $x_i(c_i)$ in what follows. For each agent $i \in S$ let:

$$x_i(z) = q_i + \frac{v_i - rz}{2v_i} \quad \text{for} \quad z \in \left[0, \frac{v_i}{r}\right]. \tag{5.12}$$

5

It is not hard to verify that $0 \leq x_i(z) \leq 1$, given the chosen parameters q and q_i for $i \in S$ as above. Furthermore, by property (i) of Lemma 5.4.2, the cost c_i of each agent $i \in S$ is at most $v_i/r \leq B$ and, therefore, $x_i(c_i)$ is well-defined.

MECHANISM 12: PRUNE-AND-ASSIGN(I)

Input: An instance $I = (N, \boldsymbol{c}, (v_i)_{i \in N}, B)$ with $\frac{v_1}{c_1} \geq \cdots \geq \frac{v_n}{c_n}$.
Output: An allocation vector $\boldsymbol{x}$ and a payment vector $\boldsymbol{p}$.

let $(\bar{\boldsymbol{x}}, r) = \text{PRUNING}(I)$
set $S = \{i \in N \mid \bar{x}_i = 1\},\ i^* \in \arg\max\{v_i \mid i \in S\},\ T = S \setminus \{i^*\}$
determine $x_i(c_i)$ for all $i \in S$ according to (5.12)
set $x_i(c_i) = 0$ for all $i \in N \setminus S$
compute the payments $\boldsymbol{p}$ according to (5.1)
return $(\boldsymbol{x}, \boldsymbol{p})$

Theorem 5.4.3. *Consider the class of instances of the divisible model with linear valuation functions. Then,* PRUNE-AND-ASSIGN *is a truthful, individually rational and budget-feasible mechanism that is* 2*-approximate.*

Given the lower bound of $1 + \sqrt{2} \approx 2.41$ by Chen et al. [2011] for this setting but with indivisible agents, Theorem 5.4.3 establishes a separation between the indivisible and divisible model in the case of linear valuation functions. The theorem follows by the following three lemmas.

Lemma 5.4.4. *PRUNE-AND-ASSIGN is monotone.*

Proof: Let $\boldsymbol{c}$ be a cost profile. As mentioned above, PRUNING is monotone and it suffices to show that the allocation scheme in (5.12) is monotone by the composability property of Gravin et al. [2020]. Fix an arbitrary agent $i \in N$ and let $c_i' < c_i$. We will prove that $x_i(c_i', \boldsymbol{c}_{-i}) \geq x_i(\boldsymbol{c})$.

Let $\bar{\boldsymbol{x}}(\boldsymbol{c})$ be the provisional allocation obtained from PRUNING. If $\bar{x}_i(\boldsymbol{c}) = 0$ the claim follows trivially. Therefore, assume that $\bar{x}_i(\boldsymbol{c}) = 1$. Because PRUNING is monotone, we have that $\bar{x}_i(c_i', \boldsymbol{c}_{-i}) = 1$, and therefore by Lemma 5.4.1 the output of PRUNING remains the same, i.e., $(\bar{\boldsymbol{x}}(\boldsymbol{c}), r(\boldsymbol{c})) = (\bar{\boldsymbol{x}}(c_i', \boldsymbol{c}_{-i}), r(c_i', \boldsymbol{c}_{-i}))$. In particular, this implies that the set S of agents remains the same, as well as the defined constants q and q_i for all $i \in S$, and the rate computed by PRUNING. And so, the allocation functions in (5.12) remain the same and the lemma follows by observing that the allocation functions are monotone non-increasing in z. $\square$

Lemma 5.4.5. *PRUNE-AND-ASSIGN is* 2*-approximate.*

Proof: Let $\boldsymbol{c}$ be a cost profile. By property *(iii)* of Lemma 5.4.2, we have:

$$\texttt{opt}_{\text{F}}(\boldsymbol{c}) \leq v(S) + r \cdot (B - c(S)) = \sum_{i \in S} (v_i - rc_i) + rB\,.$$

By the allocation rule of Prune-and-Assign, as defined in (5.12), we have for every $i \in S$:

$$x_i(z) = q_i + \frac{v_i - rz}{2v_i}, \quad \text{or, equivalently,} \quad v_i - rz = 2v_i(x_i(z) - q_i)\,.$$

Combining the two inequalities above, we obtain:

$$\begin{aligned}\mathtt{opt}_{\mathrm{F}}(\boldsymbol{c}) &\leq 2\sum_{i\in S} v_i x_i(c_i) - 2\sum_{i\in S} v_i q_i + rB \\ &= 2v(\boldsymbol{x}(\boldsymbol{c})) - 2(q_{i^*}v_{i^*} + (1 - q_{i^*} - q)v(T)) + rB = 2v(\boldsymbol{x}(\boldsymbol{c}))\,,\end{aligned}$$

where the first equality follows as $x_i(\boldsymbol{c}) = 0$ for all $i \notin S$ and by definition of q_i. The last equality follows from (5.11). This concludes the proof. □

Lemma 5.4.6. *Prune-and-Assign is budget-feasible.*

Proof: Let $\boldsymbol{c}$ be a cost profile. The payment of agent $i \in S$ is defined as:

$$p_i(\boldsymbol{c}) = c_i x_i(\boldsymbol{c}) + \int_{c_i}^{B} x_i(z, \boldsymbol{c}_{-i})dz = c_i x_i(c_i) + \int_{c_i}^{v_i/r} x_i(z)dz \leq \int_{0}^{v_i/r} x_i(z)dz\,,$$

where the second equality holds because Pruning ensures that $x_i(z, \boldsymbol{c}_{-i}) = 0$ for $z > v_i/r$. The inequality holds because Prune-and-Assign is monotone (Lemma 5.4.4).

Consider an agent $i \in S$. We have:

$$\begin{aligned}\int_0^{v_i/r} x_i(z)dz &= \int_0^{v_i/r}\left(q_i + \frac{v_i - rz}{2v_i}\right)dz = \frac{v_i}{r}\left(q_i + \frac{1}{2}\right) - \frac{r}{2v_i}\int_0^{v_i/r} z dz \\ &= \frac{v_i}{r}\left(q_i + \frac{1}{4}\right)\,.\end{aligned}$$

Summing over all agents, we obtain that the total payment is at most:

$$\begin{aligned}\sum_{i=1}^{n} p_i(\boldsymbol{c}) &= \sum_{i\in S} p_i(\boldsymbol{c}) \leq \sum_{i\in S}\frac{v_i}{r}\left(q_i + \frac{1}{4}\right) \\ &= \frac{1}{r}\left(q_{i^*}v_{i^*} + (1 - q_{i^*} - q)v(T)\right) + \frac{v(S)}{4r} = \frac{B}{2} + \frac{v(S)}{4r}\,.\end{aligned}$$

Here, the first equality follows as $x_i(\boldsymbol{c}) = 0$ for all $i \notin S$, and so $p_i(\boldsymbol{c}) = 0$ for all $i \notin S$. The last equality uses the identity in (5.11).

We conclude the proof by showing that $v(S)/4r \leq B/2$, or, equivalently, $rB \geq v(S)/2$. Starting with (5.11) and applying the definition of q_{i^*}, we have:

$$\begin{aligned}\frac{rB}{2} = q_{i^*}v_{i^*} + (1 - q_{i^*} - q)v(T) &= \frac{1}{2}(v_{i^*} + v(T)) - q\cdot\min\{v_{i^*}, v(T)\} \\ &\geq \frac{1}{2}v(S) - qv(T)\,.\end{aligned}$$

5

The lemma follows by observing that $v(T) \leq rB$ by property (ii) of Lemma 5.4.2 and $q \leq \frac{1}{2}$ by definition. □

It is worth noting that Theorem 5.4.3 also holds when the valuation functions are linear up to some threshold and remain constant beyond this threshold, i.e., when for each $i \in N$ there exists a $\gamma_i \in (0,1]$ such that $v_i(\cdot)$ is linear for $x \in [0, \gamma_i]$ and for $x \geq \gamma_i$ it holds that $v_i(x) = v_i(1)$. Namely, consider such an instance $I = (N, \boldsymbol{c}, (v_i)_{i \in N}, B)$. As for agents $i \in N$ with $\gamma_i < 1$, allocating more than a γ_i fraction does not yield additional value, we can reduce instance I to an instance $I' = (N, \boldsymbol{c}', (v_i')_{i \in N}, B)$ such that all valuation functions are linear for the entire range of $[0,1]$. In particular, for each $i \in N$ and $x \in [0,1]$ we can define $v_i'(x) = v_i(\gamma_i \cdot x)$ and define $c_i' = \gamma_i c_i$. We can then run our mechanism Prune-and-Assign for the instance I', while still running a pre-processing step that discards agents $i \in N$ that declare a cost above B for I, i.e., above $\gamma_i \cdot B$ for I'. If we then compute an allocation $\boldsymbol{x}'$ for the instance I', this can be translated to an allocation $\boldsymbol{x}$ for I with $x_i = \gamma_i \cdot x_i'$ for all $i \in N$. Note that a mechanism including such a transformation remains truthful, as γ_i is independent of the declared cost for all $i \in N$.

We conclude the section by showing that our analysis of Prune-and-Assign is tight. Consider an instance with 2 agents with $v_1 = v_2 = 1$, $c_1 = \varepsilon, c_2 = 1 - \varepsilon$, with $\varepsilon \in (0,1)$, and the budget of the auctioneer being $B = 1$. The optimal allocation is $\boldsymbol{x}^* = (1,1)$ with a value of $v(\boldsymbol{x}^*) = 2$. Consider the outcome of Prune-and-Assign for this instance. Note that Pruning does not discard either agent since for $r = 1/B$ the while condition immediately evaluates to false. Thus, $S = \{1,2\}$ and $r = 1/B$. Then, it is easy to observe that the allocation function of Prune-and-Assign prescribes an allocation of $x_1(\varepsilon) = \frac{1}{2} - \frac{\varepsilon}{2}$ and $x_2(1-\varepsilon) = \frac{1}{2} + \frac{\varepsilon}{2}$ (by viewing bidder 1 as i^*). Therefore $v(\boldsymbol{x}(\boldsymbol{c})) = 1$, which matches the approximation guarantee of Theorem 5.4.3.

5.4.2 Concave Valuations

In this section, we improve upon our mechanism in [Amanatidis et al., 2025], that achieved the first constant factor approximation for the setting of divisible agents and concave valuation functions. Prior to the above work, we obtained the best known approximation guarantee for this setting in [Klumper and Schäfer, 2022], in which the approximation guarantee grows almost linearly with the number of agent types. Here, the number of agents types refers to the number of different valuation functions. The mechanism, called Chunk-and-Solve, that we introduce in this section is inspired on both of our previous works mentioned above. The approximation guarantee of Chunk-and-Solve depends on an input parameter k, which is to be set by the user, and improves (decreases) as k increases. We first show that if this parameter k is equal to the number of agents n, our mechanism is $2(1+\varphi)$-approximate and improves upon the known guarantee of Amanatidis et al. [2025] and of Klumper and Schäfer [2022] if there are 3 or more agents types. We then show, in a more involved proof, that the approximation guarantee of our mechanism can approach $1+\varphi$ if the parameter k is set to be a large enough multiple of the number of agents n. This improves upon the guarantee of Klumper and Schäfer [2022] if there are 2 or more agents types, and converges to the guarantee of $1+\varphi$ when there is only one agent type. In this and the next section, we sometimes refer to instances of the divisible model as well as instances of the k-level model in a mechanism or lemma. If this is the case, we slightly deviate from our initial notation as introduced in Section 5.2.2 and define an instance J of the k-level model as $J = (N, \bar{\boldsymbol{c}}, (\bar{v}_i)_{i\in N}, B, k)$, while our notation for instances of the divisible model remains unchanged.

A $2(1+\varphi)$ Approximation

Recall that in the divisible model with concave non-decreasing valuation functions, we have $v_i : [0,1] \to \mathbb{R}_{\geq 0}$ with $v_i(0) = 0$ for all $i \in N$. There is a natural correspondence between the setting with $k \geq 1$ levels of service and the setting with divisible agents: If we subdivide the $[0,1]$ interval into k subintervals of length $\frac{1}{k}$ and evaluate the $v_i(\cdot)$'s at $\frac{1}{k}, \frac{2}{k}, \ldots, \frac{k}{k}$, then this can be interpreted as the value of procuring $1, 2, \ldots, k$ levels of service, respectively. We can obtain an improved approximation guarantee by applying this discretization for $k = n$ and then using a mechanism very similar to Greedy-and-Prune from Section 5.3.1, instead of using Sort-&-Reject from Section 5.3.2 as we do in [Amanatidis et al., 2025]. 5

We first introduce our mechanism Greedy-and-Prune-F, see Mechanism 13 below. The main differences with Greedy-and-Prune are as follows.

(i) In the else part of the mechanism, the initial allocation retrieves exactly an α fraction of the optimal value of the corresponding fractional k-bounded knapsack problem by allowing one level to be allocated fractionally. This allocation is constructed greedily by ordering the levels of the agents according to decreasing marginal value-over-cost ratios, where $\sigma(i)$ and $\sigma'(i)$ denote the agent and level, respectively, of the i-th most efficient ratio with $i \in [k \cdot n]$.

(ii) As the valuation functions are concave, the marginal values-over-cost ratios are evaluated for each level separately, and the threats are level specific for all agents $i \in N$, i.e., $m_i(j) \cdot \tau_i'(\boldsymbol{c}_{-i})$ for $j \in \{1, 2, \ldots, k\}$ with:

$$\tau_i'(\boldsymbol{c}_{-i}) = \frac{B}{\alpha(1+\beta)\,\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i})}. \tag{5.13}$$

Note that for each agent $i \in N$, the definition of $\rho_i(\boldsymbol{c}_{-i})$ remains as in (5.6), i.e., $\rho_i(\boldsymbol{c}_{-i}) = v_i(k)/\,\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i})$. We first prove the following theorem for

MECHANISM 13: GREEDY-AND-PRUNE-F(I)

Input: An instance $I = (N, \boldsymbol{c}, (v_i)_{i \in N}, B, k)$.
Output: An allocation vector $\boldsymbol{x}$ and a payment vector $\boldsymbol{p}$.

let $i^* \in \arg\max_{i \in N} \rho_i(\boldsymbol{c}_{-i})$
if $\rho_{i^*}(\boldsymbol{c}_{-i}) \geq \beta$ **then**
    set $x_{i^*} = k$ and $x_i = 0$ for $i \in N \setminus \{i^*\}$
end
else
    for all $k \cdot n$ marginal value-over-cost ratios, define $\sigma(i)$ and $\sigma'(i)$ with $i \in [k \cdot n]$ such that $\frac{m_{\sigma(1)}(\sigma'(1))}{c_{\sigma(1)}} \geq \frac{m_{\sigma(2)}(\sigma'(2))}{c_{\sigma(2)}} \geq \dots \geq \frac{m_{\sigma(k \cdot n)}(\sigma'(k \cdot n))}{c_{\sigma(k \cdot n)}}$
    compute $\ell \in [k \cdot n]$ such that $\sum_{i=1}^{\ell-1} m_{\sigma(i)}(\sigma'(i)) < \alpha\, \mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) \leq \sum_{i=1}^{\ell} m_{\sigma(i)}(\sigma'(i))$
    for $i \in N$ with $\exists j < \ell$ such that $\sigma(j) = i$, set $x_i = \max_{j < \ell}\{\sigma'(j) \mid \sigma(j) = i\}$
    for $h = \sigma(\ell)$, find x'_h such that $\sum_{i=1}^{n} \sum_{j=1}^{x_i} m_i(j) + (x'_h - x_h) m_h(x_h + 1) = \alpha\, \mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$ and set $x_h = x'_h$
    for $i \in N$ *with* $x_i > 0$ **do**
        for $j = 1, 2, \dots, \lceil x_i \rceil$ **do**
            if $c_i > m_i(j) \cdot \tau'_i$ **then**
                set $x_i = j - 1$ and end for loop
            end
        end
    end
end
for $i \in N$ **do**
    compute payment p_i according to (5.1)
end
return $(\boldsymbol{x}, \boldsymbol{p})$

GREEDY-AND-PRUNE-F, after which we introduce our mechanism for instances of the divisible model.

Theorem 5.4.7. *Consider the class of instances of the k-level model with concave and non-decreasing valuations. Suppose that the allocation can have one fractional level of service. Then, for $\alpha = (1 + \varphi)^{-1}$ and $\beta = \frac{\alpha}{1-\alpha}$, GREEDY-AND-PRUNE-F is a truthful, individually rational and budget-feasible mechanism that is $(1 + \varphi)$-approximate.*

We start by showing that the allocation rule of GREEDY-AND-PRUNE-F is monotone, after which we prove budget-feasibility.

Lemma 5.4.8. *The allocation rule of* GREEDY-AND-PRUNE-F *is monotone for* $\alpha \in (0,1)$ *and* $\beta > 0$.

Proof: Let $\boldsymbol{c}$ be a cost profile. Note that the two threshold values used in GREEDY-AND-PRUNE-F also depend on the parameters α and β, which are fixed and do not depend on the declarations $\boldsymbol{c}$. As an agent has a positive allocation by either being allocated as i^* or by being allocated in the greedy order, we consider the following two cases.

Case 1: $\rho_{i^*}(\boldsymbol{c}_{-i^*}) \geq \beta$. The proof for this case is the same as in the proof of Lemma 5.3.2.

Case 2: $\rho_{i^*}(\boldsymbol{c}_{-i^*}) < \beta$. In this case, the initial allocation is computed in the else statement of GREEDY-AND-PRUNE-F, and multiple agents may have a positive allocation. (Note that if no agent has a positive allocation, we do not need to examine any unilateral deviation.) Fix an agent $i \in N$ with $x_i(\boldsymbol{c}) > 0$, which implies that $\exists j \leq k$ such that $c_i \leq m_i(j)\tau_i'$. Suppose agent i unilaterally deviates to $c_i' < c_i$ and let $\boldsymbol{c}' = (c_i', \boldsymbol{c}_{-i})$. First of all, again note that for every $j \in N \setminus \{i\}$, it holds that $\rho_j(\boldsymbol{c}'_{-j}) \leq \rho_j(\boldsymbol{c}_{-j})$, whereas $\rho_i(\boldsymbol{c}'_{-i}) = \rho_i(\boldsymbol{c}_{-i})$. Therefore, under $\boldsymbol{c}'$, the if statement in Line 2 also evaluates to false and GREEDY-AND-PRUNE-F executes the else statement. Moreover, any level j of agent i may move up in the ordering in Line 6 as $c_i' < c_i$. Therefore, agent i may have a higher initial allocation after Line 9, as $\texttt{opt}_{\mathrm{F}}^k(\boldsymbol{c}') \geq \texttt{opt}_{\mathrm{F}}^k(\boldsymbol{c})$. Finally, note that for any level $j \leq k$ the threat $m_i(j)\tau_i'$ is independent of the declaration of agent i, and so $c_i' < c_i \leq m_i(j)\tau_i'$ for $j \leq x_i(\boldsymbol{c})$. Therefore, $x_i(\boldsymbol{c}') \geq x_i(\boldsymbol{c})$, proving monotonicity of the allocation rule of GREEDY-AND-PRUNE-F. □

5

Lemma 5.4.9. *For* $\alpha \in (0,1)$ *and* $\beta > 0$, GREEDY-AND-PRUNE-F *is budget feasible.*

Proof: Let $\boldsymbol{c}$ be a cost profile. We consider the following two cases.

Case 1: $\rho_{i^*}(\boldsymbol{c}_{-i^*}) \geq \beta$. In this case, only agent i^* has a positive allocation and $x_{i^*}(\boldsymbol{c}) = k$. Therefore:

$$\sum_{i \in N} p_i(\boldsymbol{c}) = p_{i^*}(\boldsymbol{c}) \leq k \cdot \frac{B}{k} = B,$$

as the allocation rule of GREEDY-AND-PRUNE-F is monotone (Lemma 5.4.8), so we can bound the payment of agent i^* by $x_{i^*}(\boldsymbol{c}) = k$ times the threshold bid of i^*. Note that we can upper bound the threshold bid by $\frac{B}{k}$, as we run a pre-processing step that removes agents with a declaration greater than $\frac{B}{k}$.

Case 2: $\rho_{i^*}(\boldsymbol{c}_{-i^*}) < \beta$. In this case, the allocation is computed in the else statement of GREEDY-AND-PRUNE-F. We now bound the payment of an agent $i \in N$ by upper bounding each allocated level with a specific threshold bid of i. Note that by Line 12 of GREEDY-AND-PRUNE-F, for level j we can use $m_i(j) \cdot \tau_i'$ as an upper bound of the threshold bid, as τ_i' is independent of the declaration c_i. For $i \in N$ and $j \in \{1, 2, \dots, k\}$ define:

$$z_{ij} = \begin{cases} 1, & \text{if } j \leq x_i, \\ 0, & \text{if } j > \lceil x_i \rceil, \\ x_i - \lfloor x_i \rfloor, & \text{otherwise.} \end{cases}$$

Then:

$$\begin{aligned} \sum_{i \in N} p_i(\boldsymbol{c}) \leq \sum_{i=1}^{n} \sum_{j=1}^{\lceil x_i \rceil} m_i(j) \cdot \tau_i' \cdot z_{ij} &= \sum_{i=1}^{n} \sum_{j=1}^{\lceil x_i \rceil} m_i(j) \cdot z_{ij} \cdot \frac{B}{\alpha(1+\beta)\,\mathtt{opt}_{\mathrm{F}}^{k}(\mathbf{c}_{-i})} \\ &\leq \sum_{i=1}^{n} \sum_{j=1}^{\lceil x_i \rceil} m_i(j) \cdot z_{ij} \cdot \frac{B}{\alpha\,\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})} \leq B, \end{aligned}$$

where the second inequality follows as for each agent $i \in N$ it holds that $(1 + \beta)\,\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i}) > \mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i}) + v_i(k) \geq \mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$, as the if condition in Line 2 evaluated to false. Finally, the last equality follows as $v(\boldsymbol{x}) \leq \alpha\,\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$, as by construction an α fraction of $\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$ is initially selected in Line 9, after which the allocation of an agent can only remain the same or decrease. This proves that GREEDY-AND-PRUNE-F is budget feasible. □

5

In order to prove the approximation guarantee of GREEDY-AND-PRUNE-F, we will make use of the following lemma.

Lemma 5.4.10. *Given an instance $I = (N, \boldsymbol{c}, (v_i)_{i \in N}, B, k)$ of the fractional k-bounded knapsack problem with concave and non-decreasing valuations and $\alpha \in (0, 1)$. Let $n = |N|$ and for all $k \cdot n$ marginal value-over-cost ratios, define $\sigma(i)$ and $\sigma'(i)$ with $i \in [k \cdot n]$ such that:*

$$\frac{m_{\sigma(1)}(\sigma'(1))}{c_{\sigma(1)}} \geq \frac{m_{\sigma(2)}(\sigma'(2))}{c_{\sigma(2)}} \geq \dots \geq \frac{m_{\sigma(k \cdot n)}(\sigma'(k \cdot n))}{c_{\sigma(k \cdot n)}},$$

where $\sigma(i)$ denotes the agent belonging to the i-th ratio and $\sigma'(i)$ the level of agent $\sigma(i)$. Let $\ell \in [k \cdot n]$ such that $\sum_{i=1}^{\ell-1} m_{\sigma(i)}(\sigma'(i)) < \alpha\,\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) \leq \sum_{i=1}^{\ell} m_{\sigma(i)}(\sigma'(i))$. Then, for $i \in \{1, 2, \dots, \ell\}$, it holds that:

$$\frac{c_{\sigma(i)}}{m_{\sigma(i)}(\sigma'(i))}(1-\alpha)\,\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) \leq B.$$

Proof: Let $\ell \in [k \cdot n]$ such that $\sum_{i=1}^{\ell-1} m_{\sigma(i)}(\sigma'(i)) < \alpha\,\mathtt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) \leq \sum_{i=1}^{\ell} m_{\sigma(i)}(\sigma'(i))$, and let $\boldsymbol{x}^*$ be the optimal solution computed by Algorithm 8 using the same tie-breaking rule. Note that we can split $\boldsymbol{x}^*$ into $\boldsymbol{x} = (x_1, x_2, \dots, x_n)$ and $\boldsymbol{y} = (x_1^* - x_1, x_2^* - x_2, \dots, x_n^* - x_n)$ with $x_i \in \{0, 1, \dots, x_i^*\}$ such that $\boldsymbol{x}^* = \boldsymbol{x} + \boldsymbol{y}$ and:

(1) $v(\boldsymbol{x}) < \alpha\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$,

(2) $v(\mathbf{y}) > (1-\alpha)\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$, and

(3) for $i \in N$ with $\exists j < \ell$ such that $\sigma(j) = i$, let $x_i = \max_{j<\ell}\{\sigma'(j) \mid \sigma(j) = i\}$.

For $i \in N$ and $j \in \{1, 2, \ldots, k\}$ define:

$$z_{ij} = \begin{cases} 1, & \text{if } j \leq x_i^*, \\ 0, & \text{if } j > \lceil x_i^* \rceil, \\ x_i^* - \lfloor x_i^* \rfloor, & \text{otherwise.} \end{cases}$$

Then, by feasibility of $\boldsymbol{x}^*$, and thus of $\boldsymbol{y}$, it follows that for $h \in \{1, 2, \ldots, \ell\}$:

$$\begin{aligned} B \geq \sum_{i=1}^{n} c_i y_i = \sum_{i=1}^{n} \sum_{j=x_i+1}^{\lceil x_i^* \rceil} c_i z_{ij} &= \sum_{i=1}^{n} \sum_{j=x_i+1}^{\lceil x_i^* \rceil} \frac{c_i}{m_i(j)} m_i(j) z_{ij} \\ &\geq \frac{c_{\sigma(\ell)}}{m_{\sigma(\ell)}(\sigma'(\ell))} \sum_{i=1}^{n} \sum_{j=x_i+1}^{\lceil x_i^* \rceil} m_i(j) z_{ij} \\ &= \frac{c_{\sigma(\ell)}}{m_{\sigma(\ell)}(\sigma'(\ell))} v(\boldsymbol{y}) > \frac{c_{\sigma(\ell)}}{m_{\sigma(\ell)}(\sigma'(\ell))} (1-\alpha)\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) \\ &\geq \frac{c_{\sigma(h)}}{m_{\sigma(h)}(\sigma'(h))} (1-\alpha)\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}), \end{aligned}$$

where the second and last inequality follow by the construction of $\boldsymbol{x}$ and $\boldsymbol{y}$, and the third inequality follows from (2). This concludes the proof. □

5

With Lemma 5.4.10 in place, the proof of the approximation guarantee of Greedy-and-Prune-F follows easily.

Lemma 5.4.11. *Given an instance $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B, k)$ with concave and non-decreasing valuations. Let $\alpha \in (0, 1)$ and $\beta > 0$ such that $\alpha = \frac{\beta}{1+\beta}$ and $\alpha(1+\beta) \leq 1-\alpha$. Then, Greedy-and-Prune-F is $1/\alpha$-approximate.*

Proof: Let $\boldsymbol{x} = \boldsymbol{x}(\boldsymbol{c})$ be the allocation computed by Greedy-and-Prune-F. We consider the following two cases.

Case 1: $\rho_{i^*}(\boldsymbol{c}_{-i^*}) \geq \beta$. In this case only agent i^* has a positive allocation of $x_{i^*} = k$. Furthermore, it holds that:

$$v_{i^*}(k) \geq \beta\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}_{-i^*}) \geq \beta\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) - \beta v_{i^*}(k),$$

and rearranging terms leads to $v(\boldsymbol{x}) = v_{i^*}(k) \geq \beta/(1+\beta)\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c}) = \alpha\,\texttt{opt}_{\mathrm{F}}^{k}(\boldsymbol{c})$.

Case 2: $\rho_{i^*}(\boldsymbol{c}_{-i^*}) < \beta$. In this case, the allocation is computed in the else statement of Greedy-and-Prune-F. In order for Greedy-and-Prune-F to achieve the best possible approximation guarantee, we want no agent $i \in N$ with $x_i > 0$ before the for loop in Line 10 to have any level removed due to the imposed threat, i.e. we want that $c_i \leq m_i(j)\tau_i'$ for all $j \leq \lceil x_i \rceil$. As $\alpha(1+\beta) \leq 1-\alpha$ and $\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i}) \leq \mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c})$, it holds that:

$$\alpha(1+\beta)\,\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i}) \leq (1-\alpha)\,\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}),$$

which for $i \in N$ with $x_i > 0$ and $j \leq \lceil x_i \rceil$ leads to:

$$\tau_i'(\boldsymbol{c}_{-i}) = \frac{B}{\alpha(1+\beta)\,\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c}_{-i})} \geq \frac{B}{(1-\alpha)\,\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c})} \geq \frac{c_i}{m_i(j)},$$

where the last inequality follows from Lemma 5.4.10 and by construction of $\boldsymbol{x}$. Therefore, $v(\boldsymbol{x}) = \alpha\,\mathtt{opt}_{\mathrm{F}}^k(\boldsymbol{c})$ by construction, concluding the proof. ☐

Together, Lemmas 5.4.8, 5.4.9 and 5.4.11 prove Theorem 5.4.7. To see this, note that rewriting $\alpha = \frac{\beta}{1+\beta}$ leads to $\beta = \frac{\alpha}{1-\alpha}$, and for $\alpha = (1+\varphi)^{-1}$, it holds that $\alpha(1+\beta) = 1-\alpha$.

We can now introduce our mechanism Chunk-and-Solve, see Mechanism 14 below, which applies a discretization to an instance of the divisible model and then uses Greedy-and-Prune-F to compute an allocation.

MECHANISM 14: Chunk-and-Solve(I, k)

Input: An instance $I = (N, \boldsymbol{c}, (v_i)_{i \in N}, B)$ and a parameter $k \in \mathbb{N}_{>0}$.
Output: An allocation vector $\boldsymbol{x}$ and a payment vector $\boldsymbol{p}$.

1 initialize for each $i \in N$, $\bar{v}_i : \{0, \ldots, k\} \to \mathbb{R}_{\geq 0}$ with $\bar{v}_i(0) = 0$.
2 **for** $i \in N$ and $j \in \{1, \ldots, k\}$ **do**
3 | set $\bar{v}_i(j) = v_i\left(\frac{j}{k}\right)$
4 **end**
5 **for** $i \in N$ **do**
6 | set $\bar{c}_i = \frac{c_i}{k}$
7 **end**
8 let $J = (N, \bar{\boldsymbol{c}}, (\bar{v}_i)_{i \in N}, B, k)$ be the resulting discretized instance of the k-level model
9 compute $\bar{\boldsymbol{x}} =$ Greedy-and-Prune-F(J) for $\alpha = (1+\varphi)^{-1}$ and $\beta = \frac{\alpha}{1-\alpha}$
10 **for** $i \in N$ **do**
11 | set $x_i = \frac{\bar{x}_i}{k}$
12 **end**
13 compute the payments $\boldsymbol{p}$ according to (5.1)
14 **return** $(\boldsymbol{x}, \boldsymbol{p})$

Theorem 5.4.12. *Consider the class of instances of the divisible model with concave and non-decreasing valuation functions. Then, for $k = n = |N|$, the mechanism* Chunk-and-Solve *is truthful, individually rational, budget-feasible and $2(1+\varphi)$-approximate.*

The following lemma will be key in proving the approximation guarantee of the above theorem.

Lemma 5.4.13. *Let $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B)$ be an instance of the divisible model with concave and non-decreasing valuation functions and let $n = |N|$. Consider the instance $J = (N, \bar{\boldsymbol{c}}, (\bar{v}_i)_{i\in N}, B, n)$ as defined in Line 8 of* Chunk-and-Solve. *Then, it holds that $\texttt{opt}_{\mathrm{F}}(\boldsymbol{c}) \le 2 \cdot \texttt{opt}^n(\bar{\boldsymbol{c}})$.*

Proof: Let $\boldsymbol{x}^* = (x_1^*, \dots, x_n^*) \in [0,1]^n$ be an optimal solution to the (non-strategic) fractional concave knapsack problem for an instance I (see (5.4) in Section 5.2.3). Let J be the discretized instance of I as defined in Line 8 of Chunk-and-Solve. Let $y_i = \lfloor n \cdot x_i^* \rfloor$ for all $i \in N$ and let $\boldsymbol{y} = (y_1, \dots, y_n)$. Define $\mathbf{1} = (1, 1, \dots, 1)$. We have:

$$\begin{aligned}\texttt{opt}_{\mathrm{F}}(\boldsymbol{c}) = \sum_{i\in N} v_i(x_i^*) &\le \sum_{i\in N} v_i\left(\frac{\lceil n \cdot x_i^* \rceil}{n}\right) \le \sum_{i\in N} v_i\left(\frac{y_i}{n}\right) + m_i(1) \\ &= \sum_{i\in N} \bar{v}_i(y_i) + \bar{v}_i(1) = \bar{v}(\boldsymbol{y}) + \bar{v}(\mathbf{1}). \qquad (5.14)\end{aligned}$$

The first inequality follows by construction, as $x_i^* \le \lceil n \cdot x_i^* \rceil / n$ for all $i \in N$. The second inequality follows as $\lceil n \cdot x_i^* \rceil \le y_i + 1$ for all $i \in N$, and as the valuation functions are concave and non-decreasing. The second equality follows as $y_i \le n$ is an integer and $m_i(1) = \bar{v}_i(1)$ for any $i \in N$.

Next, we show that both $\boldsymbol{y}$ and $\mathbf{1}$ are feasible solutions for J. For $\boldsymbol{y}$, it holds by construction that:

$$\sum_{i\in N} \bar{c}_i y_i = \sum_{i\in N} \frac{c_i}{n} y_i = \sum_{i\in N} \frac{c_i}{n} \lfloor n \cdot x_i^* \rfloor \le \sum_{i\in N} c_i x_i^* \le B,$$

where the last inequality follows from the fact that $\boldsymbol{x}^*$ is a feasible solution for I, as it is an optimal solution.

Similarly, we can verify that $\mathbf{1}$ is a feasible solution for J by observing that:

$$\sum_{i\in N} \bar{c}_i = \sum_{i\in N} \frac{c_i}{n} = \frac{\sum_{i\in N} c_i}{n} \le \frac{n \cdot \max_{i\in N} c_i}{n} = \max_{i\in N} c_i \le B\,,$$

where the final inequality holds by the cost assumption of the divisible model (see Section 5.2.3).

Since both $\boldsymbol{y}$ and $\mathbf{1}$ are feasible solutions for J, it holds that:

$$\bar{v}(\boldsymbol{y}) \leq \mathtt{opt}^n(\bar{\boldsymbol{c}}) \quad \text{and} \quad \bar{v}(\mathbf{1}) \leq \mathtt{opt}^n(\bar{\boldsymbol{c}})\,.$$

By combining this with (5.14), the lemma follows. □

We can now prove Theorem 5.4.12.

Proof of Theorem 5.4.12: Chunk-and-Solve reduces an instance I of the divisible model to an instance J with $k = n = |N|$ levels of service and uses Greedy-and-Prune-F parameterized with $\alpha = (1+\varphi)^{-1}$ and $\beta = \frac{\alpha}{1-\alpha}$ to compute an allocation. Thus, Chunk-and-Solve inherits truthfulness and individual rationality through Lemma 5.4.8. Moreover, for the approximation guarantee we have:

$$v(\boldsymbol{x}(\boldsymbol{c})) \geq \bar{v}(\bar{\boldsymbol{x}}(\bar{\boldsymbol{c}})) \geq \alpha\, \mathtt{opt}^n_{\mathrm{F}}(\bar{\boldsymbol{c}}) \geq \alpha\, \mathtt{opt}^n(\bar{\boldsymbol{c}}) \geq \frac{\alpha}{2}\, \mathtt{opt}_{\mathrm{F}}(\boldsymbol{c})\,,$$

where the first inequality follows as the valuation functions are concave and non-decreasing and the second inequality follows from Lemma 5.4.11. The last inequality follows from Lemma 5.4.13.

What remains to be shown is that Chunk-and-Solve is budget-feasible. We will show that, for an instance $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B)$ of the divisible model, the payments computed in Line 13 of Chunk-and-Solve coincide with the payments computed by Greedy-and-Prune-F for the instance $J = (N, \bar{\boldsymbol{c}}, (\bar{v}_i)_{i\in N}, B, n)$ which is constructed in Line 8 of Chunk-and-Solve. Specifically, we will show that for each agent $i \in N$, it holds that $p_i(\boldsymbol{c}) = \bar{p}_i(\bar{\boldsymbol{c}})$. Let $w = f(x) = x/n$. For each $i \in N$, we have:

$$\begin{aligned}
p_i(\boldsymbol{c}) &= c_i x_i(\boldsymbol{c}) + \int_{c_i}^{B} x_i(z, \boldsymbol{c}_{-i})\, dz = n\bar{c}_i \frac{\bar{x}_i(\bar{\boldsymbol{c}})}{n} + \int_{n\bar{c}_i}^{B} \frac{\bar{x}_i\left(z/n, \bar{\boldsymbol{c}}_{-i}\right)}{n}\, dz \\
&= \bar{c}_i \bar{x}_i(\bar{\boldsymbol{c}}) + \int_{n\bar{c}_i}^{B} f'(z) \cdot \bar{x}_i\left(f(z), \bar{\boldsymbol{c}}_{-i}\right)\, dz = \bar{c}_i \bar{x}_i(\bar{\boldsymbol{c}}) + \int_{f(n\bar{c}_i)}^{f(B)} \bar{x}_i\left(w, \bar{\boldsymbol{c}}_{-i}\right)\, dw \\
&= \bar{c}_i \bar{x}_i(\bar{\boldsymbol{c}}) + \int_{\bar{c}_i}^{B/n} \bar{x}_i\left(w, \bar{\boldsymbol{c}}_{-i}\right)\, dw = \bar{p}_i(\bar{\boldsymbol{c}}).
\end{aligned}$$

The first equality follows from applying Theorem 5.2.1 (Myerson's payment identity) to the allocation $\boldsymbol{x}(\boldsymbol{c})$ for I in Chunk-and-Solve. The second equality follows from Lines 6 and 11, since, by definition, $x_i(\boldsymbol{b}) = \bar{x}_i\left(\frac{1}{n}\boldsymbol{b}\right)/n$ holds for all $\boldsymbol{b} \in [0, B]^n$. The subsequent three equalities are obtained through simple calculus, involving a change of the variable of the integral term. Finally, the last equality holds because Greedy-and-Prune-F also computed the payments according to Theorem 5.2.1 for the instance J. Thus, as Greedy-and-Prune-F is budget-feasible by Lemma 5.4.9, budget-feasibility of Chunk-and-Solve follows. □

5

5.4.3 A More Exhaustive Analysis for Concave Valuations

In order to improve upon the approximation guarantee of Theorem 5.4.12, we run our mechanism Chunk-and-Solve with a larger parameter k. This allows us to improve (decrease) the approximation guarantee of Lemma 5.4.13, that relates the value of an optimal solution to the fractional concave knapsack problem of an instance I, to the value of an optimal solution to the k-bounded knapsack problem of an instance J, where J is the discretized instance of I. In our proof, we compute an optimal solution $\boldsymbol{x}^*$ to the fractional concave knapsack problem, and use our algorithm k-Rounding to transform $\boldsymbol{x}^*$ into a specific solution of the fractional k-bounded knapsack problem, which we use to prove the desired approximation guarantee. As the proof is quite involved, we first state our improved approximation guarantee of Chunk-and-Solve. We then give two auxiliary lemmas which will be useful for proving correctness and specific properties of our algorithm k-Rounding. We then conclude this section by proving that an allocation that follows from k-Rounding satisfies the desired approximation guarantee.

We first state the lemma that relates the value of an optimal solution of a fractional concave knapsack problem to the value of its discretized counterpart.

Lemma 5.4.14. *Let $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B)$ be an instance of the divisible model with concave and non-decreasing valuation functions and let $\varepsilon \in (0,1)$. Define $s = \lceil 1 + \frac{1}{\varepsilon} \rceil$ and $k = n \cdot s$. Consider the discretized instance $J = (N, \bar{\boldsymbol{c}}, (\bar{v}_i)_{i\in N}, B, k)$ of I as defined in Line 8 of Chunk-and-Solve. Then, it holds that $\mathtt{opt}_{\mathrm{F}}(\boldsymbol{c}) \leq \left(1 + \frac{1}{s-2}\right) \mathtt{opt}^k(\bar{\boldsymbol{c}})$.*

5

The main result of this section is the following theorem:

Theorem 5.4.15. *Consider the class of instances of the divisible model with concave and non-decreasing valuation functions. Let $s = \lceil 1 + 4\varepsilon^{-1} \rceil$, with $\varepsilon \in (0,1)$ and n the number of agents of an instance. Then, for $k = n \cdot s$, the mechanism Chunk-and-Solve is truthful, individually rational, budget-feasible and $(1 + \varphi + \varepsilon)$-approximate.*

Proof: Consider an instance I of the divisible model with concave and non-decreasing valuation functions and a cost profile $\boldsymbol{c}$. Let $\boldsymbol{x}$ be the computed allocation, i.e., $\boldsymbol{x} = $ Chunk-and-Solve$(I, n \cdot s)$. It follows from the proof of Theorem 5.4.12 that Chunk-and-Solve is truthful, individual rationality and budget-feasible. For the approximation guarantee, it follows that:

$$\alpha^{-1}\left(1 + \frac{1}{s-2}\right) v(\boldsymbol{x}) \geq \alpha^{-1}\left(1 + \frac{1}{s-2}\right) \bar{v}(\bar{\boldsymbol{x}}(\bar{\boldsymbol{c}}))$$
$$\geq \left(1 + \frac{1}{s-2}\right) \mathtt{opt}^k_{\mathrm{F}}(\bar{\boldsymbol{c}}) \geq \mathtt{opt}_{\mathrm{F}}(\boldsymbol{c}),$$

where the first inequality follows as the valuation functions are concave and non-decreasing and the second inequality follows from Lemma 5.4.11. The last inequality follows from Lemma 5.4.14, as $\texttt{opt}^k_{\mathrm{F}}(\bar{\boldsymbol{c}}) \geq \texttt{opt}^k(\bar{\boldsymbol{c}})$. As $\alpha = (1+\varphi)^{-1}$ and $s = \lceil 1 + 4\varepsilon^{-1} \rceil$, we have that:

$$\begin{aligned}\alpha^{-1}\left(1 + \frac{1}{s-2}\right) &= (1+\varphi)\left(1 + \frac{1}{\lceil 1 + 4\varepsilon^{-1}\rceil - 2}\right) \\ &\leq 1 + \varphi + \frac{3}{\frac{4}{\varepsilon} - 1} = 1 + \varphi + \frac{3\varepsilon}{4 - \varepsilon} \\ &< 1 + \varphi + \frac{3\varepsilon}{3} = 1 + \varphi + \varepsilon,\end{aligned}$$

as $\varepsilon \in (0,1)$. This concludes the proof. □

Auxiliary Lemmas

We now present two auxiliary lemmas and a fact which will be used in the proofs regarding our algorithm k-Rounding and Lemma 5.4.14. The first lemma will be useful for proofs regarding aggregated value-over-cost ratios, which we consider for our algorithm k-Rounding.

5

Lemma 5.4.16. *Let $N = \{1, 2, \ldots, n\}$. For each $i \in N$ let $a_i \in \mathbb{R}_{\geq 0}$ and $b_i \in \mathbb{R}_{>0}$. Define $i^+ \in \arg\max_{i \in N} \frac{a_i}{b_i}$ and $i^- \in \arg\min_{i \in N} \frac{a_i}{b_i}$. Then:*

$$\frac{\sum_{i=1}^n a_i}{\sum_{i=1}^n b_i} \leq \frac{a_{i^+}}{b_{i^+}} \quad \textit{and} \quad \frac{a_{i^-}}{b_{i^-}} \leq \frac{\sum_{i=1}^n a_i}{\sum_{i=1}^n b_i}.$$

Proof: Note that, by definition of i^+ and i^-, we have that for all $i \in N$ it holds that $a_i \leq b_i \cdot \frac{a_{i^+}}{b_{i^+}}$ and $b_i \cdot \frac{a_{i^-}}{b_{i^-}} \leq a_i$, as $b_i > 0$. Adding these inequalities for all i leads to:

$$\sum_{i=1}^n a_i \leq \sum_{i=1}^n b_i \cdot \frac{a_{i^+}}{b_{i^+}} \quad \text{and} \quad \sum_{i=1}^n b_i \cdot \frac{a_{i^-}}{b_{i^-}} \leq \sum_{i=1}^n a_i.$$

As $\sum_{i=1}^n b_i > 0$, rearranging terms leads to:

$$\frac{\sum_{i=1}^n a_i}{\sum_{i=1}^n b_i} \leq \frac{a_{i^+}}{b_{i^+}} \quad \text{and} \quad \frac{a_{i^-}}{b_{i^-}} \leq \frac{\sum_{i=1}^n a_i}{\sum_{i=1}^n b_i},$$

concluding the proof. □

Naturally, the following lemma shows that there always exists an optimal solution to a fractional concave knapsack problem with non-decreasing valuation functions that spends the entire budget, if the budget is smaller than the accumulated costs of all agents. Namely, given an optimal solution that does not spend the entire budget, one can increases the allocation of the agents such that the allocations are still smaller than or equal to 1 and the entire budget is spend.

Fact 5.4.17. Let $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B)$ be an instance of the divisible model with non-decreasing valuation functions with $\sum_{i\in N} c_i > B$. Then, there exists an optimal solution $\boldsymbol{x}^*$ to the fractional concave knapsack problem of I such that $\sum_{i\in N} c_i x_i^* = B$.

Finally, the last lemma exhibits a property of optimal solutions of fractional concave knapsack problems with non-decreasing valuation functions. It shows that if an agent $j \in N$ is not fully allocated, i.e., $x_j < 1$, then redistributing budget from another agent $i \in N$ to agent j results in a lower or equal efficiency of agent j compared to agent i, with the efficiency evaluated for the redistributed budget.

Lemma 5.4.18. *Let $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B)$ be an instance of the divisible model with non-decreasing valuation functions and let $\boldsymbol{x}^*$ be an optimal solution to the fractional concave knapsack problem of I. Then for any $i, j \in N$ with $x_i^* \in (0, 1]$, $x_j^* \in (0, 1)$ and $\varepsilon \in (0, \min\{x_i^* c_i, (1 - x_j^*)c_j\}]$, it holds that:*

$$\frac{v_i(x_i^*) - v_i\left(\frac{x_i^* c_i - \varepsilon}{c_i}\right)}{\varepsilon} \geq \frac{v_j\left(\frac{x_j^* c_j + \varepsilon}{c_j}\right) - v_j(x_j^*)}{\varepsilon}.$$

Proof: Towards contradiction, assume that there exists $i, j \in N$ with $x_i^* \in (0, 1]$, $x_j^* \in (0, 1)$ and $\exists \varepsilon \in (0, \min\{x_i^* c_i, (1 - x_j^*)c_j\}]$ such that:

$$0 \leq \frac{v_i(x_i^*) - v_i\left(\frac{x_i^* c_i - \varepsilon}{c_i}\right)}{\varepsilon} < \frac{v_j\left(\frac{x_j^* c_j + \varepsilon}{c_j}\right) - v_j(x_j^*)}{\varepsilon},$$

where the first inequality follows as the valuation functions are non-decreasing. Let $\boldsymbol{x}$ be such that $x_h = x_h^*$ for $h \in N$ with $h \neq i, j$, and let $x_i = x_i^* - \frac{\varepsilon}{c_i}$ and $x_j = x_j^* + \frac{\varepsilon}{c_j}$. Note that $\boldsymbol{x}$ is a feasible solution by construction, as it holds that $x_i \in [0, 1)$, $x_j \in (0, 1]$ and $\sum_{i\in N} c_i x_i = \sum_{i\in N} c_i x_i^* \leq B$. It follows by assumption that:

$$\begin{aligned} v_i(x_i^*) - v_i(x_i) &= \varepsilon \cdot \frac{v_i(x_i^*) - v_i\left(\frac{x_i^* c_i - \varepsilon}{c_i}\right)}{\varepsilon} \\ &< \varepsilon \cdot \frac{v_j\left(\frac{x_j^* c_j + \varepsilon}{c_j}\right) - v_j(x_j^*)}{\varepsilon} = v_j(x_j) - v_j(x_j^*), \end{aligned}$$

which contradicts optimality of $\boldsymbol{x}^*$ as:

$$v(\boldsymbol{x}) - v(\boldsymbol{x}^*) = v_j(x_j) - v_j(x_j^*) - (v_i(x_i^*) - v_i(x_i)) > 0.$$

□

5

Our Algorithm: k-Rounding

We now introduce our algorithm k-Rounding, see Algorithm 15, that we will use in order to relate the value of an optimal solution of a fractional concave knapsack problem to the optimal value of its discretized counterpart. k-Rounding takes as input an instance $I = (N, \boldsymbol{c}, (v_i)_{i\in N}, B)$ and a parameter $\varepsilon \in (0,1)$, which will determine the discretization step k. The discretization step is computed in Line 1, and is always a multiple of the number of agents n, i.e., $k = s \cdot n$, where s depends on the input parameter ε. Then, an optimal solution $\boldsymbol{x}^*$ to I that spends the entire budget B is computed in Line 2. Note that if $\boldsymbol{x}^*$ is evaluated for the discretized counterpart of I, i.e., $\bar{v}(k \cdot \boldsymbol{x}^*)$, value is potentially lost for agents $i \in N$ for which x_i^* is not a multiple of $1/k$. Namely, by construction, for any multiple of $1/k$ the valuation functions coincide, i.e., $v_i(j/k) = \bar{v}_i(j)$ for $j = 0, 1, \dots, k$ for all $i \in N$. However, if x_i^* is not a multiple of $\frac{1}{k}$, it may be that $v_i(x_i^*) > \bar{v}_i(k \cdot x_i^*)$. Therefore, if this difference in value is 'large' for an agent i, we round the allocation of i up to the nearest multiple of $1/k$. However, if the allocation of agent i is already sufficiently large, there is no need to round the allocation up as this loss in value can then be compensated.

If the allocation of an agent $j \in N$ is rounded up, then the allocation of other agents needs to be reduced in order for the allocation to remain feasible. Which agents k-Rounding considers for increasing and decreasing allocations is determined in Line 3. Agents i with $x_i^* = 0$ are not considered, and agents $j \in N^S$ with a relatively small allocation, i.e., $x_j^* < \frac{s-1}{k}$, are eligible to be rounded up. This will be at the expense of agents $i \in N^L$ with a relatively large allocation, i.e., $x_i^* > \frac{s}{k}$. The rest of the agents with a medium sized allocation, i.e., $i \in N^M$ with $\frac{s-1}{k} \leq x_i^* \leq \frac{s}{k}$ are not considered by k-Rounding.

k-Rounding only considers rounding the allocation of agents $j \in N^S$ up if x_j^* is not a multiple of $\frac{1}{k}$ (see Line 4). In order to determine if the allocation of agent j is rounded up, a worst-case cost of $\frac{1}{k}c_j = \delta_j^c$ is made available by decreasing the allocations of agents in N^L (see Line 5). Decreasing the allocation of agents in N^L is done in such a way that the value lost is as small as possible, i.e., by solving the optimization problem below, where x_i is the current allocation of $i \in N^L$.

$$\begin{array}{lrcll} \min & \sum_{i\in N^L} v_i(x_i) - v_i(z_i) & & & \\ \text{s.t.} & \sum_{i\in N^L} z_i c_i & = & \sum_{i\in N^L} x_i c_i - \delta_j^c & \\ & z_i & \in & \left[\frac{s}{n\cdot s}, x_i\right] & \forall i \in N^L \end{array} \tag{5.15}$$

Which renders the same solution as the solution of the concave optimization problem below.

$$\begin{array}{lrcll} \max & \sum_{i\in N^L} v_i(z_i) & & & \\ \text{s.t.} & \sum_{i\in N^L} z_i c_i & = & \sum_{i\in N^L} x_i c_i - \delta_j^c & \\ & z_i & \in & \left[\frac{s}{n\cdot s}, x_i\right] & \forall i \in N^L \end{array} \tag{5.16}$$

Algorithm 15: K-ROUNDING(I, ε)

Input: An instance $I = (N, \boldsymbol{c}, (v_i)_{i \in N}, B)$ and a parameter $\varepsilon \in (0, 1)$.
Output: A feasible solution $\boldsymbol{x}$.

1 set $k = n \cdot s$, with $s = \lceil 1 + \frac{1}{\varepsilon} \rceil$
2 let $\boldsymbol{x}^*$ be an optimal solution of the fractional concave knapsack problem of I s.t. $\sum_{i \in N} c_i x_i^* = B$ and set $\boldsymbol{x} = \boldsymbol{x}^*$ and $y_i^* = \frac{\lfloor x_i^* \cdot k \rfloor}{k}$ $\forall i \in N$
3 let $N^S = \{i \in N \mid 0 < kx_i^* < s - 1\}$ and $N^L = \{i \in N \mid kx_i^* > s\}$
4 **for** $j \in N^S$ *with* $x_j^* > y_j^*$ **do**
5   let $\delta_j^c = \frac{1}{k} c_j$ and let $\boldsymbol{x}'_{N^L}$ be the solution of (5.15) with input $(I, s, N^L, \delta_j^c, \boldsymbol{x}_{N^L})$ and let δ_j^v be the objective value
6   **if** $\frac{v_j(x_j^*) - v_j(y_j^*)}{(x_j^* - y_j^*) c_j} > \frac{\delta_j^v}{\delta_j^c}$ **then**
7   set $x_j = y_j^* + \frac{1}{k}$ and $\boldsymbol{x}_{N^L} = \boldsymbol{x}'_{N^L}$
8   **end**
9 **end**
10 **return** $\boldsymbol{x}$

Note that in both (5.15) and (5.16), the first constraint ensures that enough budget is freed up, while at the same time ensuring that the available budget is entirely spend. The second constraint ensures that the allocation of an agent $i \in N^L$ can only decrease, but remain larger than or equal to $\frac{s}{k}$.

Then, in Line 6, K-ROUNDING compares the loss in efficiency of the agents in N^L, i.e., the objective value δ_j^v of (5.15) divided by δ_j^c, is compared to the efficiency of the 'fractional' part of the allocation of j. The allocation of j is then only rounded up to the nearest multiple of $\frac{1}{k}$ if the latter efficiency is strictly greater than the former. If this is the case, then in order for the allocation to remain feasible, the allocation of the agents in N^L is updated to the solution of (5.15) in Line 7.

The following lemma establishes that K-ROUNDING terminates and returns a feasible solution, along with additional properties.

Lemma 5.4.19. *Let $\varepsilon > 0$ and consider an instance $I = (N, \boldsymbol{c}, (v_i)_{i \in N}, B)$ of the divisible model with non-decreasing valuation functions and $\sum_{i \in N} c_i > B$. Let $\boldsymbol{x}^*$ be an optimal solution to the fractional concave knapsack problem of I computed in Line 2 of K-ROUNDING. Let $k = n \cdot s$ with $s = \lceil 1 + \frac{1}{\varepsilon} \rceil$. Then, $\boldsymbol{x} =$ K-ROUNDING(I, ε) is a feasible solution to I such that:*

$$\sum_{i \in N^L} (x_i^* - x_i)\, c_i = \sum_{i \in N_u^S} \frac{1}{k} c_i \quad \text{and} \quad \sum_{i \in N^L} v_i(x_i^*) - v_i(x_i) = \sum_{i \in N_u^S} \delta_i^v,$$

with $N^L = \{i \in N \mid kx_i^* > s\}$ and $N_u^S = \{i \in N \mid x_i^* < x_i\}$.

Proof: We begin by proving that K-ROUNDING terminates and returns a feasible solution $\boldsymbol{x}$ to I. Note that a solution $\boldsymbol{x}^*$ computed in Line 2 exists by Fact 5.4.17. If a solution $\boldsymbol{x}'_{N^L}$ in Line 5 of K-ROUNDING always exists, then K-ROUNDING trivially terminates. We first show that, together, the allocations of all i in N^L are large enough to round the allocation of all j in N^S up to the nearest multiple of $\frac{1}{k}$ in Line 7. We define $N' = \{i \in N \mid x_i^* > 0\}$ as the set of all $i \in N$ with an allocation strictly greater than 0 in $\boldsymbol{x}^*$, and partition N' into N^L, N^M and N^S, with N^L and N^S as defined in Line 3 of K-ROUNDING and $N^M = \{i \in N' \mid s-1 \leq kx_i^* \leq s\}$. If $N^S \neq \emptyset$, then as $\boldsymbol{x}^*$ uses the entire budget:

$$\begin{aligned}\sum_{i\in N^L}\left(x_i^* - \frac{s}{k}\right)c_i &= B - \sum_{i\in N^L}\frac{s}{k}c_i - \sum_{i\in N^M} x_i^* c_i - \sum_{i\in N^S} x_i^* c_i \\ &> n\cdot\frac{s}{k}B - \sum_{i\in N^L}\frac{s}{k}c_i - \sum_{i\in N^M}\frac{s}{k}c_i - \sum_{i\in N^S}\frac{k\cdot y_i^*+1}{k}c_i \\ &\geq \sum_{i\in N^S}\frac{s-(k\cdot y_i^*+1)}{k}B \geq \sum_{i\in N^S}\frac{1}{k}c_i = \sum_{i\in N^S}\delta_i^c. \qquad (5.17)\end{aligned}$$

Here, the first inequality holds as $k = n\cdot s$ and by definitions of N^M, N^S and y_i^*. The second inequality follows as $|N'| \leq |N| = n$. The third inequality follows by definition of N^S and by the cost assumption. Therefore, and by the first constraint of (5.15), a solution $\boldsymbol{x}'_{N^L}$ computed by K-ROUNDING in Line 5 always exists.

Furthermore, note that every time K-ROUNDING computes a solution $\boldsymbol{x}'_{N^L}$ to (5.15), the solution satisfies $x'_i \in [\frac{s}{k}, x_i^*]$ for every $i \in N^L$. Also note that for every sequential solution, the allocation of any $i \in N^L$ can only decrease or stay the same due to the second constraint. Therefore, and by the first constraint, it holds that:

$$\sum_{i\in N^L}(x_i^* - x_i)\,c_i = \sum_{i\in N_u^S}\delta_i^c = \sum_{i\in N_u^S}\frac{1}{k}c_i.$$

To see that $\boldsymbol{x}$ is a feasible solution, note that $\boldsymbol{x}^*$ is feasible by definition and $x_i = x_i^*$ for $i \notin N^L \cup N_u^S$. Furthermore, the decrease in costs for $i \in N^L$ is greater than the increase in costs for $i \in N_u^S$ as $x_i^* > y_i^*$ for $i \in N_u^S$, so:

$$\sum_{i\in N_u^S}(x_i - x_i^*)\,c_i = \sum_{i\in N_u^S}\left(y_i^* + \frac{1}{k} - x_i^*\right)c_i < \sum_{i\in N_u^S}\frac{1}{k}c_i.$$

Let $\ell \leq |N_u^S|$ be the number of times the if statement in Line 6 in K-ROUNDING evaluates to true and let $\boldsymbol{x}^j_{N^L}$ be the solution to (5.15) computed in Line 5 for the iteration in which the if statement evaluates to true for the j-th time. So $\boldsymbol{x}^\ell_{N^L} = \boldsymbol{x}_{N^L}$ and let $\boldsymbol{x}^0_{N^L} = \boldsymbol{x}^*_{N^L}$. It holds that:

$$\sum_{i\in N^L} v_i(x_i^*) - v_i(x_i) = \sum_{j=1}^{\ell} v(\boldsymbol{x}^{j-1}_{N^L}) - v(\boldsymbol{x}^j_{N^L}) = \sum_{i\in N_u^S}\delta_i^v,$$

where the first equality holds by Line 7 of K-ROUNDING. This concludes the proof. □

For each agent $j \in N^S$ considered in Line 4 of K-ROUNDING, the next lemma establishes an upper bound on the loss in efficiency if a cost of δ_j^c has to be deducted from the budget spent on the agents in N^L.

Lemma 5.4.20. *Let $\varepsilon > 0$ and consider an instance $I = (N, \boldsymbol{c}, (v_i)_{i \in N}, B)$ of the divisible model with concave and non-decreasing valuation functions and $\sum_{i \in N} c_i > B$. Let $\boldsymbol{x}^*$ be an optimal solution to the fractional concave knapsack problem of I computed in Line 2 of K-ROUNDING. Let $k = n \cdot s$ with $s = \lceil 1 + \frac{1}{\varepsilon} \rceil$ and let $\boldsymbol{x} =$ K-ROUNDING(I, ε). Then, for $j \in N^S$ with $x_j^* > y_j^*$ it holds that:*

$$\frac{\delta_j^v}{\delta_j^c} \leq \frac{\sum_{i \in N^L} v_i(x_i) - v_i\left(\frac{s}{k}\right)}{\sum_{i \in N^L} \left(x_i - \frac{s}{k}\right) c_i}. \tag{5.18}$$

Proof: Let ℓ be the number of times the for loop in Line 4 of K-ROUNDING is executed, i.e., $\ell = |\{i \in N^S \mid x_i^* > y_i^*\}|$. For $j = 1, 2, \ldots, \ell$ let $\boldsymbol{x}^j$ be the allocation after the j-th execution of the for loop, so $\boldsymbol{x} = \boldsymbol{x}^\ell$. Let $\boldsymbol{x}^* = \boldsymbol{x}^0$. Note that the denominator on the left-hand side in (5.18) is never 0 as $\delta_j^c = \frac{1}{k} c_j > 0$ for all $j \in N^S$. And as the preconditions of Lemma 5.4.19 are met, it follows from equation (5.17) in Lemma 5.4.19 that for $j = 0, 1, \ldots, \ell$ it holds that:

$$\sum_{i \in N^L} \left(x_i^j - \frac{s}{k}\right) c_i > 0. \tag{5.19}$$

We will first prove that for any $j \in N^S$ with $x_j^* > y_j^*$ considered in iteration $\sigma(j)$ of the for loop with $\boldsymbol{x}'_{N^L}$ the solution to (5.15) in Line 5, it holds that:

$$\frac{\delta_j^v}{\delta_j^c} = \frac{\sum_{i \in N^L} v_i\left(x_i^{\sigma(j)-1}\right) - v_i\left(x'_i\right)}{\sum_{i \in N^L} \left(x_i^{\sigma(j)-1} - x'_i\right) c_i} \leq \frac{\sum_{i \in N^L} v_i\left(x'_i\right) - v_i\left(\frac{s}{k}\right)}{\sum_{i \in N^L} \left(x'_i - \frac{s}{k}\right) c_i}. \tag{5.20}$$

Note that the equality follows by definition. We define:

$$\rho^+ \in \underset{i \in N^L : x_i^{\sigma(j)-1} > x'_i}{\arg\max} \frac{v_i\left(x_i^{\sigma(j)-1}\right) - v_i\left(x'_i\right)}{\left(x_i^{\sigma(j)-1} - x'_i\right) c_i} \quad \text{and} \quad \rho^- \in \underset{i \in N^L : x'_i > \frac{s}{k}}{\arg\min} \frac{v_i\left(x'_i\right) - v_i\left(\frac{s}{k}\right)}{\left(x'_i - \frac{s}{k}\right) c_i}.$$

Note that ρ^+ and ρ^- exist as $\delta_j^c = \frac{1}{k} c_j > 0$ for all $j \in N^S$ and by (5.19). Let

5

$\varepsilon \in \left(0, \min\left\{\left(x'_{\rho^-} - \frac{s}{k}\right) c_{\rho^-}, \left(x^{\sigma(j-1)}_{\rho^+} - x'_{\rho^+}\right) c_{\rho^+}\right\}\right)$, then:

$$\frac{v_{\rho^+}\left(x^{\sigma(j)-1}_{\rho^+}\right) - v_{\rho^+}\left(x'_{\rho^+}\right)}{\left(x^{\sigma(j)-1}_{\rho^+} - x'_{\rho^+}\right) c_{\rho^+}} \leq \frac{v_{\rho^+}\left(x'_{\rho^+} + \frac{\varepsilon}{c_{\rho^+}}\right) - v_{\rho^+}\left(x'_{\rho^+}\right)}{\varepsilon}$$
$$\leq \frac{v_{\rho^-}\left(x'_{\rho^-}\right) - v_{\rho^-}\left(x'_{\rho^-} - \frac{\varepsilon}{c_{\rho^-}}\right)}{\varepsilon} \leq \frac{v_{\rho^-}\left(x'_{\rho^-}\right) - v_{\rho^-}\left(\frac{s}{k}\right)}{\left(x'_{\rho^-} - \frac{s}{k}\right) \bar{c}_{\rho^-}},$$

where the first and third inequality follow as the valuation functions are concave and non-decreasing (Fact 5.2.3). The second inequality follows from Lemma 5.4.18. Therefore, by definition of ρ^+ and ρ^- and Lemma 5.4.16 it follows that:

$$\frac{\sum_{i \in N^L} v_i\left(x^{\sigma(j)-1}_i\right) - v_i\left(x'_i\right)}{\sum_{i \in N^L}\left(x^{\sigma(j)-1}_i - x'_i\right) c_i} \leq \frac{v_{\rho^+}\left(x^{\sigma(j)-1}_{\rho^+}\right) - v_{\rho^+}\left(x'_{\rho^+}\right)}{\left(x^{\sigma(j)-1}_{\rho^+} - x'_{\rho^+}\right) c_{\rho^+}}$$
$$\leq \frac{v_{\rho^-}\left(x'_{\rho^-}\right) - v_{\rho^-}\left(\frac{s}{k}\right)}{\left(x'_{\rho^-} - \frac{s}{k}\right) c_{\rho^-}} \leq \frac{\sum_{i \in N^L} v_i\left(x'_i\right) - v_i\left(\frac{s}{k}\right)}{\sum_{i \in N^L}\left(x'_i - \frac{s}{k}\right) c_i},$$

proving (5.20).

5

Now we first consider $j \in N^S$ with $x^*_j > y^*_j$ for which the if statement in Line 6 of k-Rounding evaluated to true. From (5.20) it follows that:

$$\frac{\delta^v_j}{\delta^c_j} = \frac{\sum_{i \in N^L} v_i\left(x^{\sigma(j)-1}_i\right) - v_i\left(x^{\sigma(j)}_i\right)}{\sum_{i \in N^L}\left(x^{\sigma(j)-1}_i - x^{\sigma(j)}_i\right) c_i} \leq \frac{\sum_{i \in N^+_1} v_i\left(x^{\sigma(j)}_i\right) - v_i\left(\frac{s}{k}\right)}{\sum_{i \in N^+_1}\left(x^{\sigma(j)}_i - \frac{s}{k}\right) c_i}, \tag{5.21}$$

as in this case $\boldsymbol{x}^{\sigma(j)}_{N^L} = \boldsymbol{x}'_{N^L}$, with $\boldsymbol{x}'_{N^L}$ the solution of (5.15) computed in Line 5 in iteration $\sigma(j)$. Secondly, consider $j \in N^S$ with $x^*_j > y^*_j$ for which the if statement in Line 6 evaluated to false. Note that in this case $x^{\sigma(j-1)}_{N^L} = x^{\sigma(j)}_{N^L}$. We will show that:

$$\frac{\delta^v_j}{\delta^c_j} = \frac{\sum_{i \in N^L} v_i\left(x^{\sigma(j)}_i\right) - v_i\left(x'_i\right)}{\sum_{i \in N^L}\left(x^{\sigma(j)}_i - x'_i\right) c_i} \leq \frac{\sum_{i \in N^L} v_i\left(x^{\sigma(j)}_i\right) - v_i\left(\frac{s}{k}\right)}{\sum_{i \in N^L}\left(x^{\sigma(j)}_i - \frac{s}{k}\right) c_i}. \tag{5.22}$$

Note that $x'_i \in \left[\frac{s}{k}, x^{\sigma(j)}_i\right]$ for all $i \in N^L$. Using (5.20) for the inequality below, we have that:

$$\begin{aligned}
&\frac{\sum_{i\in N^L} v_i\left(x_i^{\sigma(j)}\right)-v_i\left(\frac{s}{k}\right)}{\sum_{i\in N^L}\left(x_i^{\sigma(j)}-\frac{s}{k}\right)c_i}\cdot\sum_{i\in N^L}\left(x_i^{\sigma(j)}-\frac{s}{k}\right)c_i=\sum_{i\in N^L} v_i\left(x_i^{\sigma(j)}\right)-v_i\left(\frac{s}{k}\right)\\
&=\frac{\sum_{i\in N^L} v_i\left(x_i^{\sigma(j)}\right)-v_i\left(x_i'\right)}{\sum_{i\in N^L}\left(x_i^{\sigma(j)}-x_i'\right)c_i}\cdot\sum_{i\in N^L}\left(x_i^{\sigma(j)}-x_i'\right)c_i\\
&+\frac{\sum_{i\in N^L} v_i\left(x_i'\right)-v_i\left(\frac{s}{k}\right)}{\sum_{i\in N^L}\left(x_i'-\frac{s}{k}\right)c_i}\cdot\sum_{i\in N^L}\left(x_i'-\frac{s}{k}\right)c_i\\
&\geq\frac{\sum_{i\in N^L} v_i\left(x_i^{\sigma(j)}\right)-v_i\left(x_i'\right)}{\sum_{i\in N^L}\left(x_i^{\sigma(j)}-x_i'\right)c_i}\cdot\sum_{i\in N^L}\left(x_i^{\sigma(j)}-\frac{s}{k}\right)c_i,
\end{aligned}$$

proving (5.22).

We now show that for $j=1,2,\ldots,\ell-1$ it holds that:

$$\frac{\sum_{i\in N^L} v_i\left(x_i^{\sigma(j)}\right)-v_i\left(\frac{s}{k}\right)}{\sum_{i\in N^L}\left(x_i^{\sigma(j)}-\frac{s}{k}\right)c_i}\leq\frac{\sum_{i\in N^L} v_i\left(x_i^{\sigma(j)+1}\right)-v_i\left(\frac{s}{k}\right)}{\sum_{i\in N^L}\left(x_i^{\sigma(j)+1}-\frac{s}{k}\right)c_i}. \tag{5.23}$$

Note that for $j\in\{1,2,\ldots,\ell-1\}$ with $\boldsymbol{x}_{N^L}^{\sigma(j)}=\boldsymbol{x}_{N^L}^{\sigma(j)+1}$, (5.23) holds with equality. Otherwise, it holds that $x_i^{\sigma(j)}\geq x_i^{\sigma(j)+1}$ for $i\in N^L$ by Line 7 of k-Rounding and the second constraint in (5.15). We have:

$$\begin{aligned}
&\frac{\sum_{i\in N^L} v_i\left(x_i^{\sigma(j)}\right)-v_i\left(\frac{s}{k}\right)}{\sum_{i\in N^L}\left(x_i^{\sigma(j)}-\frac{s}{k}\right)c_i}\cdot\sum_{i\in N^L}\left(x_i^{\sigma(j)}-\frac{s}{k}\right)c_i=\sum_{i\in N^L} v_i\left(x_i^{\sigma(j)}\right)-v_i\left(\frac{s}{k}\right)\\
&=\frac{\sum_{i\in N^L} v_i\left(x_i^{\sigma(j)}\right)-v_i\left(x_i^{\sigma(j)+1}\right)}{\sum_{i\in N^L}\left(x_i^{\sigma(j)}-x_i^{\sigma(j)+1}\right)c_i}\cdot\sum_{i\in N^L}\left(x_i^{\sigma(j)}-x_i^{\sigma(j)+1}\right)c_i\\
&+\frac{\sum_{i\in N^L} v_i\left(x_i^{\sigma(j)+1}\right)-v_i\left(\frac{s}{k}\right)}{\sum_{i\in N^L}\left(x_i^{\sigma(j)+1}-\frac{s}{k}\right)c_i}\cdot\sum_{i\in N^L}\left(x_i^{\sigma(j)+1}-\frac{s}{k}\right)c_i\\
&\leq\frac{\sum_{i\in N^L} v_i\left(x_i^{\sigma(j)+1}\right)-v_i\left(\frac{s}{k}\right)}{\sum_{i\in N^L}\left(x_i^{\sigma(j)+1}-\frac{s}{k}\right)c_i}\cdot\sum_{i\in N^L}\left(x_i^{\sigma(j)}-\frac{s}{k}\right)c_i,
\end{aligned}$$

proving (5.23). Here, the inequality again follows from (5.20).

By (5.21), (5.22) and (5.23), and as $\boldsymbol{x}=\boldsymbol{x}^\ell$, the proof follows. □

5

Approximation

We complete this section by proving Lemma 5.4.14 which, for convenience, is repeated below:

Lemma 5.4.14. *Let* $I = (N, \boldsymbol{c}, (v_i)_{i \in N}, B)$ *be an instance of the divisible model with concave and non-decreasing valuation functions and let* $\varepsilon \in (0,1)$*. Define* $s = \lceil 1 + \frac{1}{\varepsilon} \rceil$ *and* $k = n \cdot s$*. Consider the discretized instance* $J = (N, \bar{\boldsymbol{c}}, (\bar{v}_i)_{i \in N}, B, k)$ *of* I *as defined in Line 8 of* Chunk-and-Solve*. Then, it holds that* $\texttt{opt}_{\mathrm{F}}(\boldsymbol{c}) \leq \left(1 + \frac{1}{s-2}\right) \texttt{opt}^k(\bar{\boldsymbol{c}})$*.*

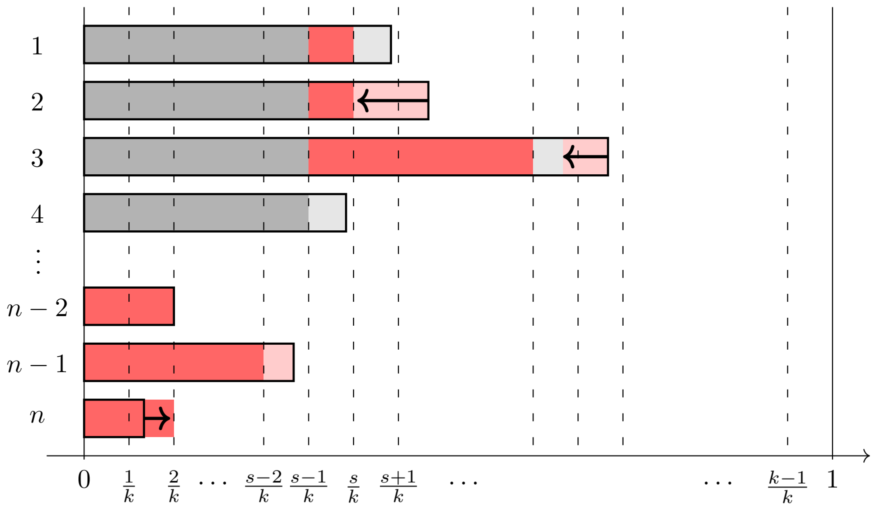

Figure 5.3: Example of an outcome of k-Rounding. The allocation of an agent i in the optimal solution of the fractional concave knapsack problem is equal to the length of the respective horizontal bar. The arrows represent the rounding decisions made by k-Rounding.

We first explain the high-level idea of the proof by means of Figure 5.3, in which the vertical axis indicates the agents, ranging from 1 to $n = |N|$, and the horizontal axis indicates the range of an allocation, ranging from 0 to 1. Figure 5.3 illustrates a possible outcome of k-Rounding. For each the agent $i \in N$, the allocation x_i^* of the optimal solution of the fractional concave knapsack problem computed in k-Rounding is equal to the full length of the respective horizontal bar. The arrows represent the rounding decisions made by k-Rounding. For example, the allocation of agent n is rounded up by k-Rounding to $\frac{2}{k}$ and the allocation of agent 2 is decreased. The allocation $\boldsymbol{x}$ constructed by k-Rounding corresponds to the highlighted dark red, dark gray and light gray parts, as well as the light red parts for the agents in N^S as defined in Line 3 of k-Rounding.

We define $\boldsymbol{y}$ by rounding all the allocations of $\boldsymbol{x}$ down to the nearest multiple of $\frac{1}{k}$. We relate the value of $\boldsymbol{x}^*$ to $\boldsymbol{y}$ in the following way. For agents $i \in N^S$ for which the allocation is rounded up, e.g., agent n in the figure, it already holds that $v_i(x_i^*) < \bar{v}_i(k \cdot y_i)$. For agents $i \in N^S$ for which x_i^* is a multiple of the discretization size of $\frac{1}{k}$, e.g., agent $n-2$ in the figure, it already holds that $v_i(x_i^*) = \bar{v}_i(k \cdot y_i)$. For all other agents, the allocation $\boldsymbol{y}$ is of equal value as the corresponding fraction of the allocation $\boldsymbol{x}^*$, e.g., all other dark red and dark gray areas in the figure. For an agent i in N^L, e.g., agents 1, 2, and 3 in the figure, or an agent i not in both N^L and N^S, e.g., agent 4 in the figure, the value of the light gray area can be related to a fraction of the value of the dark gray area, as the allocation of such an agent i is large enough. The value of $\boldsymbol{x}^*$ that is not yet accounted for is the value of the light red areas in the figure. We will relate this value to a fraction of the value of the dark red areas.

Proof of Lemma 5.4.14: First, recall that by construction, for any multiple of $\frac{1}{k}$ the valuation functions of the instances I and J coincide, i.e., $v_i(j/k) = \bar{v}_i(j)$ for $j = 0, 1, \ldots, k$ for all $i \in N$. Therefore, if $\sum_{i \in N} c_i \leq B$, $\texttt{opt}_{\mathrm{F}}(\boldsymbol{c}) = v(\mathbf{1}) = \bar{v}(\boldsymbol{k}) = \texttt{opt}^k(\bar{\boldsymbol{c}})$, proving the claim as $s \geq 3$.

So, assume that $\sum_{i \in N} c_i > B$. Let $\boldsymbol{x} = $ k-Rounding(I, ε) and let $\boldsymbol{x}^*$ be an optimal solution to the fractional concave knapsack problem of I computed in Line 2 of k-Rounding. Let $y_i = \frac{\lfloor k \cdot x_i \rfloor}{k} \geq 0$ for all $i \in N$. Note that both y_i and y_i^* (as defined in Line 2 of k-Rounding) are multiples of $\frac{1}{k}$. By Lemma 5.4.19, $k \cdot \boldsymbol{x}$ is a feasible solution to J. Therefore, $k \cdot \boldsymbol{y}$ is also a feasible solution to J and $\bar{v}(k \cdot \boldsymbol{y}) \leq \texttt{opt}^k(\bar{\boldsymbol{c}})$. We will show that $v(\boldsymbol{x}^*) \leq \left(1 + \frac{1}{s-2}\right) \bar{v}(k \cdot \boldsymbol{y})$, proving the claim.

As $v_i(0) = 0$ for all $i \in N$, in order to upper bound $v(\boldsymbol{x}^*)$ we only have to consider the agents in $N' = \{i \in N \mid x_i^* > 0\}$. Let $N^M = \{i \in N' \mid s - 1 \leq k x_i^* \leq s\}$ and let N^L and N^S be defined as in Line 3 of k-Rounding. Note that N^L, N^M and N^S form a partition of N'. We partition N^S further depending on whether x_i^* is rounded up or remains equal in k-Rounding, or if x_i^* was already a multiple of $\frac{1}{k}$, i.e., $N_u^S = \left\{i \in N^S \mid x_i^* < x_i\right\}$, $N_f^S = \left\{i \in N^S \mid y_i^* < x_i = x_i^*\right\}$ and $N_c^S = \left\{i \in N^S \mid x_i^* = y_i^*\right\}$.

We first consider N^M and N_c^S. For $i \in N^M$ it holds that:

$$\begin{aligned}
\sum_{i \in N^M} v_i(x_i^*) &= \sum_{i \in N^M} v_i(y_i^*) + \sum_{i \in N^M} v_i(x_i^*) - v_i(y_i^*) \\
&\leq \sum_{i \in N^M} \bar{v}_i(k \cdot y_i^*) + \sum_{i \in N^M} v_i\left(y_i^* + \frac{1}{k}\right) - v_i(y_i^*) \\
&\leq \sum_{i \in N^M} \bar{v}_i(k \cdot y_i) + \sum_{i \in N^M} \frac{1}{s-1} v_i\left(\frac{s-1}{k}\right) \\
&\leq \left(1 + \frac{1}{s-2}\right) \sum_{i \in N^M} \bar{v}_i(k \cdot y_i).
\end{aligned}$$

Here, the first inequality follows as $x_i^* < y_i^* + \frac{1}{k}$ and as y_i^* is a multiple of $\frac{1}{k}$ for all $i \in N$. The final two inequalities follow as $\frac{s-1}{k} \leq y_i^* = y_i$ for $i \in N^M$, as the valuations are concave and non-decreasing and as $s \geq 3$. For $i \in N_c^S$ we have that x_i^* is a multiple of $\frac{1}{k}$ and $x_i^* = x_i$ by K-ROUNDING, and therefore it holds that $v_i(x_i^*) = \bar{v}_i(k \cdot y_i)$. Combining this with the above inequality leads to:

$$\sum_{i \in N^M \cup N_c^S} v_i(x_i^*) \leq \left(1 + \frac{1}{s-2}\right) \sum_{i \in N^M \cup N_c^S} \bar{v}_i(k \cdot y_i). \tag{5.24}$$

It remains to bound the values of the optimal allocations for N^L, N_f^S and N_u^S. We start by upper bounding the value of the 'fractional' part of x_i for i in N^L or N_f^S, i.e., $v_i(x_i) - v_i(y_i)$. For $i \in N^L$ it holds by K-ROUNDING that $\frac{s}{k} \leq y_i \leq x_i < y_i + \frac{1}{k}$ and therefore:

$$\begin{aligned} \sum_{i \in N^L} v_i(x_i) - v_i(y_i) &\leq \sum_{i \in N^L} v_i\left(y_i + \frac{1}{k}\right) - v_i(y_i) \\ &\leq \sum_{i \in N^L} \frac{1}{s-1} v_i\left(\frac{s-1}{k}\right), \end{aligned} \tag{5.25}$$

again as the valuations are concave and non-decreasing. By similar reasoning, we obtain for $i \in N_f^S$ with $y_i > 0$ that:

$$v_i(x_i^*) - v_i(y_i^*) \leq v_i\left(y_i^* + \frac{1}{k}\right) - v_i(y_i^*) \leq \frac{1}{k \cdot y_i} v_i(y_i),$$

as $y_i = y_i^*$ in this case. Multiplying both sides by $k \cdot y_i$ and summing over all $i \in N_f^S$ leads to:

$$\sum_{i \in N_f^S} k \cdot y_i(v_i(x_i^*) - v_i(y_i^*)) \leq \sum_{i \in N_f^S} v_i(y_i). \tag{5.26}$$

Note that (5.26) holds for all $i \in N_f^S$ as $v_i(0) = 0$ by assumption.

We also upper bound the value that is lost by decreasing the allocations x_i^* of $i \in N^L$ in K-ROUNDING. For $i \in N_u^S$ with $y_i^* \geq \frac{1}{k}$, it holds that:

$$\frac{\delta_i^v}{\frac{1}{k} c_i} = \frac{\delta_i^v}{\delta_i^c} < \frac{v_i(x_i^*) - v_i(y_i^*)}{(x_i^* - y_i^*) c_i} \leq \frac{v_i(y_i^*) - v_i(y_i^* - \frac{1}{k})}{\frac{1}{k} c_i},$$

where the first inequality follows from if statement in Line 6 of K-ROUNDING and the second inequality follows by Fact 5.2.3 as the valuation functions are concave and non-decreasing. Multiplying the above inequality by $k \cdot y_i^* \cdot \frac{1}{k} c_i$ and summing over all $i \in N_u^S$ leads to:

$$\sum_{i \in N_u^S} k \cdot y_i^* \cdot \delta_i^v \leq \sum_{i \in N_u^S} v_i(y_i^*), \tag{5.27}$$

5

as the valuation functions are concave and non-decreasing.

We first lower bound the cost of x_i for the share above $\frac{s}{k}$ for $i \in N^L$, and use this later on to bound a related value. We have that:

$$\begin{aligned}
\sum_{i \in N^L} \left(x_i - \frac{s}{k}\right) c_i &= \sum_{i \in N^L} \left(x_i^* - \frac{s}{k}\right) c_i - \sum_{i \in N^L} (x_i^* - x_i) c_i \\
&= \sum_{i \in N^L} \left(x_i^* - \frac{s}{k}\right) c_i - \sum_{i \in N_u^S} \frac{1}{k} c_i \\
&\geq \sum_{i \in N^S} \frac{s - (k \cdot y_i^* + 1)}{k} B - \sum_{i \in N_u^S} \frac{1}{k} B \\
&\geq \sum_{i \in N_f^S} \frac{s - (k \cdot y_i + 1)}{k} B + \sum_{i \in N_u^S} \frac{s - (k \cdot y_i^* + 2)}{k} B. \qquad (5.28)
\end{aligned}$$

Here, the second equality follows from Lemma 5.4.19 and the first inequality follows by (5.17) in Lemma 5.4.19 and as $c_i \leq B$ for all $i \in N$. The second inequality follows as $y_i^* = y_i$ for $i \in N_f^S$.

We can now lower bound the value for the allocation above $\frac{s-1}{k}$ for $i \in N^L$.

$$\begin{aligned}
&\sum_{i \in N^L} v_i(y_i) - v_i\left(\frac{s-1}{k}\right) \geq \sum_{i \in N^L} v_i(x_i) - v_i\left(\frac{s}{k}\right) \\
&= \frac{\sum_{i \in N^L} v_i(x_i) - v_i\left(\frac{s}{k}\right)}{\sum_{i \in N^L} \left(x_i - \frac{s}{k}\right) c_i} \cdot \sum_{i \in N^L} \left(x_i - \frac{s}{k}\right) c_i \\
&\geq \frac{\sum_{i \in N^L} v_i(x_i) - v_i\left(\frac{s}{k}\right)}{\sum_{i \in N^L} \left(x_i - \frac{s}{k}\right) c_i} \cdot \left(\sum_{i \in N_f^S} \frac{s - (k \cdot y_i + 1)}{k} B + \sum_{i \in N_u^S} \frac{s - (k \cdot y_i^* + 2)}{k} B \right) \\
&\geq \sum_{i \in N_f^S} \frac{\delta_i^v}{\delta_i^c} \cdot \frac{s - (k \cdot y_i + 1)}{k} B + \sum_{i \in N_u^S} \frac{\delta_i^v}{\delta_i^c} \cdot \frac{s - (k \cdot y_i^* + 2)}{k} B \\
&\geq \sum_{i \in N_f^S} \frac{v_i(x_i^*) - v_i(y_i^*)}{(x_i^* - y_i^*) c_i} \cdot \frac{s - (k \cdot y_i + 1)}{k} B + \sum_{i \in N_u^S} \frac{\delta_i^v}{\frac{1}{k} B} \cdot \frac{s - (k \cdot y_i^* + 2)}{k} B \\
&\geq \sum_{i \in N_f^S} (v_i(x_i^*) - v_i(y_i^*)) \cdot (s - (k \cdot y_i + 1)) + \sum_{i \in N_u^S} \delta_i^v \cdot (s - (k \cdot y_i^* + 2)). \qquad (5.29)
\end{aligned}$$

Here, the first inequality holds as $x_i \leq y_i + \frac{1}{k}$ and as the valuation functions are concave and non-decreasing. The second inequality follows by (5.28). The third and fourth inequality follow by Lemma 5.4.20, by Line 6 of k-Rounding and by the cost assumption. The final inequality follows as $x_i^* - y_i^* \leq \frac{1}{k}$ for all $i \in N$ and by the cost assumption.

5

We continue by showing that:

$$
\begin{aligned}
&(s-2)\left(\sum_{i\in N^L} v_i(x_i^*) - v_i(x_i) + \sum_{i\in N^L} v_i(x_i) - v_i(y_i) + \sum_{i\in N_f^S} v_i(x_i^*) - v_i(y_i^*)\right) \\
&\leq (s-2)\sum_{i\in N_u^S} \delta_i^v + (s-1)\sum_{i\in N^L} v_i(x_i) - v_i(y_i) + (s-1)\sum_{i\in N_f^S} v_i(x_i^*) - v_i(y_i^*) \\
&\leq \sum_{i\in N^L\cup N_f^S} v_i(y_i) + \sum_{i\in N_u^S} v_i(y_i^*) \leq \sum_{i\in N^L\cup N_f^S\cup N_u^S} v_i(y_i),
\end{aligned}
$$

where, the first inequality follows by Lemma 5.4.19. The second inequality follows by multiplying both sides of (5.25) by $s-1$, and adding this with equations (5.29), (5.27) and (5.26). The final inequality follows as $y_i^* < y_i$ for $i \in N_u^S$. Multiplying both sides by $1/(s-2)$ leads to:

$$
\begin{aligned}
&\sum_{i\in N^L} v_i(x_i^*) - v_i(x_i) + \sum_{i\in N^L} v_i(x_i) - v_i(y_i) + \sum_{i\in N_f^S} v_i(x_i^*) - v_i(y_i^*) \\
&\leq \frac{1}{s-2}\sum_{i\in N^L\cup N_f^S\cup N_u^S} v_i(y_i). \qquad (5.30)
\end{aligned}
$$

We can now upper bound the value of x_i^* for all i in $N^L \cup N_f^S \cup N_u^S$.

5

$$
\begin{aligned}
\sum_{i\in N^L\cup N_f^S\cup N_u^S} v_i(x_i^*) &\leq \sum_{i\in N^L} v_i(x_i^*) - v_i(x_i) + \sum_{i\in N^L} v_i(x_i) - v_i(y_i) + \sum_{i\in N^L} v_i(y_i) \\
&\quad + \sum_{i\in N_f^S} v_i(x_i^*) - v_i(y_i^*) + \sum_{i\in N_f^S} v_i(y_i^*) + \sum_{i\in N_u^S} v_i(y_i) \\
&\leq \frac{1}{s-2}\sum_{i\in N^L\cup N_f^S\cup N_u^S} v_i(y_i) + \sum_{i\in N^L} v_i(y_i) + \sum_{i\in N_f^S} v_i(y_i) + \sum_{i\in N_u^S} v_i(y_i) \\
&= \left(1+\frac{1}{s-2}\right)\sum_{i\in N^L\cup N_f^S\cup N_u^S} \bar{v}_i(k\cdot y_i). \qquad (5.31)
\end{aligned}
$$

Here, the first inequality follows as $x_i^* < x_i = y_i$ for $i \in N_u^S$. The second inequality follows by (5.30) and as for $i \in N_f^S$ it holds that $y_i^* = y_i$.

Finally, combining (5.31) with (5.24) concludes the proof as:

$$
v(\boldsymbol{x}^*) = \sum_{i\in N'} v_i(x_i^*) \leq \left(1+\frac{1}{s-2}\right)\sum_{i\in N'} \bar{v}_i(k\cdot y_i) \leq \left(1+\frac{1}{s-2}\right)\bar{v}(k\cdot \boldsymbol{y}).
$$

□

5.5 Conclusion

In this chapter, we derived budget-feasible mechanisms for procurement auctions under the mild assumption of being able to afford each individual strategic agent entirely. We considered the setting in which partial allocations are allowed, and where the degree of divisibility depends on a given parameter k, i.e., the levels of service, of the instance at hand. As k grows, the granularity of the partial allocation increases, converging to the divisible agent setting as $k \to \infty$. We introduced two truthful and budget-feasible mechanisms specifically tailored to linear valuation functions. For the setting of divisible agents, our 2-approximation establishes a separation with its indivisible counterpart, i.e., $k = 1$, for which a lower bound of $1 + \sqrt{2}$ is known due to Chen et al. [2011]. For the k-level model, the approximation guarantee of our mechanism Greedy-and-Prune improves (decreases) as the granularity of the partial allocation increases, and converges to $1 + \varphi$ if $k \to \infty$.

We also considered valuation functions that are concave and non-decreasing, which we believe are both interesting and relevant for applications as such functions can model diminishing returns of the auctioneer. For the k-level model, our mechanism Sort-&-Reject greedily truncates an optimal fractional solution, which is a common approach in the literature. This leads to an approximation of $2 + \sqrt{3}$, independent of the granularity k. For the divisible model, we can run a mechanism that is similar to Greedy-and-Prune, but takes the concavity of the valuation functions into account. Running this mechanism with a large enough granularity k results in a $(1 + \varphi + \epsilon)$-approximation. It would be interesting to understand whether we can obtain mechanisms with an approximation guarantee that depends on k for the k-level model, and if the $(1 + \varphi + \epsilon)$-approximation can be improved to 2 for the divisible model. 5

Furthermore, for linear valuation functions, two important open problems remain for the divisible setting and its indivisible counterpart, for which the best-possible approximation ratio is in $[e/(e-1),\ 2]$ (due to Anari et al. [2018] and our Theorem 5.4.3) and in $[1 + \sqrt{2},\ 3]$ (due to Chen et al. [2011] and Gravin et al. [2020]), respectively. Any progress on these fronts may give rise to novel techniques, which may also be applicable to problems in richer environments. Finally, a natural direction for future research is to consider additional combinatorial constraints, which can capture interactions among the services of the agents.

Part Three

Inefficiency of Equilibria

Chapter 6

Price of Anarchy Bounds for Simultaneous First-Price Auctions

6.1 Introduction

In recent years, the online advertising landscape has undergone a significant transformation with the emergence of autobidding. Traditionally, advertisers were required to submit fine-grained bids to bidding platforms. Autobidding on the other hand enables advertisers to delegate complex bidding strategies to automated agents, by defining high-level goals and constraints so that the autobidding agent can bid on their behalf. A growing number of platforms are offering the functionality of autobidding, and the vast majority of advertisers are adopting it. According to recent estimates, more than 80% of online advertising is now managed through autobidding agents [Deng et al., 2024]. This motivates the recent line of work on the fundamental aspects of autobidding, which includes among other things the design of optimal bidding strategies and analysis of inefficiency of equilibria.

6

In addition to this shift in bidding strategies, the online advertising industry has also seen a change in auction formats, particularly with the move from second-price to first-price auctions (FPAs). Motivated by the status quo in the online advertising industry, and specifically the recent adoption of FPAs by Google Ad Exchange [Paes Leme et al., 2020], we address the following question in this chapter: Can we derive bounds on the inefficiency of equilibria of simultaneous FPAs in the presence of autobidding agents?

Autobidding. As previously mentioned, autobidding allows an advertiser to specify high-level constraints. Two constraints that have been significantly studied, both separately and together, are budget and return-on-investment (ROI) constraints. With a budget constraint the advertiser can limit their expenses, as this specifies the maximum payment that the autobidding agent can make. We solely consider ROI constraints which can be interpreted as the willingness to pay of the advertiser and specify the maximum payment that the autobidding agent can make for the acquired value.

Beyond constraints, autobidding allows the advertiser to specify the objective that they want their autobidding agent to maximize. The objective can range from utility maximization to value maximization, which have been the two most prominent objectives considered in the autobidding literature. Value maximizers focus on maximizing the acquired value regardless of the price, while satisfying the given constraints. On the other hand, utility maximizers want to maximize the difference between the acquired value and price paid. Value and utility maximizers are the extremes of a spectrum defining the behavior of agents. We study the whole range of this spectrum and also allow for hybrid agents, who have an objective function in which the payments are discounted by a factor in $(0, 1)$.

Price of Anarchy. To analyze the efficiency of FPAs in the autobidding world, we consider the price of anarchy (POA) introduced by Koutsoupias and Papadimitriou [1999]. This metric provides a quantitative assessment of the efficiency loss due to selfish behavior. Given a valuation and equilibrium class, we measures the ratio of the optimal liquid welfare against the liquid welfare achieved in the worst equilibrium. The notion of liquid welfare was first introduced by Dobzinski and Leme [2014] and is an efficiency metric for scenarios in which agents are subject to constraints. The optimal liquid welfare can be interpreted as the maximum payment that the seller can extract. Specifically, an agent's liquid welfare is defined as the minimum of the value they derive from an outcome and their available budget for this outcome, which for us is the ROI constraint. This metric provides a meaningful benchmark, as it acknowledges the limited financial resources of agents.

The liquid welfare POA of FPAs in the autobidding world was recently studied by Deng et al. [2024], both for the model with only value maximizers and the *mixed model* consisting of both value and utility maximizers. Surprisingly, they established a tight bound of 2.18 on the POA for mixed Nash equilibria in the mixed agent model, by leveraging structural insights of equilibria.

Smoothness. A powerful general technique for bounding the POA of games is the *smoothness technique*. Originally introduced by Roughgarden [2015] for strategic games and later extended by Syrgkanis and Tardos [2013] to composable auctions and mechanism design, its strength lies in its *extension theorems*. In the context of composable mechanisms [Syrgkanis and Tardos, 2013], these theorems allow simple POA proofs for pure Nash equilibria in a 'base mechanism' (e.g., a single-item first-price auction) to extend to more complex compositions (e.g., simultaneous or sequential auctions) and broader equilibrium concepts (e.g., mixed and correlated equilibria). This, in combination with the ability to provide strong general bounds, makes the technique particularly appealing. As a result, the smoothness technique has been successfully applied in various settings, but its application to the autobidding world remains unexplored.

6.1.1 Our Contributions

We derive a smoothness framework to extend the study of the price of anarchy of simultaneous first-price auctions in the autobidding world in various ways.

1. We consider the general *hybrid agent model* in which agents have different *types* depending on their payment sensitivity and ROI constraint. This includes the previously studied mixed agent model as a special case, in which agents can either be value or utility maximizers.

2. We significantly broaden the class of valuation functions considered in autobidding by studying fractionally subadditive (or equivalently, *XOS*) valuations (see [Syrgkanis and Tardos, 2013]). In particular, monotone submodular functions form an important special case within this class. To the best of our knowledge, all previous work on autobidding has solely focused on additive valuation functions.

3. We extend our POA bounds to more general equilibrium concepts, including correlated (CE) and coarse correlated equilibria (CCE).

4. We extend our results to simultaneous first-price auctions with reserve prices. This is particularly relevant when the seller can leverage predictions to approximate buyers' valuations (see also the framework by Balseiro et al. [2021a]). Using these predictions to set reserve prices leads to an improved bound on the POA.

Our Smoothness Framework. We derive our results for the aforementioned extensions through a unifying framework, adapting the smoothness framework of Syrgkanis and Tardos [2013] for utility maximizers to our setting in a non-trivial way. However, unlike the original framework, a key technical challenge is the incorporation of different agent types. To the best of our knowledge, prior works leveraging the smoothness technique assume agents are 'alike' in their strategic behavior. Consequently, the (λ, μ)-smoothness parameters established for the base mechanism can be directly lifted to the composed mechanism. In contrast, our setting allows for different agent types, each with its own smoothness parameters (λ_t, μ_t) that must be integrated. A crucial element of our approach is introducing *calibration vectors* that enable us to balance the different smoothness parameters. An upper bound on the POA is then determined by a *POA-revealing mathematical program* (*POA-RMP*). While solving this program presents its own analytical challenges, we are able to find feasible (sometimes optimal) solutions and bound the objective of the POA-RMP, thereby deriving bounds on the POA.

Agent Types	**POA Upper Bound**	**Theorem**
$\sigma_{\max} \leq 0.79$	2	Thm. 6.4.6
$\sigma = \sigma_{\max} > 0.79$	$1 + \sigma(1 + W_0(-e^{-\sigma-1}))^{-1}$	Thm. 6.4.6
$\sigma = \sigma_{\min} \geq 0.741$	$\frac{e^{1/\sigma}}{\sigma(e^{1/\sigma}-1)}$	Thm. 6.4.9

Table 6.1: Overview of POA bounds for CCE of simultaneous FPAs with fractionally subadditive valuations. Here, $\sigma_{\min}$ and $\sigma_{\max}$ refer to the smallest (0 for value maximizers) and largest (1 for utility maximizers) agent types respectively.

Once the above machinery is established, we can focus on proving the smoothness of a single-item first-price auction with reserve price for the different agent types. Similar to the original smoothness approach, our Extension Theorem then allows us to derive (sometimes tight) bounds on the POA of coarse correlated equilibria for simultaneous single-item first-price auctions with reserve prices and fractionally subadditive valuation functions.

Overview of Results. An overview of the POA bounds that we derive for CCE of simultaneous FPAs with fractionally subadditive valuations is provided in Table 6.1. Here, $W_0(\cdot)$ denotes the principal branch of the Lambert W function, introduced at the end of Section 6.2. Note that the upper bound on the POA is equal to 2 for settings in which all agents have a payment sensitivity ($\sigma_{\max}$) of at most 0.79. If $\sigma_{\max} > 0.79$, then the upper bound on the POA increases from 2 to 2.18 as $\sigma_{\max}$ increases to 1. Note that, among other things, this generalizes the result of Deng et al. [2024] who, for mixed Nash equilibria and additive valuations, prove an upper bound on the POA of 2 when only considering value maximizers and of 2.18 in the mixed agent model, to CCE and fractionally subadditive valuations.

6

For simultaneous FPAs with feasible reserve prices and additive valuations, we prove upper bounds on the POA for specific agent types as a function of η. Here, $\eta \in [0, 1)$ indicates the quality of the reserve prices where a larger η indicates a better quality. More specifically, η measures the worst relative gap between the reserve price of an auction and the maximum payment that could be extracted. For the homogeneous case of utility maximizers and value maximizers, we prove an upper bound on the POA of ${}^{e}/_{e-1+\eta}$ and $2 - \eta$ respectively (see Table 6.2). Additionally, we prove that these bounds are tight in Theorems 6.5.4 and 6.5.5. To see how the POA decreases as η improves, we refer to Figure 6.1. Note that when the quality of the reserve prices is almost perfect, i.e., $\eta \to 1$, the POA approaches 1. We also provide bounds for the mixed agent model and for the homogeneous case of $\sigma \in (0, {}^{e-1}/_{e})$ (see also Table 6.2).

We note that, in our setting with reserve prices, it is required that all items

Agent Types $\forall i$	POA Upper Bound	Theorem
$\sigma_i = 1$	$\frac{e}{e-1+\eta}$	Thm. 6.4.15
$\sigma_i = \sigma \in \left(0, \frac{e-1}{e}\right)$	$1 + \frac{1}{\sigma}(\ln(\frac{1-\sigma\eta}{1-\sigma})(1-\sigma))$	Thm. 6.4.15
$\sigma_i = 0$	$2 - \eta$	Thm. 6.4.15
$\sigma_i \in \{0, 1\}$	$(1-\eta) \cdot \frac{2-\eta+W_0(-(1-\eta)^2 e^{\eta-2})}{1-\eta+W_0(-(1-\eta)^2 e^{\eta-2})}$	Thm. 6.4.14

Table 6.2: Overview of POA bounds for well-supported CCE of simultaneous FPAs with feasible reserve prices and additive valuations, with $\eta \in [0, 1)$ indicating the quality of the reserve prices.

are sold under equilibrium outcomes (which we refer to as *well-supported equilibria*). As we demonstrate, for additive valuations this property holds up to and including CE, but fails for the class of CCE. However, building on the work of Kolumbus and Nisan [2022], we show that for a repeated single-item first-price auction with feasible reserve price, CCE arising from regret-minimizing agents are well-supported, regardless of the combination of agent types.

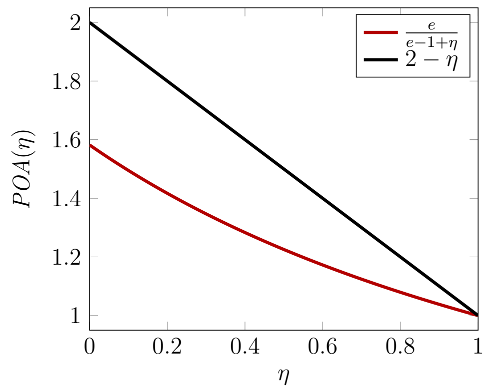

Figure 6.1: Tight POA as a function of η for the homogeneous case of value (black) and utility (red) maximizers.

6.1.2 Related Work

Autobidding. Motivated by the rise of autobidding in online advertising, recent research has focused on value-maximizing agents in an autobidding environment, considering questions such as the existence and efficiency of equilibria, and improving inefficiency. Aggarwal et al. [2019] consider a set of constraints modeling different autobidding products, and among other things, analyze optimal bidding strategies in truthful auctions and derive a liquid welfare price of anarchy bound. Deng et al. [2021] and Balseiro et al. [2021a] explore how auctions with boosts and reserve prices can improve the welfare guarantee. Additionally, Balseiro et al. [2021b, 2022] characterize revenue-optimal auctions under various informational settings, while Aggarwal et al. [2025b] focus on finding optimal bidding strategies in a multi-platform auction setting under budget and ROI constraints. Mehta [2022] and Liaw et al. [2023] considered randomized auctions, were Liaw et al. [2023] considered non-truthful auctions and also showed that for deterministic bids, the POA of the first-price auction is 2. Expanding on this, Deng et al. [2024] analyze more general equilibria with randomized bidding strategies, demonstrating that the POA remains 2 for value-maximizing agents but declines to approximately 2.18 when considering both value-maximizing and utility-maximizing agents. For further details, we refer the interested reader to the survey of Aggarwal et al. [2024].

Price of Anarchy of First-Price Auction. Under the standard model of quasi-linear utilities, the inefficiency of the first-price auction has been considered in the literature since the seminal work of Vickrey [1961]. The POA of CCE was shown by Syrgkanis and Tardos [2013] to be at most $\frac{e}{e-1}$ for simultaneous auctions with XOS valuations. This bound was shown to be tight even for a single auction by Syrgkanis [2014] and for MNE in simultaneous auctions with submodular valuations by Christodoulou et al. [2016]. Beyond the smoothness framework, Feldman et al. [2020] showed an upper bound of 2 on the Bayesian POA for subadditive valuations, and Christodoulou et al. [2016] provided a matching lower bound. More recently, Jin and Lu [2022] showed that the Bayesian POA for the single-item first-price auction is exactly $\frac{e^2}{e^2-1}$.

6

Repeated Auctions. Kolumbus and Nisan [2022] investigate a setting where two utility-maximizing agents employ regret-minimizing algorithms in a repeated auction game. Their findings show that for first-price auctions, the dynamics convergence to co-undominated CCE when agents use mean-based regret-minimizing algorithms [Braverman et al., 2018]. Our work extends these insights to the setting with reserve price and arbitrary agent types. Finally, regret minimization has also been considered in the context of autobidding by other works (see, e.g., [Aggarwal et al., 2025a, Lucier et al., 2024, Balseiro and Gur, 2019]).

6.2 Preliminaries

We consider the setting in which $n \geq 2$ agents simultaneously participate in $m \geq 1$ single-item auctions. We use $N = [n]$ and $M = [m]$ to denote the set agents and auctions, respectively, and $\mathcal{M}_j$ to denote the mechanism that is used to sell the item of auction $j \in M$. Each agent $i \in N$ submits to each auction $j \in M$ a non-negative bid $b_{ij} \in \mathbb{R}_{\geq 0}$. The corresponding bid profile of agent i is denoted by $\boldsymbol{b}_i = (b_{ij})_{j \in M}$. Let $D_i = \mathbb{R}^m_{\geq 0}$ be the set of all bid profiles of agent

i. The aggregated bid profile of all agents is denoted by $\boldsymbol{b} = (\boldsymbol{b}_i)_{i \in N} \in D$, where $D = \times_{i \in N} D_i$.

First-Price Auctions with Reserve Prices. We solely focus on first-price auctions with reserve prices. So, given reserve prices $\boldsymbol{r} = (r_j)_{j \in M} \in \mathbb{R}^m_{\geq 0}$, each mechanism $\mathcal{M}_j$ implements a first-price auction with reserve price r_j, denoted by $FPA(r_j)$: Given the bid profile $\boldsymbol{b}_j = (b_{ij})_{i \in N}$ submitted to auction j, $FPA(r_j)$ only considers the bids that are at least as large as the reserve price r_j, i.e., $b_{ij} \geq r_j$. Among these bids, and using a fixed tie-breaking rule, $\mathcal{M}_j$ allocates the item to the highest bidder, also referred to as the *actual winner* and denoted by $\mathsf{aw}(j) \in \arg\max_{i \in N : b_{ij} \geq r_j} b_{ij}$. The price of the actual winner is equal to their bid $b_{\mathsf{aw}(j)j}$. Note that the item can remain unassigned if $b_{ij} < r_j$ for all $i \in N$.

Let $\boldsymbol{x}_j(\boldsymbol{b}) = (x_{ij}(\boldsymbol{b}))_{i \in N}$ and $\boldsymbol{p}_j(\boldsymbol{b}) = (p_{ij}(\boldsymbol{b}))_{i \in N}$ refer to the respective allocation and payments of $FPA(r_j)$, i.e., for $i = \mathsf{aw}(j)$ we have that $x_{ij}(\boldsymbol{b}) = 1$ and $p_{ij}(\boldsymbol{b}) = b_{ij}$, and for $i \neq \mathsf{aw}(j)$ we have that $x_{ij}(\boldsymbol{b}) = 0$ and $p_{ij}(\boldsymbol{b}) = 0$. Note that both $\boldsymbol{x}_j$ and $\boldsymbol{p}_j$ only depend on the input profile $\boldsymbol{b}_j$; we sometime only use $\boldsymbol{b}_j$ as the argument for notational convenience.

Our global mechanism implements the above mechanisms with reserve prices $\boldsymbol{r} = (r_j)_{j \in M}$ simultaneously. That is, given a bid profile $\boldsymbol{b}$, the outcome $\mathcal{M}(\boldsymbol{r}, \boldsymbol{b}) = (\boldsymbol{x}(\boldsymbol{b}), \boldsymbol{p}(\boldsymbol{b}))$ is determined by the allocation $\boldsymbol{x}(\boldsymbol{b}) = (\boldsymbol{x}_j(\boldsymbol{b}))_{j \in M}$ and the payments $\boldsymbol{p}(\boldsymbol{b}) = (\boldsymbol{p}_j(\boldsymbol{b}))_{j \in M}$ obtained by running m first-price auctions with reserve prices $\boldsymbol{r}$ simultaneously. We define $p_i(\boldsymbol{b}) = \sum_{j \in M} p_{ij}(\boldsymbol{b})$ as the total payment of an agent i over all m auctions and $P(\boldsymbol{b}) = \sum_{i \in N} p_i(\boldsymbol{b})$ as the sum of all payments. Further, we define $\boldsymbol{x}_i(\boldsymbol{b}) = (x_{ij}(\boldsymbol{b}))_{j \in M} \subseteq \{0,1\}^m$ to denote the allocation of an agent $i \in N$ and $\boldsymbol{X}$ to denote the set of feasible allocations, i.e., $\boldsymbol{X} = \{\boldsymbol{x}(\boldsymbol{b}) = (\boldsymbol{x}_i(\boldsymbol{b}))_{i \in N} \mid \sum_{i \in N} x_{ij}(\boldsymbol{b}) \leq 1, \forall j \in M\}$. We slightly overload notation and use $\boldsymbol{x}_i(\boldsymbol{b})$ also to refer to the set of items assigned to an agent i, i.e., $\boldsymbol{x}_i(\boldsymbol{b}) = \{j \in M \mid x_{ij}(\boldsymbol{b}) = 1\} \subseteq M$. Additionally, we sometimes omit the argument $\boldsymbol{b}$ when it is irrelevant or clear from context.

For ease of notation, we use $FPA(m, \boldsymbol{r})$ and $FPA(m)$, respectively, to refer to m simultaneous first-price auctions with reserve prices $\boldsymbol{r}$ and without reserve prices. We use $FPA(r)$ to indicate that we consider a single-item first-price auction with reserve price r. If $m = 1$, we drop the auction index $j = 1$ from all our notation.

Valuation Functions. Each agent $i \in N$ has a valuation function $v_i : 2^M \to \mathbb{R}_{\geq 0}$ over the subsets of the auctioned items, where $v_i(S)$ is the value that agent i obtains when receiving the items in $S \subseteq M$. We assume w.l.o.g. that $v_i(\emptyset) = 0$. Also, we assume that v_i is *monotone*, i.e., $v_i(S) \leq v_i(T)$ for all $S \subseteq T \subseteq M$. We use $\mathcal{V}_i$ to denote the class of valuation functions of agent i and let $\mathcal{V} = \times_{i \in N} \mathcal{V}_i$ be the set of all valuation functions of the agents. We use $\boldsymbol{v} = (v_i)_{i \in N} \in \mathcal{V}$ to refer to the profile of valuation functions of the agents. We consider different classes of valuation functions:

Definition 6.2.1. Let $v_i : 2^M \to \mathbb{R}_{\geq 0}$ be a valuation function.

- v_i is *additive* if there exist additive valuations $(v_{ij})_{j \in M} \in \mathbb{R}^m_{\geq 0}$ such that for every subset $S \subseteq M$, it holds that $v_i(S) = \sum_{j \in S} v_{ij}$.
- v_i is *submodular* if $v_i(S \cup \{j\}) - v_i(S) \geq v_i(T \cup \{j\}) - v_i(T)$ for all $S \subseteq T \subseteq M$.
- v_i is *fractionally subadditive* (or, *XOS*), if there exists a class $\mathcal{L}_i = \{(v^\ell_{ij})_{j \in M} \in \mathbb{R}^m_{\geq 0}\}$ of additive valuations such that for every subset $S \subseteq M$, it holds that $v_i(S) = \max_{\ell \in \mathcal{L}_i} \sum_{j \in S} v^\ell_{ij}$.

Let $\mathcal{V}_{\text{ADD}}$, $\mathcal{V}_{\text{SUB}}$ and $\mathcal{V}_{\text{XOS}}$ refer to the set of additive, submodular and fractionally subadditive (XOS) valuation functions, respectively.

Randomized Bid Profiles. Let π be a probability distribution over the set of bid profiles in D. Throughout the paper, we use $\boldsymbol{B} \sim \pi$ to denote a random bid profile that is drawn from π. For ease of notation, we often omit the explicit reference to π and identify $\boldsymbol{B}$ with π. Further, we define Δ_i as the space of randomized bid profiles of an agent i and use $\boldsymbol{b} \in D$ to refer to a deterministic (or pure) bid profile. For a randomized bid profile $\boldsymbol{B}$, we use $f_{\boldsymbol{B}}$ and $F_{\boldsymbol{B}}$ to refer to the probability density function (PDF) and cumulative distribution function (CDF) of $\boldsymbol{B}$, respectively. The support of $\boldsymbol{B}$, denoted by $\text{supp}(\boldsymbol{B})$, refers to the set of bid profiles that have positive density under $\boldsymbol{B}$, i.e., $\text{supp}(\boldsymbol{B}) = \{\boldsymbol{b} \in D \mid f_{\boldsymbol{B}}(\boldsymbol{b}) > 0\}$. If $\text{supp}(\boldsymbol{B}) = \{\boldsymbol{b}\}$ then $\boldsymbol{B}$ chooses $\boldsymbol{b}$ deterministically; we write $\boldsymbol{B} = \boldsymbol{b}$ to indicate this. We use $\text{supp}_i(\boldsymbol{B})$ to refer to the set of bid profiles $\boldsymbol{b}_i$ of agent i that have positive density under $\boldsymbol{B}$. The *marginal* $\boldsymbol{B}_{-i}$ of $\boldsymbol{B}$ is defined by the following PDF:

$$\forall \boldsymbol{b}_{-i} \in D_{-i} : \qquad f_{\boldsymbol{B}_{-i}}(\boldsymbol{b}_{-i}) = \int_{D_i} f_{\boldsymbol{B}}(\boldsymbol{b}_i, \boldsymbol{b}_{-i}) d\boldsymbol{b}_i.$$

Given a bid profile $\boldsymbol{b}'_i$ of agent i, we denote by $(\boldsymbol{b}'_i, \boldsymbol{B}_{-i})$ the random bid profile that we obtain from $\boldsymbol{B}$ when agent i bids $\boldsymbol{b}'_i$ deterministically and the other agents bid according to the marginal $\boldsymbol{B}_{-i}$.

Hybrid Agent Model and Return-On-Investment. Each agent $i \in N$ is associated with a *gain function* $g_i : D \to \mathbb{R}$ that i wants to maximize; for each bid profile $\boldsymbol{b}$, g_i is defined as:

$$g_i(\boldsymbol{b}) = v_i(\boldsymbol{x}_i(\boldsymbol{b})) - \sigma_i \cdot p_i(\boldsymbol{b}). \tag{6.1}$$

Here, $\sigma_i \in [0, 1]$ is the *payment sensitivity* (or, *sensitivity*) parameter of agent i. Intuitively, a larger sensitivity σ_i indicates that agent i is more sensitive towards the payments (in the sense of being negatively affected). In particular, agent i is a value maximizer if $\sigma_i = 0$, and a utility maximizer if $\sigma_i = 1$. Our model thus allows us to capture a large spectrum of agents' sensitivities, ranging from value maximizers to utility maximizers. This general model is also termed the *hybrid agent model* (see, e.g., [Balseiro et al., 2021a, Aggarwal et al., 2024]). Most models studied in the literature consider a mix of value and utility maximiers only, i.e., $\sigma_i \in \{0, 1\}$ for all $i \in N$; we refer to this special case as the *mixed agent model.*

Each agent i has a *return-on-investment (ROI)* constraint that must be satisfied (see, e.g., [Aggarwal et al., 2024]). Given a random bid profile $\boldsymbol{B}$, the ROI constraint of an agent i enforces that the expected total payment of i is at most a factor $\tau_i \in \mathbb{R}_{>0}$ of their expected valuation for the received items:

$$\mathbb{E}\left[p_i(\boldsymbol{B})\right] \leq \tau_i \cdot \mathbb{E}[v_i(\boldsymbol{x}_i(\boldsymbol{B}))]. \tag{6.2}$$

Here, τ_i is also called the *target* parameter of agent i. We define $\mathcal{R}_i$ as the set of random bid profiles $\boldsymbol{B}$ that satisfy the ROI constraint (6.2) for agent i. Note that for a deterministic bid profile $\boldsymbol{B} = \boldsymbol{b}$, the ROI constraint (6.2) reduces to $p_i(\boldsymbol{b}) \leq \tau_i v_i(\boldsymbol{x}_i(\boldsymbol{b}))$.

We assume w.l.o.g. that $\tau_i\sigma_i \leq 1$ for all agents $i \in N$. Basically, if an agent with sensitivity $\sigma_i > 0$ has a target parameter $\tau_i > 1/\sigma_i$, then i prefers to withdraw whenever they are asked to pay more than $1/\sigma_i$ their valuation. In other words, there are no equilibria in which the expected payment of agent i exceeds $1/\sigma_i$ times their expected value, as agent i would have a negative expected gain and be better off by bidding 0 deterministically in each auction. In essence, this has the same effect as capping the target parameter at $1/\sigma_i$. Note that this argument applies only to agents with sensitivity $\sigma_i > 0$. In particular, the target parameter τ_i of each value maximizer i with $\sigma_i = 0$ remains unconstrained as $\tau_i\sigma_i \leq 1$ is always satisfied.

Equilibrium Notions. The objective of each agent i is to determine a random bid profile $\boldsymbol{B}_i$ that maximizes their gain, subject to their ROI constraint and the strategy $\boldsymbol{B}_{-i}$ of the other agents. That is, each agent i solves the following optimization problem, which is often defined as the *agent's optimization problem*:

6

$$\max_{\boldsymbol{B}_i} \ \mathbb{E}[g_i(\boldsymbol{B}_i, \boldsymbol{B}_{-i})] \quad \text{subject to} \quad (\boldsymbol{B}_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i.$$

The resulting bid profile $\boldsymbol{B}$ can be interpreted as an equilibrium, when an important refinement is made to the equilibrium notions as defined in Chapter 2. That is, for each agent i only deviations that satisfy the ROI constraint are considered.

Definition 6.2.2. Let $\boldsymbol{B}$ be a random bid profile satisfying the ROI constraint of each agent, i.e., $\boldsymbol{B} \in \mathcal{R}_i$ for all $i \in N$.

- $\boldsymbol{B}$ is a *CCE* if for every agent $i \in N$ we have:

$$\mathbb{E}[g_i(\boldsymbol{B})] \geq \mathbb{E}[g_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})] \qquad \forall (\boldsymbol{B}'_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i. \tag{6.3}$$

- $\boldsymbol{B}$ is a *CE* if for every agent $i \in N$ and every swapping function h : $\text{supp}_i(\boldsymbol{B}) \to \Delta_i$ we have:

$$\mathbb{E}[g_i(\boldsymbol{B})] \geq \mathbb{E}[g_i(h(\boldsymbol{B}_i), \boldsymbol{B}_{-i})] \qquad \forall (h(\boldsymbol{B}_i), \boldsymbol{B}_{-i}) \in \mathcal{R}_i. \tag{6.4}$$

- $\boldsymbol{B}$ is a *MNE* if $\boldsymbol{B} = \prod_{i \in [n]} \boldsymbol{B}_i$ and for every agent $i \in N$ we have:

$$\mathbb{E}[g_i(\boldsymbol{B})] \geq \mathbb{E}[g_i(\boldsymbol{B}_i', \boldsymbol{B}_{-i})] \qquad \forall (\boldsymbol{B}_i', \boldsymbol{B}_{-i}) \in \mathcal{R}_i. \tag{6.5}$$

Note that for a correlated equilibrium $\boldsymbol{B}$, $(h(\boldsymbol{B}_i), \boldsymbol{B}_{-i}) \in \mathcal{R}_i$ implies that:

$$\mathbb{E}\left[\mathbb{E}\left[p_i(h(\boldsymbol{b}_i), \boldsymbol{B}_{-i}) \mid \boldsymbol{B}_i = \boldsymbol{b}_i\right]\right] \leq \tau_i \cdot \mathbb{E}\left[\mathbb{E}\left[v_i(x_i(h(\boldsymbol{b}_i), \boldsymbol{B}_{-i})) \mid \boldsymbol{B}_i = \boldsymbol{b}_i\right]\right],$$

where the outer expectation is for $\boldsymbol{B}_i$ and the inner expectation is for $h(\boldsymbol{b}_i)$ and $\boldsymbol{B}_{-i}$. Also note that the gain in the definition of a *CE* $\boldsymbol{B}$ is evaluated for the entire support of $\boldsymbol{B}$. Note that it holds that *MNE* $\subseteq$ *CE* $\subseteq$ *CCE*, where *CCE*, *CE* and *MNE* refer to the sets of equilibria. Further details on the equilibrium hierarchy and why pure Nash equilibria are not considered[1] can be found in Appendix A.1. We say that an equilibrium $\boldsymbol{B}$ is *well-supported* if all the items are always sold. This is an important property that we need to exploit in our smoothness lemmas for auctions with reserve prices. As we prove in Section 6.5, for most equilibrium notions and additive valuations this is not a restriction. We write $\boldsymbol{x}_j(\boldsymbol{b}_j) \neq \mathbf{0}$ to indicate that the item is sold in auction j, and we write $\boldsymbol{x}(\boldsymbol{b}) \neq \mathbf{0}$ to indicate that all items are sold under $\boldsymbol{b}$, i.e., $\boldsymbol{x}_j(\boldsymbol{b}_j) \neq \mathbf{0}$ for all $j \in M$.

Liquid Price of Anarchy. We use *liquid welfare* as the social welfare objective which is also the standard benchmark in the autobidding literature [Aggarwal et al., 2024]. Intuitively, the liquid welfare measures the maximum amount of payments one can extract from the agents. Given an allocation $\boldsymbol{x}$, the liquid welfare of agent i with $\sigma_i > 0$ is $\min(v_i(\boldsymbol{x}_i)/\sigma_i, \tau_i v_i(\boldsymbol{x}_i))$, where the first term reflects that the gain g_i of agent i always remains non-negative and the second term is due to the ROI constraint. Note that because $\tau_i \sigma_i \leq 1$ (as discussed above), the liquid welfare of i is $\tau_i v_i(\boldsymbol{x}_i)$. The liquid welfare of agent i with $\sigma_i = 0$ is $\tau_i v_i(\boldsymbol{x}_i)$, i.e., due to the ROI constraint only. As a consequence, the liquid welfare of an agent evaluates to $\tau_i v_i(\boldsymbol{x}_i)$, independently of their type. So given a feasible allocation $\boldsymbol{x} \in \boldsymbol{X}$, the liquid welfare of $\boldsymbol{x}$ is defined as:

$$LW(\boldsymbol{x}) = \sum_{i \in N} \tau_i v_i(\boldsymbol{x}_i).$$

[1] As the ROI constraints must be satisfied, it is not true that a pure Nash equilibrium w.r.t. deterministic deviations survives when randomization is allowed.

We use $\boldsymbol{\tau v}$ to refer to the vector of τ_i-adjusted agent valuations, i.e., $\boldsymbol{\tau v} = (\tau_i v_i)_{i \in N}$. An optimal allocation $\boldsymbol{x}^* = \boldsymbol{x}^*(\boldsymbol{\tau v}) \in \boldsymbol{X}$ maximizes the liquid welfare over all feasible allocations. We use $OPT(\boldsymbol{\tau v}) = LW(\boldsymbol{x}^*)$ to denote the maximum liquid welfare with respect to $\boldsymbol{\tau v}$.

For additive valuation functions $\boldsymbol{v} = (v_{ij})_{i \in N, j \in M}$, we say that an agent $i \in N$ is the *rightful winner* of auction $j \in M$ if i is an agent with maximum τ_i-adjusted valuation. We denote this agent by $\mathsf{rw}(j) \in \arg\max_{i \in N} \tau_i v_{ij}$ and in case of ties, we let $\mathsf{rw}(j)$ denote the winner of auction j in the considered optimal allocation. Thus, in the additive case, the optimal liquid welfare is $OPT(\boldsymbol{\tau v}) = \sum_{j \in M} \tau_{\mathsf{rw}(j)} v_{\mathsf{rw}(j)j}$.

The *price of anarchy (POA)* with respect to a given class of valuation functions and an equilibrium notion is defined as the worst-case ratio of the optimal liquid welfare to the expected liquid welfare of any equilibrium. More formally, given a class of valuation functions $\mathcal{V}_{\text{VAL}}$ with $\text{VAL} \in \{\text{ADD}, \text{SUB}, \text{XOS}\}$, let $I = (N, M, \boldsymbol{\sigma}, \boldsymbol{\tau}, \boldsymbol{v}) \in \mathcal{I}_{\text{VAL}}$ denote an instance with $\boldsymbol{v} \in \mathcal{V}_{\text{VAL}}$. Then, given an equilibrium notion $EQ \in \{PNE, MNE, CE, CCE\}$, we define the POA of EQ with respect to $\mathcal{V}_{\text{VAL}}$ as:

$$EQ\text{-}POA(\mathcal{V}_{\text{VAL}}) = \sup_{I \in \mathcal{I}_{\text{VAL}}} \sup_{\boldsymbol{B} \in EQ(I)} \frac{\sum_{i \in N} \tau_i v_i(x_i^*(\boldsymbol{\tau v}))}{\mathbb{E}[\sum_{i \in N} \tau_i v_i(x_i(\boldsymbol{B}))]}.$$

Reserve Prices. In auctions with reserve prices, an item might not be sold simply because the reserve price is too high. As in previous works [Balseiro et al., 2021a, Deng et al., 2024], we therefore make the assumption that for additive valuations the reserve price r_j is at most the $\tau_{\mathsf{rw}(j)}$-adjusted value of the rightful winner $\mathsf{rw}(j)$ for each auction $j \in M$, i.e., $\tau_{\mathsf{rw}(j)} v_{\mathsf{rw}(j)j}$. We introduce a parameter η_j that measures the relative gap between r_j and $\tau_{\mathsf{rw}(j)} v_{\mathsf{rw}(j)j}$. More formally, we define $\eta_j \in [0, 1)$ such that $r_j = \eta_j \tau_{\mathsf{rw}(j)} v_{\mathsf{rw}(j)j}$ and define $\eta = \min_{j \in N} \eta_j$ as the smallest relative gap. Note that the price of anarchy could be unbounded if $\eta = 1$, as it is not guarantee that the optimal liquid welfare is achieved in each equilibrium. Specifically, consider a single auction with $r = \tau_{\mathsf{rw}} v_{\mathsf{rw}}$. Suppose that agent rw is a utility maximizer and the only agent that can meet the reserve price. Then rw deterministically bidding 0 is a *MNE* with a liquid welfare of 0, while the optimal liquid welfare is $\tau_{\mathsf{rw}} v_{\mathsf{rw}} > 0$. This would lead to an unbounded *POA*. We therefore only consider reserve prices $\boldsymbol{r}$ that satisfy $\eta_j \in [0, 1)$ for each auction $j \in M$ and call such call reserve prices *feasible*. For fractionally subadditive valuations $\boldsymbol{v} \in \mathcal{V}_{\text{XOS}}$, we say that reserve prices are feasible if, for all agents $i \in N$, they are feasible for the *additive representatives* $(v_{ij}^*)_{j \in M}$ with respect to their optimal allocation $\boldsymbol{x}_i^*$ (see Section 6.3.2 for more details). As observed by Balseiro et al. [2021a], a reserve price r_j might be interpreted as a prediction of the $\tau_{\mathsf{rw}(j)}$-adjusted value of the rightful winner. In this context, η is an indicator of the quality of the prediction, and as η increases so does the quality of the prediction. For a simultaneous FPA with reserve prices $\boldsymbol{r}$, we use $\mathcal{I}_{\text{VAL}^+}$ to denote the class of instances $I = (N, M, \boldsymbol{\sigma}, \boldsymbol{\tau}, \boldsymbol{v}, \boldsymbol{r}) \in \mathcal{I}_{\text{VAL}^+}$ with $\boldsymbol{v} \in \mathcal{V}_{\text{VAL}}$, and define the *EQ-POA*$(\mathcal{V}_{\text{VAL}^+})$ accordingly.

Uniform Target Parameter. All previous works studying the mixed agent model assume that the ROI constraint is imposed on the value maximizers only, while the utility maximizers remain unconstrained. As argued in [Deng et al., 2024], this is equivalent to assuming a uniform ROI target parameter of $\tau_i = 1$ for all agents $i \in N$. Without loss of generality, we can extend the assumption of a uniform target parameter to our hybrid agent model. To see this, consider an instance $I = (N, M, \boldsymbol{\sigma}, \boldsymbol{\tau}, \boldsymbol{v})$ with arbitrary valuations, sensitivities and target parameters. As mentioned earlier, we can assume w.l.o.g. that $\sigma_i \tau_i \leq 1$ for every agent $i \in N$. Now consider the instance $I' = (N, M, \boldsymbol{\sigma}', \boldsymbol{\tau}', \boldsymbol{v}')$ that is just a scaled version of I with $\boldsymbol{v}' = \boldsymbol{\tau v}$, $\boldsymbol{\sigma}' = \boldsymbol{\tau\sigma}$ and $\boldsymbol{\tau}' = \mathbf{1}$. Note that $\sigma_i' \leq 1$ for each agent $i \in N$. For the instance I', the gain g_i' of an agent i for a bid profile $\boldsymbol{b}$ is $g_i'(\boldsymbol{b}) = v_i'(x_i(\boldsymbol{b})) - \sigma_i' p_i(\boldsymbol{b}) = \tau_i v_i(x_i(\boldsymbol{b})) - \tau_i \sigma_i p_i(\boldsymbol{b}) = \tau_i g_i(\boldsymbol{b})$, which is just the gain function for the instance I scaled by the target parameter of agent i. And for the instance I', the ROI constraint of agent i is $p_i(\boldsymbol{b}) \leq \tau_i' v_i'(x_i(\boldsymbol{b})) = \tau_i v_i(x_i(\boldsymbol{b}))$, which coincides with the ROI constraint for the instance I. Therefore, when considering the instance I', agents still want to maximize $g_i(\boldsymbol{b})$ and their bidding space remains unchanged compared to I. And so, the set of equilibria for the instances I and I' coincide. Additionally, note that the liquid welfare also coincides, as for the instance I' the liquid welfare of an agent i given a bid profile $\boldsymbol{b}$ is $\tau_i' v_i'(x_i(\boldsymbol{b})) = \tau_i v_i(x_i(\boldsymbol{b}))$, which is equal to the liquid welfare of agent i for the instance I.

In the subsequent sections, we will therefore consider a uniform target parameter $\tau_i = 1$ for all agents $i \in N$. In this case, agents will behave differently depending on their sensitivity parameter. Depending on the sensitivity parameter σ_i of an agent i, we define the *type* of agent i. The set of different agent types is defined as $T = \{t \mid \exists i \in N \text{ with } \sigma_i = t\}$. We use $N_t \subseteq N$ to refer to the subset of agents having type t, i.e., $N_t = \{i \in N \mid \sigma_i = t\}$. Also, given a type $t \in T$, we use σ_t to refer to the sensitivity of type t for clarity.

The Lambert W Function. In order to derive bounds on the POA analytically, we use the Lambert W function, which is the multi-valued inverse of $f(z) = ze^z$. We exclusively use the primary real branch of the W function, also called the *principal branch* and commonly denoted by W_0, that we define below. We refer the interested reader to [Corless et al., 1996].

Definition 6.2.3. The principal branch of the Lambert W function is the real function $W_0 : [-\frac{1}{e}, \infty) \to [-1, \infty)$ such that $z = W_0(z)e^{W_0(z)}$.

Fact 6.2.4. For every $z > -\frac{1}{e}$, the derivative of $W_0(z)$ is $\frac{W_0(z)}{z(1+W_0(z))}$.

For readability, the proofs of some auxiliary lemmas are omitted and can be found in Appendix A.

6.3 Smoothness and the Extension Theorem

We introduce our smoothness framework for simultaneous first-price auctions with reserve prices under the hybrid agent model. The advantage of the smoothness framework is that we can focus on proving smoothness for a single auction only. Our Extension Theorem then enables us to derive bounds on the POA for simultaneous auctions with fractionally subadditive valuation functions. In contrast to the original framework by Syrgkanis and Tardos [2013], the technical challenge is to incorporate different agent types.

We first need to introduce the notion of ROI-restricted bid profiles.

Definition 6.3.1. Let $\boldsymbol{B}_i'$ be a random bid profile of agent i over D_i. We say that $\boldsymbol{B}_i'$ is *ROI-restricted* if for each $\boldsymbol{b}_{-i} \in D_{-i}$, it holds that:

$$\mathbb{E}[p_i(\boldsymbol{B}_i', \boldsymbol{b}_{-i})] \leq \mathbb{E}[v_i(\boldsymbol{x}_i(\boldsymbol{B}_i', \boldsymbol{b}_{-i}))].$$

We can now introduce our smoothness notion.

Definition 6.3.2: (Smoothness). Consider a first-price auction $FPA(r)$ with a feasible reserve price. Let the rightful winner i be of type $t \in T$. Then, $FPA(r)$ is (λ_t, μ_t)*-smooth for type* t with $\lambda_t, \mu_t > 0$, if there exists a random bid $B_i' = B_i'(\boldsymbol{v})$ such that B_i' is ROI-restricted, and for every bid profile $\boldsymbol{b}$ such that $\boldsymbol{x}(\boldsymbol{b}) \neq \boldsymbol{0}$ it holds that:

$$\mathbb{E}\left[g_i\left(B_i', \boldsymbol{b}_{-i}\right)\right] \geq \lambda_t v_i - \mu_t p_{\mathsf{aw}(\boldsymbol{b})}(\boldsymbol{b}). \tag{6.6}$$

We remark that, crucially, the random deviation B_i' of the rightful winner i may depend on the valuations $\boldsymbol{v}$ but *not* on the bid profile $\boldsymbol{b}$. Note that (6.6) needs to hold only for bid profiles $\boldsymbol{b}$ under which the item is sold; this condition is redundant if the reserve price is $r = 0$.

6.3.1 Smoothness Lemmas

We present two smoothness lemmas for the general model of a first-price auction with a feasible reserve price and hybrid agents. We start with a simple smoothness lemma for value maximizers.

Lemma 6.3.3. *Consider a first-price auction* $FPA(r)$ *with a feasible reserve price. Let the rightful winner be of type* $t \in T$ *with* $\sigma_t = 0$. *Then, for every* $\mu \in (0, (1-\eta)^{-1}]$, $FPA(r)$ *is* (λ_t, μ_t)*-smooth for type* t *with* $\lambda_t = \mu_t = \mu$.

Proof: Assume that the rightful winner i is of type $t \in T$ with $\sigma_t = 0$. We need to show that there exists a ROI-restricted random bid B_i' such that for every bid profile $\boldsymbol{b}$ with $\boldsymbol{x}(\boldsymbol{b}) \neq \mathbf{0}$ and $\mathsf{aw} = \mathsf{aw}(\boldsymbol{b})$ the actual winner, it holds that:

$$\mathbb{E}\left[g_i\left(B_i', \boldsymbol{b}_{-i}\right)\right] \geq \mu v_i - \mu p_{\mathsf{aw}}(\boldsymbol{b}). \tag{6.7}$$

Let $B_i' = B_i'(\boldsymbol{v})$ be a random unilateral deviation of i drawn from $[\eta v_i, v_i]$ with CDF $F(z) = F_{B_i'}(z) = \mu z / v_i + 1 - \mu$. Note that the domain is well-defined as $\eta \in [0, 1)$, and it is easy to verify that $F(\cdot)$ is non-negative and increasing over $[\eta v_i, v_i]$ and $F(v_i) = 1$. Also, B_i' is ROI-restricted as the condition is even pointwise satisfied, i.e., for every $z \in [\eta v_i, v_i]$ it holds that $p_i(z, \cdot) \leq v_i(\boldsymbol{x}_i(z, \cdot))$.

It remains to show that B_i' satisfies (6.7). Note that the expected gain of i is always non-negative, as i bids above v_i with 0 probability. Thus, (6.7) holds trivially if $v_i \leq p_{\mathsf{aw}}(\boldsymbol{b})$. Therefore, assume that $v_i > p_{\mathsf{aw}}(\boldsymbol{b})$ and define $\theta_i := \max(\eta v_i, \max_{j \neq i} b_j)$. For every $z \geq \theta_i$, i wins the item under the bid profile $(z, \boldsymbol{b}_{-i})$ and pays $p_i(z, \boldsymbol{b}_{-i}) = z$. As the item is sold under the bid profile $\boldsymbol{b}$ by assumption, the actual winner under $\boldsymbol{b}$ either pays the reserve price or their maximum bid, i.e., $p_{\mathsf{aw}}(\boldsymbol{b}) = \max(\eta v_i, \max_j b_j)$. We obtain:

$$\theta_i = \max\left(\eta v_i, \max_{j \neq i} b_j\right) \leq \max\left(\eta v_i, \max_j b_j\right) = p_{\mathsf{aw}}(\boldsymbol{b}) < v_i.$$

This leads to the desired result as:

$$\begin{aligned}\mathbb{E}\left[g_i\left(B_i', \boldsymbol{b}_{-i}\right)\right] &= v_i(1 - F(\theta_i)) = v_i\left(1 - \left(\frac{\mu\theta_i}{v_i} + 1 - \mu\right)\right) \\ &= \mu v_i - \mu\theta_i \geq \mu v_i - \mu p_{\mathsf{aw}}(\boldsymbol{b}).\end{aligned}$$

Note that the first equality holds because the sensitivity of i is $\sigma_t = 0$. □

6

In comparison to the existing smoothness proofs without reserve price, the intuition is that a reserve price allows us to increase the probability mass of the random deviation for larger bids. This provides a better trade-off in terms of the smoothness parameters. Our smoothness lemma for agent types t with $\sigma_t > 0$ follows the same approach, but is technically more involved.

Lemma 6.3.4. *Consider a first-price auction FPA(r) with a feasible reserve price. Let the rightful winner* rw *be of type $t \in T$ with $\sigma = \sigma_t > 0$, so $r = \eta v_{\mathsf{rw}}$ with $\eta \in [0, 1)$. Then, for every μ satisfying:*

$$\begin{cases} \mu \geq \sigma\left(\ln\left(\frac{1-\sigma\eta}{1-\sigma}\right)\right)^{-1}, & \text{if } \sigma < 1, \\ \mu > 0, & \text{if } \sigma = 1, \end{cases} \tag{6.8}$$

FPA(r) is (λ_t, μ_t)-smooth for type t with $\mu_t = \mu$ and:

$$\lambda_t = \frac{\mu}{\sigma}\left(1 - \frac{1 - \sigma\eta}{e^{\sigma/\mu}}\right).$$

We elaborate on the expression in (6.8) in Lemma 6.3.4. First, note that $\sigma \in (0,1]$, as $\sigma > 0$ and $\sigma \leq 1$ by assumption. The $\ln(\cdot)$ expression decreases as η increases and converges to 0 (from above) as $\eta \to 1$; the lower bound restriction on μ thus increases as η increases. As $\eta < 1$, the $\ln(\cdot)$ expression is well-defined for all combinations of σ and η, except when $\sigma = 1$. In the latter case, we only impose the restriction that $\mu > 0$.

Given μ, we define a parameter $\gamma = \gamma(\mu)$ as follows:

$$\gamma(\mu) = \frac{1}{\sigma}\left(1 - \frac{1-\sigma\eta}{e^{\sigma/\mu}}\right), \tag{6.9}$$

which will be useful in the proof of Lemma 6.3.4. Note that γ is well-defined because $\sigma > 0$ by assumption.

The following corollary is an immediate consequence of the definitions above.

Corollary 6.3.5. *Let μ satisfy* (6.8) *and let γ be defined as in* (6.9)*. Then,* $\gamma \in [\eta, 1]$.

Proof: Note that the interval $[\eta, 1]$ is well-defined because $\eta \in [0,1)$ by assumption. We first prove the lower bound on γ. Note that $e^{\sigma/\mu} > 1$ as $\mu > 0$ and $\sigma > 0$, and therefore:

$$\gamma = \frac{1}{\sigma}\left(1 - \frac{1-\sigma\eta}{e^{\sigma/\mu}}\right) > \frac{1}{\sigma}\cdot\sigma\eta = \eta.$$

For the upper bound on γ, we have that:

$$\gamma = \frac{1}{\sigma}\left(1 - \frac{1-\sigma\eta}{e^{\sigma/\mu}}\right) \leq \frac{1}{\sigma}\left(1 - \frac{1-\sigma\eta}{e^{\ln\left(\frac{1-\sigma\eta}{1-\sigma}\right)}}\right) = \frac{1}{\sigma}\cdot\sigma = 1,$$

where the inequality follows from (6.8) and because e^x is non-decreasing in x. □ 6

We now continue with the proof of Lemma 6.3.4.

Proof of Lemma 6.3.4: Assume that the rightful winner i is of type $t \in T$ with $\sigma = \sigma_t > 0$. We need to show that there exists a ROI-restricted random bid B_i' such that for every bid profile $\boldsymbol{b}$ with $\boldsymbol{x}(\boldsymbol{b}) \neq \boldsymbol{0}$ and $\mathsf{aw} = \mathsf{aw}(\boldsymbol{b})$ the actual winner, it holds that:

$$\mathbb{E}\left[g_i\left(B_i', \boldsymbol{b}_{-i}\right)\right] \geq \mu\gamma v_i - \mu p_{\mathsf{aw}}(\boldsymbol{b}). \tag{6.10}$$

Let $B_i' = B_i'(\boldsymbol{v})$ be a random unilateral deviation of i drawn from $[\eta v_i, \gamma v_i]$ with PDF $f(z) = f_{B_i'}(z) = \mu/(v_i - \sigma z)$. Note that the domain is well-defined as $\gamma \in [\eta, 1]$ by Corollary 6.3.5, and that $f(\cdot)$ is non-negative. Also note that:

$$\begin{aligned}\int_{\eta v_i}^{\gamma v_i} f(z)dz &= \int_{\eta v_i}^{\gamma v_i} \frac{\mu}{v_i - \sigma z}dz = \mu\int_{\eta v_i}^{\gamma v_i}\left(\frac{-\ln\left(v_i - \sigma z\right)}{\sigma}\right)' dz \\ &= \frac{\mu}{\sigma}\ln\left(\frac{1-\sigma\eta}{1-\sigma\gamma}\right) = \frac{\mu}{\sigma}\ln\left(e^{\sigma/\mu}\right) = 1,\end{aligned}$$

where the fourth equality follows from the definition of γ in (6.9). Furthermore, B'_i is ROI-restricted as the condition is even pointwise satisfied, i.e., for $z \in [\eta v_i, \gamma v_i]$ it holds that $p_i(z, \cdot) \leq \gamma v_i(\boldsymbol{x}_i(z, \cdot)) \leq v_i(\boldsymbol{x}_i(z, \cdot))$, as $\gamma \in [\eta, 1]$ by Corollary 6.3.5.

It remains to show that B'_i satisfies (6.10). Note that the expected gain of i is always non-negative, as i bids above v_i with 0 probability and $\sigma \leq 1$. Thus, (6.10) holds trivially if $\gamma v_i \leq p_{\mathsf{aw}}(\boldsymbol{b})$. Therefore, assume that $\gamma v_i > p_{\mathsf{aw}}(\boldsymbol{b})$ and define $\theta_i := \max(\eta v_i, \max_{j \neq i} b_j)$. Then, for every $z \geq \theta_i$, i wins the item under bid profile $(z, \boldsymbol{b}_{-i})$ and pays $p_i(z, \boldsymbol{b}_{-i}) = z$. As the item is sold under bid profile $\boldsymbol{b}$ by assumption, the actual winner under $\boldsymbol{b}$ either pays the reserve price or their maximum bid, i.e., $p_{\mathsf{aw}}(\boldsymbol{b}) = \max(\eta v_i, \max_j b_j)$. We obtain:

$$\theta_i = \max\left(\eta v_i, \max_{j \neq i} b_j\right) \leq \max\left(\eta v_i, \max_{j} b_j\right) = p_{\mathsf{aw}}(\boldsymbol{b}) < \gamma v_i.$$

This leads to the desired result as:

$$\begin{aligned}\mathbb{E}\left[g_i\left(B'_i, \boldsymbol{b}_{-i}\right)\right] &= \int_{\theta_i}^{\gamma v_i} (v_i - \sigma p_i(z, \boldsymbol{b}_{-i})) f(z) dz = \int_{\theta_i}^{\gamma v_i} (v_i - \sigma z) f(z) dz \\ &= \int_{\theta_i}^{\gamma v_i} \mu dz = \mu \gamma v_i - \mu \theta_i \geq \mu \gamma v_i - \mu p_{\mathsf{aw}}(\boldsymbol{b}).\end{aligned}$$

Note that the first equality holds because the sensitivity of i is σ. □

6.3.2 Extension Theorem

6

In this subsection we present our Extension Theorem that uses our smoothness definition for a single auction and specific agent type, in order to derive bounds on the POA of simultaneous auctions with feasible reserve prices and fractionally subadditive valuation functions.

We first introduce some additional notation. Given fractionally subadditive valuation functions $\boldsymbol{v} = (v_i)_{i \in N} \in \mathcal{V}_{\text{XOS}}$ and an optimal liquid welfare allocation $\boldsymbol{x}^* = \boldsymbol{x}^*(\boldsymbol{v})$. As each valuation function v_i satisfies the XOS property as defined in Definition 6.2.1, we can define $(v^*_{ij})_{j \in M}$ as the *additive representatives* of $i \in N$ with respect to their optimal allocation $\boldsymbol{x}^*_i$.[2] For each $i \in N$, the following two properties follow directly from Definition 6.2.1:

XOS1: $v_i(\boldsymbol{x}^*_i) = \sum_{j \in M} v^*_{ij} x^*_{ij}$.

XOS2: for any allocation $x_i \subseteq M$, it holds that $v_i(\boldsymbol{x}_i) \geq \sum_{j \in M} v^*_{ij} x_{ij}$.

To see the latter, recall that the additive representatives $(v_{ij})_{j \in M}$ of i with respect to $\boldsymbol{x}_i$ are chosen as maximizers from the class $\mathcal{L}_i \ni (v^*_{ij})_{j \in M}$, and thus $v_i(\boldsymbol{x}_i) =$

[2]Note that for additive valuations these representatives simply coincide with the input valuations.

$\sum_{j \in M} v_{ij} x_{ij} \geq \sum_{j \in M} v^*_{ij} x_{ij}$. We will use $\boldsymbol{v}^* = (v^*_i)_{i \in N}$ to refer to the additive valuation functions obtained from the representatives, i.e., we define $v^*_i(\boldsymbol{x}_i) = \sum_{j \in M} v^*_{ij} x_{ij}$ for any allocation $\boldsymbol{x}_i$. Given these representatives, we can basically regard each valuation function v_i as being additive in the proof of the Extension Theorem. To this end, when we refer to an auction $\mathcal{M}_j$ as (λ_t, μ_t)-smooth for each type $t \in T$, this is with respect to the valuations $\boldsymbol{v}^*$ in the remainder of this section.

Before stating our Extension Theorem, we introduce the notion of *calibration vectors*, which will be crucial in the proof.

Definition 6.3.6: (Calibration Vectors). Let T be a set of different agent types with sensitivities $\boldsymbol{\sigma} = (\sigma_t)_{t \in T} \in [0,1]^{|T|}$ and let $\boldsymbol{\mu} = (\mu_t)_{t \in T}$ be such that $\mu_t > 0$ for each $t \in T$. We define the set of *calibration vectors* $\mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma})$ as follows:

$$\mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma}) = \left\{ \boldsymbol{\delta} \in (0,1]^{|T|} \;\middle|\; \max_{t \in T}(\delta_t \mu_t) + \max_{t \in T}(\delta_t(1 - \sigma_t)) \leq 1 \right\}. \qquad (6.11)$$

We later show in Lemma 6.4.1 in Section 6.4 below, that $\mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma})$ is always non-empty. We can now state the main result of this section:

Theorem 6.3.7: (Extension Theorem). *Consider the class of simultaneous first-price auctions with fractionally subadditive valuations $\boldsymbol{v} \in \mathcal{V}_{\text{XOS}}$ and feasible reserve prices. Let T be the set of different agent types with sensitivities $\boldsymbol{\sigma}$. Assume that, for each type $t \in T$, each auction $\mathcal{M}_j$ of $\mathcal{M}$ is (λ_t, μ_t)-smooth. Then, the price of anarchy of well-supported coarse correlated equilibria is:*

$$\textit{CCE-POA}(\mathcal{V}_{\text{XOS}}) \leq \left(\max_{\boldsymbol{\delta} \in \mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma})} \min_{t \in T} \delta_t \lambda_t \right)^{-1}.$$

Note that the reserve price r_j of an auction j is assumed to be feasible with respect to the additive representatives, i.e., $r_j = \eta_j v^*_{\mathsf{rw}(j)j}$ for some $\eta_j \in [0, 1)$. 6

The following corollary and the subsequent lemma will be useful in the proof of the Extension Theorem.

Corollary 6.3.8. *Consider a random bid profile $\boldsymbol{B}'_i$ of agent i that is ROI-restricted. Let $\boldsymbol{B}_{-i}$ be an arbitrary probability distribution over D_{-i} of the other agents. Then, $(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i$.*

Proof: We have:

$$\begin{aligned}
\mathbb{E}\left[p_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})\right] &= \int_{D_{-i}} f_{\boldsymbol{B}_{-i}}(\boldsymbol{b}_{-i}) \cdot \mathbb{E}[p_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i})] d\boldsymbol{b}_{-i} \\
&\leq \int_{D_{-i}} f_{\boldsymbol{B}_{-i}}(\boldsymbol{b}_{-i}) \cdot \mathbb{E}[v_i(x_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i}))] d\boldsymbol{b}_{-i} = \mathbb{E}[v_i(\boldsymbol{x}_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}))],
\end{aligned}$$

where, the inequality follows because $\boldsymbol{B}'_i$ is ROI-restricted. $\square$

Lemma 6.3.9. *Consider the simultaneous first-price auctions with fractionally subadditive valuations $\boldsymbol{v} \in \mathcal{V}_{\text{XOS}}$ and feasible reserve prices. Let T be the set of different agent types. Assume that, for each type $t \in T$, each auction $\mathcal{M}_j$ of $\mathcal{M}$ is (λ_t, μ_t)-smooth. Let $\boldsymbol{B}$ be a well-supported coarse correlated equilibrium. Then, for each $t \in T$ it holds that:*

$$\sum_{i \in N_t} \mathbb{E}\left[g_i(\boldsymbol{B})\right] \geq \sum_{j \in M : \mathsf{rw}(j) \in N_t} \lambda_t v^*_{\mathsf{rw}(j)j} - \mu_t \mathbb{E}\left[p_{\mathsf{aw}(j)j}(\boldsymbol{B})\right]. \tag{6.12}$$

Proof: Consider some agent $i \in N_t$ of type $t \in T$ and assume that i is the rightful winner of auction j, i.e., $i = \mathsf{rw}(j)$. There exists a ROI-restricted random bid B'_{ij} such that for each bid profile $\boldsymbol{b}_j$ with $x_j(\boldsymbol{b}_j) \neq \boldsymbol{0}$, we have:

$$\mathbb{E}\left[g_i\left(B'_{ij}, (\boldsymbol{b}_j)_{-i}\right)\right] \geq \lambda_t v^*_{\mathsf{rw}(j)j} - \mu_t p_{\mathsf{aw}(j)j}(\boldsymbol{b}_j), \tag{6.13}$$

as $\mathcal{M}_j$ is (λ_t, μ_t)-smooth by assumption. We define a random deviation $\boldsymbol{B}'_i$ of agent i for the global mechanism $\mathcal{M}$ by simply drawing a bid b'_{ij} for each auction $j \in M$ independently according to B'_{ij} if $i = \mathsf{rw}(j)$ is the rightful winner of j, and let $B'_{ij} = b'_{ij} = 0$ deterministically if $i \neq \mathsf{rw}(j)$. For each bid profile $\boldsymbol{b}_{-i}$, we have:

$$\begin{aligned}\mathbb{E}\left[p_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i})\right] &= \mathbb{E}\left[\sum_{j \in M} p_{ij}(B'_{ij}, (\boldsymbol{b}_j)_{-i})\right] \leq \mathbb{E}\left[\sum_{j \in M} v^*_{ij} x_{ij}(B'_{ij}, (\boldsymbol{b}_j)_{-i})\right] \\ &= \mathbb{E}\left[v^*_i(x_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i}))\right] \leq \mathbb{E}[v_i(x_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i}))],\end{aligned}$$

where the inequality holds because B'_{ij} is ROI-restricted for each j (which holds trivially also for all auctions j with $i \neq \mathsf{rw}(j)$). Note that the second equality holds by the definition of v^*_i, and the last inequality holds by the property XOS2. We conclude that $\boldsymbol{B}'_i$ is ROI-restricted for each agent i.

6

Given an arbitrary bid profile $\boldsymbol{b}$ with $\boldsymbol{x}(\boldsymbol{b}) \neq \boldsymbol{0}$, from the above we conclude that:

$$\begin{aligned}\mathbb{E}[g_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i})] &= \sum_{j \in M} \mathbb{E}\left[g_i\left(B'_{ij}, (\boldsymbol{b}_j)_{-i}\right)\right] \\ &\geq \sum_{j \in M : \mathsf{rw}(j) = i} \lambda_t v^*_{\mathsf{rw}(j)j} - \mu_t p_{\mathsf{aw}(j)j}(\boldsymbol{b}).\end{aligned} \tag{6.14}$$

Here, the inequality holds by applying (6.13) for all auctions j in which $i = \mathsf{rw}(j)$, and using that the expected gain of i is non-negative for all auctions j in which $i \neq \mathsf{rw}(j)$.

Let $\boldsymbol{B}$ be a well-supported CCE. As argued above, the random deviation $\boldsymbol{B}'_i$ is ROI-restricted for every $i \in N_t$. By Corollary 6.3.8, we have that $(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}) \in \mathcal{R}_i$, and so the CCE condition (6.3) of agent i with respect to the deviation $\boldsymbol{B}'_i$ can be applied. Using this, we obtain:

$$\begin{aligned}\mathbb{E}\left[g_i(\boldsymbol{B})\right] &\geq \mathbb{E}\left[g_i(\boldsymbol{B}_i', \boldsymbol{B}_{-i})\right] = \mathbb{E}\left[\mathbb{E}_{\boldsymbol{B}_i'}\left[g_i(\boldsymbol{B}_i', \boldsymbol{B}_{-i})\right]\right] \\ &\geq \sum_{j \in M : \mathsf{rw}(j) = i} \lambda_t v^*_{\mathsf{rw}(j)j} - \mu_t \mathbb{E}\left[p_{\mathsf{aw}(j)j}(\boldsymbol{B})\right].\end{aligned}$$

Here, the second inequality follows because (6.14) holds for every $\boldsymbol{b} \in \operatorname{supp}(\boldsymbol{B})$ (note that $\boldsymbol{x}(\boldsymbol{b}) \neq \boldsymbol{0}$ because $\boldsymbol{B}$ is well-supported), and we can thus take the expectation with respect to $\boldsymbol{B}$ on both sides. Finally, summing the above inequality over all agents $i \in N_t$ proves the claim. □

We can now complete the proof of the Extension Theorem.

Proof of Theorem 6.3.7: Let $\boldsymbol{B}$ be a well-supported CCE. Let $\boldsymbol{\delta} \in \mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma}) \subseteq (0,1]^{|T|}$ be an arbitrary calibration vector. Note that such a vector is guaranteed to exist by Lemma 6.4.1. We have:

$$\begin{aligned}\mathbb{E}[LW(\boldsymbol{B})] &= \sum_{t \in T} \sum_{i \in N_t} \delta_t \mathbb{E}\left[v_i(x_i(\boldsymbol{B}))\right] + (1 - \delta_t) \mathbb{E}\left[v_i(x_i(\boldsymbol{B}))\right] \\ &\geq \sum_{t \in T} \sum_{i \in N_t} \delta_t \mathbb{E}\left[v_i(x_i(\boldsymbol{B}))\right] + (1 - \delta_t) \mathbb{E}\left[p_i(\boldsymbol{B})\right] \\ &= \sum_{t \in T} \sum_{i \in N_t} \delta_t \mathbb{E}\left[g_i(\boldsymbol{B})\right] + (1 - \delta_t + \delta_t \sigma_t) \mathbb{E}\left[p_i(\boldsymbol{B})\right] \\ &\geq \sum_{t \in T} \sum_{i \in N_t} \delta_t \mathbb{E}\left[g_i(\boldsymbol{B})\right] + \left(1 - \max_{t \in T}(\delta_t(1 - \sigma_t))\right) \sum_{t \in T} \sum_{i \in N_t} \mathbb{E}\left[p_i(\boldsymbol{B})\right] \\ &= \sum_{t \in T} \sum_{i \in N_t} \delta_t \mathbb{E}\left[g_i(\boldsymbol{B})\right] + \left(1 - \max_{t \in T}(\delta_t(1 - \sigma_t))\right) \mathbb{E}\left[\sum_{i \in N} p_i(\boldsymbol{B})\right].\end{aligned}$$

Here, the first inequality is obtained by using the ROI constraint (6.2) for each agent $i \in N$. The last equality holds because $\{N_t\}_{t \in T}$ is a partition of N.

We continue with lower bounding the first part of the above inequality. Using Lemma 6.3.9, we obtain that:

$$\begin{aligned}\sum_{t \in T} \sum_{i \in N_t} \delta_t \mathbb{E}\left[g_i(\boldsymbol{B})\right] &\geq \sum_{t \in T} \sum_{\substack{j \in M: \\ \mathsf{rw}(j) \in N_t}} \delta_t \lambda_t v^*_{\mathsf{rw}(j)j} - \delta_t \mu_t \mathbb{E}\left[p_{\mathsf{aw}(j)j}(\boldsymbol{B})\right] \\ &\geq \sum_{t \in T} \sum_{\substack{j \in M: \\ \mathsf{rw}(j) \in N_t}} \min_{t \in T}(\delta_t \lambda_t) v^*_{\mathsf{rw}(j)j} - \max_{t \in T}(\delta_t \mu_t) \mathbb{E}\left[p_{\mathsf{aw}(j)j}(\boldsymbol{B})\right] \\ &= \min_{t \in T}(\delta_t \lambda_t) LW(\boldsymbol{x}^*) - \max_{t \in T}(\delta_t \mu_t) \mathbb{E}\left[\sum_{i \in N} p_i(\boldsymbol{B})\right]. \qquad (6.15)\end{aligned}$$

6

were the last equality follows from property XOS1. Using (6.15), we finally obtain that:

$$\begin{aligned}\mathbb{E}[LW(\boldsymbol{B})] &\geq \min_{t\in T}(\delta_t\lambda_t)OPT(\boldsymbol{v}) \\ &\quad + \left(1 - \max_{t\in T}(\delta_t\mu_t) - \max_{t\in T}(\delta_t(1-\sigma_t))\right)\mathbb{E}\left[\sum_{i\in N} p_i(\boldsymbol{B})\right] \\ &\geq \min_{t\in T}(\delta_t\lambda_t)OPT(\boldsymbol{v}),\end{aligned}$$

where the second inequality holds because $\boldsymbol{\delta} \in \mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma})$ (see Definition 6.3.6). By rearranging terms and selecting a calibration vector $\boldsymbol{\delta} \in \mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma})$ that maximizes $\min_{t\in T}(\delta_t\lambda_t)$, the proof follows. □

Note that in the proof above, the whole purpose of our calibration vector was to eventually lower bound the total payments in the final expression by 0. This also explains the specific definition of the feasibility constraint of $\mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma})$ in (6.11).

6.4 A POA-Revealing Mathematical Program

We present our approach to determine the price of anarchy through the Extension Theorem. A crucial building block for deriving our POA-revealing mathematical program is the characterization of optimal calibration vectors given below.

Lemma 6.4.1. *Let T be a set of different agent types with sensitivities $\boldsymbol{\sigma} = (\sigma_t)_{t\in T} \in [0,1]^{|T|}$. Let $\boldsymbol{\mu} = (\mu_t)_{t\in T}$ and $\boldsymbol{\lambda} = (\lambda_t)_{t\in T}$ be such that $\mu_t > 0$ and $\lambda_t > 0$ for each $t \in T$. Then, $\max_{\boldsymbol{\delta}\in\mathcal{C}(\boldsymbol{\mu},\boldsymbol{\sigma})} \min_{t\in T} \lambda_t\delta_t = O$, where:*

$$O = \min\left\{\min_{t\in T}\lambda_t, \left(\max_{t\in T}\left(\frac{\mu_t}{\lambda_t}\right) + \max_{t\in T}\left(\frac{1-\sigma_t}{\lambda_t}\right)\right)^{-1}\right\}. \tag{6.16}$$

Proof: Define $\boldsymbol{\delta}'$ such that $\delta'_t = O/\lambda_t$ for each $t \in T$. We first show that $\boldsymbol{\delta}' \in \mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma})$. First, note that $\boldsymbol{\delta}' \in (0,1]^{|T|}$, since for each $t \in T$ it holds that:

$$\delta'_t = \frac{O}{\lambda_t} \leq \frac{\min_{t\in T}\lambda_t}{\lambda_t} \leq 1,$$

and $\delta'_t > 0$ as $O > 0$. Furthermore, $\boldsymbol{\delta}'$ satisfies:

$$\begin{aligned}\max_{t\in T}(\delta'_t\mu_t) + \max_{t\in T}(\delta'_t(1-\sigma_t)) &= \max_{t\in T}\left(\frac{O}{\lambda_t}\mu_t\right) + \max_{t\in T}\left(\frac{O}{\lambda_t}(1-\sigma_t)\right) \\ &= O\cdot\left(\max_{t\in T}\left(\frac{\mu_t}{\lambda_t}\right) + \max_{t\in T}\left(\frac{1-\sigma_t}{\lambda_t}\right)\right) \leq 1,\end{aligned}$$

where the last inequality follows from (6.16). Hence, $\boldsymbol{\delta}' \in \mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma})$ and therefore $\min_{t \in T} \lambda_t \delta'_t = O$.

To complete the proof, we need to show that $\max_{\boldsymbol{\delta} \in \mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma})} \min_{t \in T} \lambda_t \delta_t = O$. Towards a contradiction, assume that there exists a $\bar{\boldsymbol{\delta}} \in \mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma})$ with:

$$\min_{t \in T} \lambda_t \bar{\delta}_t > O. \tag{6.17}$$

We distinguish two cases for the value of O.
<u>Case 1:</u> $O = \min_{t \in T} \lambda_t$. In this case we get the following contradiction:

$$\min_{t \in T} \lambda_t = O < \min_{t \in T} \lambda_t \bar{\delta}_t \leq \min_{t \in T} \lambda_t,$$

where the equality and first inequality hold by assumption. The second inequality follows as $\bar{\boldsymbol{\delta}} \in \mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma})$, so $\bar{\delta}_t \leq 1$ for all $t \in T$.
<u>Case 2:</u> $O = \left(\max_{t \in T} \left(\frac{\mu_t}{\lambda_t}\right) + \max_{t \in T} \left(\frac{1-\sigma_t}{\lambda_t}\right)\right)^{-1}$. Let $\hat{t} = \arg\max_{t \in T} \mu_t / \lambda_t$ and $\tilde{t} = \arg\max_{t \in T} (1 - \sigma_t)/\lambda_t$. In this case we get the following contradiction:

$$\begin{aligned} 1 < \left(\frac{\mu_{\hat{t}}}{\lambda_{\hat{t}}} + \frac{1-\sigma_{\tilde{t}}}{\lambda_{\tilde{t}}}\right) \cdot \min_{t \in T} \lambda_t \bar{\delta}_t &= \frac{\mu_{\hat{t}}}{\lambda_{\hat{t}}} \cdot \min_{t \in T} \lambda_t \bar{\delta}_t + \frac{1-\sigma_{\tilde{t}}}{\lambda_{\tilde{t}}} \cdot \min_{t \in T} \lambda_t \bar{\delta}_t \\ &\leq \mu_{\hat{t}} \bar{\delta}_{\hat{t}} + (1 - \sigma_{\tilde{t}})\, \bar{\delta}_{\tilde{t}} \\ &\leq \max_{t \in T} \left(\mu_t \bar{\delta}_t\right) + \max_{t \in T} \left((1 - \sigma_t)\, \bar{\delta}_t\right) \\ &\leq 1. \end{aligned}$$

Here, the first inequality holds by assumption and the last inequality holds since $\bar{\boldsymbol{\delta}} \in \mathcal{C}(\boldsymbol{\mu}, \boldsymbol{\sigma})$. This concludes the proof. □

We use the above lemma together with our smoothness lemmas to derive a mathematical program that captures a bound on the POA. We refer to this program as the *POA-revealing mathematical program (POA-RMP)*. Let T be the set of different agent types with sensitivities $\boldsymbol{\sigma} \in [0, 1]^{|T|}$. Lemmas 6.3.3 and 6.3.4 then imply the following smoothness parameters, depending on the respective agent type $t \in T$ and assuming that $\eta = \min_{j \in M} \eta_j \in [0, 1)$. 6

$$\lambda_t = \mu_t \left(1 - \tfrac{1-\eta}{e^{1/\mu_t}}\right) \qquad \mu_t > 0 \qquad \forall t \in T\colon\ \sigma_t = 1 \tag{6.18}$$

$$\lambda_t = \tfrac{\mu_t}{\sigma_t} \left(1 - \tfrac{1-\sigma_t \eta}{e^{\sigma_t/\mu_t}}\right) \qquad \mu_t \geq \sigma_t \left(\ln\left(\tfrac{1-\sigma_t \eta}{1-\sigma_t}\right)\right)^{-1} \qquad \forall t \in T\colon\ \sigma_t \in (0, 1) \tag{6.19}$$

$$\lambda_t = \mu_t \qquad \mu_t \in \left(0, \tfrac{1}{1-\eta}\right] \qquad \forall t \in T\colon\ \sigma_t = 0 \tag{6.20}$$

To obtain a bound on the POA, we determine a vector $\boldsymbol{\mu} = (\mu_t)_{t \in T}$ that maximizes the expression in (6.16) subject to the constraints above. Thus, our POA-RMP is:

$$\begin{aligned} \text{POA-RMP} \;=\; \max\ \min \left\{ \min_{t \in T} \lambda_t, \left(\max_{t \in T} \left(\frac{\mu_t}{\lambda_t}\right) + \max_{t \in T} \left(\frac{1-\sigma_t}{\lambda_t}\right)\right)^{-1} \right\} \\ \text{subject to (6.18)–(6.20).} \end{aligned}$$

6.4.1 Solving POA-RMP by Partitioning Agent Types

We are left to determine a lower bound on the optimal value of POA-RMP. The main technical challenge is to find a solution analytically. Below, we provide a policy to define a solution $\boldsymbol{\mu}$ of POA-RMP. We first focus on the case without reserve prices, i.e., $\eta = 0$ in (6.18)–(6.20).

The main idea is to partition the agent types T of the rightful winners into two classes H_ω (high) and L_ω (low), where ω is a separation parameter. We then define μ_t depending on the class t belongs to. Intuitively, H_ω contains agent types that are structurally close to utility maximizers, while L_ω contains agent types that are structurally close to value maximizers.

Definition 6.4.2. Given $\omega \in (0,1)$, define $H_\omega = \{t \in T \mid \sigma_t \geq \omega\}$ and $L_\omega = \{t \in T \mid \sigma_t < \omega\}$. Define $\boldsymbol{\mu}^*(\omega) \in \mathbb{R}^{|T|}_{>0}$ such that

$$\mu_t^*(\omega) = \begin{cases} \frac{\sigma_t}{-\ln(1-\omega)}, & \text{if } t \in H_\omega, \\ \frac{\sigma_t}{-\ln(1-\sigma_t)}, & \text{if } t \in L_\omega \text{ and } \sigma_t > 0, \\ 1, & \text{if } t \in L_\omega \text{ and } \sigma_t = 0. \end{cases} \tag{6.21}$$

The following corollary is easy to verify.

Corollary 6.4.3. $\boldsymbol{\mu}^*(\omega)$ *is a feasible solution of POA-RMP.*

Proof: Clearly, (6.18) holds for all utility maximizing types $t \in H_\omega$ with $\sigma_t = 1$. Also, (6.19) is satisfied for each $t \in H_\omega$ with $\sigma_t < 1$ because $\sigma_t \geq \omega$ (as $t \in H_\omega$) and because the function $f(z) = -\ln(1-z)$ is non-negative and non-decreasing on $(0,1)$. Further, (6.19) is satisfied with equality for all types $t \in L_\omega$ with $\sigma_t > 0$. Finally, (6.20) holds for all value maximizing types $t \in L_\omega$ with $\sigma_t = 0$. □

6

We introduce some more notation. For a given set T of agent types, we define $\sigma_{\max} = \max_{t\in T} \sigma_t$ and $\sigma_{\min} = \min_{t\in T} \sigma_t$. If $\{t \in T \mid \sigma_t > 0\} \neq \emptyset$, we define $\sigma_{\min_+} = \min_{t\in T:\sigma_t>0} \sigma_t$. We also define $\lambda_t^*(\omega) = \lambda_t$ with λ_t as defined in (6.18)–(6.20) with $\mu_t = \mu_t^*(\omega)$ as in Definition 6.4.2. For these definitions of $\lambda_t^*(\omega)$ and $\mu_t^*(\omega)$, we state several properties which will be useful in the proofs below. The proof of the properties can be found in Appendix A.2.1

Lemma 6.4.4. *The following properties hold for every set of agent types T with $\sigma_{\max} > 0$ and every $\omega \in (0, \sigma_{\max}] \cap (0,1)$.*

(i) $\lambda_t^*(\omega) = \begin{cases} \frac{\omega}{-\ln(1-\omega)}, & \text{if } t \in H_\omega, \\ \frac{\sigma_t}{-\ln(1-\sigma_t)}, & \text{if } t \in L_\omega \text{ and } \sigma_t > 0, \\ 1, & \text{if } t \in L_\omega \text{ and } \sigma_t = 0. \end{cases}$

(ii) $\min_{t\in T:\sigma_t>0} \lambda_t^*(\omega) = \frac{\omega}{-\ln(1-\omega)}$.

(iii) $\max_{t\in T:\sigma_t>0} \frac{1-\sigma_t}{\lambda_t^*(\omega)} = \begin{cases} \frac{-\ln(1-\omega)(1-\sigma_{\min_+})}{\omega}, & \textit{if } \sigma_{\min_+} \geq \omega, \\ \frac{-\ln(1-\sigma_{\min_+})(1-\sigma_{\min_+})}{\sigma_{\min_+}}, & \textit{if } \sigma_{\min_+} < \omega. \end{cases}$

(iv) $\max_{t\in T:\sigma_t>0} \frac{\mu_t^*(\omega)}{\lambda_t^*(\omega)} = \frac{\sigma_{\max}}{\omega}$.

The following lemma will be useful when proving the bounds on the POA of Theorem 6.4.6 below. The proof is given in Appendix A.2.2.

Lemma 6.4.5. *Let T be a set of different agent types with sensitivities $\boldsymbol{\sigma} \in [0,1]^{|T|}$. If $\sigma_{\max} > 0$, then for every $\omega \in (0, \sigma_{\max}] \cap (0,1)$ the optimal value of POA-RMP is at least:*

$$\min\left\{\frac{\omega}{-\ln(1-\omega)}, \frac{\omega}{\omega+\sigma_{\max}}\right\}.$$

With the help of Lemma 6.4.5, it is not hard to derive the following bounds on the POA. Recall that $W_0(\cdot)$ is the principal branch of the Lambert function.

Theorem 6.4.6. *Consider the class of simultaneous first-price auctions with fractionally subadditive valuations $\boldsymbol{v} \in \mathcal{V}_{\text{XOS}}$. Then:*

$$\textit{CCE-POA}(\mathcal{V}_{\text{XOS}}) \leq \begin{cases} 1 + \frac{\sigma_{\max}}{1+W_0(-e^{-\sigma_{\max}-1})}, & \textit{if } \sigma_{\max} > 1 + \frac{W_0(-2e^{-2})}{2}, \\ 2, & \textit{otherwise.} \end{cases}$$

Theorem 6.4.6 reveals an intriguing threshold phenomenon: the POA stays at 2 for all $\sigma_{\max} < 0.79$, and increases from 2 to 2.18 as $\sigma_{\max}$ increases to 1. The following two mathematical claims will also be useful in the proof of Theorem 6.4.6. The proofs of the claims can be found in Appendix A.2.3.

6

Claim 6.4.7. Consider the function $f(z) = 1 + W_0\left(-e^{-z-1}\right)$. For every $z \in \left(1 + {W_0(-2e^{-2})}/{2}, 1\right]$, it holds that $f(z) < z$.

Claim 6.4.8. Consider the function $f(z) = 1 + W_0\left(-e^{-z-1}\right)$. For every $z \in (1 + {W_0(-2e^{-2})}/{2}, 1]$, it holds that $f(z) + z = -\ln(1 - f(z))$.

We can now prove Theorem 6.4.6.

Proof of Theorem 6.4.6: If $\sigma_{\max} = 0$, then all agents are value maximizers. Assume that the value maximizers are of type t. Then, for every $\omega \in (0,1)$, using (6.20) and (6.21) we get that the value of the objective of POA-RMP is at least:

$$\min\left\{\lambda_t, \left(\frac{\mu_t}{\lambda_t} + \frac{1}{\lambda_t}\right)^{-1}\right\} = \min\left\{1, (1+1)^{-1}\right\} = \frac{1}{2}.$$

Now assume that $\sigma_{\max} > 0$. We distinguish the following two cases.

Case 1: $\sigma_{\max} \leq 1 + {}^{W_0(-2e^{-2})}/_2$. Let $\omega = \sigma_{\max}$. Then $\omega \in (0, \sigma_{\max}] \cap (0,1)$ as $\sigma_{\max} < 1$, and we can invoke Lemma 6.4.5 with $\omega = \sigma_{\max}$. Hence, we obtain that the value of the objective function of POA-RMP is at least:

$$\min\left\{\frac{\omega}{-\ln(1-\omega)}, \frac{\omega}{\omega+\sigma_{\max}}\right\} = \min\left\{\frac{\sigma_{\max}}{-\ln(1-\sigma_{\max})}, \frac{1}{2}\right\} = \frac{1}{2}. \tag{6.22}$$

Here, the first and second equality follow by choice of $\omega = \sigma_{\max}$ as $\frac{z}{-\ln(1-z)} \geq \frac{1}{2}$ for all $z \leq 1 + {}^{W_0(-2e^{-2})}/_2$.

Case 2: $\sigma_{\max} > 1 + {}^{W_0(-2e^{-2})}/_2$. Let $\omega = 1 + W_0\left(-e^{-\sigma_{\max}-1}\right)$. Using Claim 6.4.7, we get that $\omega = f(\sigma_{\max}) < \sigma_{\max}$ whenever $\sigma_{\max} \in (1 + {}^{W_0(-2e^{-2})}/_2, 1]$. Therefore, $\omega \in (0, \sigma_{\max}] \cap (0,1)$ as $\sigma_{\max} \leq 1$ and we can invoke Lemma 6.4.5 with $\omega = 1 + W_0\left(-e^{-\sigma_{\max}-1}\right)$. Hence, we obtain that the value of the objective function of POA-RMP is at least:

$$\begin{aligned}\min\left\{\frac{\omega}{-\ln(1-\omega)}, \frac{\omega}{\omega+\sigma_{\max}}\right\} &= \frac{\omega}{\omega+\sigma_{\max}} \\ &= \frac{1+W_0\left(-e^{-\sigma_{\max}-1}\right)}{1+W_0\left(-e^{-\sigma_{\max}-1}\right)+\sigma_{\max}}.\end{aligned} \tag{6.23}$$

Here, the first equality holds by Claim 6.4.8 and the second equality follows by the definition of ω.

The proof follows by applying Theorem 6.3.7 and Lemma 6.4.1, together with the lower bounds on POA-RMP from (6.22) or (6.23), depending on the above case distinction. □

6

The bounds on the POA in Theorem 6.4.6 depend on the maximum sensitivity of the agents considered ($\sigma_{\max}$), whereas the minimum sensitivity can be arbitrary. Contrarily, we can consider settings for which we derive bounds on the POA depending on the minimum sensitivity of the agents considered ($\sigma_{\min}$), whereas the maximum sensitivity can be arbitrary. In this case, we derive bounds on the POA for $\sigma_{\min} \geq 0.741$ in Theorem 6.4.9 below. Note that for $\sigma_{\min} = 1$, we retrieve the bound of $e/(e-1)$.

Theorem 6.4.9. *Let β be the solution to the equation $\beta = 1 - e^{-\frac{1}{\beta}}$, i.e., $\beta \approx 0.741$. Consider the class of simultaneous first-price auctions with fractionally subadditive valuations $\boldsymbol{v} \in \mathcal{V}_{\text{XOS}}$. If $\sigma_{\min} \geq \beta$, then:*

$$\textit{CCE-POA}\left(\mathcal{V}_{\text{XOS}}\right) \leq \frac{e^{\frac{1}{\sigma_{\min}}}}{\sigma_{\min}\left(e^{\frac{1}{\sigma_{\min}}}-1\right)} \in \left[\frac{e}{e-1}, 1.83\right].$$

Proof: Set $\omega = 1 - e^{-1/\sigma_{\min}}$. Note that $\omega \leq \sigma_{\min}$. Namely, consider $h(z) = z - 1 + e^{-\frac{1}{z}}$. For all $z \in [\beta, 1]$, it holds that $h(z) \geq 0$. To see this, note that

$h(\beta) = \beta - 1 + e^{-\frac{1}{\beta}} = 0$ by definition of β. Also, the function h is non-decreasing for all $z \geq \beta$, as $h'(z) = 1 + \frac{e^{-\frac{1}{z}}}{z^2} > 0$ in this case. Therefore, $h(\sigma_{\min}) = \sigma_{\min} - \omega \geq 0$.

For $\boldsymbol{\mu} := \boldsymbol{\mu}^*(\omega)$ and $\boldsymbol{\lambda} := \boldsymbol{\lambda}^*(\omega)$ as defined in (6.21) and *(i)* of Lemma 6.4.4, respectively, it holds that:

$$\begin{aligned}
\max_{t \in T} \frac{\mu_t}{\lambda_t} + \max_{t \in T} \frac{1-\sigma_t}{\lambda_t} &= \max_{\substack{t \in T: \\ \sigma_t > 0}} \frac{\mu_t}{\lambda_t} + \max_{\substack{t \in T: \\ \sigma_t > 0}} \frac{1-\sigma_t}{\lambda_t} = \frac{\sigma_{\max} - \ln(1-\omega)(1-\sigma_{\min})}{\omega} \\
&\leq \frac{1 - \ln(1-\omega)(1-\sigma_{\min})}{\omega} = \frac{-\ln(1-\omega)}{\omega} \\
&+ \frac{1 + \ln(1-\omega)\sigma_{\min}}{\omega} = \frac{-\ln(1-\omega)}{\omega}.
\end{aligned} \tag{6.24}$$

Here, the first equality holds as there are no value maximizers by assumption. The second equality follows from properties *(iii)* and *(iv)* of Lemma 6.4.4, as $\omega \leq \sigma_{\min}$. Finally, the last equality follows as $1 + \ln(1-\omega)\sigma_{\min} = 0$ by definition of ω. We conclude that:

$$\begin{aligned}
\text{POA-RMP} &\geq \min\left\{ \min_{t \in T} \lambda_t, \left(\max_{t \in T} \frac{\mu_t}{\lambda_t} + \max_{t \in T} \frac{1-\sigma_t}{\lambda_t} \right)^{-1} \right\} \\
&= \min\left\{ \frac{\omega}{-\ln(1-\omega)}, \left(\max_{t \in T} \frac{\mu_t}{\lambda_t} + \max_{t \in T} \frac{1-\sigma_t}{\lambda_t} \right)^{-1} \right\} \\
&\geq \frac{\omega}{-\ln(1-\omega)} = \left(1 - e^{-\frac{1}{\sigma_{\min}}}\right) \sigma_{\min},
\end{aligned}$$

where the first equality follows from property *(ii)* of Lemma 6.4.4 and the second inequality follows from (6.24). The proof follows by applying Theorem 6.3.7 and Lemma 6.4.1. □

6

6.4.2 Improving Inefficiency via Reserve Prices

In this section, we show that our smoothness framework allows us to capture how inefficiency improves in simultaneous first-price auctions with feasible reserve prices. Similarly to the existing literature (see, e.g., [Balseiro et al., 2021a]), we present the POA upper bounds as functions of the parameter η. Recall that η measures the minimum relative gap between the reserve price and the value of the rightful winner of an auction. Naturally, our results show that the bound on the POA improves as η increases. We only focus on additive valuations in this section, as all correlated equilibria are well-supported in this case, which we later proof in Section 6.5. Additionally, for reasons of analytical tractability, we focus on two specific compositions of agent types. First, we consider the mixed agent model. Recall that in this case agents are either value maximizers, i.e., agents of type t with $\sigma_t = 0$, or utility maximizers, i.e., agents of type t with $\sigma_t = 1$. This generalizes a result of Deng et al. [2024] (Theorem 3) to well-supported coarse correlated equilibria. In the second setting, we consider the scenario of homogeneous agents, i.e., all agents share one common type t.

As there are only two agent types in the mixed agent model, we naturally define μ_t for each agent type.

Definition 6.4.10. Let $\zeta = 2-\eta+W_0\left(-\frac{(1-\eta)^2}{e^{2-\eta}}\right)$. Let T be a set of different agent types containing value and utility maximizers, i.e., $|T| = 2$. Define $\bar{\boldsymbol{\mu}}(\eta) \in \mathbb{R}_{>0}^{|T|}$ such that:

$$\bar{\mu}_t(\eta) = \begin{cases} \frac{1}{1-\eta}, & \text{if } \sigma_t = 0, \\ \frac{1}{\zeta}, & \text{if } \sigma_t = 1. \end{cases} \tag{6.25}$$

The following fact will be helpful when showing that $\bar{\boldsymbol{\mu}}(\eta)$ is a feasible solution.

Fact 6.4.11. Let $h(z) = 2 - z + W_0\left(-\frac{(1-z)^2}{e^{2-z}}\right)$. Then, $h(z) > 0$ for all $z \in [0, 1)$.

Note that Fact 6.4.11 holds as $-\frac{(1-z)^2}{e^{2-z}}$ is non-decreasing for $z \in [0, 1)$, $W_0(\cdot)$ is non-decreasing and $W_0(-e^{-2}) \geq -0.16$. The following corollary stating that $\bar{\boldsymbol{\mu}}(\eta)$ is a feasible solution is now easy to verify.

Corollary 6.4.12. *$\bar{\boldsymbol{\mu}}(\eta)$ is a feasible solution of POA-RMP.*

Namely, for agents of type t with $\sigma_t = 0$, the constraint in (6.20) is satisfied by construction. For agents of type t with $\sigma_t = 1$, the constraint in (6.18) holds as $\zeta > 0$ for all $\eta \in [0, 1)$ by Fact 6.4.11. We also define $\bar{\lambda}_t(\eta) = \lambda_t$ with λ_t as defined in (6.18) and (6.20) with $\mu_t = \bar{\mu}_t(\eta)$. The following lemma will be useful when proving our bound on the POA for the mixed agent model.

Lemma 6.4.13. *Let $\zeta = 2 - \eta + W_0\left(-\frac{(1-\eta)^2}{e^{2-\eta}}\right)$. Let T be a set of different agent types containing value and utility maximizers, i.e., $|T| = 2$. Then, the following properties hold.*

(i) $\min_{t \in T} \bar{\lambda}_t(\eta) \geq \frac{1-(1-\eta)e^{-\zeta}}{\zeta}$.

(ii) $\max_{t \in T} \frac{\bar{\mu}_t(\eta)}{\bar{\lambda}_t(\eta)} \leq \frac{1}{1-(1-\eta)e^{-\zeta}}$.

(iii) $\max_{t \in T} \frac{1-\sigma_t}{\bar{\lambda}_t(\eta)} \leq 1 - \eta$.

(iv) $\zeta = 2 - \eta - (1-\eta)^2 e^{-\zeta}$.

6

The proof of Lemma 6.4.13 can be found in Appendix A.2.4. We can now prove Theorem 6.4.14.

Theorem 6.4.14. *Consider the class of simultaneous first-price auctions with feasible reserve prices with additive valuations $\boldsymbol{v} \in \mathcal{V}_{\text{ADD}}$ in the mixed agent model, i.e., $|T| = 2$ and for all $t \in T$ it holds that $\sigma_t \in \{0, 1\}$. Let:*

$$Q(\eta) = (1-\eta) \cdot \frac{2 - \eta + W_0\left(-(1-\eta)^2 e^{\eta-2}\right)}{1 - \eta + W_0\left(-(1-\eta)^2 e^{\eta-2}\right)}.$$

Then, for well-supported coarse correlated equilibria, CCE-POA$(\mathcal{V}_{\text{ADD}}) \leq Q(\eta)$. Furthermore, $Q(\eta)$ is non-increasing in $[0, 1)$ with $\lim_{z\to 1} Q(z) = 1$.

Proof of Theorem 6.4.14: Let $\zeta = 2 - \eta + W_0\left(-\frac{(1-\eta)^2}{e^{2-\eta}}\right)$ and define $\bar{\boldsymbol{\mu}} = \bar{\boldsymbol{\mu}}(\eta)$ as in (6.25) and $\bar{\boldsymbol{\lambda}} = \bar{\boldsymbol{\lambda}}(\eta)$. Note that $\bar{\boldsymbol{\mu}}$ is a feasible solution by Corollary 6.4.12. We proceed as in Section 6.4.1 and use $\bar{\boldsymbol{\mu}}$ to obtain the following lower bound on the value of the objective function of POA-RMP:

$$\begin{aligned}
&\min\left\{\min_{t\in T} \bar{\lambda}_t, \left(\max_{t\in T} \frac{\bar{\mu}_t}{\bar{\lambda}_t} + \max_{t\in T} \frac{1-\sigma_t}{\bar{\lambda}_t}\right)^{-1}\right\} \\
&\geq \min\left\{\frac{1-(1-\eta)e^{-\zeta}}{\zeta}, \left(\frac{1}{1-(1-\eta)e^{-\zeta}} + 1 - \eta\right)^{-1}\right\} \\
&= \left(1-(1-\eta)e^{-\zeta}\right) \cdot \min\left\{\frac{1}{\zeta}, \frac{1}{2-\eta-(1-\eta)^2 e^{-\zeta}}\right\} \\
&= \frac{1-(1-\eta)e^{-\zeta}}{\zeta}.
\end{aligned} \tag{6.26}$$

Here, the second inequality follows from properties *(i)–(iii)* of Lemma 6.4.13. The final equality holds due to property *(iv)* of the same lemma. We conclude the proof with:

$$\begin{aligned}
\mathit{POA\text{-}CCE}(\mathcal{V}_{\text{ADD}}) &\leq \frac{\zeta}{1-(1-\eta)e^{-\zeta}} = (1-\eta) \cdot \frac{\zeta}{1-\eta-(1-\eta)^2 e^{-\zeta}} \\
&= (1-\eta) \cdot \frac{\zeta}{\zeta - 1} = Q(\eta),
\end{aligned}$$

where the inequality follows from Theorem 6.3.7, Lemma 6.4.1 and (6.26). The second equality holds by property *(iv)* of Lemma 6.4.13.

Finally, we show that $\lim_{z\to 1} Q(z) = 1$. Let $h(z) = -(1-z)^2 e^{z-2}$. Then by using l'Hôpital's rule and Fact 6.2.4, it holds that:

$$\lim_{z\to 1} Q(z) = \lim_{z\to 1} \frac{-2 + z - W_0(h(z)) + (1-z)\left(\frac{(1-z^2)W_0(h(z))}{-(1-z)^2(1+W_0(h(z)))} - 1\right)}{\frac{(1-z^2)W_0(h(z))}{-(1-z)^2(1+W_0(h(z)))} - 1} = \frac{-2+1}{-1},$$

6

as $W_0(h(1)) = 0$. This concludes the proof. □

Finally, we consider the setting of homogeneous agents, i.e., all agents share one common type t.

Theorem 6.4.15. *Consider the class of simultaneous first-price auctions with feasible reserve prices and additive valuations $\boldsymbol{v} \in \mathcal{V}_{\text{ADD}}$ in which all agents have a common type t with $\sigma = \sigma_t$. Let:*

$$P_\sigma(\eta) = \begin{cases} \frac{e}{e-1+\eta}, & \text{if } \sigma = 1, \\ 1 + \frac{1}{\sigma}\left(\ln\left(\frac{1-\sigma\eta}{1-\sigma}\right)(1-\sigma)\right), & \text{if } \sigma \in \left(0, \frac{e-1}{e}\right), \\ 2-\eta, & \text{if } \sigma = 0. \end{cases}$$

Then, for well-supported coarse correlated equilibria, CCE-POA($\mathcal{V}_{\text{ADD}}$) $\leq P_\sigma(\eta)$. Furthermore, P_σ is non increasing in $[0, 1)$ with $P_\sigma(1) = 1$ for all types t.

Proof: As all agents have a common type t with $\sigma = \sigma_t$, lower bounding the objective function of POA-RMP reduces to determining a scalar $\mu = \mu_t$ that satisfies (6.18), (6.19) or (6.20), depending on the type t. Then, for every feasible choice of μ and corresponding $\lambda = \lambda(\mu)$, we get the following lower bound:

$$\min\left\{\lambda, \frac{\lambda}{\mu + 1 - \sigma}\right\}. \tag{6.27}$$

We distinguish three cases.

6

Case 1: $\sigma = 1$. We set $\mu = \sigma$. Note that this choice of μ is feasible as the non-negativity constraint in (6.18) is satisfied. Furthermore, the right hand side of (6.27) is equal to $\lambda = \mu\left(1 - (1-\eta)e^{-\frac{1}{\mu}}\right) = 1 - (1-\eta)e^{-1} = (P_\sigma(\eta))^{-1}$.

Case 2: $\sigma \in (0, \frac{e-1}{e})$. We set $\mu = \sigma\left(\ln\left(\frac{1-\sigma\eta}{1-\sigma}\right)\right)^{-1}$, satisfying the constraint in (6.19) with equality. We have that:

$$\lambda = \frac{\mu}{\sigma}\left(1 - \frac{1-\sigma\eta}{e^{\frac{\sigma}{\mu}}}\right) = \sigma \cdot \ln\left(\frac{1-\sigma\eta}{1-\sigma}\right)^{-1}, \tag{6.28}$$

and since $\eta \in [0, 1)$ and $\sigma < \frac{e-1}{e}$, this leads to:

$$\mu + 1 - \sigma = 1 + \sigma\left(\frac{1}{\ln\left(\frac{1-\sigma\eta}{1-\sigma}\right)} - 1\right) > 1. \tag{6.29}$$

Therefore, using (6.27) we can lower bound the objective of POA-RMP with:

$$\min\left\{\lambda, \frac{\lambda}{\mu+1-\sigma}\right\} = \frac{\lambda}{\mu+1-\sigma} = \frac{\sigma}{\sigma + \ln\left(\frac{1-\sigma\eta}{1-\sigma}\right)(1-\sigma)} = (P_\sigma(\eta))^{-1},$$

where the first and second equality follow by (6.29) and (6.28), respectively.

Case 3: $\sigma = 0$. We set $\mu = (1-\eta)^{-1}$, which clearly satisfies the constraint in (6.20) and prescribes that $\lambda = (1-\eta)^{-1}$. We obtain:

$$\min\left\{\lambda, \frac{\lambda}{\mu+1-\sigma}\right\} = \min\left\{\frac{1}{1-\eta}, \frac{(1-\eta)^{-1}}{1+(1-\eta)^{-1}}\right\} = \min\left\{\frac{1}{1-\eta}, \frac{1}{2-\eta}\right\}$$
$$= \frac{1}{2-\eta} = (P_\sigma(\eta))^{-1}.$$

As the lower bound on the objective value of POA-RMP is greater than or equal to $(P_\sigma(\eta))^{-1}$ for all three cases, it holds by Theorem 6.3.7 and Lemma 6.4.1 that for well-supported coarse-correlated equilibria:

$$\textit{CCE-POA}(\mathcal{V}_{ADD}) \leq P_\sigma(\eta),$$

concluding the proof. □

6.5 Well-Supported Equilibria and Lower Bounds

When deriving bounds on the POA for simultaneous first-price auctions with feasible reserve prices, we crucially only consider equilibria that are well-supported. Recall that an equilibrium $\boldsymbol{B}$ is well-supported if for each auction $j \in M$ the item is always sold, i.e., for each auction $j \in M$ it holds that $\mathbb{P}[x_j(\boldsymbol{B}) \neq \mathbf{0}] = 1$. Even though we consider feasible reserve prices, not all equilibria are necessarily well-supported; this depends on the equilibrium notion and the type of valuation functions at hand. In this section, we first consider additive valuations and show that, up to correlated equilibria, all equilibria are well-supported. We then give three lower bounds on the POA for specific settings with additive valuations. We conclude this section by showing that for submodular valuations, there exist equilibria that are not well-supported. In the homogeneous case of utility maximizers, this leads to an unbounded POA.

6.5.1 Additive Valuations

In the case of additive valuations, it intuitively does not make sense for the rightful winner i of an auction j to bid below the reserve price r_j. As we consider feasible reserve prices, agent i could potentially achieve some positive gain from winning auction j, unless some other agent wins the item at a price of v_{ij}. We use this intuition in the proof below and show that, up to CE, all equilibria are well-supported.

Lemma 6.5.1. *Consider the class of simultaneous first-price auctions with additive valuations $\boldsymbol{v} \in \mathcal{V}_{\text{ADD}}$ and feasible reserve prices. Then, every CE is well-supported.*

Proof: For contradiction, suppose that $\boldsymbol{B}$ is a CE that is not well-supported, i.e., there exists an auction $j \in M$ such that:

$$\mathbb{P}[x_j(\boldsymbol{B}) \neq \boldsymbol{0}] < 1. \tag{6.30}$$

Let agent $i \in N$ be the rightful winner of auction j, i.e., $i = \mathsf{rw}(j)$. We define the deviation of agent i as follows. For all auctions $k \neq j$, agent i follows the recommendation. However, for auction j, agent i follows the recommendation only if the recommendation is to bid above or equal to the reserve price r_j. Otherwise, agent i ignores the recommendation and bids exactly the reserve price r_j for auction j. More formally, for any $\boldsymbol{b}_i$ in the support of $\boldsymbol{B}$, we define:

$$h(\boldsymbol{b}_i) = \boldsymbol{b}'_i = (b_{i1}, \ldots, b_{i(j-1)}, \max\{b_{ij}, r_j\}, b_{i(j+1)}, \ldots, b_{im}). \tag{6.31}$$

Note that by (6.30), there exist recommendation(s) $\boldsymbol{b}_i$ in the support of $\boldsymbol{B}$ that recommend agent i to bid below the reserve price r_j for auction j.

We first show that the deviation as defined in (6.31) satisfies the ROI constraint for agent i. The conditional PDF of $\boldsymbol{B}_{-i}$ given $\boldsymbol{B}_i = \boldsymbol{b}_i$ is denoted as $f_{\boldsymbol{B}_{-i}|\boldsymbol{B}_i}(\boldsymbol{b}_{-i} \mid \boldsymbol{b}_i)$. Note that $\boldsymbol{B} \in \mathcal{R}_i$ as $\boldsymbol{B}$ is a correlated equilibrium, i.e., the ROI constraint is satisfied for agent i:

$$\mathbb{E}\left[\mathbb{E}\left[\sum_{k\in M} p_{ik}(\boldsymbol{b}_i, \boldsymbol{B}_{-i}) \,\middle|\, \boldsymbol{B}_i = \boldsymbol{b}_i\right]\right] \leq \mathbb{E}\left[\mathbb{E}\left[\sum_{k\in M} v_{ik}(x_{ik}(\boldsymbol{b}_i, \boldsymbol{B}_{-i})) \,\middle|\, \boldsymbol{B}_i = \boldsymbol{b}_i\right]\right]. \tag{6.32}$$

Recall that the outer expectation is for $\boldsymbol{B}_i$ and the inner expectation is for $\boldsymbol{B}_{-i}$. The ROI constraint for the deviation of agent i is:

$$\mathbb{E}\left[\mathbb{E}\left[\sum_{k\in M} p_{ik}(h(\boldsymbol{b}_i), \boldsymbol{B}_{-i}) \,\middle|\, \boldsymbol{B}_i = \boldsymbol{b}_i\right]\right] \leq \mathbb{E}\left[\mathbb{E}\left[\sum_{k\in M} v_{ik}(x_{ik}(h(\boldsymbol{b}_i), \boldsymbol{B}_{-i})) \,\middle|\, \boldsymbol{B}_i = \boldsymbol{b}_i\right]\right]. \tag{6.33}$$

Note that for auctions $k \neq j$, agent i always follows the recommendation, so the terms for these auctions k on the left- and right-hand side of equations (6.32) and (6.33) coincide. Agent i also follows the recommendation for auction j if $b_{ij} \geq r_j$, so in this case the terms also coincide. Finally, consider recommendations for auction j with $b_{ij} < r_j$. In this case the terms for auction j on the left- and right-hand side of equation (6.32) are equal to 0, as agent i never wins auction j when bidding below the reserve price r_j. For the deviation, the terms on the left- and right-hand side in (6.33) satisfy:

$$\begin{aligned} &\mathbb{E}\left[\mathbb{E}\left[p_{ij}(h(\boldsymbol{b}_i), \boldsymbol{B}_{-i}) \mid \boldsymbol{B}_i = \boldsymbol{b}_i \wedge b_{ij} < r_j\right]\right] \\ &\leq \mathbb{E}\left[\mathbb{E}\left[v_{ij}(x_{ij}(h(\boldsymbol{b}_i), \boldsymbol{B}_{-i})) \mid \boldsymbol{B}_i = \boldsymbol{b}_i \wedge b_{ij} < r_j\right]\right], \end{aligned}$$

6

as $r_j < v_{ij}$ (reserve prices are feasible) and agent i deviates to the reserve price r_j for auction j if $b_{ij} < r_j$. Combining this with the reasoning above, it follows that the deviation of agent i satisfies the ROI constraint in (6.33), i.e., $(h(\boldsymbol{B}_i), \boldsymbol{B}_{-i}) \in \mathcal{R}_i$.

We now show that agent i is strictly better off with this deviation. First, note that we have:

$$\begin{aligned}
\mathbb{E}[g_i(\boldsymbol{b}_i, \boldsymbol{B}_{-i}) \mid \boldsymbol{B}_i = \boldsymbol{b}_i] &= \int g_i(\boldsymbol{b}_i, \boldsymbol{b}_{-i}) \cdot f_{\boldsymbol{B}_{-i}|\boldsymbol{B}_i}(\boldsymbol{b}_{-i} \mid \boldsymbol{b}_i) d\boldsymbol{b}_{-i} \\
&= \int \left[\sum_{k \in M; k \neq j} v_{ik} x_{ik}(\boldsymbol{b}_i, \boldsymbol{b}_{-i}) - \sigma_i p_{ik}(\boldsymbol{b}_i, \boldsymbol{b}_{-i}) \right] \cdot f_{\boldsymbol{B}_{-i}|\boldsymbol{B}_i}(\boldsymbol{b}_{-i} \mid \boldsymbol{b}_i) d\boldsymbol{b}_{-i} \\
&+ \int [v_{ij} x_{ij}(\boldsymbol{b}_i, \boldsymbol{b}_{-i}) - \sigma_i p_{ij}(\boldsymbol{b}_i, \boldsymbol{b}_{-i})] \cdot f_{\boldsymbol{B}_{-i}|\boldsymbol{B}_i}(\boldsymbol{b}_{-i} \mid \boldsymbol{b}_i) d\boldsymbol{b}_{-i}.
\end{aligned}$$

Next we have that:

$$\begin{aligned}
\mathbb{E}[g_i(h(\boldsymbol{b}_i), \boldsymbol{B}_{-i}) \mid \boldsymbol{B}_i = \boldsymbol{b}_i] &= \int g_i(h(\boldsymbol{b}_i), \boldsymbol{b}_{-i}) \cdot f_{\boldsymbol{B}_{-i}|\boldsymbol{B}_i}(\boldsymbol{b}_{-i} \mid \boldsymbol{b}_i) d\boldsymbol{b}_{-i} \\
&= \int \left[\sum_{k \in M; k \neq j} v_{ik} x_{ik}(\boldsymbol{b}_i, \boldsymbol{b}_{-i}) - \sigma_i p_{ik}(\boldsymbol{b}_i, \boldsymbol{b}_{-i}) \right] \cdot f_{\boldsymbol{B}_{-i}|\boldsymbol{B}_i}(\boldsymbol{b}_{-i} \mid \boldsymbol{b}_i) d\boldsymbol{b}_{-i} \\
&+ \int [v_{ij} x_{ij}(h(\boldsymbol{b}_i), \boldsymbol{b}_{-i}) - \sigma_i p_{ij}(h(\boldsymbol{b}_i), \boldsymbol{b}_{-i})] \cdot f_{\boldsymbol{B}_{-i}|\boldsymbol{B}_i}(\boldsymbol{b}_{-i} \mid \boldsymbol{b}_i) d\boldsymbol{b}_{-i},
\end{aligned}$$

where the second equality follows by definition of the deviation. Subtracting the former conditional expectation from the latter leads to:

$$\begin{aligned}
&\mathbb{E}[g_i(h(\boldsymbol{b}_i), \boldsymbol{B}_{-i}) \mid \boldsymbol{B}_i = \boldsymbol{b}_i] - \mathbb{E}[g_i(\boldsymbol{b}_i, \boldsymbol{B}_{-i}) \mid \boldsymbol{B}_i = \boldsymbol{b}_i] \\
&= \int \Big[v_{ij} x_{ij}(h(\boldsymbol{b}_i), \boldsymbol{b}_{-i}) - \sigma_i p_{ij}(h(\boldsymbol{b}_i), \boldsymbol{b}_{-i}) - v_{ij} x_{ij}(\boldsymbol{b}) + \sigma_i p_{ij}(\boldsymbol{b}) \Big] \\
&\cdot f_{\boldsymbol{B}_{-i}|\boldsymbol{B}_i}(\boldsymbol{b}_{-i} \mid \boldsymbol{b}_i) d\boldsymbol{b}_{-i}. \qquad (6.34)
\end{aligned}$$

Note that the expression in brackets in (6.34) evaluates to 0 for $\boldsymbol{b}_i$ with $b_{ij} \geq r_j$, as in that case $b'_{ij} = b_{ij}$. Otherwise, if $b_{ij} < r_j$, then agent i never wins auction j under $\boldsymbol{b}_i$ and if agent i wins auction j under $(h(\boldsymbol{b}_i), \boldsymbol{b}_{-i})$, the expression evaluates to:

$$v_{ij} x_{ij}(h(\boldsymbol{b}_i), \boldsymbol{b}_{-i}) - \sigma_i p_{ij}(h(\boldsymbol{b}_i), \boldsymbol{b}_{-i}) = v_{ij} - \sigma_i r_j > 0,$$

as $b'_{ij} = r_j < v_{ij}$ (reserve prices are feasible) and $\sigma_i \leq 1$. Furthermore, it follows from (6.30) that there is some positive probability that agent i wins auction j in this case. This leads to:

$$\begin{aligned}
\mathbb{E}[g_i(h(\boldsymbol{B}_i), \boldsymbol{B}_{-i})] &= \mathbb{E}\left[\mathbb{E}[g_i(h(\boldsymbol{b}_i), \boldsymbol{B}_{-i}) \mid \boldsymbol{B}_i = \boldsymbol{b}_i]\right] \\
&> \mathbb{E}\left[\mathbb{E}[g_i(\boldsymbol{b}_i, \boldsymbol{B}_{-i}) \mid \boldsymbol{B}_i = \boldsymbol{b}_i]\right] = \mathbb{E}[g_i(\boldsymbol{B})],
\end{aligned}$$

6

contradicting that $\boldsymbol{B}$ is a CE. □

It is no longer true that all equilibria are well-supported beyond correlated equilibria. Since we consider feasible reserve prices, this phenomenon is not due to reserve prices exceeding the agents' willingness to pay. Inspired by Syrgkanis [2014], we show that already for a single first-price auction, there exist CCE for which the item is not always sold. Even though the rightful winner i can ensure that the item is always sold, agent i might have to pay more in some cases or even not win the item when deviating from the joint probability distribution, causing the gain of agent i to not strictly improve.

Lemma 6.5.2. *Consider the class of simultaneous first-price auctions with additive valuations $\boldsymbol{v} \in \mathcal{V}_{\text{ADD}}$ and feasible reserve prices. Then, there exist CCE that are not well-supported.*

Proof: Consider a single first-price auction with a feasible reserve price r among 2 agents with values $v_1 = 1$ and $v_2 = 0$. Let $\sigma_1 = 1$. The type of agent 2 can be arbitrary. Assume that ties are broken in favor of agent 1 and note that $r < 1$, as the reserve price is feasible. Let X be a random variable with the following CDF and PDF:

$$F_X(x) = \frac{1}{e} \cdot \frac{1-r}{1-x} \quad \text{and} \quad f_X(x) = \frac{1}{e} \cdot \frac{1-r}{(1-x)^2} \quad \text{with} \quad x \in \left[0, \frac{e-1+r}{e}\right].$$

Note that the domain is well-defined as $0 \le x < 1$, and it is easy to verify that $F_X(\cdot)$ is non-negative and increasing over the domain and $F_X((e-1+r)/e) = 1$. Now consider the randomized bid profile $\boldsymbol{B}$ for which the agents always bid identically according to X, i.e., draw some value x from X and let both agents bid x. We prove that $\boldsymbol{B}$ is a CCE.

Since ties are broken in favor of agent 1, the expected gain of agent 2 under $\boldsymbol{B}$ is 0 and agent 2 satisfies the ROI constraint. Moreover, as $v_2 = 0$, no unilateral deviation of agent 2 leads to a positive expected gain, and so agent 2 satisfies the equilibrium condition. Note that agent 1 satisfies the ROI constraint for $\boldsymbol{B}$ as $x \le (e-1+r)/e < 1 = v_1$, so agent 1 never bids above v_1. As agent 1 does not win when bidding below the reserve price, the expected gain of agent 1 is:

$$\mathbb{E}[g_1(\boldsymbol{B})] = \int_r^{\frac{e-1+r}{e}} (1-x) f_X(x) dx = \frac{1-r}{e}\Big[-\ln(1-x)\Big]_r^{\frac{e-1+r}{e}} = \frac{1-r}{e}.$$

First, consider any deterministic unilateral deviation b_1' of agent 1. If $b_1' \in [0, r)$, then agent 1 will never win the auction and always have a gain of 0. Secondly, consider a bid $b_1' \in [r, \frac{e-1+r}{e}]$. Let $B_2 = \boldsymbol{B}_{-1}$, then:

$$\mathbb{E}[g_1(b_1', B_2)] = (1-b_1')F_X(b_1') = (1-b_1')\frac{1-r}{e(1-b_1')} = \frac{1-r}{e},$$

where the first equality holds as agent 1 wins and pays b_1' if agent 2 bids at most b_1'. Finally, if $b_1' > \frac{e-1+r}{e}$, agent 1 always wins the auction but this leads to an expected gain of $1 - b_1'$ that is strictly smaller than $\frac{1-r}{e}$. Therefore, we can conclude that for any randomized unilateral deviation Y with PDF f_Y of agent 1 satisfying the ROI constraint, it holds that:

$$\mathbb{E}[g_1(Y, B_2)] = \int_0^\infty \mathbb{E}[g_1(y, B_2)] \cdot f_Y(y)dy \leq \int_0^\infty \frac{1-r}{e} f_Y(y)dy = \frac{1-r}{e}.$$

Therefore, agent 1 also satisfies the equilibrium condition, proving that $\boldsymbol{B}$ is a CCE that is not well-supported. Note that this holds for any feasible reserve price r. Additionally, for any reserve price $r > 0$, the CCE $\boldsymbol{B}$ is not well-supported as in this case the probability that the item is not sold is greater than 0.

Additionally, note that this can be extended to hold for any simultaneous first-price auction, i.e., *FPA*$(m, \boldsymbol{r})$ with $m > 1$, simply by considering $\boldsymbol{B}$ for an auction $j \in \{1, 2, \ldots, m\}$ and letting the randomized bid profiles of the auctions $i \neq j$ be independent of $\boldsymbol{B}$. □

As the item is not always sold under the CCE $\boldsymbol{B}$ in the proof of Lemma 6.5.2, a lower bound on the POA can be implied. This results is stated below in Theorem 6.5.3 and shows that, for the homogeneous case of utility maximizers, the *CCE-POA* of simultaneous first-price auctions *with* feasible reserve prices is not better than that of simultaneous first-price auctions *without* reserve prices.

Theorem 6.5.3. *Consider the class of simultaneous first-price auctions with additive valuations* $\boldsymbol{v} \in \mathcal{V}_{\text{ADD}}$ *and feasible reserve prices. Then, it holds that:*

$$\textit{CCE-POA}(\mathcal{V}_{\text{ADD}}) \geq \frac{e}{e-1}.$$

Proof: Consider the *FPA*(r) among two agents with feasible reserve price and the CCE $\boldsymbol{B}$ as in the proof of Lemma 6.5.2. It is easy to see that the optimal liquid welfare in this instance is $1 = v_{\mathsf{rw}}$. The expected liquid welfare of $\boldsymbol{B}$ is $\mathbb{E}[\sum_{i \in N} v_i(x_i(\boldsymbol{B}))] = \mathbb{P}[X \geq r] = \frac{e-1}{e}$. Therefore, *CCE-POA* $\geq \frac{e}{e-1}$, concluding the proof. □

However, as shown in Theorem 6.4.15, the presence of feasible reserve prices does improve the POA when only considering well-supported CCE. The following result shows that this bound on the POA is even tight for all $\eta \in (0, 1)$ in the case of homogeneous utility maximizers.

Theorem 6.5.4. *Consider the class of simultaneous first-price auctions with additive valuations* $\boldsymbol{v} \in \mathcal{V}_{\text{ADD}}$ *and feasible reserve prices among a set of* N *homogeneous agents with* $\sigma_i \geq \frac{e-1}{e-\eta}$ *for all* $i \in N$*. Then, for well-supported CCE, it holds that:*

$$\textit{CCE-POA}(\mathcal{V}_{\text{ADD}}) \geq \frac{e}{e-1+\eta}.$$

6

Proof: Consider a first-price auction with reserve price r among $n = 2$ agents of the same type, i.e., $\sigma := \sigma_1 = \sigma_2$. Let $\boldsymbol{v} = (1, \eta)$, with $\eta \in [0, 1)$, and let $\sigma = \frac{e-1}{e-\eta}$. The feasible reserve price is $r = \eta < v_1$. We assume that ties are broken in favor of agent 1 for every bid profile $\boldsymbol{b} = (b, b)$ with $b > \eta$, and in favor of agent 2 for $\boldsymbol{b} = (\eta, \eta)$. Let X be a random variable with the following CDF and PDF:

$$F_X(x) = \frac{1}{e} \cdot \frac{1-\sigma\eta}{1-\sigma x} \quad \text{and} \quad f_X(x) = \frac{\sigma}{e} \cdot \frac{1-\sigma\eta}{(1-\sigma x)^2} \quad \text{with} \quad x \in \left[\eta, \frac{e-1+\sigma\eta}{e\sigma}\right].$$

Note that the domain is well-defined as $x \geq 0$, and it is easy to verify that $F_X(\cdot)$ is non-negative and increasing over the domain and $F_X((e-1+\sigma\eta)/e\sigma) = 1$. Now consider the randomized bid profile $\boldsymbol{B}$ in which the agents always bid identically according to X, i.e., draw some value x from X and let both agents bid x. We prove that $\boldsymbol{B}$ is a CCE.

First, we verify that $\boldsymbol{B}$ satisfies the ROI constraints of the two agents. Note that this is trivially true for agent 2 as, by the tie-breaking rule, agent 2 only wins the item for the bid profile (η, η) and $\eta = v_2$. Note that $\boldsymbol{B} \in \mathcal{R}_1$ since:

$$\mathbb{E}[p_1(\boldsymbol{B})] \leq \frac{e-1+\sigma\eta}{e\sigma} \cdot \mathbb{E}[x_1(\boldsymbol{B})] \leq \mathbb{E}[x_1(\boldsymbol{B})] = \mathbb{E}[v_1(x_1(\boldsymbol{B}))],$$

where the first inequality holds as agent 1 never bids above $\frac{e-1+\sigma\eta}{e\sigma}$, while the second inequality holds by assumption.

Secondly, we show that no agent can improve by unilaterally deviating. First, note that under $\boldsymbol{B}$ agent 2 only wins when bidding bid η, and that the gain of agents 2 is non-negative in this case, i.e., $v_2 - \sigma\eta = \eta(1-\sigma) \geq 0$. We consider deterministic deviations of agent 2. Agent 2 would never win when bidding $b_2 < \eta = r$. If agent 2 wins when bidding $b_2 > \eta$, agent 2 violates their ROI constraint. And if agents 2 wins when bidding $b_2 = \eta$, their ROI constraint is tight. Therefore, agent 2 has no randomized unilateral deviation that is beneficial and satisfies the ROI constraint.

6

As ties are broken in favor of agent 1 for every bid profile $\boldsymbol{b} = (b, b)$ with $b > \eta$, the expected gain of agent 1 under $\boldsymbol{B}$ is:

$$\begin{aligned}
\mathbb{E}[g_1(\boldsymbol{B})] &= \int_{\eta}^{\frac{e-1+\sigma\eta}{e\sigma}} (1-\sigma x) f_X(x) dx = \int_{\eta}^{\frac{e-1+\sigma\eta}{e\sigma}} \frac{\sigma}{e} \frac{1-\sigma\eta}{1-\sigma x} dx \\
&= \frac{1-\sigma\eta}{e} \Big[-\ln(1-\sigma x) \Big]_{\eta}^{\frac{e-1+\sigma\eta}{e\sigma}} = \frac{1-\sigma\eta}{e}.
\end{aligned}$$

Consider a deterministic deviation b_1' of agent 1. Clearly, $b_1' \leq \eta$ leads to an expected gain of $\mathbb{E}[g_1(b_1', X)] = 0$, as agent 2 always wins the item at (η, η). Moreover, note that bidding $b_1' > \frac{e-1+\sigma\eta}{e\sigma}$ is not beneficial. In this case agent 1 would always win the item but:

$$\mathbb{E}[g_1(b_1', X)] = 1 - \sigma b_1' < 1 - \sigma \frac{e-1+\sigma\eta}{e\sigma} = \frac{1-\sigma\eta}{e} = \mathbb{E}[g_1(\boldsymbol{B})].$$

Finally, consider $b'_1 \in (\eta, \frac{e-1+\sigma\eta}{e\sigma}]$. We obtain that:

$$\mathbb{E}[g_1(b'_1, X)] = \mathbb{P}[b'_1 \geq X](1-\sigma b'_1) = F_X(b'_1)(1-\sigma b'_1) = \frac{1-\sigma\eta}{e} = \mathbb{E}[g_1(\boldsymbol{B})].$$

Therefore, we can conclude that for any randomized unilateral deviation Y with PDF f_Y of agent 1 satisfying the ROI constraint, it holds that:

$$\mathbb{E}[g_1(Y, X)] = \int_0^\infty \mathbb{E}[g_1(y, X)] \cdot f_Y(y)dy \leq \int_0^\infty \frac{1-\sigma\eta}{e} f_Y(y)dy = \frac{1-\sigma\eta}{e},$$

proving that $\boldsymbol{B}$ is a CCE.

We conclude the proof by showing how inefficient the CCE $\boldsymbol{B}$ is. Clearly, the optimal liquid welfare is $v_1 = 1$, and therefore:

$$\textit{CCE-POA} \geq \frac{1}{\mathbb{E}[v_1(x_1(\boldsymbol{B})) + v_2(x_2(\boldsymbol{B}))]} = \frac{1}{1 - F_X(\eta) + \eta F_X(\eta)} = \frac{e}{e-1+\eta}.$$

□

Note that the lower bound on the POA for homogeneous agents in Theorem 6.5.4 holds for a certain range of η, depending on the sensitivity of the agents. Given the sensitivity of the homogeneous agents, Figure 6.2 shows the maximum value of η for which the precondition of Theorem 6.5.4 is satisfied. Note that Theorem 6.5.4 only holds for sensitivities greater than or equal to $\frac{e-1}{e}$. If all agents are value maximizers, we can prove a lower bound on the POA of $2-\eta$ for the whole range of η, matching the upper bound in Theorem 6.4.15.

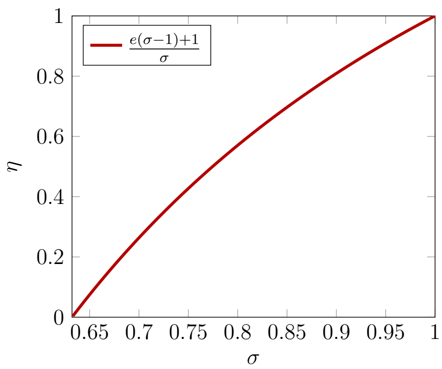

Figure 6.2: Maximum value of η (vertical axis) for which the precondition of Theorem 6.5.4 is satisfied for a certain sensitivity of the homogeneous agents (horizontal axis).

Theorem 6.5.5. *Consider the class of simultaneous first-price auctions with additive valuations $\boldsymbol{v} \in \mathcal{V}_{ADD}$ and feasible reserve prices among value maximizers. Then, for $\eta \in [0, 1)$, it holds that:*

$$\textit{MNE-POA}(\mathcal{V}_{ADD}) \geq 2 - \eta.$$

Proof: Consider a $FPA(2, \boldsymbol{r})$ among two value maximizing agents with additive valuations and feasible reserve prices. Assume that ties are broken in favor of agent 1 for both auctions. For auctions i and j, consider a small $\varepsilon > 0$ and let $\eta_i \in [0, 1-\varepsilon]$ and $\eta_j = 1 - \varepsilon \geq 0$. Note that $\eta = \eta_i$. Let $v_{1i} = 1$ and $v_{1j} = 0$ and let $v_{2i} = 0$ and $v_{2j} = 1 - \eta_i$. Note that the reserve prices are feasible and equal to $r_i = \eta_i$ and $r_j = (1-\varepsilon)(1-\eta_i)$. Consider a bid profile $\boldsymbol{b}$ with $\boldsymbol{b}_1 = (\eta_i, 1-\eta_i)$ and $\boldsymbol{b}_2 = (0, 1-\eta_i)$. We show that $\boldsymbol{b}$ is a MNE.

First note that under $\boldsymbol{b}$ the ROI constraints are satisfied for both agents, as agent 1 wins both items with a total payment of 1. Also note that agent 2 cannot win anything by deviating without violating their ROI constraint. This is because agent 1 deterministically bids the value v_{2j} of agent 2 for item j and ties are broken in favor of agent 1, and agent 2 has 0 value for item i. As agent 1 wins both items under $\boldsymbol{b}$, agent 1 cannot improve their gain with any (random) unilateral deviation. Therefore, $\boldsymbol{b}$ is a MNE.

It is easy to see that the optimal liquid welfare in this instance is achieved when agent 1 is allocated item i and agent 2 is allocated item j, and is equal to $2 - \eta_i = 2 - \eta$. The liquid welfare of the MNE $\boldsymbol{b}$ is equal to $v_{1i} + v_{1j} = 1$. Therefore, the *MNE-POA* of this instance is to $2 - \eta$, concluding the proof. □

6.5.2 Submodular Valuations

If we go beyond additive valuations and consider simultaneous first-price auctions with submodular valuations and feasible reserve prices, then there already exist MNE that are not well-supported, as we show in the lemma below. Using this result, we show that omitting the precondition of equilibria being well-supported results in an unbounded *MNE-POA* in the presence of utility maximizers.

6

Lemma 6.5.6. *Consider the class of simultaneous first-price auctions with submodular valuations* $\boldsymbol{v} \in \mathcal{V}_{\text{SUB}}$ *and feasible reserve prices. Then, there exist MNE that are not well-supported.*

Proof: Consider a $FPA(2, \boldsymbol{r})$ among two agents with feasible reserve prices. Let $\sigma_1 = 1$ for agent 1 and for agent 2, let $\sigma_2 \leq 1$ be arbitrary. Consider the submodular valuations with $v_1(\{i\}) = \varepsilon$, $v_1(\{j\}) = 1$ and $v_1(\{i,j\}) = 1 + \frac{\varepsilon}{2}$. Let the values of agent 2 be $v_2(\{i\}) = v_2(\{j\}) = v_2(\{i,j\}) = \frac{\varepsilon}{4}$. Now consider the feasible reserve prices $r_i = \frac{\varepsilon}{2}$ and $r_j = 1 - \frac{\varepsilon}{2}$. The bid profile $\boldsymbol{b}$ with $\boldsymbol{b}_1 = (\frac{\varepsilon}{2}, 0)$ and $\boldsymbol{b}_2 = (0,0)$ is a MNE. First, note that under $\boldsymbol{b}$ the ROI constraints are satisfied for both agents. Also note that agent 2 cannot win anything without violating their ROI constraint, as, for both items, the reserve price exceeds their value. Agent 1 only wins item i under $\boldsymbol{b}$ and thus has a gain of $g_1(\boldsymbol{b}) = \frac{\varepsilon}{2}$. Note that for any $\boldsymbol{b}^1$ with $b^1_i \geq \frac{\varepsilon}{2}$ and $b^1_j < 1 - \frac{\varepsilon}{2}$, $\boldsymbol{b}^2$ with $b^2_i < \frac{\varepsilon}{2}$ and $b^2_j \geq 1 - \frac{\varepsilon}{2}$, $\boldsymbol{b}^3$ with $b^3_i \geq \frac{\varepsilon}{2}$ and $b^3_j \geq 1 - \frac{\varepsilon}{2}$ and $\boldsymbol{b}^4$ with $b^4_i < \frac{\varepsilon}{2}$ and $b^4_j < 1 - \frac{\varepsilon}{2}$, it holds that $g_1(\boldsymbol{b}^k, \boldsymbol{b}_2) \leq \frac{\varepsilon}{2}$

for $k = 1, 2, 3, 4$. Therefore, agent 1 also satisfies the equilibrium conditions, as, regardless of the ROI constraint, it holds for any randomized bid profile B_1' that:

$$\mathbb{E}[g_1(B_1', \boldsymbol{b}_2)] = \int_0^\infty g_1(\boldsymbol{b}', \boldsymbol{b}_2) f_{\boldsymbol{B}_1'}(\boldsymbol{b}') d\boldsymbol{b}' \leq \frac{\varepsilon}{2} \int_0^\infty f_{\boldsymbol{B}_1'}(\boldsymbol{b}') d\boldsymbol{b}' = \frac{\varepsilon}{2} = g_1(\boldsymbol{b}).$$

Therefore, the bid profile $\boldsymbol{b}$ is a MNE. Furthermore, it holds that $\boldsymbol{b}$ is not well-supported, as the item of auction j is not sold. □

Theorem 6.5.7. *Consider the class of simultaneous first-price auctions with submodular valuations $\boldsymbol{v} \in \mathcal{V}_{\text{SUB}}$ and feasible reserve prices. Then, it holds that the MNE-POA$(\mathcal{V}_{\text{SUB}})$ is unbounded.*

Proof: Consider the $FPA(2, \boldsymbol{r})$ among two agents with feasible reserve prices and the MNE $\boldsymbol{b}$ as in the proof of Lemma 6.5.6. It is easy to see that the optimal liquid welfare in this instance is achieved when agent 1 is allocated both items and is equal to $1 + \frac{\varepsilon}{2}$. The liquid welfare of the MNE $\boldsymbol{b}$ is equal to ε. Therefore, the *MNE-POA* of this instance is equal to $\frac{1+\frac{\varepsilon}{2}}{\varepsilon}$, and as $\varepsilon \to 0$, this becomes unbounded $\frac{1+\frac{\varepsilon}{2}}{\varepsilon} \to \infty$. □

6.6 Learning Outcomes

In this section, we consider our model in the framework of a repeated auction setting as in [Kolumbus and Nisan, 2022]. Specifically, we consider 2 agents[3] with arbitrary payment sensitivities that repeatedly participate in a first-price auction with a feasible reserve price r in which ties are broken uniformly at random. It is assumed that each agent i cannot overbid, i.e., $b_i \leq v_i$ (as it still holds w.l.o.g. that $\tau = 1$). The item sold as well as the values and types of the agents are assumed to be the same at each repetition of the auction. It is assumed that the value of both agents, the bids and the reserve price r are multiples of a minimum value $\varepsilon > 0$. The objective of the agents is additive over the auctions, and each agent i wants to maximize their gain g_i. We note that in the homogeneous case of utility maximizers and no reserve price ($r = 0$), our setting and results coincide with [Kolumbus and Nisan, 2022]. The agents use regret-minimization algorithms, with the regret of agent i at time T given bids $(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_T)$ defined as $R_i^T = \sum_{t=1}^T \max_b g_i(b, \boldsymbol{b}_{-i}^t) - g_i(b_i^t, \boldsymbol{b}_{-i}^t)$, with b the optimal fixed action (bid) in hindsight. In their work, Kolumbus and Nisan [2022] define a specific class of CCE, called co-undominated. In a co-undominated CCE, there is no action in the support of an agent that is weakly dominated in the support of the other agents.

[3]We consider 2 agents for simplicity, but the results can be extended to more than 2 agents.

6

Definition 6.6.1. Let A_1 and A_2 denote subsets of the action spaces of two agents in a two-agent game. Action $i \in A_1$ of agent 1 is weakly dominated in A_2 by action i' of agent 1 if $\forall j \in A_2 : g_1(i,j) \leq g_1(i',j)$ and $\exists j \in A_2 : g_1(i,j) < g_1(i',j)$.

Definition 6.6.2: ([Kolumbus and Nisan, 2022]). Let $\boldsymbol{B}$ be a CCE of a (finite) two-agent game with action spaces I_1 and I_2. Denote its support (A_1, A_2) by $A_1 = \{i \in I_1 \mid \exists j \in I_2 \text{ with } B_{ij} > 0\}$ and $A_2 = \{j \in I_2 \mid \exists i \in I_1 \text{ with } B_{ij} > 0\}$. A CCE is called co-undominated if for every $i \in A_1$, and every $i' \in I_1$, action i is not weakly dominated in A_2 by i', and similarly for A_2.

If mean-based regret-minimizing algorithms converge, Kolumbus and Nisan [2022] prove that they converge to co-undominated CCE. This convergence result also holds for our setting of a first-price auction with feasible reserve price and different agent types. Additionally, it turns out co-undominated CCE have the desired property of being well-supported.

Lemma 6.6.3. *Consider a first-price auction with feasible reserve price among two agents with arbitrary types and action spaces I_1 and I_2. Then, any co-undominated CCE is well-supported.*

Proof: Let $\boldsymbol{B}$ be a co-undominated CCE of a $FPA(r)$ with feasible reserve price among two agents with arbitrary types and action spaces I_1 and I_2. For contradiction, assume that $\boldsymbol{B}$ is not well-supported, i.e., the item is not always sold. Denote the support (A_1, A_2) of $\boldsymbol{B}$ by $A_1 = \{i \in I_1 \mid \exists j \in I_2 \text{ with } B_{ij} > 0\}$ and $A_2 = \{j \in I_2 \mid \exists i \in I_1 \text{ with } B_{ij} > 0\}$. Let agent 1 have support A_1 and assume w.l.o.g. that $v_1 \geq v_2$. As the item is not always sold, $\exists i \in A_1$ and $\exists j \in A_2$ such that $i < r$, $j < r$ and $B_{ij} > 0$. However, in this case action i of agent 1 is weakly dominated in A_2 by action r as $\forall j' \in A_2$ it holds that $0 = g_1(i,j') \leq g_1(r,j')$ and $0 = g_1(i,j) < g_1(r,j) = v_1 - \sigma_1 r$, as $r < v_1$ and $\sigma_1 \leq 1$. □

6

Therefore, if the dynamics converge to a co-undominated CCE $\boldsymbol{B}$, the POA results of the previous sections apply as $\boldsymbol{B}$ is well-supported. Furthermore, we can extend the result of Kolumbus and Nisan [2022] in case of convergence of the dynamics to a single distribution. In our setting, the wining bid doesn't converge to the second highest value but (almost always) converges to the maximum of the reserve prices and the second highest value. If $v_1 > v_2$ and $\sigma_1 = 0$, the winning bid doesn't necessarily converge as agent 1 is not sensitive to the price. In this case the support of agent 1 in a co-undominated CCE is in $[\max\{r, v_2\}, \max\{r, v_2 + \varepsilon\}, \ldots, \max\{r, v_1\}]$. Otherwise, if $v_1 \geq v_2$, the support of agent 1 in a co-undominated CCE is in $[\max\{r, v_2 - 2\varepsilon\}, \max\{r, v_2 - \varepsilon\}, \ldots, \max\{r, v_2 + \varepsilon\}]$ and the wining bid indeed converges to the maximum of the reserve price and the second highest value. This results follows from the three lemmas below. For their proofs, we refer to Appendix A.3.

Lemma 6.6.4. *Consider a first-price auction with feasible reserve price among two agents with arbitrary types and values* $v_1 > v_2 > 2\varepsilon$*. Let the discrete bid levels be multiples of* $\varepsilon \leq \frac{1}{2}(v_1 - v_2)$*. Then, in any co-undominated CCE, if* $\sigma_1 > 0$ *the support of agent 1 is in* $[\max\{r, v_2\}, \max\{r, v_2 + \varepsilon\}]$.

Lemma 6.6.5. *Consider a first-price auction with feasible reserve price among two agents with arbitrary types and values* $v_1 = v_2 = v > 2\varepsilon$*. Let the discrete bid levels be multiples of* ε*. Then, in any co-undominated CCE, if* $\sigma_1 > 0$ *and* $\sigma_2 > 0$ *there is an agent with their support in* $[\max\{r, v - 2\varepsilon\}, \max\{r, v - \varepsilon\}, v] = [\max\{r, v - 2\varepsilon\}, v - \varepsilon, v]$.

Lemma 6.6.6. *Consider a first-price auction with feasible reserve price among two agents with arbitrary types and values* $v_1 \geq v_2 \geq \varepsilon$*. Let the discrete bid levels be multiples of* ε*. Then, in any co-undominated CCE, if* $\sigma_1 = 0$ *and* $v_1 = v_2$ *agent 1 deterministically bids* v_1*, otherwise, if* $\sigma_1 = 0$ *and* $v_1 > v_2$ *the support of agent 1 is in* $[\max\{r, v_2\}, \max\{r, v_2 + \varepsilon\}, \dots, \max\{r, v_1\}]$.

6.7 Conclusion

In this chapter, we derived bounds on the price of anarchy for simultaneous single-time first-price auctions under the hybrid agent model, in which the negative effect of payments on the utility (gain) of an agent depends on the agent's type. This is modeled by payment sensitivities, which range from 0 in the case of value maximizers to 1 in the case of utility maximizers. Additionally, each agent has a return-on-investment constraint, enforcing that the expected payment does not exceed the expected acquired value. Here, the POA measures the efficiency loss due to strategic behavior in terms of liquid welfare, a metric that takes the limited financial resources (or willingness to pay) of agents into account. This model is especially relevant in online advertising, as advertisers have recently transitioned to delegating bidding decisions to automated agents.

We derived bounds on the POA by extending the smoothness framework of Syrgkanis and Tardos. A crucial element of our approach is introducing calibration vectors that enable us to balance the different smoothness parameters of each agent type in our Extension Theorem. We can then obtain bounds on the POA through our POA-revealing mathematical program. The generality and extendability of our approach allows us to derive bounds for fractionally subadditive valuation functions, coarse correlated equilibria, and any composition of agent types, extending and adding to known bounds Additionally, we derive bounds for well-supported equilibria of simultaneous FPAs with additive valuation functions and feasible reserve prices, and prove that any CE is well-supported in this case. For the homogeneous setting of utility or value maximizers, we derive the tight bounds of $e/(e-1+\eta)$ and $2-\eta$, where η measures the worst relative gap between the reserve price of an auction and the maximum payment that could be extracted.

An interesting direction for future work would be to incorporate reserve prices beyond additive valuation functions, as we show that for the homogeneous setting of utility maximizers, there are MNE that are not well-supported, resulting in an unbounded POA. Furthermore, a natural direction is to investigate if our framework can be applied to other auction formats, such as ad-auctions and multi-unit auctions, and even more general models, e.g., models that can additionally incorporate budget constraint of agents.

Final Remarks and Appendix

Chapter 7

Final Remarks

In the first part of this thesis, we explored the domain of mechanism design without monetary transfers. In Chapter 3, we contributed to the emerging line of research on learning-augmented mechanism design. We derived new results for the generalized assignment problem in the private graph model augmented with output predictions. For the bipartite matching problem, the approximation guarantee of our deterministic weakly group-strategyproof mechanism Boost smoothly interpolates between the optimal consistency-robustness trade-off and if the confidence parameter γ is large, the approximation guarantee is near-optimal for all other values of the error measure. Boost is inspired by the well-known deferred acceptance algorithm and makes tentative matches based on boosted valuations, where only the values of predicted edges are boosted by a factor of $\gamma \geq 1$. Boost can also be applied in the non-augmented setting by not using the prediction, or, said differently, by setting $\gamma = 1$ representing no confidence in the prediction. In that case Boost is always 2-approximate, just as the optimal greedy mechanism of Dughmi and Ghosh [2010] that matches edges irrevocably in the non-augmented setting. We believe that tentative matches are required in order to achieve improved approximation guarantees in the learning-augmented setting, and thus mechanisms beyond tweaking the known greedy mechanism had to be designed. Boost also proved to be a crucial component for more general variants of the problem, for which randomized mechanisms were known in the non-augmented setting. Namely, our mechanism randomizes over Boost, Trust (that simply adheres to the prediction), and a greedy mechanism, and achieves a superior consistency-robustness trade-off compared to a mechanism randomizing between the prediction and the best-known non-augmented mechanism. We suspect that further improving this consistency-robustness trade-off requires the development of a novel mechanism, that might additionally give rise to improvements in the non-augmented setting. 7

In Chapter 4, we introduced the perspective of mechanism design with outliers and considered single facility location on the real line as a test case, with the goal of achieving improved approximation guarantees by leveraging outliers. For the two most-studied objectives of utilitarian and egalitarian social cost, we provided a complete overview by deriving deterministic strategyproof mechanisms achieving the best-possible approximation guarantees. As the non-strategic benchmark got stronger and a strategyproof mechanism could not leverage outliers under the utilitarian objective, this led to the counter-intuitive phenomenon of the approximation guarantee degrading as the number of outliers increased. Furthermore, we derived a negative result for $n \leq 2z$ and showed that the problem is inapproximable under both objectives. It is not yet clear if the above-mentioned phenomenon and strong negative result persist for other classes of mechanism design problems, and it would be interesting to see if there are classes for which outliers can be leveraged to achieve improved approximation guarantees. If, at first glance, no mechanism design problem can be identified for which an improved approximation guarantee can be achieved, alternative approaches that incorporate outliers could be considered. We believe that gaining insights in the impact of outliers is theoretically interesting and relevant in applications involving agents with extreme or atypical preferences.

Budget-feasible mechanism design was the second central theme of this thesis. A comprehensive treatment of knapsack procurement auctions was given in Chapter 5, under the mild assumption of being able to afford each individual agent entirely. We connected the perspectives of the indivisible and divisible setting by considering partial allocations of granularity k, converging to the divisible setting as k increases. We introduced two truthful and budget-feasible mechanisms specifically tailored to linear valuation functions. For the divisible setting, our 2-approximate mechanism Prune-and-Assign, inspired by Gravin et al. [2020], establishes a separation with its indivisible counterpart for which a lower bound of $1 + \sqrt{2}$ is known due to Chen et al. [2011]. For the k-level model, the approximation guarantee of our mechanism improves as the granularity increases, dropping below 3 for $k \geq 8$ and converging to $1 + \varphi \approx 2.62$ as $k \to \infty$. It would be interesting to see if a suitable pruning step for the k-level model would give rise to an interpolation between our 2-approximation and the 3-approximation of Gravin et al. [2020] for the indivisible setting. We also considered valuation functions that are concave and non-decreasing, modeling diminishing returns of the auctioneer. The approximation guarantee of our mechanism for the k-level model is independent of the granularity k and equal to $2 + \sqrt{3}$. For the divisible setting, our mechanism achieves an approximation guarantee of $1 + \varphi$ by using a large granularity in order to upper bound the payments in a fine-grained manner. It would be interesting to understand whether we can obtain mechanisms with approximation guarantees equal to the ones derived for the settings with linear valuations, i.e., a guarantee depending on k for the k-level model. Given the similar ideas behind existing mechanisms and the scarcity of lower bounds, we believe that progress on these fronts will require novel techniques.

In the third part of this thesis, we changed our perspective from the design of mechanisms to the analysis of strategic interactions within an already predefined environment. In Chapter 6, we considered simultaneous single-item first-price auctions (FPAs) under the hybrid agent model with return-on-investment constraints. By combining the perspectives of smoothness and autobidding, we derived a general and extendable framework that allowed us to derive bounds on the price of anarchy (POA) for fractionally subadditive valuation functions, coarse correlated equilibria, and any composition of agent types. Additionally, we derived bounds via our framework for well-supported equilibria when considering additive valuation functions and feasible reserve prices, and proved that any correlated equilibrium is well-supported in this case. For the homogeneous setting of utility or value maximizers, our framework allowed us to derive tight bounds as a function of a measure of the quality of the reserve prices. An interesting follow-up question from both a theoretical and practical perspective is how, or even if, reserve prices can be incorporated in our framework beyond additive valuation functions.

Ultimately, this thesis has contributed towards the field of mechanism design by closing and improving gaps of achievable approximation guarantees of strategyproof mechanisms for a variety of problems. The recurring theme throughout the chapters has been the adoption of new perspectives to uncover novel insights, which is a key ingredient in the advancement of research.

Chapter A

Missing Material of Chapter 6

A.1 Equilibrium Notions

A.1.1 Hierarchy of Equilibrium Notions

To see that *MNE* $\subseteq$ *CE*, consider a *MNE* $\boldsymbol{B} = \prod_{i\in[n]} \boldsymbol{B}_i$. Note that $\boldsymbol{B}$ also satisfies the *CE* requirements:

$$\mathbb{E}[\mathbb{E}[g_i(h(\boldsymbol{b}_i), \boldsymbol{B}_{-i}) \mid \boldsymbol{B}_i = \boldsymbol{b}_i]] = \mathbb{E}[g_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})] \leq \mathbb{E}[g_i(\boldsymbol{B})],$$

where the equality holds as $\boldsymbol{B}_i$ and $\boldsymbol{B}_{-i}$ are independent random bid profiles and the swapping function h can be written as a randomized bid vector $\boldsymbol{B}'_i$ in this case. The randomized deviations of i that satisfy the ROI constraint also coincide for both equilibrium notions by the same reasoning.

To see that *CE* $\subseteq$ *CCE*, for contradiction, consider a *CE* $\boldsymbol{B}$ that is not a *CCE*. Then, there is an agent $i \in N$ with randomized bid profile $\boldsymbol{B}'_i$ satisfying the ROI constraint as in (6.2) such that $\mathbb{E}[g_i(\boldsymbol{B})] < \mathbb{E}[g_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})]$. However, this would contradict that $\boldsymbol{B}$ is a *CE*, as a swapping function h for agent i that maps all bid profiles to the randomized bid vector $\boldsymbol{B}'_i$ is equivalent to playing $\boldsymbol{B}'_i$ independent of $\boldsymbol{B}_{-i}$. This contradicts that $\boldsymbol{B}$ is a *CE*, as deviating to $\boldsymbol{B}'_i$ would strictly improve the gain of agent i and $\boldsymbol{B}'_i$ satisfies the ROI constraint as it coincides with the satisfied ROI constraint for *CCE* in this case.

A.1.2 Equilibrium Notions under ROI

As mentioned in Section 6.2, it is not true that a *PNE* w.r.t. deterministic deviations survives when randomization is allowed due to the ROI constraints. Therefore, given the following definition of a *PNE* for our setting, it is not true that *PNE* $\subseteq$ *MNE*. We illustrate this with an example after introducing the definition of a *PNE*.

Definition A.1.1. Let $\boldsymbol{b}$ be a deterministic bid profile satisfying the ROI constraint of each agent, i.e., $\boldsymbol{b} \in \mathcal{R}_i$ for all $i \in N$. Then, $\boldsymbol{b}$ is a *PNE* if for every agent $i \in N$ it holds that:

$$g_i(\boldsymbol{b}) \geq g_i(\boldsymbol{b}'_i, \boldsymbol{b}_{-i}) \quad \forall (\boldsymbol{b}'_i, \boldsymbol{b}_{-i}) \in \mathcal{R}_i.$$

Example A.1.2. Consider a *FPA*(2) with items i and j among 2 agents with additive valuations. In case of ties, they are always broken in favor of agent 1. Let $\boldsymbol{v}_1 = (v_{1i}, v_{1j}) = (0.5, 0.3)$ and let $\boldsymbol{v}_2 = (v_{2i}, v_{2j}) = (0.7, 0)$. Further, let $\sigma_1 = \sigma_2 = 0$ and recall that we can assume w.l.o.g. that $\tau_1 = \tau_2 = 1$. Consider the deterministic bid profile $\boldsymbol{b} = ((0, 0.2), (0.7, 0.2))$.

It is not hard to see that $\boldsymbol{b}$ is a *PNE*. Agent 2 always wins item i and never wins item j, which it has 0 value for, and agent 2 satisfies the ROI constraint. Agent 1 always wins item j and never wins item i, and satisfies the ROI constraint. However, agent 1 does have a positive value for item i, but agent 1 must at least bid 0.7 to win item i, which would violate the ROI constraint, regardless of agent 1 winning item j or not.

If we consider randomized deviations, we will see that $\boldsymbol{b}$ is not a *MNE*. Namely, consider the randomized bid profile $\boldsymbol{B}'_1$ for agent 1 which draws both bid profiles $\boldsymbol{b}_1^1 = (0, 0.2)$ and $\boldsymbol{b}_1^2 = (0.7, 0.2)$ with probability $\frac{1}{2}$. This leads to an expected gain for agent 1 of $\frac{1}{2} \cdot 0.3 + \frac{1}{2}(0.5 + 0.3) = 0.55$, which is strictly better than 0.3, which is the gain of agent 1 under $\boldsymbol{b}$. Further, $\boldsymbol{B}'_1$ satisfies the ROI constraint as $\frac{1}{2} \cdot 0.2 + \frac{1}{2}(0.7 + 0.2) = 0.55$.

What Example A.1.2 also shows is that, given some bid profile, it might be beneficial for an agent i to sometimes pay more than v_{ij} for an item j if this still leads to a positive gain, and compensate this 'overpaying' by 'underpaying' for another item so that the ROI constraint is still satisfied. To also account for this behavior in the definition of a *CE*, the ROI constraint for a deviation is evaluated for the entire support, and not only for the conditioned bid profile. Further, a deviation may require multiple recommendations to be mapped to a different strategy, i.e., the swapping function is $h : \text{supp}_i(\boldsymbol{B}) \to \Delta_i$. And finally, the gain of a deviation of an agent i is evaluated for the entire support, not only for the conditioned bid profile. The following example illustrates why these three elements are relevant.

A

Example A.1.3. Consider a *FPA*(2) with items i and j among 2 agents with additive valuations. In case of ties, they are broken in favor of agent 1. Further, let $\sigma_1 = \sigma_2 = 0$ (and $\tau_1 = \tau_2 = 1$).

First, consider the valuation $\boldsymbol{v}$ with $\boldsymbol{v}_1 = (v_{1i}, v_{1j}) = (0.2, 0.5)$ and $\boldsymbol{v}_2 = (v_{2i}, v_{2j}) = (0.2, 0)$. Consider $\boldsymbol{B}$ for which bid profiles $\boldsymbol{b}^1 = ((0, 0.7), (0.2, 0.7))$ and $\boldsymbol{b}^2 = ((0.2, 0.3), (0.2, 0.3))$ are both drawn with probability $\frac{1}{2}$. Note that both agents satisfy the ROI constraint under $\boldsymbol{B}$. If $\boldsymbol{b}^1$ is recommended, agent 1 only wins item j, but by deviating to 0.2 for item i, agent 1 could win both items and strictly improve their gain. However, this does not satisfy the ROI constraint when evaluated only for $\boldsymbol{b}^1$, but does satisfy the ROI constraint when evaluated for the entire support of $\boldsymbol{B}$.

Secondly, consider valuation $\boldsymbol{v}$ with $\boldsymbol{v}_1 = (v_{1i}, v_{1j}) = (0.4, 0.2)$ and $\boldsymbol{v}_2 = (v_{2i}, v_{2j}) = (0.6, 0)$. Consider $\boldsymbol{B}$ for which both bid profiles $\boldsymbol{b}^1 = ((0, 0.2), (0.5, 0.2))$ and $\boldsymbol{b}^2 = ((0, 0.2), (0.7, 0.1))$ are both drawn with probability $\frac{1}{2}$. Note that both agents satisfy the ROI constraint for $\boldsymbol{B}$. If $\boldsymbol{b}^1$ is recommended, agent 1 only wins item j, but by deviating to 0.5 for item i, agent 1 could win both items and strictly improve their gain. However, this would not satisfy the ROI constraint. But if additionally, when $\boldsymbol{b}^2$ is recommended, agent 1 deviates to 0.1 for item j, agent 1 will still win item j and the ROI will be satisfied.

Finally, consider valuation $\boldsymbol{v}$ with $\boldsymbol{v}_1 = (v_{1i}, v_{1j}) = (0.5, 0.1)$ and $\boldsymbol{v}_2 = (v_{2i}, v_{2j}) = (0, 0.2)$. Consider $\boldsymbol{B}$ for which bid profiles $\boldsymbol{b}^1 = ((0.4, 0), (0.4, 0.2))$ and $\boldsymbol{b}^2 = ((0.6, 0), (0.6, 0.2))$ are both drawn with probability $\frac{1}{2}$. Note that both agents satisfy the ROI constraint for $\boldsymbol{B}$. If $\boldsymbol{b}^1$ is recommended, agent 1 only wins item i, but by deviating to 0.2 for item j, agent 1 could win both items and strictly improve their gain. However, this would not satisfy the ROI constraint. But if additionally, when $\boldsymbol{b}^2$ is recommended, agent 1 deviates to 0 for item i, agent 1 will satisfy the ROI constraint. If in this case, the gain of agent 1 is only evaluated for the conditioned bid profile $\boldsymbol{b}^1$, the gain of agent 1 increases from 0.5 to 0.6. However, if the gain of agent 1 is evaluated for the entire support, the total gain of agent 1 decreases from 0.5 to 0.3.

A.2 Missing Material of Section 6.4.1

A.2.1 Proof of Lemma 6.4.4

Proof of Lemma 6.4.4: For brevity, let $\lambda_t = \lambda_t^*(\omega)$ and $\mu_t = \mu_t^*(\omega)$ for each $t \in T$, and let $L = L_\omega$ and $H = H_\omega$. Note that H is guaranteed to be non-empty by the range of ω, i.e., there is a type $t \in H$ with $\sigma_t = \sigma_{\max}$. All statements follow from the definition of $\mu_t^*(\omega)$ in (6.21) and elementary calculus.

(i) For every $t \in L$ with $\sigma_t = 0$, $\lambda_t = \mu_t = 1$ holds. Then, for all $t \in L$ with $\sigma_t > 0$, we have that:

$$\lambda_t = \frac{\mu_t}{\sigma_t}\left(1 - e^{-\frac{\sigma_t}{\mu_t}}\right) = \frac{1 - e^{\ln(1-\sigma_t)}}{-\ln(1-\sigma_t)} = \frac{\sigma_t}{-\ln(1-\sigma_t)}.$$

Similarly, for all $t \in H$ we have that:

$$\lambda_t = \frac{\mu_t}{\sigma_t}\left(1 - e^{-\frac{\sigma_t}{\mu_t}}\right) = \frac{1 - e^{\ln(1-\omega)}}{-\ln(1-\omega)} = \frac{\omega}{-\ln(1-\omega)}.$$

(ii) First, assume that $\{t \in L \mid \sigma_t > 0\} = \emptyset$. Then $L = \{t \in T \mid \sigma_t = 0\}$, and by identity *(i)* we have:

$$\min_{t \in T: \sigma_t > 0} \lambda_t = \min_{t \in H} \lambda_t = \frac{\omega}{-\ln(1-\omega)}.$$

Next, assume that that $\{t \in L \mid \sigma_t > 0\} \neq \emptyset$. By the definition of L, the fact that $f(x) = \frac{x}{-\ln(1-x)}$ is non-increasing, and statement *(i)*, we obtain:

$$\min_{t \in L: \sigma_t > 0} \lambda_t = \min_{t \in L: \sigma_t > 0} \frac{\sigma_t}{-\ln(1-\sigma_t)} \geq \frac{\omega}{-\ln(1-\omega)} = \min_{t \in H} \lambda_t.$$

Therefore:

$$\min_{t \in T: \sigma_t > 0} \lambda_t = \min\left\{\min_{t \in H} \lambda_t, \min_{t \in L: \sigma_t > 0} \lambda_t\right\} = \min_{t \in H} \lambda_t = \frac{\omega}{-\ln(1-\omega)}.$$

(iii) First, consider the case that $\sigma_{\min_+} \geq \omega$. Then $\{t \in L \mid \sigma_t > 0\} = \emptyset$. By statement *(i)* and since $\min_{t \in H} \sigma_t = \sigma_{\min_+}$, it holds that:

$$\max_{t \in T: \sigma_t > 0} \frac{1-\sigma_t}{\lambda_t} = \max_{t \in H} \frac{1-\sigma_t}{\lambda_t} = \frac{-\ln(1-\omega)(1-\sigma_{\min_+})}{\omega}.$$

Next, assume that $\sigma_{\min_+} < \omega$. Then $\{t \in L \mid \sigma_t > 0\} \neq \emptyset$. We have that:

$$\begin{aligned}
\max_{t \in H} \frac{1-\sigma_t}{\lambda_t} &= \frac{-\ln(1-\omega)\max_{t \in H}(1-\sigma_t)}{\omega} = \frac{-\ln(1-\omega)\,(1-\min_{t \in H} \sigma_t)}{\omega} \\
&\leq \frac{-\ln(1-\omega)\,(1-\omega)}{\omega} \leq \max_{\substack{t \in L: \\ \sigma_t > 0}} \frac{-\ln(1-\sigma_t)\,(1-\sigma_t)}{\sigma_t} \\
&= \max_{\substack{t \in L: \\ \sigma_t > 0}} \frac{1-\sigma_t}{\lambda_t}. \qquad \text{(A.1)}
\end{aligned}$$

The first inequality holds because $\sigma_t \geq \omega$ for all $t \in H$. The second inequality holds because $\sigma_t < \omega$ for all $t \in L$ with $\sigma_t > 0$ (note that there is at least one such type), and due to the fact that the function $h(x) = \frac{-\ln(1-x)(1-x)}{x}$ is non-increasing in $(0, 1)$. Using (A.1), we conclude that:

$$\begin{aligned}
\max_{\substack{t \in T: \\ \sigma_t > 0}} \frac{1-\sigma_t}{\lambda_t} &= \max\left\{\max_{t \in H} \frac{1-\sigma_t}{\lambda_t}, \max_{\substack{t \in L: \\ \sigma_t > 0}} \frac{1-\sigma_t}{\lambda_t}\right\} = \max_{\substack{t \in L: \\ \sigma_t > 0}} \frac{1-\sigma_t}{\lambda_t} \\
&= \max_{\substack{t \in L: \\ \sigma_t > 0}} \frac{-\ln(1-\sigma_t)(1-\sigma_t)}{\sigma_t} = \frac{-\ln(1-\sigma_{\min_+})(1-\sigma_{\min_+})}{\sigma_{\min_+}}.
\end{aligned}$$

The last equality holds because the function $h(x)$, as defined above, is non-increasing in $(0, 1)$, and thus the maximum is attained for $\min_{t \in L: \sigma_t > 0} \sigma_t = \sigma_{\min_+}$.

(iv) If $L = \emptyset$, the statement follows by statement *(i)*. Otherwise, we have that:

$$\max_{\substack{t \in T: \\ \sigma_t > 0}} \frac{\mu_t}{\lambda_t} = \max\left\{\max_{t \in H} \frac{\mu_t}{\lambda_t}, \max_{\substack{t \in L: \\ \sigma_t > 0}} \frac{\mu_t}{\lambda_t}\right\} = \max\left\{\max_{t \in H} \frac{\sigma_t}{\omega}, 1\right\} = \frac{\sigma_{\max}}{\omega},$$

where, the last equality follows from the definition of H and because $\omega \le \sigma_{\max}$ by assumption. □

A.2.2 Proof of Lemma 6.4.5

Proof of Lemma 6.4.5: Let $\mu_t^*(\omega)$ and $\lambda_t^*(\omega)$ be as previously defined in (6.21) and Lemma 6.4.4, respectively. For brevity, let $\lambda_t = \lambda_t^*(\omega)$ and $\mu_t = \mu_t^*(\omega)$ for each $t \in T$. As argued in Corollary 6.4.3, $(\mu_t)_{t \in T}$ is a feasible solution to POA-RMP. We bound the two terms of the min expression in the objective function of POA-RMP separately.

First, we show that:

$$\min_{t \in T} \lambda_t = \frac{\omega}{-\ln(1-\omega)}. \tag{A.2}$$

When $\sigma_{\min} > 0$, this immediately follows from property *(ii)* of Lemma 6.4.4. Otherwise, when $\sigma_{\min} = 0$, we have that:

$$\min_{t \in T} \lambda_t = \min\left\{\min_{t \in T: \sigma_t > 0} \lambda_t, \min_{t \in T: \sigma_t = 0} \lambda_t\right\} = \min\left\{\frac{\omega}{-\ln(1-\omega)}, 1\right\} = \frac{\omega}{-\ln(1-\omega)}.$$

The last equality holds since $\frac{z}{-\ln(1-z)} < 1$, for all $z \in (0,1)$.

Next, we upper bound the second term of the min expression. We prove that:

$$\max_{t \in T} \frac{\mu_t}{\lambda_t} + \max_{t \in T} \frac{1-\sigma_t}{\lambda_t} \le \frac{\sigma_{\max}}{\omega} + 1. \tag{A.3}$$

When $\sigma_{\min} > 0$, we have:

$$\begin{aligned}
\max_{t \in T} \frac{\mu_t}{\lambda_t} + \max_{t \in T} \frac{1-\sigma_t}{\lambda_t} &= \max_{t \in T: \sigma_t > 0} \frac{\mu_t}{\lambda_t} + \max_{t \in T: \sigma_t > 0} \frac{1-\sigma_t}{\lambda_t} = \frac{\sigma_{\max}}{\omega} + \max_{t \in T: \sigma_t > 0} \frac{1-\sigma_t}{\lambda_t} \\
&= \frac{\sigma_{\max}}{\omega} + \max\left\{\frac{-\ln(1-\omega)\,(1-\omega)}{\omega}, \max_{\substack{t \in L: \\ \sigma_t > 0}} \frac{-\ln(1-\sigma_t)\,(1-\sigma_t)}{\sigma_t}\right\} \\
&\le \frac{\sigma_{\max}}{\omega} + 1,
\end{aligned}$$

where the second and third equality follow from properties *(iv)* and *(i)* of Lemma 6.4.4, respectively. Then, the inequality holds by exploiting that $\frac{-\ln(1-z)(1-z)}{z} < 1$

for all $z \in (0,1)$. Similarly, when $\sigma_{\min} = 0$, we obtain:

$$\begin{aligned}\max_{t\in T}\frac{\mu_t}{\lambda_t} + \max_{t\in T}\frac{1-\sigma_t}{\lambda_t} &= \max\left\{\max_{\substack{t\in T:\\ \sigma_t>0}}\frac{\mu_t}{\lambda_t}, \max_{\substack{t\in T:\\ \sigma_t=0}}\frac{\mu_t}{\lambda_t}\right\} \\ &\quad + \max\left\{\max_{\substack{t\in T:\\ \sigma_t>0}}\frac{1-\sigma_t}{\lambda_t}, \max_{\substack{t\in T:\\ \sigma_t=0}}\frac{1-\sigma_t}{\lambda_t}\right\} \\ &= \max\left\{\frac{\sigma_{\max}}{\omega}, 1\right\} + \max\left\{\max_{\substack{t\in T:\\ \sigma_t>0}}\frac{1-\sigma_t}{\lambda_t}, 1\right\} = \frac{\sigma_{\max}}{\omega} + 1,\end{aligned}$$

where the second equality holds by properties *(i)* and *(iv)* of Lemma 6.4.4, i.e., for each $t \in T$ with $\sigma_t = 0$, $\mu_t = \lambda_t = 1$. The last equality holds because $\omega \le \sigma_{\max}$ and by property *(iii)* of Lemma 6.4.4, as $\frac{-\ln(1-z)(1-z)}{z} < 1$ for all $z \in (0,1)$.

Finally, combining (A.2) and (A.3), we obtain that the optimal value of POA-RMP is at least:

$$\min\left\{\min_{t\in T}\lambda_t, \left(\max_{t\in T}\frac{\mu_t}{\lambda_t} + \max_{t\in T}\frac{1-\sigma_t}{\lambda_t}\right)^{-1}\right\} \ge \min\left\{\frac{\omega}{-\ln(1-\omega)}, \frac{\omega}{\omega+\sigma_{\max}}\right\}.$$

□

A.2.3 Proof of Claims 6.4.7 and 6.4.8

Proof of Claim 6.4.7: Let $z_0 = 1 + W_0(-2e^{-2})/2$. Consider the function $h(z) = z - 1 - W_0\left(-e^{-z-1}\right)$ in $[z_0, 1]$. By Definition 6.2.3, we have:

$$e^{W_0(-2e^{-2})}W_0(-2e^{-2}) = -2e^{-2} \quad \Leftrightarrow \quad -e^2 = e^{W_0(-2e^{-2})}\frac{W_0(-2e^{-2})}{2}.$$

Using this, we obtain:

$$\begin{aligned}W_0\left(-e^{-\frac{W_0(-2e^2)}{2}-2}\right) &= W_0\left(-e^{-\frac{W_0(-2e^2)}{2}}e^{-2}\right) \\ &= W_0\left(e^{\frac{W_0(-2e^2)}{2}}\frac{W_0(-2e^2)}{2}\right).\end{aligned} \tag{A.4}$$

We now compute $h(z_0)$:

$$\begin{aligned}h(z_0) = z_0 - 1 - W_0(-e^{-z_0-1}) &= \frac{W_0\left(-2e^{-2}\right)}{2} - W_0\left(-e^{-\frac{W_0(-2e^2)}{2}-2}\right) \\ = \frac{W_0\left(-2e^{-2}\right)}{2} - W_0\left(e^{\frac{W_0(-2e^2)}{2}}\frac{W_0(-2e^2)}{2}\right) &= \frac{W_0\left(-2e^{-2}\right)}{2} - \frac{W_0\left(-2e^{-2}\right)}{2} = 0.\end{aligned}$$

A

Here, the third equality follows from (A.4) and the fourth equality follows by Definition 6.2.3. Having established that $h(z_0) = 0$, we now show that the derivative $h'(z) > 0$ for all $z > z_0$. By Fact 6.2.4, we have:

$$h'(z) = 1 - \left(W_0\left(-e^{-z-1}\right)\right)' = 1 + \frac{W_0\left(-e^{-z-1}\right)}{e^{-z-1}\left(1 + W_0\left(-e^{-z-1}\right)\right)} \cdot \left(-e^{-z-1}\right)',$$

and simplifying leads to:

$$h'(z) = \frac{1 + 2W_0\left(-e^{-z-1}\right)}{1 + W_0\left(-e^{-z-1}\right)}.$$

The strict inequality $h'(z) > 0$ holds because W_0 is the principal branch of the Lambert function, and since W_0 is non-decreasing on $\left[-\frac{1}{e}, \infty\right)$, we have:

$$1 + 2W_0(-e^{-z-1}) > 1 + 2W_0(-e^{-z_0-1}) > 0.59,$$

where the last inequality is verified numerically. The argument for the denominator is identical i.e., we numerically verify that $1 + W_0(-e^{-z_0-1}) > 0.79$.

Since $h(z)$ is strictly increasing in $(z_0, 1]$ and $h(z_0) = 0$, it holds that $h(z) > 0$ for all $z > z_0$ or, equivalently, that $f(z) < z$. □

Proof of Claim 6.4.8: Since $1 - f(z) = -W_0(-e^{-z-1}) > 0$ (as $-e^{-z-1} > -\frac{1}{e}$ holds for $z > 0$), $f(z) + z = -\ln(1 - f(z))$ is equivalent to:

$$e^{f(z)+z} = \frac{1}{1 - f(z)} \quad \Longleftrightarrow \quad e^{f(z)+z}\left(1 - f(z)\right) = 1.$$

Substituting $f(z)$ in this expression leads to:

$$\begin{aligned} e^{f(z)+z}\left(1 - f(z)\right) &= -e^{z+1}\left(W_0\left(-e^{-z-1}\right)e^{W_0\left(-e^{-z-1}\right)}\right) \\ &= -e^{z+1}\left(-e^{-z-1}\right) = 1, \end{aligned}$$

where the second equality follows by Definition 6.2.3. □

A

A.2.4 Proof of Lemma 6.4.13

Proof of Lemma 6.4.13: By (6.18) and (6.20), it holds that:

$$\bar{\lambda}_t(\eta) = \begin{cases} \frac{1}{1-\eta}, & \text{if } \sigma_t = 0, \\ \frac{1-(1-\eta)e^{-\zeta}}{\zeta}, & \text{if } \sigma_t = 1. \end{cases}$$

For brevity, let $\bar{\lambda}_t = \bar{\lambda}_t(\eta)$ and $\bar{\mu}_t = \bar{\mu}_t(\eta)$ for each $t \in T$.

(i) For type t with $\sigma_t = 1$, we have:

$$\begin{aligned}
\bar{\lambda}_t &= \frac{1-(1-\eta)e^{-\zeta}}{\zeta} = \frac{1}{1-\eta} \cdot \frac{1-\eta-(1-\eta)^2 e^{-\zeta}}{\zeta} \\
&= \frac{1}{1-\eta} \cdot \frac{1-\eta-(1-\eta)^2 e^{-W_0\left(-\frac{(1-\eta)^2}{e^{2-\eta}}\right)} e^{-2+\eta}}{\zeta} \\
&= \frac{1}{1-\eta} \cdot \frac{1-\eta+W_0\left(-\frac{(1-\eta)^2}{e^{2-\eta}}\right)}{\zeta} \\
&< \frac{1}{1-\eta} \cdot \frac{2-\eta+W_0\left(-\frac{(1-\eta)^2}{e^{2-\eta}}\right)}{\zeta} = \frac{1}{1-\eta}\frac{\zeta}{\zeta} = \frac{1}{1-\eta},
\end{aligned}$$

where the fourth equality follows by Definition 6.2.3. On the other hand, for t with $\sigma_t = 0$, we have that $\bar{\lambda}_t = \frac{1}{1-\eta}$. As these are the only two types by assumption, this leads to $\min_{t \in T} \bar{\lambda}_t \geq \frac{1-(1-\eta)e^{-\zeta}}{\zeta}$.

(ii) For type t with $\sigma_t = 1$, by our choice of smoothness parameters we have that:

$$\frac{\bar{\mu}_t}{\bar{\lambda}_t} = \frac{1}{1-(1-\eta)e^{-\zeta}} > 1,$$

where the inequality follows by definition of ζ, as $\eta \in [0,1)$ and by Claim 6.4.11. For t with $\sigma_t = 0$, we have that $\frac{\bar{\mu}_t}{\bar{\lambda}_t} = 1$, and the property follows.

(iii) For type t with $\sigma_t = 1$, we have that $\frac{1-\sigma_t}{\bar{\lambda}_t} = 0 < 1-\eta$ as $\eta \in [0,1)$. On the other hand, for type t with $\sigma_t = 0$, we have that $\frac{1-\sigma_t}{\bar{\lambda}_t} = 1-\eta$, and the property follows.

(iv) By the definition of ζ and Definition 6.2.3 we obtain that:

$$\begin{aligned}
2-\eta-(1-\eta)^2 e^{-\zeta} &= 2-\eta-(1-\eta)^2 e^{-W_0\left(-\frac{(1-\eta)^2}{e^{2-\eta}}\right)} e^{-2+\eta} \\
&= 2-\eta+W_0\left(-\frac{(1-\eta)^2}{e^{2-\eta}}\right) = \zeta.
\end{aligned}$$

A

□

A.3 Missing Material of Section 6.6

The following lemma will be useful in the proof of Lemmas 6.6.4 and 6.6.6.

Lemma A.3.1. *Let $\boldsymbol{B}$ be a co-undominated CCE of a FPA(r) with feasible reserve price among two agents with arbitrary types and action spaces I_1 and I_2.*

Denote its support (A_1, A_2) by $A_1 = \{i \in I_1 \mid \exists j \in I_2$ with $B_{ij} > 0\}$ and $A_2 = \{j \in I_2 \mid \exists i \in I_1$ with $B_{ij} > 0\}$. Then, if $v_1 > v_2$ there is no strategy $i \in A_1$ with $i < r$ and if $v_1 = v_2$ there is an agent j with no strategy $i \in A_j$ with $i < r$.

Proof: Let ℓ_2 be the lowest bid in A_2 and note that $\ell_2 \leq v_2$ as agent 2 does not overbid. Consider the following two cases.

Case 1: $v_1 > v_2$. We show that any action $b_1 \in A_1$ of agent 1 with $b_1 < r$ is weakly dominated in A_2 by action $b_1' = \max\{r, \ell_2\} < v_1$. Note that $v_1 - \sigma_1 b_1' > (1 - \sigma_1)v_1 \geq 0$, as $\sigma_1 \leq 1$. For all $b_2 \in A_2$ we have:

$$0 = g_1(b_1, b_2) \leq g_1(b_1', b_2) = \begin{cases} v_1 - \sigma_1 b_1', & \text{if } b_1' > b_2, \\ \frac{1}{2}(v_1 - \sigma_1 b_1'), & \text{if } b_1' = b_2, \\ 0, & \text{otherwise.} \end{cases}$$

For $\ell_2 \in A_2$ we have:

$$0 = g_1(b_1, \ell_2) < g_1(b_1', \ell_2) = \begin{cases} v_1 - \sigma_1 b_1', & \text{if } b_1' > \ell_2, \\ \frac{1}{2}(v_1 - \sigma_1 b_1'), & \text{if } b_1' = \ell_2. \end{cases}$$

Therefore, any action $b_1 < r$ of agent 1 is weakly dominated in A_2 by action $b_1' = \max\{r, \ell_2\} \geq r$.

Case 2: $v_1 = v_2$. Let ℓ_1 be the lowest bid in A_1 and, for contradiction, assume that both $\ell_1 < r$ and $\ell_2 < r$. We show that action ℓ_1 of agent 1 is weakly dominated in A_2 by action r. Recall that $v_1 - \sigma_1 r > 0$. For all $b_2 \in A_2$ we have:

$$0 = g_1(\ell_1, b_2) \leq g_1(r, b_2) = \begin{cases} v_1 - \sigma_1 r, & \text{if } r > b_2, \\ \frac{1}{2}(v_1 - \sigma_1 r), & \text{if } r = b_2, \\ 0, & \text{otherwise.} \end{cases}$$

For $\ell_2 \in A_2$ we have:

$$0 = g_1(\ell_1, \ell_2) < g_1(r, \ell_2) = v_1 - \sigma_1 r,$$

concluding the proof. □

A

Proof of 6.6.4: We adapt the proof in [Kolumbus and Nisan, 2022] to fit our setting. Consider agents 1 and 2 with $v_1 > v_2 > 2\varepsilon$. We will prove that bids that are not in $[\max\{r, v_2\}, \max\{r, v_2 + \varepsilon\}]$ cannot be in the support of agent 1 in a co-undominated CCE.

As agent 2 does not overbid, for any support A_2 of agent 2 in a co-undominated CCE, any bid $b_1' > b_1 = \max\{r, v_2 + \varepsilon\}$ of agent 1 is weakly dominated by b_1, as agent 1 will always win at a lower price, i.e., for any b_2 in the support of A_2 it holds that $g_1(b_1', b_2) < g_1(b_1, b_2)$. So in a co-undominated CCE, agent 1 does not bid above $\max\{r, v_2 + \varepsilon\}$.

Let ℓ_1 and ℓ_2 be the lowest bids in the support of agent 1 and 2 in a co-undominated CCE. For contradiction, assume that $\ell_1 < v_2$. Note that $v_2 = \max\{r, v_2\}$ as $r \leq \ell_1$ by Lemma A.3.1. Consider the following cases.

Case 1: $\ell_1 = \ell_2$. In this case ℓ_1 is weakly dominated by $\ell_1 + \varepsilon$, as agent 1 loses with ℓ_1 unless agent 2 also bids ℓ_2 and $g_1(\ell_1, \ell_2) = \frac{1}{2}(v_1 - \sigma_1 \ell_1) < v_1 - \sigma_1(\ell_1 + \epsilon) = g_1(\ell_1 + \epsilon, \ell_2)$, as $\ell_1 < v_2 \leq v_1 - 2\varepsilon \leq \frac{1}{\sigma_1} v_1 - 2\varepsilon$.

Case 2: $\ell_1 < \ell_2$. In this case ℓ_1 is weakly dominated by v_2, as agent 1 always loses with ℓ_1 and $g_1(v_2, \ell_2) \geq \frac{1}{2}(v_1 - \sigma_1 v_2) > \frac{1}{2}(v_1 - \sigma_1 v_1) \geq 0$, as agent 2 does not overbid ($\ell_2 \leq v_2$).

Case 3: $\ell_1 > \ell_2$. In this case ℓ_2 is weakly dominated by ℓ_1, as agent 2 always loses with ℓ_2 and $g_2(\ell_1, \ell_1) \geq \frac{1}{2}(v_2 - \sigma_2(v_2 - \varepsilon)) > 0$.

We conclude that if $r < v_2$, the support of agent 1 in a co-undominated CCE is in $[\max\{r, v_2\}, \max\{r, v_2 + \varepsilon\}]$. For $r \geq v_2$, it follows from the fact that agent 1 does not underbid the reserve price in a co-undominated CCE (Lemma A.3.1) or overbid $\max\{r, v_2 + \varepsilon\}$, that the support of agent 1 in a co-undominated CCE is in $[\max\{r, v_2\}, \max\{r, v_2 + \varepsilon\}]$. □

Proof of 6.6.5: Let ℓ_1 and ℓ_2 be the lowest bids and let h_1 and h_2 be the highest bids in the support of agents 1 and 2 in a co-undominated CCE, respectively. As we assume that agents do not overbid, h_1 and h_2 are smaller than or equal to v. Consider the following cases.

Case 1: $\sigma_1 < 1$ and $\sigma_2 < 1$. Let $\ell_1 < \ell_2 \leq v$ (or symmetrically $\ell_2 < \ell_1 \leq v$). Note that in this case $\ell_2 \geq r$, as otherwise r weakly dominates $\ell_2 < r$ as agent 2 would always lose with ℓ_2 and sometimes win with $r < v$. In this case, ℓ_1 is weakly dominated by ℓ_2, as agent 1 always loses with ℓ_1 and $g_1(\ell_2, \ell_2) = \frac{1}{2}(v - \sigma_1 \ell_2) > 0$, as $r \leq \ell_2 \leq v$ and $\sigma_1 < 1$. Furthermore, for $\ell_1 = \ell_2 \leq v - 2\varepsilon$, ℓ_1 is weakly dominated by $\max\{r, \ell_1 + \varepsilon\}$, as $g_1(\ell_1, \ell_2) \leq \frac{1}{2}(v - \sigma_1 \ell_1) < v - \sigma_1(\max\{r, \ell_1 + \varepsilon\}) = g_1(\max\{r, \ell_1 + \varepsilon\}, \ell_2)$. Therefore, in this case, the support of each agent in a co-undominated CCE is in $[v - \varepsilon, v] = [\max\{r, v - \varepsilon\}, v]$.

Case 2: $\sigma_1 = 1$ and $\sigma_2 = 1$. Let $\ell_1 < \ell_2 < v$ (or symmetrically $\ell_2 < \ell_1 < v$). Again note that $\ell_2 \geq r$. In this case, ℓ_1 is weakly dominated by ℓ_2, as agent 1 always loses with ℓ_1 and $g_1(\ell_2, \ell_2) = \frac{1}{2}(v - \ell_2) > 0$, as $\ell_2 < v$. Furthermore, for $\ell_1 = \ell_2 < v - 2\varepsilon$, ℓ_1 is weakly dominated by $\max\{r, \ell_1 + \varepsilon\}$, as $g_1(\ell_1, \ell_2) \leq \frac{1}{2}(v - \ell_1) < v - (\max\{r, \ell_1 + \varepsilon\}) = g_1(\max\{r, \ell_1 + \varepsilon\}, \ell_2)$ if $r \leq v - 2\varepsilon$ and $g_1(\ell_1, \ell_2) = 0 < v - (v - \varepsilon) = g_1(\max\{r, \ell_1 + \varepsilon\}, \ell_2)$ if $r = v - \varepsilon$. Additionally, if $\ell_1 = h_1 = v$, then agent 1 always has a utility of 0. For the co-undominated CCE condition to hold, it must be that also $\ell_2 = h_2 = v$. Therefore, in this case, the support of each agent in a co-undominated CCE is in $[\max\{r, v - 2\varepsilon\}, \max\{r, v - \varepsilon\}, v]$.

Case 3: $\sigma_1 = 1$ and $\sigma_2 < 1$ (or $\sigma_1 < 1$ and $\sigma_2 = 1$). The previous reasoning holds for $\ell_1 < \ell_2 < v$ for agent 1, for $\ell_2 < \ell_1 \leq v$ for agent 2 and for $\ell_2 = \ell_1 \leq v - 2\varepsilon$

A

for agent 2. Therefore, in this case, the support of agent 2 in a co-undominated CCE is in $[v - \varepsilon, v] = [\max\{r, v - \varepsilon\}, v]$. □

Proof of 6.6.6: First, consider the case that $v_1 > v_2 > 0$. For contradiction, suppose that there exists a co-undominated CCE for which the lowest bid ℓ_1 in the support of agent 1 is smaller than $\max\{r, v_2\}$. As agent 1 does not bid below the reserve price r by Lemma A.3.1, it must be that $r \leq \ell_1 < v_2$. In this case, any bid $b_2 < \ell_1$ of agent 2 is weakly dominated by ℓ_1, as agent 2 always loses with b_2 and $g_2(\ell_1, \ell_1) = \frac{1}{2}(v_2 - \sigma_2 \ell_1) > 0$. Now let ℓ_2 be the lowest bid in the support of agent 2 in the co-undominated CCE. It must be that $\ell_2 \geq \ell_1$. But in this case, ℓ_1 is weakly dominated by $\ell_2 + \varepsilon > \ell_1$, as agent 1 always loses when agent 2 bids above ℓ_2 and agent 1 is strictly better off when agent 2 bids ℓ_2: $g_1(\ell_1, \ell_2) \leq \frac{1}{2} v_1 < v_1 = g_1(\ell_2 + \varepsilon, \ell_2)$, leading to a contradiction. Note that this holds for any $r \leq \ell_1 < v_2$. Therefore, and as agent 1 does not overbid, the support of agent 1 in a co-undominated CCE is in $[\max\{r, v_2\}, \max\{r, v_2 + \varepsilon\}, \ldots, \max\{r, v_1\}]$. (Note that if the highest bid h_2 in the support of agent 2 in the co-undominated CCE is $h_2 = v_2$, the lowest bid ℓ_1 in the support of agent 1 in the co-undominated CCE will be $\ell_1 > v_2$.)

Secondly, consider the case that $v_1 = v_2 > 0$. For contradiction, suppose that there exists a co-undominated CCE for which the lowest bid ℓ_1 in the support of agent 1 is smaller than v_1. In this case, any bid $b_2 < \ell_1$ of agent 2 is weakly dominated by $\max\{r, \ell_1\}$, as agent 2 always loses with b_2 and $g_2(\ell_1, \max\{r, \ell_1\}) \geq \frac{1}{2}(v_2 - \sigma_2 \max\{r, \ell_1\}) > 0$. Now let $\ell_2 < v_2$ be the lowest bid in the support of agent 2 in the co-undominated CCE. It must be that $\ell_2 \geq \max\{r, \ell_1\}$. But in that case, ℓ_1 is weakly dominated by $\ell_2 + \varepsilon > \ell_1$, as agent 1 always loses when agent 2 bids above ℓ_2 and agent 1 is strictly better off when agent 2 bids ℓ_2 as $g_1(\ell_1, \ell_2) \leq \frac{1}{2} v_1 < v_1 = g_1(\ell_2 + \varepsilon, \ell_2)$, leading to a contradiction. For $\ell_2 = v_2$, note that any bid $b_1 < \ell_2$ of agent 1 is weakly dominated by $v_1 = \ell_2$, as $0 = g_1(b_1, \ell_2) < \frac{1}{2} v_1 = g_1(v_1, \ell_2)$. Therefore, and as agent 1 does not overbid and $r < v_1 = v_2$, the support of agent 1 in a co-undominated CCE is in $[v_1]$. □

A

Bibliography

J. Aerts, J. Korst, F. Spieksma, W. Verhaegh, and G. Woeginger. Random redundant storage in disk arrays: complexity of retrieval problems. *IEEE Transactions on Computers*, 52(9):1210–1214, 2003.

G. Aggarwal, A. B. Varadaraja, and A. Mehta. Autobidding with constraints. In *Proceedings of the International Conference on Web and Internet Economics*, pages 17–30, 2019.

G. Aggarwal, A. Badanidiyuru, S. R. Balseiro, K. Bhawalkar, Y. Deng, Z. Feng, G. Goel, C. Liaw, H. Lu, M. Mahdian, et al. Auto-bidding and auctions in online advertising: A survey. *ACM SIGecom Exchanges*, 22(1):159–183, 2024.

G. Aggarwal, G. Fikioris, and M. Zhao. No-regret algorithms in non-truthful auctions with budget and roi constraints. In *Proceedings of the ACM on Web Conference*, pages 1398–1415, 2025a.

G. Aggarwal, A. Gupta, X. Tan, and M. Zhao. Multi-platform autobidding with and without predictions. In *Proceedings of the ACM on Web Conference*, pages 2850–2859, 2025b.

A. Agrawal, T. Inamdar, S. Saurabh, and J. Xue. Clustering what matters: Optimal approximation for clustering with outliers. *Journal of Artificial Intelligence Research*, 78:143–166, 2023.

P. Agrawal, E. Balkanski, V. Gkatzelis, T. Ou, and X. Tan. Learning-augmented mechanism design: Leveraging predictions for facility location. *Mathematics of Operations Research*, 49(4):2626–2651, 2024.

M. Almanza, A. Epasto, A. Panconesi, and G. Re. k-clustering with fair outliers. In *Proceedings of the ACM International Conference on Web Search and Data Mining*, page 5–15, 2022.

G. Amanatidis, G. Birmpas, and E. Markakis. Coverage, matching, and beyond: new results on budgeted mechanism design. In *Proceedings of the International Conference on Web and Internet Economics*, pages 414–428, 2016.

G. Amanatidis, G. Birmpas, and E. Markakis. On budget-feasible mechanism design for symmetric submodular objectives. In *Proceedings of the International Conference on Web and Internet Economics*, pages 1–15, 2017.

G. Amanatidis, P. Kleer, and G. Schäfer. Budget-feasible mechanism design for non-monotone submodular objectives: Offline and online. In *Proceedings of the ACM Conference on Economics and Computation*, pages 901–919, 2019.

G. Amanatidis, S. Klumper, E. Markakis, G. Schäfer, and A. Tsikiridis. Partial allocations in budget-feasible mechanism design: bridging multiple levels of service and divisible agents. *ACM Transactions on Economics and Computation*, 13(2):1–28, 2025.

N. Anari, G. Goel, and A. Nikzad. Budget feasible procurement auctions. *Operations Research*, 66(3):637–652, 2018.

A. Antoniadis, C. Coester, M. Eliáš, A. Polak, and B. Simon. Online metric algorithms with untrusted predictions. *ACM Transactions on Algorithms*, 19 (2):1–34, 2023a.

A. Antoniadis, T. Gouleakis, P. Kleer, and P. Kolev. Secretary and online matching problems with machine learned advice. *Discrete Optimization*, 48:100778, 2023b.

K. R. Apt and J. Heering. Characterization of incentive compatible single-parameter mechanisms revisited. *The Journal of Mechanism and Institution Design*, 7(1):113–129, 2022.

A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*, pages 482–491, 2001.

G. Auricchio, H. J. Clough, and J. Zhang. On the capacitated facility location problem with scarce resources. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 2024a.

G. Auricchio, Z. Wang, and J. Zhang. Facility location problems with capacity constraints: Two facilities and beyond. *arXiv preprint arXiv:2404.13566*, 2024b.

Y. Azar, S. Leonardi, and N. Touitou. Flow time scheduling with uncertain processing time. In *Proceedings of the ACM SIGACT Symposium on Theory of Computing*, pages 1070–1080, 2021.

Y. Azar, D. Panigrahi, and N. Touitou. Online graph algorithms with predictions. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 35–66, 2022.

H. Aziz, H. Chan, B. Lee, B. Li, and T. Walsh. Facility location problem with capacity constraints: Algorithmic and mechanism design perspectives. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1806–1813, 2020a.

H. Aziz, H. Chan, B. E. Lee, and D. C. Parkes. The capacity constrained facility location problem. *Games and Economic Behavior*, 124:478–490, 2020b.

E. Balkanski, P. Garimidi, V. Gkatzelis, D. Schoepflin, and X. Tan. Deterministic budget-feasible clock auctions. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 2940–2963, 2022.

E. Balkanski, V. Gkatzelis, and G. Shahkarami. Randomized strategic facility location with predictions. *Advances in Neural Information Processing Systems*, 37:35639–35664, 2024a.

E. Balkanski, V. Gkatzelis, X. Tan, and C. Zhu. Online mechanism design with predictions. In *Proceedings of the ACM Conference on Economics and Computation*, pages 1184–1184, 2024b.

S. Balseiro, Y. Deng, J. Mao, V. Mirrokni, and S. Zuo. Robust auction design in the auto-bidding world. *Advances in Neural Information Processing Systems*, 34:17777–17788, 2021a.

S. Balseiro, Y. Deng, J. Mao, V. Mirrokni, and S. Zuo. Optimal mechanisms for value maximizers with budget constraints via target clipping. *Columbia Business School Research Paper*, (4081457), 2022.

S. R. Balseiro and Y. Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 65(9):3952–3968, 2019.

S. R. Balseiro, Y. Deng, J. Mao, V. S. Mirrokni, and S. Zuo. The landscape of auto-bidding auctions: Value versus utility maximization. In *Proceedings of the ACM Conference on Economics and Computation*, pages 132–133, 2021b.

S. Banerjee, V. Gkatzelis, A. Gorokh, and B. Jin. Online nash social welfare maximization with predictions. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 1–19, 2022.

Z. Barak, A. Gupta, and I. Talgam-Cohen. Mac advice for facility location mechanism design. *Advances in Neural Information Processing Systems*, 37:129564–129604, 2024.

X. Bei, N. Chen, N. Gravin, and P. Lu. Worst-case mechanism design via bayesian analysis. *SIAM Journal on Computing*, 46(4):1428–1448, 2017.

B. Berger, M. Feldman, V. Gkatzelis, and X. Tan. Learning-augmented metric distortion via (p,q)-veto core. In *Proceedings of the ACM Conference on Economics and Computation*, page 984, 2024.

K. C. Border and J. S. Jordan. Straightforward elections, unanimity and phantom voters. *The Review of Economic Studies*, 50(1):153–170, 1983.

M. Braverman, J. Mao, J. Schneider, and M. Weinberg. Selling to a no-regret buyer. In *Proceedings of the ACM Conference on Economics and Computation*, pages 523–538, 2018.

I. Caragiannis and G. Kalantzis. Randomized learning-augmented auctions with revenue guarantees. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 2687–2694, 2024.

D. Chakrabarty and G. Goel. On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and gap. *SIAM Journal on Computing*, 39(6):2189–2211, 2010.

H. Chan and J. Chen. Truthful multi-unit procurements with budgets. In *Proceedings of the International Conference on Web and Internet Economics*, pages 89–105, 2014.

H. Chan, A. Filos-Ratsikas, B. Li, M. Li, and C. Wang. Mechanism design for facility location problems: A survey. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 4356–4365, 2021.

H. Chan, Z. Gong, M. Li, C. Wang, and Y. Zhao. Facility location games with ordinal preferences. *Theoretical Computer Science*, 979:114208, 2023.

M. Charikar, S. Khuller, D. M. Mount, and G. Narasimhan. Algorithms for facility location problems with outliers. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, page 642–651, 2001.

K. Chen. A constant factor approximation algorithm for k-median clustering with outliers. In *Proceedings of the ACM-SIAM symposium on Discrete algorithms*, pages 826–835, 2008.

N. Chen, N. Gravin, and P. Lu. On the approximability of budget feasible mechanisms. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 685–699, 2011.

N. Chen, N. Gravin, and P. Lu. Truthful generalized assignments via stable matching. *Mathematics of Operations Research*, 39(3):722–736, 2014.

Q. Chen, N. Gravin, and S. Im. Strategic facility location via predictions. *arXiv preprint arXiv:2410.07497*, 2024.

Z. Chen, K. C. K. Fong, M. Li, K. Wang, H. Yuan, and Y. Zhang. Facility location games with optional preference. *Theoretical Computer Science*, 847:185–197, 2020.

Y. Cheng, W. Yua, and G. Zhang. Strategy-proof approximation mechanisms for an obnoxious facility game on networks. *Theoretical Computer Science*, 497: 154–163, 2013.

G. Christodoulou, A. Kovács, A. Sgouritsa, and B. Tang. Tight bounds for the price of anarchy of simultaneous first-price auctions. *ACM Transactions on Economics and Computation*, 4(2):1–33, 2016.

G. Christodoulou, A. Sgouritsa, and I. Vlachos. Mechanism design augmented with output advice. *Advances in Neural Information Processing Systems*, 37: 41934–41953, 2024.

I. Cohen, A. Eden, T. Eden, and A. Vasilyan. Plant-and-steal: Truthful fair allocations via predictions. *Advances in Neural Information Processing Systems*, 37:110057–110096, 2024.

R. Colini-Baldeschi, S. Klumper, G. Schäfer, and A. Tsikiridis. To trust or not to trust: Assignment mechanisms with predictions in the private graph model. In *Proceedings of the ACM Conference on Economics and Computation*, pages 1134–1154, 2024.

R. M. Corless, G. H. Gonnet, D. E. Hare, D. J. Jeffrey, and D. E. Knuth. On the lambert w function. *Advances in Computational Mathematics*, 5:329–359, 1996.

R. Dabas and N. Gupta. On variants of facility location problem with outliers. *arXiv preprint arXiv:2107.00403*, 2021.

M. Dawande, J. Kalagnanam, P. Keskinocak, F. S. Salman, and R. Ravi. Approximation algorithms for the multiple knapsack problem with assignment restrictions. *Journal of Combinatorial Optimization*, 4:171–186, 2000.

A. Deligkas, A. Filos-Ratsikas, and A. A. Voudouris. Heterogeneous facility location with limited resources. *Games and Economic Behavior*, 139:200–215, 2023.

Y. Deng, J. Mao, V. Mirrokni, and S. Zuo. Towards efficient auctions in an auto-bidding world. In *Proceedings of the Web Conference*, pages 3965–3973, 2021.

Y. Deng, J. Mao, V. Mirrokni, H. Zhang, and S. Zuo. Efficiency of the first-price auction in the autobidding world. *Advances in Neural Information Processing Systems*, 37:139270–139293, 2024.

M. Dinitz, S. Im, T. Lavastida, B. Moseley, and S. Vassilvitskii. Algorithms with prediction portfolios. *Advances in Neural Information Processing Systems*, 35: 20273–20286, 2022.

S. Dobzinski and S. Dughmi. On the power of randomization in algorithmic mechanism design. *SIAM Journal on Computing*, 42(6):2287–2304, 2013.

S. Dobzinski and R. P. Leme. Efficiency guarantees in auctions with budgets. In *International Colloquium on Automata, Languages, and Programming*, pages 392–404, 2014.

S. Dobzinski, C. H. Papadimitriou, and Y. Singer. Mechanisms for complement-free procurement. In *Proceedings of the ACM conference on Electronic Commerce*, pages 273–282, 2011.

L. E. Dubins and D. A. Freedman. Machiavelli and the gale-shapley algorithm. *The American Mathematical Monthly*, 88(7):485–494, 1981.

S. Dughmi and A. Ghosh. Truthful assignment without money. In *Proceedings of the ACM Conference on Electronic commerce*, pages 325–334, 2010.

S. Fadaei and M. Bichler. A truthful mechanism for the generalized assignment problem. *ACM Transactions on Economics and Computation*, 5(3):1–18, 2017a.

S. Fadaei and M. Bichler. Generalized assignment problem: Truthful mechanism design without money. *Operations Research Letters*, 45(1):72–76, 2017b.

U. Feige and J. Vondrak. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*, pages 667–676, 2006.

I. Feigenbaum and J. Sethuraman. Strategyproof mechanisms for one-dimensional hybrid and obnoxious facility location models. In *AAAI Workshop on Incentive and Trust in E-Communities*, volume 15, page 08, 2015.

M. Feldman, A. Fiat, and I. Golomb. On voting and facility location. In *Proceedings of the ACM Conference on Economics and Computation*, pages 269–286, 2016.

M. Feldman, H. Fu, N. Gravin, and B. Lucier. Simultaneous auctions without complements are (almost) efficient. *Games and Economic Behavior*, 123:327–341, 2020.

Q. Feng, Z. Zhang, Z. Huang, J. Xu, and J. Wang. Improved algorithms for clustering with outliers. In *Proceedings of the International Symposium on Algorithms and Computation*, 2019.

A. Filos-Ratsikas, P. Kanellopoulos, A. A. Voudouris, and R. Zhang. The distortion of distributed facility location. *Artificial Intelligence*, 328:104066, 2024.

C. K. K. Fong, M. Li, P. Lu, T. Todo, and M. Yokoo. Facility location games with fractional preferences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1039–1046, 2018.

D. Fotakis and C. Tzamos. On the power of deterministic mechanisms for facility location games. *ACM Transactions on Economics and Computation*, 2(4):15:1–15:37, 2014.

Z. Friggstad, K. Khodamoradi, M. Rezapour, and M. R. Salavatipour. Approximation schemes for clustering with outliers. *ACM Transactions on Algorithms*, 15(2):1–26, 2019.

L. Gai, M. Liang, and C. Wang. Two-facility-location games with mixed types of agents. *Applied Mathematics and Computation*, 466:128479, 2024.

D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

D. Gale and M. Sotomayor. Some remarks on the stable matching problem. *Discrete Applied Mathematics*, 11(3):223–232, 1985.

V. Gkatzelis, K. Kollias, A. Sgouritsa, and X. Tan. Improved price of anarchy via predictions. In *Proceedings of the ACM Conference on Economics and Computation*, pages 529–557, 2022.

V. Gkatzelis, D. Schoepflin, and X. Tan. Clock auctions augmented with unreliable advice. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 2629–2655, 2025.

G. Goel, A. Nikzad, and A. Singla. Allocating tasks to workers with matching constraints: truthful mechanisms for crowdsourcing markets. In *Proceedings of the International Conference on World Wide Web*, pages 279–280, 2014.

D. Goyal, R. Jaiswal, and A. Kumar. Fpt approximation for constrained metric k-median/means. In *Proceedings of the International Symposium on Parameterized and Exact Computation*, volume 180, pages 14:1–14:19, 2020.

N. Gravin, Y. Jin, P. Lu, and C. Zhang. Optimal budget-feasible mechanisms for additive valuations. *ACM Transactions on Economics and Computation*, 8(4): 1–15, 2020.

S. Gupta, R. Kumar, K. Lu, B. Moseley, and S. Vassilvitskii. Local search methods for k-means with outliers. *Proceedings of the VLDB Endowment*, 10(7):757–768, 2017.

K. Han, Y. Wu, H. Huang, and S. Cui. Triple eagle: Simple, fast and practical budget-feasible mechanisms. *Advances in Neural Information Processing Systems*, 36:33894–33911, 2023.

J. W. Hatfield and P. R. Milgrom. Matching with contracts. *American Economic Review*, 95(4):913–935, 2005.

D. S. Hochbaum. A nonlinear knapsack problem. *Operations Research Letters*, 17(3):103–110, 1995.

H. Huang, K. Han, S. Cui, and J. Tang. Randomized pricing with deferred acceptance for revenue maximization with submodular objectives. In *Proceedings of the ACM Web Conference*, pages 3530–3540, 2023.

G. Istrate and C. Bonchis. Mechanism design with predictions for obnoxious facility location. *arXiv preprint arXiv:2212.09521*, 2022.

P. Jalaly and É. Tardos. Simple and efficient budget feasible mechanisms for monotone submodular valuations. In *Proceedings of the International Conference on Web and Internet Economics*, pages 246–263, 2018.

B. Jin and W. Ma. Online bipartite matching with advice: Tight robustness-consistency tradeoffs for the two-stage model. *Advances in Neural Information Processing Systems*, 35:14555–14567, 2022.

Y. Jin and P. Lu. First price auction is $1 - 1/e^2$ efficient. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*, pages 179–187, 2022.

P. Kanellopoulos, A. A. Voudouris, and R. Zhang. Truthful two-facility location with candidate locations. In *Proceedings of the International Symposium on Algorithmic Game Theory*, 2023a.

P. Kanellopoulos, A. A. Voudouris, and R. Zhang. On discrete truthful heterogeneous two-facility location. *SIAM Journal on Discrete Mathematics*, 37: 779–799, 2023b.

S. Klumper and G. Schäfer. Budget feasible mechanisms for procurement auctions with divisible agents. In *Proceedings of the International Symposium on Algorithmic Game Theory*, pages 78–93, 2022.

Y. Kolumbus and N. Nisan. Auctions between regret-minimizing agents. In *Proceedings of the ACM Web Conference*, pages 100–111, 2022.

E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.

R. Krishnaswamy, S. Li, and S. Sandeep. Constant approximation for k-median and k-means with outliers via iterative rounding. In *Proceedings of the ACM SIGACT Symposium on Theory of Computing*, pages 646–659, 2018.

H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955.

T. Lavastida, B. Moseley, R. Ravi, and C. Xu. Learnable and instance-robust predictions for online matching, flows and load balancing. In *Annual European Symposium on Algorithms*, volume 204, page 59, 2021a.

T. Lavastida, B. Moseley, R. Ravi, and C. Xu. Using predicted weights for ad delivery. In *Proceedings of the SIAM Conference on Applied and Computational Discrete Algorithms*, pages 21–31, 2021b.

S. Leonardi, G. Monaco, P. Sankowski, and Q. Zhang. Budget feasible mechanisms on matroids. In *Proceedings of the International Conference on Integer Programming and Combinatorial Optimization*, pages 368–379, 2017.

M. Li, P. Lu, Y. Yao, and J. Zhang. Strategyproof mechanism for two heterogeneous facilities with constant approximation ratio. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 238–245, 2020.

M. Li, C. Wang, and M. Zhang. Budget feasible mechanisms for facility location games with strategic facilities. *Autonomous Agents and Multi-Agent Systems*, 36(2):35, 2022.

C. Liaw, A. Mehta, and A. Perlroth. Efficiency of non-truthful auctions in autobidding: The power of randomization. In *Proceedings of the ACM Web Conference*, page 3561–3571, 2023.

X. Liu, H. Chan, M. Li, and W. Wu. Budget feasible mechanisms: A survey. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 8132–8141, 2024.

P. Lu, X. Sun, Y. Wang, and Z. A. Zhu. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 315–324, 2010.

P. Lu, Z. Wan, and J. Zhang. Competitive auctions with imperfect predictions. In *Proceedings of the ACM Conference on Economics and Computation*, pages 1155–1183, 2024.

B. Lucier, S. Pattathil, A. Slivkins, and M. Zhang. Autobidders with budget and roi constraints: Efficiency, regret, and pacing dynamics. In *Proceedings of the Conference on Learning Theory*, pages 3642–3643, 2024.

T. Lykouris and S. Vassilvitskii. Competitive caching with machine learned advice. *Journal of the ACM*, 68(4):1–25, 2021.

B. Maity, S. Das, and A. Dasgupta. Linear programming based approximation to individually fair k-clustering with outliers. *arXiv preprint arXiv:2412.10923*, 2024.

S. Martello and P. Toth. *Knapsack problems: algorithms and computer implementations*. John Wiley & Sons, Inc., 1990.

A. Mehta. Auction design in an auto-bidding setting: Randomization improves efficiency beyond vcg. In *Proceedings of the ACM Web Conference*, page 173–181, 2022.

M. Mitzenmacher and S. Vassilvitskii. Algorithms with predictions. *Communications of the ACM*, 65(7):33–35, 2022.

H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4): 437–455, 1980.

R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1): 58–73, 1981.

J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.

R. Neogi, K. Pashkovich, and C. Swamy. Budget-feasible mechanism design: Simpler, better mechanisms and general payment constraints. In *Proceedings of the Innovations in Theoretical Computer Science Conference*, 2024.

N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proceedings of the ACM symposium on Theory of Computing*, pages 129–140, 1999.

N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.

Z. Nutov, I. Beniaminy, and R. Yuster. A (1–1/e)-approximation algorithm for the generalized assignment problem. *Operations Research Letters*, 34(3):283–288, 2006.

R. Paes Leme, B. Sivan, and Y. Teng. Why do competitive markets converge to first-price auctions? In *Proceedings of the Web Conference*, pages 596–605, 2020.

H. Peters, H. van der Stel, and T. Storcken. Range convexity, continuity, and strategy-proofness of voting schemes. *Zeitschrift für Operations Research*, 38 (2):213–229, 1993.

S. Prasad, M.-F. F. Balcan, and T. Sandholm. Bicriteria multidimensional mechanism design with side information. *Advances in Neural Information Processing Systems*, 36:40832–40852, 2023.

A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. *ACM Transactions on Economics and Computation*, 1(4):1–26, 2013.

M. Purohit, Z. Svitkina, and R. Kumar. Improving online algorithms via ml predictions. *Advances in Neural Information Processing Systems*, 31, 2018.

A. E. Roth. The economics of matching: Stability and incentives. *Mathematics of Operations Research*, 7(4):617–628, 1982.

T. Roughgarden. Algorithmic game theory. *Communications of the ACM*, 53(7): 78–86, 2010.

T. Roughgarden. Intrinsic robustness of the price of anarchy. *Journal of the ACM*, 62(5):1–42, 2015.

T. Roughgarden. Beyond worst-case analysis. *Communications of the ACM*, 62 (3):88–96, 2019.

T. Roughgarden. *Beyond the worst-case analysis of algorithms*. Cambridge University Press, 2021.

A. Rubinstein and J. Zhao. Beyond worst-case budget-feasible mechanism design. In *Proceedings of the Innovations in Theoretical Computer Science Conference*, 2023.

J. Schummer and R. V. Vohra. Mechanism design without money. *Algorithmic Game Theory*, 10:243–299, 2007.

P. Serafino and C. Ventre. Heterogeneous facility location without money. *Theoretical Computer Science*, 636:27–46, 2016.

Y. Singer. Budget feasible mechanisms. In *Proceedings of the Symposium on Foundations of Computer Science*, pages 765–774, 2010.

V. Syrgkanis. *Efficiency of mechanisms in complex markets*. Cornell University, 2014.

V. Syrgkanis and E. Tardos. Composable and efficient mechanisms. In *Proceedings of the ACM Symposium on Theory of Computing*, page 211–220, 2013.

W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

C. Xu and P. Lu. Mechanism design with predictions. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 571–577, 2022.

X. Xu, B. Li, M. Li, and L. Duan. Two-facility location games with minimum distance requirement. *Journal of Artificial Intelligence Research*, 70:719–756, 2021.

Z. Zhang, Q. Feng, J. Huang, Y. Guo, J. Xu, and J. Wang. A local search algorithm for k-means with outliers. *Neurocomputing*, 450:230–241, 2021.

Samenvatting

In deze thesis onderzoeken we optimalisatieproblemen in een strategische setting. Met strategisch bedoelen we dat een deel van de input van het probleem alleen bekend is bij individuele strategische spelers. Om het optimalisatieprobleem op te lossen wordt deze ontbrekende informatie bij de spelers opgevraagd. Dit proces noemen we een mechanisme. De spelers maken strategische keuzes over de informatie die ze delen, en zullen niet eerlijk de ontbrekende informatie opgeven als dat in hun voordeel is. Dit soort strategisch gedrag kan onwenselijk zijn vanuit een maatschappelijk oogpunt en kan leiden tot inefficiënte oplossingen. Daarom zijn we op zoek naar mechanismen die niet gemanipuleerd kunnen worden: dit soort mechanismen noemen we *strategyproof*. Bij een strategyproof mechanisme is het voor een speler het beste om eerlijk te zijn, ongeacht wat de andere spelers doen. De eerste twee delen van deze thesis gaan over het ontwerpen van strategyproof mechanismen.

In het eerste deel behandelen we twee problemen waarbij geen geld wordt uitgewisseld. Verder zijn de problemen verrijkt met een voorspelling van de optimale oplossing. Het doel is om een mechanisme te ontwerpen dat een optimale approximatie realiseert wanneer de voorspelling feilloos is (*consistency*), en een degelijke approximatie behoudt wanneer de voorspelling niet accuraat is (*robustness*). Alle mechanismen die we ontwerpen hebben een approximatie die afhangt van de nauwkeurigheid van de voorspelling en geleidelijk interpoleert tussen de twee uitersten van consistency en robustness. Daarnaast beschikken sommige mechanismen over een parameter waarmee kan worden ingesteld hoe zwaar de voorspelling meeweegt in het oplossingsproces, wat bepaalt hoe geleidelijk de interpolatie verloopt.

Hoofdstuk 3 bestudeert het algemene allocatieprobleem. In dit probleem moet een allocatie van spelers aan objecten gevonden worden van maximale waarde. Voor elk koppel van een speler en een object is het algemeen bekend wat de waarde en het gewicht zijn. Echter is niet elk koppel compatibel en beschikt alleen de speler over deze informatie. Verder heeft elk object een capaciteit en mag in een allocatie het totale gewicht van gealloceerde spelers aan het object niet de capaciteit overschrijden. Voor de variant van het probleem waarin alle gewichten en capaciteiten gelijk zijn aan 1 heeft ons mechanisme de optimale balans tussen consistency en robustness. Voor algemenere varianten van het probleem ontwerpen we probabilistische mechanismen.

In Hoofdstuk 4 introduceren we eerst het begrip *outliers* binnen het ontwerpen van mechanismen, en passen we dit toe op het fundamentele probleem van *facility location* op de reële lijn. Het doel is hier om de locatie te bepalen van een nieuwe publieke voorziening die wordt gebruikt door n spelers, waarbij elke speler een private voorkeurslocatie voor de voorziening heeft. We richten ons op het minimaliseren van ofwel de totale kosten ofwel de maximale kosten, dat wil zeggen: respectievelijk de som of het maximum van de kosten over alle spelers. We minimaliseren deze doelfuncties voor slechts een deel van de spelers, wat relevant is wanneer z spelers, de zogenaamde outliers, voorkeurslocaties hebben die significant afwijken van de rest. We bestuderen problemen met maximaal $z < n/2$ outliers, omdat we aantonen dat er anders geen strategyproof mechanisme met een begrensde approximatie bestaat. Ons mechanisme voor het minimaliseren van de maximale kosten behaalt een 2-approximatie, en we laten zien dat dit niet verbeterd kan worden, zelfs niet als het probleem wordt verrijkt met voorspellingen, in tegenstelling tot het probleem zonder outliers. Voor het minimaliseren van de totale kosten tonen we aan dat strategyproof mechanismen niet effectief gebruik kunnen maken van outliers, waardoor de approximatie verslechtert naarmate het aantal outliers toeneemt. In dit geval leidt het verrijken met voorspellingen wel tot betere approximaties, en heeft ons mechanisme een optimale balans tussen consistency en robustness.

Het tweede deel van deze thesis behandelt mechanismen waarbij geldtransacties beperkt zijn toegestaan. In Hoofdstuk 5 bestuderen we een inkoopveiling: spelers bieden bij een veilingmeester goederen (of diensten) aan met een algemeen bekende waarde en een private kostprijs. De veilingmeester moet vervolgens beslissen welke spelers worden ingehuurd en hoeveel ze betaald worden, waarbij de veilingmeester niet meer kan uitgeven dan zijn budget. We richten ons op problemen met waardefuncties die concaaf en niet-dalend zijn, en waarbij gedeeltelijke toewijzingen zijn toegestaan; dat wil zeggen, een speler kan elk aantal in $0, 1, \ldots, k$ toegewezen krijgen voor $k \geq 1$. Dit overbrugt de situatie met ondeelbare goederen ($k = 1$) met de situatie waarbij elke fractionele toewijzing mogelijk is ($k \to \infty$). We behalen verbeterde resultaten voor het speciale geval van lineaire waardefuncties en tonen een onderscheid aan tussen de fractionele en ondeelbare setting. Ondanks deze verbeteringen blijft het bepalen van de best mogelijke approximatie een open onderzoeksvraag.

In het derde deel verschuift de focus van het ontwerpen en analyseren van strategyproof mechanismen naar het analyseren van stabiele uitkomsten die voortkomen uit strategische interacties binnen een reeds vastgesteld mechanisme. Hoofdstuk 6 onderzoekt de efficiëntie van stabiele uitkomsten *(price of anarchy (POA))* van simultane eersteprijsveilingen, waarbij een veilingmeester meerdere goederen tegelijk verkoopt aan een groep spelers. Per veiling wint de speler met het hoogste bod en is de prijs gelijk aan deze hoogste bieding. Het model dat we bestuderen heeft verschillende soorten spelers: de impact van de prijzen op de nutsfunctie van een speler wordt verdisconteerd met een parameter tussen 0 (prijzen hebben geen invloed) en 1. Daarnaast moet elke speler aan de voorwaarde voldoen dat hun verwachte totale betaling niet hoger is dan hun verwachte totale waarde van de gewonnen goederen. We introduceren een nieuwe methode waarmee een grens op de POA berekend kan worden voor verschillende soorten waardefuncties, verschillende definities van een stabiele uitkomst, en voor elke samenstelling van spelers. Als de veilingmeester geschikte minimumprijzen invoert kan onze methode ook gebruikt worden als de waardefuncties additief zijn, wat vooral relevant is wanneer de veilingmeester voorspellingen kan gebruiken die de waarderingen van spelers benaderen.

Summary

The main focus of this thesis is on game-theoretic versions of optimization problems. In contrast to standard optimization problems, a part of the input is unknown and distributed over multiple strategic agents. In order to compute a feasible solution of a certain quality, the unknown input is retrieved by requesting the agents to declare it. However, as the agents are strategic and their goal not necessarily aligns with the optimization objective, their declaration will not be truthful if this is in their own best interest. This strategic behavior can be undesirable from a societal perspective and may lead to inefficient solutions. A way to address this is the design of strategyproof (SP) mechanisms, which retrieve the declarations and compute a feasible solution such that it is always in an agent's best interest to declare their input truthfully, regardless of the actions of the other agents. This is the main focus of the first two parts of this thesis.

Part One considers two mechanism design problems without monetary transfers in the learning-augmented framework. This emerging new line of research bridges worst-case analysis and machine learning by augmenting problems with a prediction regarding the unknown input. The goal is to achieve an optimal approximation guarantee when the prediction is perfect (consistency), while maintaining the worst-case approximation guarantee when the prediction is imperfect (robustness). Both mechanism design problems are augmented with a prediction of the optimal solution, and all mechanisms in Part One provide error-dependent approximation guarantees that smoothly interpolate between the consistency and robustness guarantees. Additionally, some mechanisms are parameterized by a confidence parameter that can be used to specify how much the mechanism should trust the prediction, which determines the slope of the interpolation.

Chapter 3 considers the generalized assignment problem (GAP) in the private graph model, where the goal is to find a feasible assignment of agents to resources of maximum value. In this problem, each agent has a value and size for each resource which is publicly known, but the agent-resource compatibilities are private information. Each resource has a publicly known capacity, which cannot be exceeded by the total size of the allocated agents. For the GAP variant where all sizes and capacities are equal to 1, i.e., the bipartite matching problem, our mechanism achieves the optimal consistency-robustness trade-off. For

more general variants of GAP we derive randomized mechanisms with improved approximation guarantees in expectation.

In Chapter 4, we first introduce the notion of outliers within mechanism design, and apply it to the fundamental problem of single facility location on the real line. Here, the goal is to determine the location of a new public facility used by n agents who each have a private preferred location of the facility. We focus on either minimizing the utilitarian or egalitarian objective, that is, minimize the total cost or the maximum cost over all agents. We revisit the problem by minimizing the objective for only a subset of agents, a natural approach when z agents—referred to as outliers—have preferred locations that significantly differ from the rest. We consider settings with at most $z < n/2$ outliers, as we show that no SP mechanism can achieve a bounded approximation guarantee otherwise. For the egalitarian objective, our mechanism achieves a tight 2-approximation and we show that this cannot be improved when augmenting the problem with various types of predictions, contrary to the problem without outliers. For the utilitarian objective, we show that SP mechanisms cannot effectively leverage outliers which causes the approximation guarantee to deteriorate as the number of outliers increases. In this case, augmenting with a prediction does lead to improved guarantees, and our mechanism achieves the optimal consistency-robustness trade-off.

Part Two considers budget-feasible mechanism design in which monetary transfers are allowed, but limited. More specifically, Chapter 5 considers knapsack procurement auctions: agents offer a service with a publicly known valuation and a private cost to the auctioneer, and the auctioneer must decide which agents to hire and how much to pay them, while being constrained by a budget on the total payment. We focus on settings with concave and non-decreasing valuations and in which partial allocations are allowed, i.e., the allocation of an agent can be any number in $\{0, 1, \ldots, k\}$ for $k \geq 1$. This bridges the indivisible setting of $k = 1$ with the divisible setting which allows any fractional allocation as $k \to \infty$. We derive improved results for the special case of linear valuations, and establish a separation between the divisible and indivisible setting. Even though we provide significant improvements, determining the best possible approximation guarantee remains an open question.

In Part Three, we shift our focus from the design and analysis of strategyproof mechanisms to the analysis of equilibria resulting from strategic interactions within an already predefined environment. Chapter 6 considers the efficiency of equilibrium outcomes (price of anarchy (POA)) of simultaneous single-item first-price auctions (FPAs) under the hybrid agent model, in which the impact of payments on the utility function of an agent are discounted by a parameter that ranges from 1 for classical utility maximizers to 0 for value maximizers. Additionally, each agent has a return-on-investment constraint which enforces that their expected payment does not exceed their expected acquired value. We introduce a new smoothness framework through which we derive bounds on the POA for coarse correlated equilibria, fractionally subadditive valuations, and any composition of agent types. For additive valuations and feasible reserve prices, bounds on the POA for well-supported equilibria can also be derived through our framework, which is especially relevant when the seller can leverage predictions to approximate agents' valuations.

Titles in the ILLC Dissertation Series:

ILLC DS-2020-12: **Bastiaan van der Weij**
Experienced listeners: Modeling the influence of long-term musical exposure on rhythm perception

ILLC DS-2020-13: **Thom van Gessel**
Questions in Context

ILLC DS-2020-14: **Gianluca Grilletti**
Questions & Quantification: A study of first order inquisitive logic

ILLC DS-2020-15: **Tom Schoonen**
Tales of Similarity and Imagination. A modest epistemology of possibility

ILLC DS-2020-16: **Ilaria Canavotto**
Where Responsibility Takes You: Logics of Agency, Counterfactuals and Norms

ILLC DS-2020-17: **Francesca Zaffora Blando**
Patterns and Probabilities: A Study in Algorithmic Randomness and Computable Learning

ILLC DS-2021-01: **Yfke Dulek**
Delegated and Distributed Quantum Computation

ILLC DS-2021-02: **Elbert J. Booij**
The Things Before Us: On What it Is to Be an Object

ILLC DS-2021-03: **Seyyed Hadi Hashemi**
Modeling Users Interacting with Smart Devices

ILLC DS-2021-04: **Sophie Arnoult**
Adjunction in Hierarchical Phrase-Based Translation

ILLC DS-2021-05: **Cian Guilfoyle Chartier**
A Pragmatic Defense of Logical Pluralism

ILLC DS-2021-06: **Zoi Terzopoulou**
Collective Decisions with Incomplete Individual Opinions

ILLC DS-2021-07: **Anthia Solaki**
Logical Models for Bounded Reasoners

ILLC DS-2021-08: **Michael Sejr Schlichtkrull**
Incorporating Structure into Neural Models for Language Processing

ILLC DS-2021-09: **Taichi Uemura**
Abstract and Concrete Type Theories

ILLC DS-2021-10: **Levin Hornischer**
Dynamical Systems via Domains: Toward a Unified Foundation of Symbolic and Non-symbolic Computation

ILLC DS-2021-11: **Sirin Botan**
Strategyproof Social Choice for Restricted Domains

ILLC DS-2021-12: **Michael Cohen**
Dynamic Introspection

ILLC DS-2021-13: **Dazhu Li**
Formal Threads in the Social Fabric: Studies in the Logical Dynamics of Multi-Agent Interaction

ILLC DS-2021-14: **Álvaro Piedrafita**
On Span Programs and Quantum Algorithms

ILLC DS-2022-01: **Anna Bellomo**
Sums, Numbers and Infinity: Collections in Bolzano's Mathematics and Philosophy

ILLC DS-2022-02: **Jan Czajkowski**
Post-Quantum Security of Hash Functions

ILLC DS-2022-03: **Sonia Ramotowska**
Quantifying quantifier representations: Experimental studies, computational modeling, and individual differences

ILLC DS-2022-04: **Ruben Brokkelkamp**
How Close Does It Get?: From Near-Optimal Network Algorithms to Suboptimal Equilibrium Outcomes

ILLC DS-2022-05: **Lwenn Bussière-Carae**
No means No! Speech Acts in Conflict

ILLC DS-2022-06: **Emma Mojet**
Observing Disciplines: Data Practices In and Between Disciplines in the 19th and Early 20th Centuries

ILLC DS-2022-07: **Freek Gerrit Witteveen**
Quantum information theory and many-body physics

ILLC DS-2023-01: **Subhasree Patro**
Quantum Fine-Grained Complexity

ILLC DS-2023-02: **Arjan Cornelissen**
Quantum multivariate estimation and span program algorithms

ILLC DS-2023-03: **Robert Paßmann**
Logical Structure of Constructive Set Theories

ILLC DS-2023-04: **Samira Abnar**
Inductive Biases for Learning Natural Language

ILLC DS-2023-05: **Dean McHugh**
Causation and Modality: Models and Meanings

ILLC DS-2023-06: **Jialiang Yan**
Monotonicity in Intensional Contexts: Weakening and: Pragmatic Effects under Modals and Attitudes

ILLC DS-2023-07: **Yiyan Wang**
Collective Agency: From Philosophical and Logical Perspectives

ILLC DS-2023-08: **Lei Li**
Games, Boards and Play: A Logical Perspective

ILLC DS-2023-09: **Simon Rey**
Variations on Participatory Budgeting

ILLC DS-2023-10: **Mario Giulianelli**
Neural Models of Language Use: Studies of Language Comprehension and Production in Context

ILLC DS-2023-11: **Guillermo Menéndez Turata**
Cyclic Proof Systems for Modal Fixpoint Logics

ILLC DS-2023-12: **Ned J.H. Wontner**
Views From a Peak: Generalisations and Descriptive Set Theory

ILLC DS-2024-01: **Jan Rooduijn**
Fragments and Frame Classes: Towards a Uniform Proof Theory for Modal Fixed Point Logics

ILLC DS-2024-02: **Bas Cornelissen**
Measuring musics: Notes on modes, motifs, and melodies

ILLC DS-2024-03: **Nicola De Cao**
Entity Centric Neural Models for Natural Language Processing

ILLC DS-2024-04: **Ece Takmaz**
Visual and Linguistic Processes in Deep Neural Networks: A Cognitive Perspective

ILLC DS-2024-05: **Fatemeh Seifan**
Coalgebraic fixpoint logic Expressivity and completeness result

ILLC DS-2024-06: **Jana Sotáková**
Isogenies and Cryptography

ILLC DS-2024-07: **Marco Degano**
Indefinites and their values

ILLC DS-2024-08: **Philip Verduyn Lunel**
Quantum Position Verification: Loss-tolerant Protocols and Fundamental Limits

ILLC DS-2024-09: **Rene Allerstorfer**
Position-based Quantum Cryptography: From Theory towards Practice

ILLC DS-2024-10: **Willem Feijen**
Fast, Right, or Best? Algorithms for Practical Optimization Problems

ILLC DS-2024-11: **Daira Pinto Prieto**
Combining Uncertain Evidence: Logic and Complexity

ILLC DS-2024-12: **Yanlin Chen**
On Quantum Algorithms and Limitations for Convex Optimization and Lattice Problems

ILLC DS-2024-13: **Jaap Jumelet**
Finding Structure in Language Models

ILLC DS-2025-01: **Julian Chingoma**
On Proportionality in Complex Domains

ILLC DS-2025-02: **Dmitry Grinko**
Mixed Schur-Weyl duality in quantum information

ILLC DS-2025-03: **Rochelle Choenni**
Multilinguality and Multiculturalism: Towards more Effective and Inclusive Neural Language Models

ILLC DS-2025-04: **Aleksi Anttila**
Not Nothing: Nonemptiness in Team Semantics

ILLC DS-2025-05: **Niels M. P. Neumann**
Adaptive Quantum Computers: decoding and state preparation

ILLC DS-2025-06: **Alina Leidinger**
Towards Language Models that benefit us all: Studies on stereotypes, robustness, and values

ILLC DS-2025-07: **Zhi Zhang**
Advancing Vision and Language Models through Commonsense Knowledge, Efficient Adaptation and Transparency

ILLC DS-2025-08: **Sophie Klumper**
The Gap and the Gain: Improving the Approximate Mechanism Design Frontier in Constrained Environments