TECHNISCHE UNIVERSITÄT MÜNCHEN

TUM School of Computation, Information and Technology

Auction Design in the Presence of Financial Constraints

Eleni Batziou

Vollständiger Abdruck der von der TUM School of Computation, Information and Technology der Technischen Universität München zur Erlangung einer

Doktorin der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitz:	Prof. Dr. Stefanie Rinderle-Ma
Prüfende der Dissertation:	1. Prof. Dr. Martin Bichler 2. Prof. Dr. Susanne Albers

Die Dissertation wurde am 09.04.2024 bei der Technischen Universität München eingereicht und durch die TUM School of Computation, Information and Technology am 02.10.2024 angenommen.

Abstract

Traditional market design studies rules that lead to efficient allocations of goods to a set of participants and minimize strategic manipulation. Along with deriving an allocation rule, a set of prices denotes the monetary transfer agents must submit for this exchange. A fundamental challenge is the computation of market equilibria. In the perfect competition model, with fully divisible items, where bidders are price-takers, Walrasian equilibria constitute a design concept of choice due to their desirable properties, including maximizing welfare and being core-stable. Assuming the presence of hard budget limits, such equilibria are not guaranteed to exist. Ignoring financial constraints has been proven to result in lower efficiency and instability.

The introduction of exogenous budgets poses additional complications to market analysis: payoff-maximizing equilibria might be impossible to attain and computation of core-stable, welfare-maximizing outcomes quickly becomes intractable. Addressing these complexities, in the first project of this dissertation, we examine a simple market model with restricted preferences, the assignment market, and design an ascending auction that always converges in a core-stable outcome while considering each bidder's budget constraint. Additionally, under appropriate conditions and decisions made by an auctioneer, the resulting outcome is simultaneously welfare-maximizing. When these conditions cease to hold, one can no longer hope for incentive-compatible mechanisms that satisfy the desired properties. Our research reveals the underlying hardness of combining core stability and welfare maximization in the presence of financial constraints: even under complete information and with access to value queries, determining such an outcome is an NP-hard problem.

While the first project of this dissertation examines the impact of financial constraints on the side of bidders, the second publication introduces a budget limit for the auctioneer. The setting can be understood as a reverse auction for conservation: where a principal aims to purchase land from a set of sellers. Each seller owns private land parcels, and the principal offers monetary incentives to encourage their participation. In this case, the objective is maximizing biodiversity gains while implementing an auction, which is measured by the number of sellers not developing their parcels. Drawing from the literature on budget-feasible mechanism design, we experimentally demonstrate the power of approximation algorithms. The proposed mechanism remains within a predefined budget for a slight loss in efficiency compared to the optimal solution. It reaches a solution in polynomial time while being incentive-compatible and addressing the challenge of eliciting sellers' opportunity costs. Governments or authorities possess fixed funds targeted to conservation, and thus, as shown in our work, various standard auction mechanisms are deemed unattractive due to the arbitrarily high budget violation. The budget-feasible clock auction suggested in this work is an attractive alternative for policymakers in biodiversity conservation.

Shifting the focus to eliciting preference in the complex combinatorial auction domain, where agents seek to purchase mutually disjoint subsets of items, in further work conducted in the course of this doctoral degree, we addresse the challenge posed by the exponential growth in the number of bundles in these auctions. Recognizing the computational burden of traditional value queries, we introduce a novel approach that leverages reinforcement learning techniques for combinatorial optimization. Our parametric randomized algorithm predicts value queries, optimizing parameters through gradient descent based on estimated gradients concerning the auction objective of welfare maximization. Experimental validation on real-world combinatorial auction scenarios demonstrates a significant reduction in auction runtime while maintaining comparable allocative efficiency to state-of-the-art mechanisms. Our findings pave the way for the integration of reinforcement learning in preference elicitation for large combinatorial markets and broader systems using recommendation.

Acknowledgements

First, I would like to thank my advisor, Martin Bichler, for his guidance and support throughout these years. I also want to express my gratitude to my other coauthors and consider myself fortunate to have had the opportunity to work alongside them. I sincerely thank the examination committee, for their valuable insights and their dedicated time in evaluating my dissertation.

I would like to express my appreciation to my colleagues at the Chair of Decision Sciences and Systems, as well as the members of the AdONE community, for the exciting discussions that have helped broaden my research horizons and the fun working environment.

Since moving to a new country during a pandemic has been quite a challenge, I would like to thank all my friends, from Munich to Athens and beyond, for their continuous support and the memories we've shared all this time.

I am always grateful for the love and support of my family; it is thanks to them that I'm always determined to pursue my ambitions. Last, my deepest gratitude goes to Π for always being there and believing in me even in challenging times. I am thankful we've shared this journey and can't wait for the next steps.

Contents

1 Introduction

1.1 Motivation

In traditional economic theory, market design is characterized by the use of mathematical optimization, game theory, and mechanism design to facilitate the exchange of goods, services, or resources among market participants. The primary goal is to establish rules that ensure efficient allocations, maximize welfare, and minimize strategic manipulation. Market design seeks to enhance market efficiency, promote competition, and achieve desirable socio-economic outcomes, extending beyond financial markets to include labor, auction, and resource allocation markets.

Significant contributions in the field of market design have been made by economists who have addressed various challenges and analyzed market models. The power of market design has been employed in solving important, real-life problems. Roth and Sotomayor (1989) applied matching algorithms to address the college admissions problem, as well as the assignment of medical graduates to residency programs (NRMP) (Roth, 1984). Following this work, Roth et al. (2004) developed an algorithm that matched multiple donor-patient pairs in kidney transplant exchanges, effectively saving human lives. Another notable work by Myerson (1981) revealed the connection between monetary transfers to agents that incentivize truthful reporting of their private information and the resulting allocation of the mechanism. His work focused on designing mechanisms that encourage truthful revelation of participants' private information, a crucial property that leads to efficiency and fairness in various allocation problems. Focusing on auctions, Paul Milgrom and Robert B. Wilson made significant contributions to simultaneous, multi-item auctions (SMRA) (Milgrom, 2000) and introduced innovative auction formats with applications in various real-world scenarios. The most prominent application is the auctioning of electromagnetic spectrum licenses, where the Federal Communications Commission (FCC) in the United States used their auction design to allocate permits to telecommunication companies, resulting in billions of dollars saved in bidder payments. The Royal Swedish Academy of Sciences has celebrated the importance of these works by awarding the authors the Nobel Prize in Economics Sciences in 2007, 2012, and 2020.

A core challenge in economics is analyzing market equilibria, which represent states where supply equals demand in an economy. A valuable tool in understanding the functioning of markets and implications for economic outcomes is a competitive equilibrium. This refers to a point where, in the context of perfect competition, the quantity demanded by consumers is equal to the amount supplied by producers. At the resulting equilibrium price, there is no incentive for price changes, and the market is stable because no excess demand or supply can alter prices. Market participants, either buyers or sellers, have no incentive to deviate from the competitive equilibrium. Competitive equilibria are often associated with allocative efficiency, maximizing social welfare among the agents in the market.

The seminal work of Arrow and Debreu (1954) demonstrated that a set of competitive equilibrium prices always exists given convex preferences, perfect competition and demand independence. In their model, each participant has an endowment of goods and money, and utilities are considered cardinal. In this market, participants are price-takers seeking to maximize their total value. The Arrow-Debreu model has also been the origin of the important welfare theorems.

Assuming divisibility of goods simplifies the analysis and allows for the use of continuous mathematics to derive solutions and insights. Divisibility implies that goods can be split into any fraction, facilitating the exploration of optimal allocations and equilibria without the complexity introduced by indivisible items, where allocation might require discrete optimization techniques. A prominent example is that of Fisher markets, where each participant possesses a monetary budget, while the total quantity of each good available in the market is also fixed, making them ideal for studying resource allocation under budget constraints. Considerable progress has been made towards developing algorithms determining allocations and prices in Fisher markets (Vazirani, 2007; Vazirani and Yannakakis, 2011; Cole et al., 2016). The objective in these markets is the maximization of the Nash social welfare function, and a common assumption in this literature is the divisibility of goods.

Contrary to these models, in real-world markets, perfect divisibility of goods ceases to be a realistic assumption. A significant line of work (Kelso and Crawford, 1982; Bikhchandani and Mamer, 1997; Gul et al., 2000; Leme, 2017; Baldwin and Klemperer, 2019b) examines necessary conditions for the existence of competitive equilibria in markets where goods are indivisible and utilities are quasilinear. Under quasilinear utilities, buyers aim to maximize their payoff, namely the difference between value and paid price, and have no limit on the amount they are allowed to spend.

In practical market settings, the assumption of quasilinearity often breaks down due to the presence of budget constraints, which significantly influence buyer behavior and market dynamics. Budget constraints emerge as a critical factor, influencing strategies and outcomes in various market scenarios, from housing markets to spectrum auctions. This dissertation delves into the complexities of designing and analyzing mechanisms within these exogenous constraints, challenging the conventional equilibrium models and proposing novel solutions to ensure core stability and efficiency. Taking a housing market as an example, buyers seeking to purchase real estate enter the market facing a spending limit. This limit can dictate how aggressive their offers can be on properties, which defines the final assignment. Bidding strategies of low-budget participants are effectively constrained, while the ones enjoying higher budgets can easily outbid them. Moving away from quasi-linearity, an extensive line of research work has examined the effect of private budget constraints in mechanism design and auctions (Che and Gale, 2000; Benoît and Krishna, 2001; Borgs et al., 2005; Dobzinski et al., 2008; Colini-Baldeschi et al., 2011; Dütting et al., 2016). From a practical perspective, budgets play a prominent role in internet advertising (Borgs et al., 2007; Conitzer et al., 2019, 2022) and spectrum auctions (Bichler and Goeree, 2017).

1.2 Contributions

Addressing the complexities introduced by financial constraints on the side of market participants, this dissertation embarks on a comprehensive exploration of auction mechanisms that aim at balancing efficiency, stability, and budget feasibility. Central to this investigation are two pivotal contributions: the design of an auction mechanism that assures core-stability in assignment markets with financially constrained buyers and the application of budget-feasible mechanisms in the context of environmental conservation efforts. The existence of exogenous budget constraints challenges the traditional paradigms of market design and our research proposes novel frameworks that seek to accommodate the financial limitations of market participants, while maintaining important mechanisms properties and achieving high welfare guarantees. The contributions made in this work not only advance our understanding of market mechanisms in the presence of financial constraints but also offer practical frameworks that can be applied across a variety of domains, from environmental conservation to spectrum allocation.

1.2.1 Part I: Core-Stability in Assignment Markets with Financially Constrained Buyers

One of the most essential notions defining stability in market design is that of the *core*. The core comprises the set of outcomes where no subset of participants can increase their collective payoff by forming a separate coalition and reallocating resources among themselves. In markets where agents are subject to binding budget constraints, a core outcome is not guaranteed to be welfare-maximizing, underscoring the challenges posed by budgets in auction design. A core outcome is not necessarily equivalent to a competitive equilibrium, and various outcomes from the core could lead to different levels of social welfare. Centered around the question of determining core-stable outcomes that simultaneously maximize welfare among market participants under binding financial constraints, Bichler and Waldherr (2022) demonstrate that the problem becomes intractable even for small instances in combinatorial exchanges. This work poses the natural question of whether there is hope for polynomial-time algorithms that result in such outcomes in simpler, more restrictive market models.

In the first contribution of this dissertation (Batziou et al., 2022a), drawing upon this question, we shift our focus on assignment markets, with unit-demand bidders that are constrained by a hard budget limit. Each bidder is interested in acquiring at most one item in these markets. Using solely demand queries, we develop an iterative auction algorithm that determines core-stable outcomes without knowing bidders' true valuations. This ascending auction extends the well-known DGS algorithm of Demange et al. (1986) under the influence of budget limits. When the objective involves choosing the welfare-maximizing, among all core outcomes, the auctioneer faces a challenging problem: excluding the appropriate buyers from expressing demand for items in subsequent rounds. Yet, accessing only the demand sets of each bidder, the resulting welfare may be arbitrarily low. Given the restrictive nature of unit-demand valuations, a naturally arising question is whether this condition is enough to guarantee the existence of a mechanism that satisfies incentive compatibility, welfare maximization, and core stability. We answer this question negatively: no auction mechanism can satisfy these properties simultaneously under budget constraints. Our result can be seen as an extension of Ausubel (2006) that demonstrates the clash between incentive compatibility and core stability in general quasilinear utility models.

Under full access to bidders' private information, contrary to expectations of polynomial time solutions, our paper reveals a novel result: determining core-stable, welfare-maximizing outcomes with financially constrained buyers is an NP-complete optimization problem. The complexity stems from an intricate reduction from the maximum independent set problem. These results underscore that, even in the most straightforward multi-item assignment markets, the presence of budgets renders the computation of desirable outcomes quickly intractable.

1.2.2 Part II: Budget-Feasible Market Design for Biodiversity Conservation: Considering Incentives and Spatial Coordination

The settings examined in the two publications of this dissertation consist of budget-constrained agents on either side of a market: in the first work (see Chapter 3), buyers face a financial limit, while in the second work (see Chapter 4), the constraint lies on the side of the auctioneer.

Prior to the seminal paper of Singer (2010), research on budget-feasible mechanisms had been scarce: the large majority of work assumed a principal has access to infinite resources that can be sacrificed in exchange for incentive compatibility. Incentive compatibility is a fundamental objective in market design, guaranteeing that no agent misreports their values to the auctioneer with the prospect of receiving higher utility. In a procurement auction, typically used by governments or organizations, a principal wishes to purchase services from sellers who submit bids based on their private costs. Budget feasible procurement is a practically motivated problem, with various applications (Roth and Schoenebeck, 2012; Singer, 2012; Singer and Mittal, 2013; Horel et al., 2014; Goel et al., 2014), where an auctioneer seeks to maximize a social value function on subsets of items, and the sum of payments is constrained by the budget. The problem can be cast as a variant of the famous *knapsack* problem: given a budget and set of items with costs and values, what is the optimal set of items that maximizes value under the budget?

The work of Singer sparked a line of research around budget feasible auctions, intending to achieve improved approximations of the optimal solution. Inspired by the well-known greedy algorithm in submodular maximization, that selects items based on marginal contribution to value and achieves a $1 - 1/e$ approximation (Nemhauser et al., 1978), a prominent stream of research on budget feasible auctions has assumed submodular auctioneer values. Given the relative hardness in achieving good approximations for non-monotone submodular values (Amanatidis et al., 2019; Bei et al., 2017), several works have introduced the additional constraint of monotonicity on the valuations (Chen et al., 2011; Anari et al., 2014; Jalaly and Tardos, 2021), thus achieving more promising, even constant approximation ratios. Beyond submodular, additive (Chen et al., 2011; Anari et al., 2014; Gravin et al., 2020) as well as subadditive (Dobzinski et al., 2011; Bei et al., 2017) value functions have been further examined.

Most work on budget-feasible auctions can be described as sealed-bid auctions, where participants' private information is made known to the auctioneer. A crucial limitation of such auctions is that, in practice, studies (Kagel et al., 1987) show that participants tend to misreport despite the theoretical promise of incentive compatibility. Bypassing this shortcoming, Milgrom and Segal (2020) introduced *clock auctions*, a novel family of mechanisms that progresses over rounds where a gradually decreasing clock price is offered, and bidders might accept and continue or exit the auction. The process ends when the budget limit has been reached. Posted price schemes, where an auctioneer broadcasts a fixed take-it-or-leave-it price, have been employed in budget-feasible mechanism design (Badanidiyuru et al., 2012; Balkanski and Hartline, 2016) and can be considered as a particular form of clock auctions. Contrary to prior work, Balkanski et al. (2022) propose a clock auction mechanism that respects budgets and, for the special case of monotone submodular valuations, achieves the best-known approximation of the optimal, a constant factor of 4.75. The algorithm is based on a backward greedy method for submodular maximization, with bidders being eliminated in each round while prices offered are scaled to fit the set budget limit. Being a clock auction at its core, this algorithm has desirable properties, such as resistance to collusion, incentive compatibility, and transparency, while being deterministic with a polynomial runtime.

In the second contribution of this dissertation (Batziou and Bichler, 2023), we adopt a mechanism design perspective to address a relevant environmental emergency: biodiversity conservation. With the rapid decline in wildlife populations observed in recent years, the need to incentivize private landowners to conserve becomes pressing. Thus, the challenge of providing appropriate incentives to sellers through monetary payments to prevent agricultural use arises. Governing bodies concerned with modeling such markets are subject to natural financial limitations: a fraction of public funds is reserved for environmental purposes and can thus be utilized as payments. In this work, we examine the clock auction algorithm (Balkanski et al., 2022), which satisfies incentive compatibility and respects budgets while achieving a close approximation of the optimal solution in practice. Indeed, our experimental analysis reveals that this approximation algorithm achieves high welfare compared to the optimal solution while remaining feasible under monetary constraints, thus becoming a promising alternative for policymakers. Our findings indicate that the celebrated VCG auction algorithm (Vickrey, 1961; Clarke, 1971; Groves, 1973), which has been employed as the state-of-the-art method in conservation auctions (Polasky et al., 2014), may result in payments that arbitrarily exceed budget, making it an undesirable option in practice.

1.2.3 Other Contributions

In a market, goods can be traded as single items or combined in bundles. In combinatorial auctions, bidders can express preferences and place bids on bundles of items, thus allowing for complex bidding strategies. A crucial challenge in such auctions is *preference elicitation*: given the exponential number of bundles being auctioned, representing all possible combinations of individual items into bundles, querying bidders for their true valuations for each possible combination becomes intractable. To overcome the exponential amount of communication required, iterative combinatorial auctions (ICAs) (Parkes, 2005) have emerged as an efficient model where a set of demand queries is sent to the bidders in each round, and prices can be computed based on the (ascending) clock round (Ausubel et al., 2004). An auctioneer seeking to generate an optimal outcome requires access to the full valuation profiles, which could, in turn, prove time-consuming, as bidders enter the auction with a predetermined set of bundle-value pairs, beyond which they might be unable to respond accurately to demand queries, as computation can become cumbersome (Scheffel et al., 2012; Bichler et al., 2013). Considering the inefficiencies present in combinatorial auctions, a recent stream of work has employed learning algorithms that estimate the valuation functions of bidders using a selected, small set of targeted queries and achieve close approximation of the optimal solution under full access to preference profiles (Brero et al., 2018, 2019; Weissteiner and Seuken, 2020; Beyeler et al., 2021; Weissteiner et al., 2022b,a, 2023). Using Support Vector Regressors and neural networks, among others, a candidate estimate valuation function is built based on bidders' true values, as responded to in the sent queries. In each round, the set of queries is determined by solving the winner determination problem (WDP), traditionally solved by a Mixed Integer Program (MIP). However, even if the iterative query process significantly decreases the amount of required communication, the computational bottleneck of solving the MIP remains.

Alleviating the computational burden, in work conducted in the course of this doctoral program (Batziou et al., 2022b), we build on the deep learning-inspired algorithm for iterative combinatorial auctions (Weissteiner and Seuken, 2020), solving the optimization problem using the well-known *Reinforce* algorithm (Williams, 1992). Our reinforcement learning framework computes a randomized allocation in the form of a distribution over queries, which will constitute the query set of the next round, and the algorithm parameters are optimized following traditional policy gradient methodology. Contributing to a line of research revolving around using reinforcement learning techniques in combinatorial optimization (Bello et al., 2016; Kool et al., 2019; Deudon et al., 2018), our work can be seen as a first step in incorporating reinforcement learning in preference elicitation mechanisms. Experimental results on the spectrum auction test suite (SATS) (Weiss et al., 2017) demonstrate that our framework achieves comparable efficiency in a fraction of the computation time. Thus, sacrificing a small fraction of efficiency, our framework can be seen as a promising alternative for implementing large-scale auctions.

While this dissertation has revolved mainly around auction design, the author has also conducted work in the field of computational complexity, in the context of problems in fair division (Batziou et al., 2021). Fair division is a study at the intersection of economics and computer science and is centered around distributing a limited set of goods or resources among agents in a fair, equitable manner. A notable problem within fair division is the *consensus halving* where a given set of resources is divided between two groups of agents, such that every agent values the two resulting parts equally. The resources in question are often represented as a [0,1] interval, and the existence of a solution can be determined using the Borsuk-Ulam theorem from algebraic topology. In this work, we prove that computation of an approximate solution to the consensus halving problem is polynomial time equivalent to computing an approximate solution to the Borsuk-Ulam search problem, as well as define a novel complexity class, *BBU*, that refers to the computational version of the Borsuk-Ulam theorem, and explore its properties and relation to FIXP.

1.3 Outline

The structure of this dissertation is as follows: First, in Chapter 2, we introduce necessary definitions and notation with regard to auctions. We explore desiderata for auctions, various payment rules, and notable auction algorithms, as well as a high-level idea of computational complexity classes relevant to market design problems. In Chapter 3, the first project that deals with computing core-stable outcomes in assignment markets, where buyers are constrained by a budget, is presented. Chapter 4 includes the work on mechanism design for biodiversity conservation, where a principal is concerned with remaining budget-feasible. Finally, Chapter 5 concludes this dissertation and presents directions for further research work.

2 Theoretical Background

This chapter forms the foundation of the publications included in the remaining part of this dissertation. Key notions and theorems are introduced, providing the reader with insight on core concepts of market design, and aid in understanding the context of this work.

We draw on a multitude of topics, ranging from designing mechanisms to algorithms and complexity, explaining the main concepts required to understand the diverse topics examined throughout the course of this dissertation.

The notation employed in this chapter may differ from that used in the included publications. Nonetheless, we deliberately adjust notation to simplify the comprehension of fundamental notions.

2.1 Preliminaries

An auction can be described as the process where a set of buyers places competitive bids for acquiring goods or services provided by a set of sellers. The auction is defined by a set of n *bidders* (or buyers) $i \in \mathcal{I}$, a set of sellers $j \in \mathcal{J}$, with $\mathcal{I} \cap \mathcal{J} = \emptyset$, and a finite set of indivisible *items* $k \in \mathcal{K}$. We define the set of *bundles* $S \in 2^{\mathcal{K}}$ as the set of all possible combinations of the $\mathcal{K}$ items. Bidders submit bids that a seller can reject or accept, resulting in a transaction of goods in exchange for a monetary payment.

The final result in an auction is defined as *outcome* $o = (x, p) \in O$, and consists of an *allocation* $x = (S_1, S_2, \ldots, S_n)$ and *payment* vectors $p = (p_1, p_2, \ldots, p_n)$. The allocation rule determines which bundle is won by each bidder and payments define the monetary transfer the bidder has to commit to purchase the allocated bundle. A hard *budget* constraint $b^i \in \mathbb{R}_{\geq 0}$ is a financial limit that dictates the upper bound on the payments of each bidder.

Bidders' preferences are quantified by a *valuation function*, which indicates how much agent i values each good or alternative, and is defined as a mapping over the set of all bundles:

$$v_i : 2^{\mathcal{K}} \to \mathbb{R}_{\geq 0}$$

In a similar manner, each seller is characterized by a *reserve price* for each bundle $Z \in 2^{\mathcal{K}}$ with $v_j(Z) \in \mathbb{R}_{\geq 0}$. Sellers do not accept bids below their reserve price for any bundle, and therefore the sale cannot take place.

Without the presence of appropriate incentives, bidders might misreport in an auction seeking to achieve an individually more profitable outcome. The bidding function of each bidder is defined as:

$$b_i : 2^{\mathcal{K}} \rightarrow \mathbb{R}_{\geq 0}$$

Placing bids, agents express their willingness to pay for a certain item or bundle. Under *truthful* bidding, $b_i(S) = v_i(S)\ \forall i \in I,\ S \subseteq 2^{\mathcal{K}}$, and thus bids correspond to true valuations.

Important properties of the valuation functions of all bidders that are valid throughout this dissertation are the following:

- **Empty bundle:** The value of the empty bundle is always 0, $v_i(0) = 0$.
- **Monotonicity:** For sets $S' \subseteq S$, $v_i(S') \leq v_i(S)$, namely the valuation of a bidder for larger sets is weakly increasing.
- **Independent private values:** The value of a bidder is private information and depends solely on the received bundle, i.e. does not change even if information on other bidders' values or allocation is known.

In simple, non-combinatorial models, such as *assignment markets*, bidders only express preferences on single items and are interested in winning solely one item (*unit demand*). In such markets, a common assumption is that the set of *sellers* coincides with the set of items - namely each seller is selling a single item.

Valuation functions can be subject to various mathematical rules that dictate how combinations of sets are evaluated by bidders. Common types of valuations in auction theory and mechanism design literature are:

- *Additive:* For sets $S, T \subseteq 2^{\mathcal{K}}$: $v(S) + v(T) = v(S \cup T)$
- *Subadditive:* For sets $S, T \subseteq 2^{\mathcal{K}}$: $v(S) + v(T) \geq v(S \cup T)$
- *Submodular:* For sets $S, T \subseteq 2^{\mathcal{K}}$: $v(S) + v(T) \geq v(S \cup T) + v(S \cap T)$
- *Superadditive:* For sets $S, T \subseteq 2^{\mathcal{K}}$: $v(S) + v(T) \leq v(S \cup T)$

A hierarchy between the first three of the above classes is shown in the diagram of Figure 2.1. Superadditive functions can be seen as the reverse of subadditive ones.

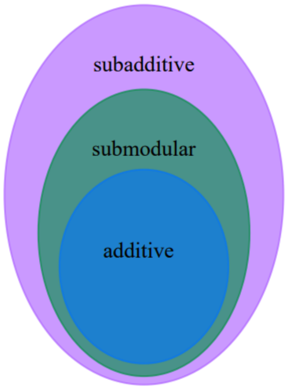

Figure 2.1: Value function hierarchy.

A standard assumption in auction theory regarding bidders' utility function is that of *quasi-linearity*. Thus payoff π_i is defined as the difference between value and price.

$$\pi_i(S, p) = v_i(S) - p_i(S)$$

Equivalently, the payoff of sellers is defined as $\pi_j(Z, p) = p_j(Z) - v_j(Z)$.

Quasilinear functions are used to describe games where utility is transferable. An important observation relative to the underlying topic of this dissertation is that, under the presence of budget constraints, quasilinearity ceases to hold from some price on: if prices exceed budget b^i for a bidder, utility is non transferable. Payoffs are expressed as the difference between value and payment, only when prices are feasible given b^i.

2.2 Mechanism Design Criteria

Mechanism design is a field in economics that focuses on designing good mechanisms and rules in settings where self-interested agents, that possess private information about their preferences, strategically interact. Since mechanism design aims to achieve desirable system-wide outcomes, it is often viewed as the reverse of game theory, focusing on creating rules that guide strategic interactions among agents. Auctions are a particular class of mechanisms where a set of bids are mapped to an outcome.

In the realm of mechanism design, the goal is to construct rules that guide participants towards outcomes that are beneficial both on an individual level and for the market as a whole. This balancing involves ensuring that mechanisms are not only efficient and fair but also robust against strategic manipulation by self-interested parties. The selection criteria discussed in this section, including Pareto optimality, budget balance, and various forms of incentive compatibility, serve as essential building blocks in this process. However, the key difficulty lies in managing the trade-offs required to satisfy these criteria simultaneously. For instance, requiring budget balance might restrict the mechanism's ability to yield optimal allocations, whereas prioritising allocative efficiency could encourage strategic bidding, compromising the auction's integrity. Furthermore, the design of incentive compatible mechanisms often includes complex rules that can be impractical in real-world applications, where simplicity and transparency are crucial to assure participants' trust and engagement. Thus, the discussion in this section not only highlights the theoretical ideals mechanism designers aim for but also underscores the considerations that must be balanced in developing mechanisms tailored for the dynamics of actual markets.

Approaching mechanism design from the perspective of auctions, the goal is defined by a *social choice function*, that, given agents' bids, selects an optimal outcome. A social choice function aggregates individual information among a group of participants to reach a collective decision, essentially translating individuals' preferences into a single choice that reflects the welfare of the group.

Definition 2.1 (SOCIAL CHOICE FUNCTION)**.** *A social choice function* $f : b_1 \times \cdots \times b_n \to O$ *receives the set of all bid functions* $b_i : 2^{\mathcal{K}} \to \mathbb{R}_{\geq 0}$ *for each bidder* $i \in \mathcal{I}$ *as input and selects an outcome among all possible ones in* O*.*

An auction mechanism M is said to *implement* a social choice function f that results in outcome o and can be described as a method to select an outcome based on agents preferences as reported.

To highlight the underlying difficulty of designing good mechanisms, we address the problem of deciding on an appropriate solution of the social choice function given agents' preference reports. The resulting outcome is selected based on these reported preferences. However, as agents are payoff-maximizing, the optimal choice of reporting strategy need not necessarily be equal to their true preference: they might profit from misreporting by placing lower bids and thus paying less under the same allocation rule. The majority of literature in mechanism design assumes *utilitarian social welfare* functions, where social welfare can be described as the sum of agents' valuations. In the case where agents report their true preferences, an auction may indeed reach the outcome that maximizes social welfare. Such an allocation is considered *efficient*.

We proceed by outlining important properties of social choice functions and mechanisms that lead to desirable outcomes.

Pareto optimality (or Pareto efficiency) characterizes an outcome of a social choice function where no alternative outcome increases one participant's payoff without reducing that of another. Formally:

Definition 2.2 (PARETO OPTIMALITY)**.** *The outcome* o *of a social choice function* f *is (weakly) Pareto optimal if there is no other outcome* $o' \in O$ *such that:*

- $\pi_i(o') > \pi_i(o)$ *for all* $i \in \mathcal{I}$*, and*

- $\pi_j(o') > \pi_j(o)$ *for some* $j \in \mathcal{I}$

The Pareto efficiency criterion does not necessarily imply an equitable distribution of resources, but focuses on maximizing the efficiency of resource allocation. A fundamental benchmark in evaluating economic policies, mechanisms, and market outcomes, it serves as a guide to understand whether resources in a system are allocated in the most effective way possible, without explicitly considering the fairness of such allocations.

Since a mechanism implementing a social choice function decides on an allocation of items/ bundles to agents, a measure of quality of the resulting allocation is *allocative efficiency*. It represents an economic state that optimizes the allocation of resources in order to maximize the aggregate valuation across all market participants. This optimization criterion is fundamental in welfare economics for evaluating the efficacy of market mechanisms. Deciding on an efficient allocation in combinatorial auctions given a set of valuations is typically done by solving the Winner Determination Problem (WDP), introduced in Section 2.3.

Definition 2.3 (ALLOCATIVE EFFICIENCY)**.** *Assume* $x = (S_1, \ldots, S_n)$ *is an allocation maximizing social welfare and* $x' = (S'_1, \ldots, S'_n)$ *is the resulting allocation of the auction mechanism. Allocative efficiency of allocation* x' *is the metric defined by:*

$$\frac{\sum_{i=1}^{n} v_i(S'_i)}{\sum_{i=1}^{n} v_i(S_i)}$$

Allocative efficiency measures the quality of distribution of resources within an economy, compared to the optimal allocation of goods that maximize overall social welfare based on consumer preferences. Transitioning from this concept, *budget balance* takes a complementary perspective by emphasizing the financial sustainability of economic mechanisms, ensuring that the total inflows match the outflows to maintain the system's equilibrium without generating surplus or deficit. This property helps ensure fairness by preventing financial imbalances, ensuring no net transfers into or out of the system, as such imbalances could deter buyer participation in the market.

Definition 2.4 (BUDGET BALANCE)**.** *A mechanism is budget balanced if the total payments of buyers equals the amount received by sellers.*

An efficient mechanisms that simultaneously satisfies budget balance is considered to also be Pareto optimal. Pareto optimality is crucial in mechanism design as it guarantees that resources are allocated in a way that maximizes the overall efficiency of the system without disadvantaging any participant, promoting equitable and efficient outcomes in economic interactions.

Since the outcome of a mechanism is decided based on reports of agents' private information, participants may engage in strategic bidding behaviors in order to increase their payoff and steer the mechanism towards a more favorable result. Misreporting preferences may be observed in various ways: declaring lower bids to minimize payments, placing additional bids to increase competitors' prices etc. Incentives in a mechanism should be designed in such a manner that dissuades agents from seeking individual profit by falsely declaring their private information. This strategic interaction, where individuals act based on private information and the asymmetry of knowledge, leads to the concept of a *Nash equilibrium* - a state where no participant can improve their payoff by unilaterally changing, assuming others' strategies remain unchanged. Extending this concept to the domain of Bayesian games, where players have incomplete information but form beliefs about unknown factors, a *Bayes Nash Equilibrium (BNE)* represents as a strategy profile in which players, guided by their beliefs and the strategies of others, select strategies that maximize their expected payoff. This equilibrium concept is particularly relevant in settings where strategic misreporting of private information and the resulting impact on mechanism outcomes are considered.

Definition 2.5 (BAYES NASH INCENTIVE COMPATIBILITY)**.** *A mechanism is Bayes Nash incentive compatible if truthful reporting is a Bayes Nash equilibrium in the game produced by the mechanism.*

In mechanisms that adhere to the principle of Bayes Nash incentive compatibility, agents achieve the highest expected utility by truthfully reporting their preferences.

A primary objective in the design of such mechanisms is to deter strategic bidding by creating systems that are *strategyproof* (also known as *dominant strategy incentive compatible*). This means that participants are disincentivized from misrepresenting their private information. This concept represents a robust form of incentive compatibility, ensuring that truthfulness is not only the best policy for maximizing utility but also the most rational course of action for all agents involved.

Definition 2.6 (STRATEGYPROOFNESS)**.** *A mechanism is strategyproof if truthful reporting of valuations is the optimal strategy for each bidder, namely, for all bidders $i \in \mathcal{I}$:*

$$\pi_i(f(v_i, v_{-i})) \geq \pi_i(f(v_i', v_{-i}))$$

where v_i corresponds to a truthful report for bidder i, v_i' is an alternative report, v_{-i} the reports of all other agents excluding i and $\pi_i(f(v))$ the utility of the bidder under outcome $f(v) \in O$ of the social choice function.

Strategyproofness is a fundamental property in mechanisms, enabling agents to determine their optimal strategies without needing to know the strategies and valuations of others, thereby minimizing computational costs.

This principle is not limited to individual participants: it expands to include groups through the concept of *group-strategyproofness*. This broader principle ensures that no group can gain an advantage by jointly misrepresenting their preferences. This aspect is especially important because it guards against collusion in auctions, preventing bidders from collaborating in a way that would result in a higher collective payoff and reduced contributions.

Definition 2.7 (GROUP-STRATEGYPROOFNESS). *A mechanism is group-strategyproof if for all groups of bidders $G \subset \mathcal{I}$, there is no v'_G such that all agents in G can improve their payoff by misreporting. Formally, for all $G \subset \mathcal{I}$, there exists $i \in \mathcal{I}$ such that:*

$$\pi_i(f(v_i, v_{-i})) \geq \pi_i(f(v'_G, v_{-G}))$$

where v_i corresponds to a truthful report for bidder i, v'_G is an alternative report by group G, v_{-G} the reports of all other agents excluding G and $\pi_i(f(v))$ the utility of the bidder under outcome $f(v) \in O$ of the social choice function.

Furthermore, it is essential for designers to create conditions where individuals are motivated to willingly engage in a mechanism. This necessity is captured in the concept of *individual rationality*, which dictates that the outcomes for agents should be at least as favorable as any alternative outside option, thereby ensuring that participants do not incur loss. This concept is important in ensuring that every agent's involvement is not only strategic but also beneficial, aligning with the broader goals of fair and efficient mechanism design.

Definition 2.8 (INDIVIDUAL RATIONALITY). *A mechanism is individually rational if for any bid profile $b = (b_1, \ldots, b_n)$, the outcome of the mechanism is such that no bidder pays more than their bid $b_i(S_i) \geq p_i(S_i)$ for their allocated bundle S_i for all $i \in \mathcal{I}$.*

Individually rational mechanisms encourage participation, since agents value expected gains from partaking higher than abstaining, given their beliefs on others' preferences.

The previously outlined properties summarize critical goals that mechanism designers strive to achieve. These objectives include ensuring strategyproofness to prevent manipulation through strategic bidding, extending strategyproofness to group contexts to avoid collusion, and maintaining individual rationality to guarantee that participation is voluntarily beneficial based on expected outcomes. However, literature in the field indicates a significant challenge: it is not always possible to develop mechanisms that simultaneously fulfill these properties. This limitation stems from the inherent complexities and trade-offs involved in designing systems that are both efficient and fair under varying conditions and assumptions.

The challenge of reconciling these properties is well-documented in economic theory and game theory literature. A series of impossibility theorems (Gibbard, 1973; Satterthwaite, 1975; Hurwicz et al., 1975; Green and Laffont, 1977; Myerson and Satterthwaite, 1983) show that for different categories of preferences (restricted or general), one has to give up on a certain subset of properties in order to design feasible mechanisms.

Further exploration into mechanism design has led to the development of concepts such as approximate mechanism design, discussed in Chapter 2.8 where the focus shifts to achieving as close as possible to the desired properties under practical constraints. This line of work suggests that while the perfect mechanism may be unattainable, significant advancements can be made in understanding the balance between efficiency, fairness, and incentive compatibility.

In summary, while the goals of mechanism design are clear and well-defined, the existing stream of literature acknowledges the complex interplay between these objectives, indicating that the search for mechanisms that simultaneously satisfy all desired properties remains an ongoing and challenging area of research. A key objective of this dissertation is to contribute to the understanding of the interplay between incentive compatibility, budget feasibility, and efficiency in the context of auctions with budget constraints.

2.3 Winner Determination Problem

Building on the foundational criteria for mechanism design, the subsequent discussion delves into the practical application of these principles within the context of auction design, focusing particularly on the challenge of allocation. The underlying question is thus how to efficiently distribute goods or resources among a group of participants. A fundamental objective in auctions is the maximization of social welfare, quantified by aggregating the utilities that bidders derive from the items or bundles they are allocated in the final assignment, in addition to the utilities sellers obtain from the transactions. This aggregation assumes a state of budget balance, implying that the payments of bidders are directly reflective of the values accrued by sellers, thereby cancelling out these financial exchanges in the welfare calculation. Consequently, welfare is computed solely based on the valuations of assigned bundles. Within this framework, bidder truthfulness becomes critical; it is assumed that bids accurately reflect the actual valuations of bidders, i.e., $b_i(S) = v_i(S)$. This transition from theoretical principles to their application in addressing the allocation problem highlights the importance of truthful reporting in achieving optimal social welfare outcomes.

Finding an allocation that maximizes social welfare, coined as the *Winner Determination Problem (WDP)* in literature, is described as a combinatorial optimization problem in the following manner:

$$
\begin{aligned}
\text{maximize} \quad & \sum_{i \in \mathcal{I}} \sum_{S \subseteq \mathcal{K}} v_i(S) x_i(S) & \\
\text{subject to} \quad & \sum_{i \in \mathcal{I}} \sum_{k \in S : S \subseteq \mathcal{K}} x_i(S) \leq 1 & \forall\, k \in \mathcal{K} \\
& \sum_{S \subseteq \mathcal{K}} x_i(S) \leq 1 & \forall\, i \in \mathcal{I} \\
& x_i(S) \in \{0, 1\} & \forall\, i \in \mathcal{I},\ S \subseteq \mathcal{K}
\end{aligned}
$$

Adopting the above formulation, an efficient outcome can be derived as the optimal solution to the WDP. A variable $x_i(S)$ equals 1 if bundle S is assigned to bidder i in the final allocation. The first constraint ensures that each item in a bundle is allocated at most once. The second constraint places a restriction on the amount of bundles that each bidder can obtain: at most one. This constraint expresses the XOR bidding language: bidders submit multiple bids on bundles but may only win one. While this is a common assumption under certain bid languages, this constraint might be omitted allowing for expressing more general preferences. Such an example is the OR bidding language, where bidders may win any combination of bundles if this leads to an increase in utility. The last constraint reflects the indivisibility of items in the auction.

Lehmann et al. (2005) conclude that no general purpose algorithm exists for efficiently (i.e. in polynomial time) solving the WDP for every problem instance, due to the exponential size of the representation in the number of items. Hence, the problem is classified as complete for the complexity class NP, defined in detail in Section 2.7. Circumventing the negative complexity result, they provide approximation algorithms as well as identify conditions that lead to polynomial time solutions. Sandholm (2002) provides a comprehensive overview of algorithms designed to solve the WDP efficiently, exploring both exact and heuristic solutions and highlighting the trade-offs involved.

2.4 Competitive Equilibria and the Core

In this section, the discussion transitions to the broader economic framework of competitive markets. A central challenge in such markets lies in determining a balance between the quantities of goods demanded and supplied in the economy. This state is captured by a *competitive equilibrium*: an outcome where demand equals supply, ensuring market efficiency and participant satisfaction. In the idealistic model of perfect competition, Arrow and Debreu (1954) prove that, when agents' preferences are convex and there is a universal demand potential for all goods, every exchange economy admits a competitive equilibrium.

In the landscape of auctions, especially under the presence of financial constraints, the *demand set* of a bidder becomes a focal point. Defined as the collection of bundles that maximize a bidder's utility for a given vector of prices, the demand set is critical in constructing auction mechanisms that navigate towards competitive equilibria. The significance of this notion is further magnified when considering auctions within financially constrained environments, where the alignment of bidder preferences with feasible budgetary outcomes is crucial for sustaining market equilibrium. Following, we provide a formal definition of the concept:

Definition 2.9 (DEMAND SET)**.** *For bidder i with utility function π_i and price function p_i, the demand set includes all bundles that maximize utility and is defined as:*

$$D_i(p_i) = \{S \mid \pi_i(S, p_i) \geq \max_{T \subseteq 2^{\mathcal{K}}} \pi_i(T, p_i),\ \pi_i(S, p_i) \geq 0,\ S \subseteq 2^{\mathcal{K}}\}$$

A bundle $S \subseteq 2^{\mathcal{K}}$ is demanded at prices p if $T \in D_i(p_i)$ for at least one bidder i. An outcome where each bidder is assigned a bundle from their demand set is considered *envy-free*. Once all items have been allocated to bidders based on their demand sets, and the market is cleared, the resulting outcome corresponds to a *competitive equilibrium*.

Definition 2.10 (COMPETITIVE EQUILIBRIUM). *An allocation $x = (S_1, \ldots, S_n)$ and a set of nonnegative prices $p = (p_1, \ldots, p_n)$ are a competitive equilibrium if $S_i \in D_i(p_i)$ and $\cup_{i \in \mathcal{I}} S_i = \mathcal{K}$.*

In addition to market clearing, a competitive equilibrium defines the prices where preferences of consumers match costs of producers. From a state of competitive equilibrium, no agent has incentive to deviate and resources are allocated in an efficient manner. Necessary conditions for an outcome to be a competitive equilibrium are envy-freeness, individual rationality as well as budget balance.

The notion of competitive equilibrium is applicable in the general setting where prices can be personalized, as the price function of each bidder is defined as $p_i : 2^{\mathcal{K}} \rightarrow \mathbb{R}_{\geq 0}$ and dictates the price to be paid for each bundle. Assuming non-linear and personalized prices, a competitive equilibrium is maximizing social welfare in a market (Bikhchandani and Ostroy, 2002).

An equivalent definition has been introduced by Walras (1874) for markets where prices are linear and anonymous, and therefore guarantee fairness among participants. In this paradigm, an auctioneer adjusts prices in such a way that all bidders witness the exact same price for the same bundle, and the bundle price is determined by simply summing all items within it, i.e. $p(S) = \sum_{k \in S} p(k)$. This design concept is defined as a *Walrasian equilibrium*, and corresponds to the market clearing point under linear and anonymous pricing rules.

Revolving around the notion of Walrasian equilibria are the two fundamental theorems of welfare economics (Arrow and Debreu, 1954; McKenzie, 1959).

Theorem 2.1 (THE FIRST WELFARE THEOREM). *Let allocation $x = S_1, \ldots, S_n$ and prices p constitute a Walrasian equilibrium. Then allocation x maximizes welfare over all feasible allocations.*

Blumrosen and Nisan (2007) introduce a stronger version of the aforementioned theorem, where allocations need not necessarily be fractional, but can be integral, capturing the case of indivisible items.

Theorem 2.2 (The second welfare theorem)**.** *Let allocation $x = S_1, \ldots, S_n$ be Pareto efficient. Then there exist a vector of prices p such that the outcome defined by x and p defines a Walrasian equilibrium.*

Due to their attractive properties, Walrasian equilibria have been a prominent solution concept in combinatorial auctions. A long line of research has been involved with determining conditions that guarantee existence of Walrasian equilibria in auctions (Bikhchandani and Mamer, 1997; Fujishige and Yang, 2003; Baldwin and Klemperer, 2019a; Leme, 2017; Shioura and Tamura, 2015). For the restricted case of assignment markets, with quasilinear utilities and unit-demand preferences, Shapley and Shubik (1971) proved existence of a Walrasian equilibria, as well as uniqueness of minimum Walrasian prices.

Another measure of stability of an allocation is the *core*, that is extensively examined in the course of this dissertation, and which can be defined as the set of feasible allocations that cannot be improved upon by any coalition of bidders with sellers or the auctioneer. A *core outcome* $o = (x, p)$ is one where no coalition of bidders can strictly increase payoff in an alternative allocation x' by proposing a counter offer to the seller or auctioneer. The notion of the core is better understood under the lens of a *coalitional* game, and a concept in the definition is that of *coalitional value*, which encompasses the objective value of the WDP for a subset M of bidders.

Definition 2.11 (Coalitional value)**.** *Let $\mathcal{I}_0 = \mathcal{I} \cup \{0\}$ denote the set of buyers including the auctioneer as the 0-th agent, and X_C the set of feasible allocations including subset $C \subseteq \mathcal{I}_0$. For coalition C, the coalitional value is defined as:*

$$w(C) = \begin{cases} \max_{x \in X_C} \sum_{i \in C} \pi_i(S_i) & \text{if } 0 \in C \\ 0 & \text{otherwise} \end{cases}$$

The coalitional value corresponds to the maximal utility that the set of coalition agents can generate working with the auctioneer, and is equal to 0 if the auctioneer does not participate in the coalition as there are no items remaining to be distributed. Two necessary conditions for an outcome $o = (x, p)$ to be *core stable* are the following: first, the payoff of any coalition $C \subseteq \mathcal{I}_0$ under allocation x must be at least as high as the social welfare that the agents in C can generate on their own. The second condition states that the sum of all payoffs under allocation x cannot exceed the social welfare $w(\mathcal{I}_0)$ that can be generated by the combination of all agents.

Definition 2.12 (Core)**.** *A core payoff vector π can be described as:*

$$Core(\mathcal{I}_0, w) = \left\{ \pi \in \mathbb{R}^{n+1}_{\geq 0} : \pi \geq w(C) \text{ and } \sum_{i \in \mathcal{I}_0} \pi = w(\mathcal{I}_0) \text{ for all } C \subseteq \mathcal{I}_0 \right\}$$

In a core outcome $o = (x, p)$ we refer to x as a core allocation and p as core prices. The concept of the core is inherently connected to stability and fairness in an allocation, even in the face of strategic behavior by participants, and therefore outcomes outside of the core are considered impractical, as they may lead to lower auctioneer revenue or be prone to collusion (Day and Raghavan, 2007; Day and Milgrom, 2008).

The seminal work of Shapley and Shubik (1971) pioneered the study of core stability in the context of assignment markets, a simple market model introduced in the next section of this dissertation. Reaching core outcomes is an important objective in various auction designs, including multi-item (Demange et al., 1986) and spectrum auctions (Ausubel and Milgrom, 2002; Ausubel et al., 2004).

In the context of this dissertation, core stability emerges as a critical objective in the design of auction mechanisms under financial constraints, addressing the complexities that arise when participants face budget limits.

2.5 Assignment Markets

Having delved into the complexities of achieving optimal allocations and maintaining market stability under the constraints of financial limitations and strategic behaviors, we explore the model of assignment markets, to illustrate a specific, yet profound, application of these principles. This model will be the key market under consideration in the publication presented in Chapter 3.

Assignment markets impose a strong restriction on preferences: each bidder is only interested in winning at most one single item, an assumption known as unit-demand. This market model simplifies the challenge of allocating bundles of items to bidders by focusing on allocating single items. Computing an optimal allocation via the WDP can thus be formulated as an assignment problem. For this, efficient ascending auction schemes exist, with truthful bidding being an ex-post equilibrium and that are described using the duality theory of linear programming. Primal-dual algorithms are a particularly strong tool implementing such auctions, since they do not require prior knowledge of private values of bidders. Starting from these simple ascending models, research has been extended to the more complex, combinatorial auction setting, allowing for ex-post truthfulness.

The assignment problem can be understood as the problem of finding the maximum weight matching in a weighted bipartite graph. A graph is weighted if the edges connecting its vertices have weights (values) and is considered bipartite if it consists of two distinct sets of nodes (bidders and items) where edges connect nodes from one set to the other, representing potential matches or allocations. In assignment markets, this translates to matching a set of bidders to a set of items in a way that maximizes overall welfare or efficiency. The challenge lies in finding an allocation that optimally aligns the preferences of all participants. Bidders may place bids on multiple items but are only allowed to win at most one. A matching M is a subset of edges such that no two edges in M share a common vertex.

Modifying notation, we denote as $v_{i,k}$ and $x_{i,k}$ the value and assignment of bidder i for item k respectively, to differentiate from the combinatorial setting where, previously, we used $v_i(S),\ x_i(S)$ to define the value and assignment for bundle S. An illustrative example of an assignment market with bidders and items defining a bipartite graph is introduced in Figure 2.2, where the edges represent values of bidders for items and the highlighted ones constitute the final allocation.

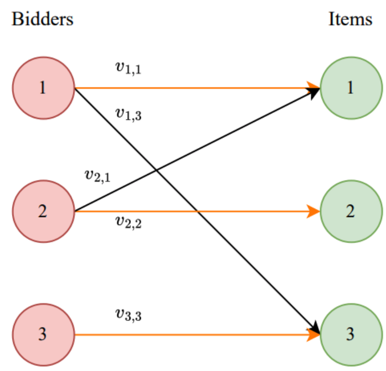

Figure 2.2: A simple assignment market consisting of three bidders and three items.

Casting the allocation problem as an integer program as follows:

$$
\begin{array}{llll}
\text{maximize} & \sum_{i\in\mathcal{I}}\sum_{k\in\mathcal{K}} v_{i,k}x_{i,k} & & \\
\text{subject to} & \sum_{i\in\mathcal{I}} x_{i,k} = 1 & \forall\, k\in\mathcal{K} & (p(k)) \\
& \sum_{k\in\mathcal{K}} x_{i,k} = 1 & \forall\, i\in\mathcal{I} & (\pi_i) \\
& x_{i,k}\in\{0,1\} & \forall\, i\in\mathcal{I},\ k\in\mathcal{K} &
\end{array}
$$

Similar to the combinatorial version of the WDP introduced in Section 2.3, the first constraint requires that every item is assigned to exactly one bidder, while the second ensures that each bidder is only assigned one item. The variables $x_{i,k}$ can take binary values 0 or 1, with 1 corresponding to bidder i being allocated item k, and 0 otherwise. On the right, we define the dual variables $p(k)$, corresponding to the item prices, and π_i representing bidder utilities. Since the constraint matrix in the assignment problem is totally unimodular, its linear relaxation always admits an integral optimal solution, which can be determined in polynomial runtime.

The first analysis of assignment markets by Shapley and Shubik (1971) concluded that the weak core is non-empty and unique minimum competitive equilibrium prices exist and are derived as the set of solutions of the dual linear program. Computationally efficient algorithms to derive this set are known under the assumptions that bidders submit sealed bids. The resulting prices correspond to VCG prices, and therefore guarantee incentive compatibility (Leonard, 1983). Applying the first welfare theorem in the context of assignment markets guarantees that a Walrasian equilibrium maximizes social welfare.

Based on the well-known Hungarian algorithm of Kuhn (1955), Demange et al. (1986) introduced a novel ascending auction algorithm that efficiently finds a competitive equilibrium. Their algorithm is based on the primal-dual design principle, and a central notion is that of the demand set of a bidder, as only demands are reported in each round. Each bidder announces their most desired item and an initial check is performed regarding whether each item k can be allocated to a bidder i that demands it. If the answer is positive, the allocation x and prices p correspond to a competitive equilibrium. Otherwise, if no such assignment exists, a set of items $\mathcal{O}$ is overdemanded, namely the number of bidders with demand only in this set is larger than its cardinality. By definition, a competitive equilibrium must have no excess demand. The algorithm determines a minimally overdemanded set: an overdemanded set $O \in \mathcal{O}$ that has no proper subset of items in $\mathcal{O}$. For all items in the chosen set O, prices are raised. Bidders report their demand sets based on the updated prices, and the iterative procedure continues until overdemand is resolved, at which point a competitive equilibrium has been reached. Given the nature of the valuations and the ascending prices, the algorithm terminates in a finite number of steps. The proposed mechanism is incentive compatible and robust against collusion. It is implementable in polynomial time and can be seen as an ascending implementation of the celebrated VCG auction mechanism.

When the unit-demand assumption in assignment markets is removed, allowing buyers to demand more than one unit, the beneficial traits typically seen in these markets cease to exist. The absence of the unit-demand constraint means buyers can seek multiple units, complicating the balance of supply and demand. This shift can result in challenges in determining efficient and stable outcomes in the market, as the straightforward dynamics associated with single-unit demand no longer apply. The section that follows discusses complex auction formats and their resulting complexities.

To mitigate these complexities, the first publication in this dissertation investigates the impact of financial constraints in the simplistic framework of assignment markets. Examining auction mechanisms that can achieve efficient allocations in the presence of budget constraints, Talman and Yang (2011) developed a dynamic auction that always results in a core allocation while van der Laan and Yang (2016) introduce an ascending auction resulting in a rationed equilibrium. Our ascending auction algorithm (Batziou et al., 2022a) always yields a core stable allocation and, under specific conditions, achieves welfare maximization. Approaching the issue from a computational complexity standpoint, we prove that the task of determining outcomes that simultaneously satisfy core stability and welfare maximization criteria is NP-hard, indicating computational infeasibility. In a predecessor of our complexity result, Bichler and Waldherr (2022), classified the equivalent problem in combinatorial auctions as complete for the complexity class Σ_p^2, a class that encompasses a broader and more complex set of problems than NP. This analysis of assignment markets, drawing from the hardness of the combinatorial setting, particularly in the context of financial constrained settings, underlines the model's utility in exploring market outcomes and attributes that are desirable within constrained environments.

2.6 Auction Formats

Auctions are mechanisms that model transactions between a set of buyers that place bids and compete for a set of items or services, and an auctioneer or a set of sellers that provide offers. Bids typically have the form of monetary payments, expressing an agent's willingness to pay for an item or bundle. Use cases of auctions range from art houses to real estate, electricity and environmental conservation. Common objectives used in auctions are maximization of allocative efficiency (*efficient auctions*) or auctioneer revenue (*optimal auctions*).

For the purpose of simplicity in the analysis, we will focus on presenting auction formats for the case of a *single item* - where only one item is available in the market and multiple participants place bids expressing their preferences. The bidding and pricing mechanisms presented can naturally be extended to multiple items.

2.6.1 Open Auctions

In *open auctions*, participants publicly announce bid prices, in a competitive setting. Each bidder has knowledge of the current highest bid and, accounting for competition, can adjust bids. Since the auction takes place in real-time, it allows for dynamic interactions between participants, and the process terminates after a specified amount of time, matching buyers to items. The most notable example is the *English auction*, where an auctioneer, starting from the reserve price, receives ascending bids. Participants start with a low opening bid and progressively increase in each bidding round, until all but one bidders drop out of the process. The highest bidder is pronounced as winner of the auction and due to its public nature, the process is transparent. English auctions have been applied in a multitude of settings, such as art auctions, fundraising events or online platforms such as eBay.

While the English auction is an attractive, dynamic format for single-item settings, scalability becomes an issue when selling multiple items. Due to the need for coordination and active participation requirement, it can be considered a less attractive format for more complex settings. In a *clock auction*, participants observe a clock price, which is gradually increased in each round, and assess their utility for the current price. If utility is positive, meaning that value exceeds clock price, bidders have incentive to remain in the auction for the next round. Unlike the English auction, it does not allow for jump bids, where a bidder might arbitrarily increase a bid between consecutive rounds. Instead, the bidding power is limited - a bidder either has positive utility and remains in the auction, or decides to permanently exit. The price is therefore dictated by the auctioneer, and the remaining, winning bidder must submit the amount indicated by the clock price, which would be the exact value of the second highest bidder, as the clock ends when only one bidder is left in the auction. Clock auctions are strategyproof, since it is in the best interest of participants to remain in the auction, while this remains a rational choice. Perhaps the most notable application of ascending clock auctions is the use for spectrum allocation by the Federal Communications Commission of the United States (Ausubel et al., 2004).

Depending on the auction format, clock prices may increase or decrease. The latter case describes *descending* auctions, where a high asking price is broadcast to the participants. Since no bidder can possibly observe a positive payoff under this price, the prices are systematically reduced in each stage, until one bidder signals that the current offer is accepted. *Dutch auctions*, originally designed for the sale of flowers and other perishable goods, use descending prices to explore participants' willingness to pay fast. The bidder that first announces their acceptance of the clock price is allocated the item.

2.6.2 Sealed-Bid Auctions

In contrast to open formats, where prices are publicly announced and known among all participants, the model of *sealed-bid* auctions promotes privacy of bids and preferences. Participants independently place bids in a sealed envelope or online platform, and confidentiality is assured since no bidder has knowledge of others' bids. The highest bidder is declared winner, and payments are dictated according to the rule of choice by the auctioneer.

Arguably the most widely used payment rule is the *first-price*, where the winner's price is set equal to the actual winning bid. An obvious limitation of the first-price rule is preserving strategyproofness: a bidder may strategically misreport preferences based on an estimate of others' private values, and submit lower bids and aiming at maximizing utility by paying less than the full value amount.

Circumventing this shortcoming, Vickrey (Vickrey, 1961) proposed the *second-price* payment rule, based on which the winner's price equals the second highest bid received. This rule limits the influence of an individual bidder in the final price and allocation. The novelty of this algorithm lies in the fact that it guarantees truthful reporting while achieving optimal efficiency. Each bidder's dominant strategy is bidding their exact value, thus a lower bid does not lead to higher payoff. This pricing rule has led to the development of the celebrated *Vickrey-Clarke-Groves (VCG)* auction mechanism (Vickrey, 1961; Clarke, 1971; Groves, 1973), that is proved to be the *unique* strategyproof and efficient mechanism in settings with independent private values (Green and Laffont, 1977; Holmström, 1979). Despite the theoretical guarantees, the VCG mechanism is not extensively used in real-world applications. The algorithm is prone to collusion, since it does not satisfy group-strategyproofness, which renders it undesirable in practice. In a combinatorial setting, VCG payments are computed as the marginal contribution of each agent on the total auction welfare, a computation that requires exponential time in worst-case. An additional practical limitation is that VCG does not account for budgets: payments are computed in a strategyproof manner but may arbitrarily exceed the constraint. Moreover, the resulting allocation is not optimal for the auctioneer, with respect to the goal of maximizing revenue.

In the context of procurement, the auction process is reversed. A set of sellers compete to provide services to a single buyer. Unlike the auction models explored above, in a *reverse auction*, a buyer (government, principal) announces a need for a certain amount of service and sellers compete to win by naming prices. Reverse auctions are common in government procurement, as well as online platforms.

The choice of auction format for a mechanism design problem is a complex task. A multitude of factors need to be accounted for, such as the number of participants, their objectives in the auction, the amount of information available as well as the budget constraints, to name but a few.

In the second publication of this dissertation, we explore the performance of different auction formats, including VCG and clock auction mechanisms, and discuss their guarantees and performance, in the presence of financial constraints. The VCG algorithm, due to the intricate payment rule design, may arbitrarily exceed the given budget, while clock prices offered to participants can be crafted in a manner that respects the constraint. Our work explores the trade-off in terms of the mechanism properties, and the implications on the practical environment of biodiversity conservation.

2.7 Complexity Theory

Given the computational hardness of allocation and pricing rules requiring an exponential amount of information, an important challenge is to understand when a problem becomes tractable and whether exact solutions can be determined in a time-efficient manner. The research field of computational complexity not only examines the theoretical limitations of achieving optimal allocations but also justifies the shift towards approximation mechanisms as a practical response to these computational barriers, a class of mechanisms examined in detail in Section 2.8, providing a framework for achieving approximately optimal solutions in a computationally efficient manner.

Computational complexity theory emerges as a useful tool to analyze the practical implications of market designs. As an example of the relative hardness of computation of an outcome, one may consider the following problems. On the one hand, sorting a set of bids, in a single-item, open auction, and on the other, determining the winning bidder in a combinatorial auction. While the first task can be efficiently and exactly performed by a sorting algorithm of choice, the second requires solving the WDP with an exponential number of variables.

Based on the nature of the solution, problems may be broadly classified into different categories. *Decision problems* model examples that can be expressed as questions on a set of (possibly infinite) input values that receive a 'yes' or 'no' answer. On the other hand, *optimization problems* revolve around the task of determining the best among a set of feasible solutions. An example of a decision version of an allocation problem is framed as 'Is there an allocation that yields revenue higher than $500?', while the equivalent optimization problem would be stated as 'Find a feasible allocation that maximizes revenue'. Computational complexity theory includes techniques for transforming problems between the two versions, and thus is focused primarily on analyzing decision variants.

Initiated in the 1960s, computational complexity theory encompasses the study of cost and resources required to solve a computational problem. The seminal work of Hartmanis and Stearns (1965) introduced the *multi-tape Turing machine*, which, to this day, represents the standard model of efficient computation. The model of a Turing machine is a central concept in complexity theory, driven by both historical and practical reasons. Its simplicity facilitates precise definitions and accurately models the behavior of real-world computers. Turing machines, functioning as finite state machines, possess the ability to access and modify symbols on tapes based on a set of rules. Tapes consist of infinite strings over a specified alphabet, and a movable tape head scans symbols within tape cells while moving left and right. In a *deterministic* Turing Machine, the transition function, given a state and symbol scanned by the tape head, dictates the symbols written to the tape, the direction of the tape head's movement, and the subsequent state.

Non-deterministic Turing Machines, in contrast to their deterministic counterparts, operate as non-deterministic finite automata. They constitute a purely theoretical concept, as they cannot be implemented as a computing device. The key distinction lies in the fact that, for a given combination of tape symbol and state, the forthcoming state and transition are not uniquely determined. Instead, there exists a path (or tree) of possibilities.

A *computational complexity class* is defined as a set encompassing all problems requiring the same amount of a specific computational resource to reach a solution. The primary criteria for classifying computational problems are time and space requirements for a Turing Machine to compute a solution. Formally, $time_M(x)$ represents the number of computation steps that a Turing Machine M performs until halting on input x, while $space_M(x)$ denotes the number of cells required until the machine halts. Unlike the standard model of a Turing machine with an unbounded tape, *bounded Turing machines* are subject to constraints on the tape space that can be used for computation and are thus limited in sizes.

Definition 2.13 (BOUNDED TURING MACHINES)**.** *Let $T : \mathbb{N} \to \mathbb{N}$ describe a function with $T(n) \geq n$. The machine M is time-bounded if for any input x on the tape alphabet, it holds that $time_M(x) \leq T(x)$. An equivalent definition applies to space-bounded Turing machines.*

Worst-case analysis involves determining an upper bound on time or space required for an algorithm to solve a problem for the most pessimistic input scenario. The *big-O notation* $O()$, the most prominent one in this context, indicates an asymptotic upper bound on the growth rate of a function, up to a constant factor, and is used to classify algorithms based on their runtime.

Definition 2.14 (BIG-O NOTATION). *Assume function $f(x)$ receiving input x describing the time complexity of an algorithm as well as comparison function g. Then $f(x) = O(g(x))$ if there exists an input size x_0 and a number $M \in \mathbb{R}_{\geq 0}$ such that $|f(x)| \leq M \cdot g(x)$ for all $x \geq x_0$.*

The formal definition of a deterministic complexity class, described in terms of deterministic Turing machines, is as follows:

Definition 2.15 (COMPLEXITY CLASS). *Let $T : \mathbb{N} \to \mathbb{N}$ describe a function with $T(n) \geq n$. The classes of languages computed by an $O(T(n))$ and $O(S(n))$ deterministic time- or space- bounded Turing machine are respectively denoted as DTIME$(T(n))$ and DSPACE$(S(n))$.*

For the case non-deterministic Turing machines, the same definition applies and the relevant classes are denoted as *NDTIME* and *NDSPACE*.

Arguably the two fundamental classes within the study of computational complexity are the classes P and NP. The class P consists of computational problems admitting efficient algorithms.

Definition 2.16 (CLASS P). *The complexity class P consists of decision problems for which a deterministic Turing machine can achieve a solution within a polynomial number of computation steps. This class comprises problems considered tractable, implying their efficient solvability. Employing the formal definition outlined earlier, the class P is defined as:*

$$P = \bigcup_{k \geq 0} \textit{DTIME}(n^k)$$

The class NP contains decision problems whose answers may be verified efficiently, but a solution cannot be reached in polynomial time.

Definition 2.17 (CLASS NP). *The complexity class NP consists of decision problems for which a non-deterministic Turing machine can achieve a solution within a polynomial number of computation steps. This class comprises problems whose solutions can be verified by a deterministic Turing machine in polynomial time, thus solutions are efficiently checked. Following the formal definition outlined earlier, the class NP is defined as:*

$$NP = \bigcup_{k \geq 0} \textit{NDTIME}(n^k)$$

Deterministic Turing machines are a subclass of non-deterministic, and therefore a natural inclusion is P $\subseteq$ NP. The conjecture P=NP remains one of the biggest open questions in the computer science community.

In complexity theory, a *reduction* is defined as an algorithm that transforms a problem into another. Efficient reductions, such as those executed in polynomial time, suggest an equivalence in difficulty between two problems.

The inaugural problem of the class NP is the well-known *Boolean satisfiability problem* (SAT) (Cook, 1971). In a seminal result, Cook proved that any problem belonging to class NP can be polynomial time reduced by a deterministic Turing machine to the problem of determining satisfiability of a Boolean formula. The notions of hardness and completeness shed light on the hierarchy of difficulty of problems in a class.

Definition 2.18 (NP-HARDNESS). *A problem π is NP-hard if every other problem π' in the class NP can be transformed into π in polynomial time.*

Definition 2.19 (NP-COMPLETENESS). *A problem π is NP-complete if it belongs to the class NP and is NP-hard.*

An NP-hard problem is at least as hard as the hardest problems within class NP, while the class of NP-complete problems includes the hardest problems of the class. Demonstrating NP-hardness is performed via a polynomial time reduction from a known NP-complete problem.

The field of computational complexity provides crucial insights into the limits of what can be computed efficiently, distinguishing tractable problems from those with exponential complexity. This distinction has significant implications for economic systems and algorithmic design, highlighting the critical balance between theoretical exploration and practical application.

2.8 Approximation Mechanisms

Given the intractability of achieving exact solutions for NP-hard problems, researchers turn to approximation algorithms as a pragmatic approach to finding solutions that are close to optimal. As discussed in Section 2.3, the WDP is an NP-complete optimization problem, thus solving the allocation and pricing problems to optimality is intractable. Commercial solvers might require days or weeks to reach an optimal outcome for problem formulations with a large number of variables and constraints. Exploring the classes of problem instances that allow for tractable solutions is a key focus in combinatorial optimization, but an alternative approach involves designing algorithms that approximate the optimal solution as closely as possible while remaining computationally efficient. In contrast to heuristic methods that are commonly employed in combinatorial optimization, approximation algorithms provide worst-case bounds on the solution quality. *Approximation mechanisms* draw on algorithms for NP-hard problems and seek to find near-optimal solutions for markets where admitting exact ones is computationally infeasible.

Typically, approximation algorithms are designed for theoretical problem classes such as knapsack or bin packing, and do not consider incentives. Approximation mechanisms, on the other hand, are designed with incentive compatibility and computational efficiency in mind. In fact, the design goal is to accommodate a large number of the properties presented in Section 2.2, while achieving a close approximation of the optimal solution in polynomial time.

Problems in (algorithmic) mechanisms design can broadly be classified into two categories. The first one involves problems where an optimal mechanism is known but computationally infeasible. Such examples have largely been focused on combinatorial auctions, with the objective of maximizing social welfare among participants (Dobzinski et al., 2012; Holzman et al., 2004; Lavi and Swamy, 2011; Lehmann et al., 2005). However, combining computational tractability with allocative efficiency and incentives for truthful reporting is challenging: the notable VCG auction mechanism becomes quickly intractable, and approximately solving the WDP no longer satisfies the truthfulness property (Nisan and Ronen, 2007). A stream of research addressing this challenge has focused on truthful approximations: designing mechanisms that admit solutions close to the optimal, in a multiplicative manner. The second class contains optimization problems for which no optimal truthful mechanism is known, but that also do not possess any proof of intractability. An important problem in this class is scheduling on unrelated machines, and the goal is to achieve an optimal approximation ratio, dropping the computational feasibility requirement.

Deterministic approximation schemes result in a fixed outcome when provided with the same input, and a notable example is the family of greedy algorithms. Aiming at satisfying truthfulness, a useful allocation rule is the *maximal-in-range (MIR)*, that determines the allocation that maximizes welfare within a restricted set of allocations, the range. Mechanisms that satisfy MIR are strategyproof. Due to the strong requirement of determinism, one cannot expect close approximations achievable while maintaining truthfulness (Lavi et al., 2003; Papadimitriou et al., 2008).

Contrary to the deterministic case, *randomized* approximation schemes produce outcomes in expectation, including an element of stochasticity in their nature. An extension of the MIR principle to the randomized realm is the property of *maximal-in-distributional range (MIDR)*, where the choice of range is replaced by a set of distributions over allocations. The allocation rule decides on the outcome that maximizes expected welfare for the bidders, once the distribution is fixed. Lavi and Swamy (2011) and Dobzinski et al. (2012) introduced a powerful tool for achieving good approximations under randomization: *black box reductions*. Based on this technique, a randomized approximation mechanism can be internally invoked multiple times and achieve truthfulness, while maintaining the same approximation factor.

A restricted assumption on the bidders in combinatorial auctions that is of interest in this dissertation is that of *single-mindedness*. Single-minded bidders are only interested in acquiring a single item or package.

Definition 2.20 (SINGLE-MINDEDNESS). *Let $v_i : 2^{\mathcal{K}} \rightarrow \mathbb{R}_{\geq 0}$ define the valuation function of bidder i. The bidder is considered single-minded if there exists a bundle T and value $v_i(T)$ such that $v_i(S_i) = v_i(T)$ for all $S_i \supseteq T$ and $v_i(S_i) = 0$ for all other bundles S.*

Mechanisms with single-minded bidders fall into the category of *single-parameter* mechanism design, a class that admits simple, strategyproof approximation designs. The following theorem, introduced by Nisan (2007), summarizes necessary conditions for strategyproofness in auctions with single-minded bidders.

Theorem 2.3. *A mechanism $\mathcal{M}$ implementing a social choice function f with single-minded bidders is strategyproof if and only if the following two conditions are met:*

1. *Monotonicity of allocation rule: f is monotone in the valuations, namely if a bidder wins bundle S_i by bidding $v_i(S_i)$, will still remain winner by placing a higher bid $v_i'(S_i) > v_i(S_i)$.*
2. *Critical payments: the price paid by a winning bidder p_i equals the infimum of all values such that bid $v_i(S_i)$ remains winner of S_i.*

Greedy-acceptance algorithms, a class of deterministic approximation mechanisms, under the single-mindedness assumption, satisfy the aforementioned criteria and can be thus characterized as strategyproof (Lehmann et al., 2002). The underlying principle of greedy algorithms lies in making the locally optimal choice in each step, and in the context of auctions, accepting at each stage the candidate optimal bid to the solution. However, in the multi-minded setting, no payment scheme exists that could render greedy mechanisms strategyproof: the unique truthful mechanism is the VCG. The idea of greedy selection is the core of *deferred acceptance auctions (DAAs)*, proposed by Milgrom and Segal (2020), which follows the principle of the *deferred acceptance* algorithm for the one-to-one stable marriage problem of Gale and Shapley (1962). Instead of sequentially accepting, DAAs iteratively reject the lowest value bids in each stage, thus employing a backward greedy technique. These auctions satisfy numerous desirable properties, such as strategyproofness, and can be implemented in a multitude of ways: as sealed-bid or in a clock manner.

Definition 2.21 (DEFERRED-ACCEPTANCE AUCTIONS). *Deferred-acceptance auctions consist of a series of steps t. We define as $\Lambda_t \subseteq \mathcal{I}$ the set of bidders that are active in each step, namely whose bids have not been rejected by the auctioneer, with $\Lambda_1 = \mathcal{I}$. The auction consists of a set of deterministic scoring rules $\sigma_i^{\Lambda_t}(b_i, b_{\mathcal{I}/\ \Lambda_t})$, non-increasing in b_i. At each step t, a decision is made as follows:*

- *If Λ_t is feasible, all bidders in Λ_t are accepted and each receives their critical payment $p_i(b_i) = sup\{b_i' \mid i \in \Lambda(b_i', b_{-i})\}$, where $\Lambda(b_i', b_{-i})$ represents the set of bidders that would have been accepted if the bids reported were b_i' instead of b_i.*
- *If Λ_t is infeasible, remove agents i from the next round, i.e. $\Lambda_{t+1} = \Lambda_t/\ \{i\}$, with $i \in argmax\{\sigma_i^{\Lambda_t}(b_i, b_{\mathcal{I}/\ \Lambda_t})\}$ being the active bidder with the lowest scores.*

Any function of the bids can be expressed as a scoring function: from simple versions returning only the bid of a bidder, to more complex where the ratio of contribution of a bid to a bundle is computed and beyond. A notable result in DAAs for combinatorial auctions with single-minded bidders by Dütting et al. (2017) presents a $O(\sqrt{m \log m})$ approximation bound (with $m = |\mathcal{K}|$), and provide additional results for the closely related family of knapsack auctions.

Resorting to approximation algorithms has long been the goal of mechanism designers concerned with maintaining desirable properties while remaining computationally tractable. Due to the attractive trade-off offered by such schemes, a central research direction focuses in improving approximation bounds, constantly decreasing the optimality gap by novel algorithmic techniques.

3 Contribution 1: Core-Stability in Assignment Markets with Financially Constrained Buyers

Peer-Reviewed Conference Paper

Title: Core-Stability in Assignment Markets with Financially Constrained Buyers

Authors: Eleni Batziou, Martin Bichler, Maximilian Fichtl

In: Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22)

Abstract: We study markets where a set of indivisible items is sold to bidders with unit-demand valuations, subject to a hard budget limit. Without financial constraints and pure quasilinear bidders, this assignment model allows for a simple ascending auction format that maximizes welfare and is incentive-compatible and core-stable. Introducing budget constraints, the ascending auction requires strong additional conditions on the unit-demand preferences to maintain its properties. We show that, without these conditions, we cannot hope for an incentive-compatible and core-stable mechanism. We design an iterative algorithm that depends solely on a trivially verifiable ex-post condition and demand queries, and with appropriate decisions made by an auctioneer, always yields a welfare-maximizing and core-stable outcome. If these conditions do not hold, we cannot hope for incentive-compatibility and computing welfare-maximizing assignments and core-stable prices is hard: Even in the presence of value queries, where bidders reveal their valuations and budgets truthfully, we prove that the problem becomes NP-complete for the assignment market model. The analysis complements complexity results for markets with more complex valuations and shows that even with simple unit-demand bidders the problem becomes intractable. This raises doubts on the efficiency of simple auction designs as they are used in high-stakes markets, where budget constraints typically play a role.

Contribution of dissertation author: Methodology, software, experimental design, investigation, visualization, joint paper management

Citation: Batziou et al. (2022a)

Comment: The article included in this dissertation is the accepted peer-reviewed manuscript presented at the ACM EC '22 conference.

Core-Stability in Assignment Markets with Financially Constrained Buyers

ELENI BATZIOU, MARTIN BICHLER, and MAXIMILIAN FICHTL, Technical University of Munich

CCS Concepts: • **Theory of computation → Computational pricing and auctions**; *Market equilibria*; • **Mathematics of computing** → Integer programming.

Additional Key Words and Phrases: core stability, budget constraints, unit demand, welfare maximization

ACM Reference Format:
Eleni Batziou, Martin Bichler, and Maximilian Fichtl. 2022. Core-Stability in Assignment Markets with Financially Constrained Buyers. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22), July 11–15, 2022, Boulder, CO, USA.* ACM, New York, NY, USA, 29 pages. https://doi.org/10.1145/3490486.3538262

We study markets where a set of indivisible items is sold to bidders with unit-demand valuations, subject to a hard budget limit. Without financial constraints and pure quasilinear bidders, this assignment model allows for a simple ascending auction format that maximizes welfare and is incentive-compatible and core-stable. Introducing budget constraints, the ascending auction requires strong additional conditions on the unit-demand preferences to maintain its properties. We show that, without these conditions, we cannot hope for an incentive-compatible and core-stable mechanism. We design an iterative algorithm that depends solely on a trivially verifiable ex-post condition and demand queries, and with appropriate decisions made by an auctioneer, always yields a welfare-maximizing and core-stable outcome. If these conditions do not hold, we cannot hope for incentive-compatibility and computing welfare-maximizing assignments and core-stable prices is hard: Even in the presence of value queries, where bidders reveal their valuations and budgets truthfully, we prove that the problem becomes NP-complete for the assignment market model. The analysis complements complexity results for markets with more complex valuations and shows that even with simple unit-demand bidders the problem becomes intractable. This raises doubts on the efficiency of simple auction designs as they are used in high-stakes markets, where budget constraints typically play a role.

1 INTRODUCTION

The idea of a market exchange automatically channeling self-interest toward welfare maximizing outcomes is a central theme in neoclassical economics. The initial conjecture of the "invisible hand" goes back to Adam Smith. Formally, the Arrow–Debreu model showed that under convex preferences and perfect competition there must be a set of Walrasian equilibrium prices [Arrow and Debreu, 1954]. In these models, market participants are price-takers, and they sell or buy divisible goods in order to maximize their total value subject to their budget or initial wealth. The more recent stream of research on competitive equilibrium theory assumes indivisible goods and

EC '22, July 11–15, 2022, Boulder, CO, USA.

ACM ISBN 978-1-4503-9150-4/22/07…$15.00
https://doi.org/10.1145/3490486.3538262

quasilinear utility functions, i.e. buyers maximize value minus the price they pay (their payoff) and there are no budget constraints [Baldwin and Klemperer, 2019, Bikhchandani and Mamer, 1997, Gul and Stacchetti, 1999, Kelso and Crawford, 1982, Leme and Wong, 2017]. The underlying question is under which conditions on the preferences markets with indivisible goods can be assumed to be core-stable[1] and welfare-maximizing. This literature focuses on larger markets where bidders are assumed to be price-takers and it emphasizes core-stability over incentive-compatibility.

For two-sided matching markets where quasilinear buyers have unit-demand, referred to as assignment markets, welfare-maximization, core-stability and even incentive-compatibility can be achieved with a polynomial-time auction algorithm [Shapley and Shubik, 1971]. These auctions can be interpreted as primal-dual algorithms, where the auctioneer specifies a price vector (a demand query) in each round and the bidders respond with their demand set, i.e. the set of goods that maximize their payoff for given prices.

Unfortunately, it is well-known that we cannot hope for such positive results with more general quasilinear preferences. Incentive-compatibility and core-stability are conflicting in general markets with quasilinear utilities [Ausubel and Milgrom, 2006]. Even if we give up on incentive-compatibility, only very restricted types of valuations (e.g., substitutes valuations) allow for Walrasian equilibria.[2] As discused in Bikhchandani and Ostroy [2002], under general valuations (allowing for substitutes and complements), competitive equilibrium prices need to be non-linear and personalized and the core can be empty.

Compared to early utility models in general equilibrium theory such as the Fisher markets for divisible items [Eisenberg and Gale, 1959, Orlin, 2010], quasilinear utility functions imply that bidders do not have budget constraints. In many markets, this is too strong an assumption [Che and Gale, 2000, Dobzinski et al., 2008, Dütting et al., 2016]. Bidders might well maximize payoff, but they need to respect budget constraints. Spectrum auctions are just one example, where bidders have general valuations with complements and substitutes and they are typically financially constrained [Bichler and Goeree, 2017]. Incentive-compatible mechanisms are known to be impossible in multi-object markets [Dobzinski et al., 2008]. It is interesting to understand how core-stable and welfare-maximizing prices can be computed in the presence of budget constraints if we assume bidders to be price-takers. This question was recently analyzed for markets that allow for general valuations where the auctioneer has complete information about values and budgets [Bichler and Waldherr, 2022]. The result is a spoiler: Computing the welfare-maximizing core-stable outcome is a Σ_2^p-hard optimization problem. Such problems are considered intractable, even for very small problem instances.

The intuition behind this result is that the allocation and pricing problems cannot be treated independently anymore. With quasilinear utility functions, the auctioneer first determines the welfare-maximizing outcome and then a corresponding price vector. If budget constraints are binding, then these constraints on the prices need to be considered when computing the welfare-maximizing outcome, which transforms the allocation and pricing problem into a bilevel integer program for general valuations. Considering that budget constraints are a reality in many markets, this casts doubts whether simple market institutions based on polynomial-time algorithms (e.g., simple ascending auctions as used in selling spectrum nowadays) can find a welfare-maximizing outcome even if bidders were price-takers. General preferences as allowed in combinatorial exchanges, might be too much to ask for. A natural question is whether we can at least hope for core-stable

[1]The core is the set of feasible outcomes that cannot be improved upon by a subset of the economy's participants.

[2]A Walrasian equilibrium describes a competitive equilibrium where supply equals demand and prices are linear (i.e., there is a price for each good) and anonymous (i.e., the price is the same for all participants and there is no price differentiation) [Baldwin and Klemperer, 2019, Bikhchandani and Mamer, 1997, Leme, 2017].

and welfare-maximizing outcomes in markets where bidders have unit-demand valuations, and seek to maximize payoff subject budget constraints.

1.1 Contributions

We study the properties that can be achieved in assignment markets with unit-demand bidders who aim to maximize their payoff but have hard budget constraints as illustrated in the previous example. The aim of this work is to compute welfare-maximizing and core-stable outcomes in the presence of such financial constraints. If we cannot achieve incentive-compatibility and core-stability with such simple valuations, also markets with more complex preferences will not satisfy these properties.

We first introduce and analyze an iterative process that always finds a core-stable outcome using only demand queries based on prices and no direct access to valuations. In contrast to Aggarwal et al. [2009], where bidders' valuations are directly queried, a distinguishing part of our algorithm is that it relies exclusively on demand queries and provides a natural generalization of the auction by Demange et al. [1986]. Moreover, we place emphasis on the question of when we can expect the outcome of the auction to not only lie in the core, but also maximize welfare among all core allocations, as welfare maximization is typically a central goal in market design. During the auction process, the auctioneer may sometimes have to make decisions on which buyer to exclude from certain items in subsequent rounds. One of our main results is that, for any market instance, if the auctioneer would be able to guess the right decisions throughout the auction, we terminate in a welfare-maximizing core-allocation. In particular, if the auctioneer does not have to make any such decisions - which is trivial to check ex-post - our result implies that the auction always finds a welfare-maximizing core-outcome. Unfortunately, we do not know ex-ante whether the condition holds, and if it does not, welfare can be arbitrarily low.

Now, it is important to understand whether we can hope for an incentive-compatible and welfare-maximizing core-selecting mechanism without additional conditions beyond the already strong restriction to unit-demand valuations. Unfortunately, the answer to this question is negative. In a novel result, we show that no auction mechanism for the assignment market can be incentive-compatible and core-stable when buyers face budget constraints. If we give up on incentive-compatibility and assume full access to the true valuations (i.e., via value queries) and buyer budgets, we can compute a core-stable and welfare-maximizing outcome.

One might expect that the problem admits a polynomial time solution, since, without the presence of budget constraints, the problem lies in complexity class P. Unfortunately, a main finding of this paper shows that determining core-stable, welfare-maximizing outcomes with financially constrained buyers is an NP-complete optimization problem, even for the assignment market with full access to valuations and budgets. This means, the existence of budget constraints renders the problem of determining welfare-maximizing, core-stable outcomes NP-hard. The hardness proof requires an involved reduction from the maximum independent set problem. One aspect that is making the reduction difficult is that prices need to be considered as continuous variables. These results show that, even for the simplest type of multi-object markets, those with only unit-demand bidders, we cannot expect core-stable and welfare-maximizing outcomes unless additional strong conditions are satisfied that are typically unknown ex-ante.

2 RELATED LITERATURE

Two-sided matching markets describe markets where buyers want to win at most one item (also known as the unit-demand model) and sellers sell only one item. Buyers and sellers are disjoint sets of agents and each buyer forms exclusive relationships with a seller. Such markets are central to the economic sciences. The well-known marriage model of Gale and Shapley [1962] assumes ordinal preferences and non-transferable utility. Shapley and Shubik [1971] analyzed such markets with quasilinear utility functions and showed that the core of this game is nonempty and encompasses all competitive equilibria. Under the quasilinear utility model, buyers maximize value minus price, while sellers maximize price minus cost. While their setting assumes access to all valuations, Demange et al. [1986] showed that an ascending auction with only demand queries results in a competitive equilibrium at the lowest possible price, i.e. at the competitive equilibrium price vector that is optimal for buyers. In such an auction, the auctioneer specifies a price vector (the demand query) in each round, and buyers respond with their demand set, i.e. the set of goods that maximize payoff at the prices.

The housing market of Shapley and Scarf [1974] is an example of a market without transferable utility or monetary funds. In this market, each agent is endowed with a good or house, and each agent is interested in one house only. The goal of this market is to redistribute ownership of the houses in accordance with the ordinal preferences of the agents. In such housing markets, the core set is nonempty. If no agent is indifferent between any two houses, then the economy has a unique competitive allocation, which is also the unique strict core allocation. An allocation belongs to the *strong* core, if no coalition of buyers and sellers can make all members as well off and at least one member better off by trading items among themselves. We assume an allocation belongs to the *weak* core if no coalition can lead to all members' utilities improved when redistributing items amongst themselves.

Quinzii [1984] generalizes the model of Shapley and Scarf [1974] to one with multiple agents with unit-demand and transferable but non-quasilinear utility. Buyers derive utility from at most one good and a transfer of money. Sellers aim at obtaining the highest possible price above a reservation level. She proved the general existence of the core in her model, and its equivalence to competitive equilibria. In a closely related model, Gale [1984] shows that a competitive equilibrium always exists. These models allow for budgets, but differ from hard budget constraints as examined in our work, where bidders are not permitted to spend more than a certain amount of money. Alaei et al. [2016] provide a structural characterization of utilities in competitive equilibria and a mechanism that is group-strategyproof. These non-quasilinear models assume utility functions that are not necessarily quasilinear, but where small changes of prices do not lead to a discontinuous change of the bidders' utilities as is the case with hard budget constraints.

Closer to this paper is another line of research that focuses on assignment markets where buyers maximize payoff subject to a *hard* budget constraint. Aggarwal et al. [2009] show that an extension of the Hungarian algorithm is incentive-compatible and bidder-optimal if the auction is in general position, a rather specific condition that is usually unknown ex-ante and hard to check. Typically, ascending auctions use only demand queries, namely the auctioneer specifies a price and the bidders respond with their demand set. The auction additionally requires value and budget queries, thereby asking for the value of a specific good to a bidder and their budget during the auction. This is quite different from the ascending auctions based on price-based demand queries only, as described in Demange et al. [1986] and the subsequent literature or compared to ascending auctions used in the field. Fujishige and Tamura [2007] consider two-sided markets with budget-constrained bidders whose valuation functions are more general than unit-demand. Their results imply that in the unit-demand setting, there always exists a core allocation. These prior results aim exclusively for core-stability, but do not attempt to maximize welfare as done in this work.

In contrast to competitive equilibrium theory, Henzinger and Loitzenbauer [2015] and its predecessor Dütting et al. [2013] do not aim for core-stability. Note that with hard budget constraints, core-stability does not imply envy-freeness. Consider for example a market with two buyers and one seller selling a single good. If both buyers have the same budget and the same valuation for the good, which exceeds their budget, the only possibility such that no bidder envies the other one is that the good remains unsold. Such an outcome is clearly not in the core, because there is a coalition of buyer and seller who want to deviate. Note that such types of envy cannot arise without binding budget constraints, because if the price is at the value of two bidders, they are indifferent between getting the object or the empty set. It depends on the considered market, whether envy-freeness and bidder-optimality or core-stability should be preferred. While their model appears to be reasonable in cases where all items are sold by one large seller, like ad auctions, it may not seem reasonable for individual sellers to participate in an auction where items remain unsold for the sake of envy-freeness. Finally, van der Laan and Yang [2016] propose an ascending auction for the assignment market that results in an *equilibrium under allotment*, which is in general not a core-stable outcome.

Core-stability and incentive-compatibility are arguably the most important axioms in market design. Whether one can design assignment markets that satisfy these axioms in the presence of hard budget constraints has not yet been answered. We show that, without strong additional assumptions, this is not possible. Importantly, even with access to all valuations and budgets, the problem is computationally intractable for large market instances.

3 PRELIMINARIES

A two-sided matching market $M = (\mathcal{B}, \mathcal{S}, v, b, r)$ consists of two disjoint sets of agents $\mathcal{B}$ and $\mathcal{S}$, representing bidders $i \in \mathcal{B} = \{1, 2, \ldots, n\}$ and goods $j \in \mathcal{S} = \{1, 2, \ldots, m\} \cup \{0\}$. We identify good j with the seller owning it, i.e. each seller owns one good. The 0-item is a dummy item and does not have value to any bidder, meaning that receiving good 0 corresponds no real good. Additionally, the market is defined by each bidder i's valuation $v_i : \mathcal{S} \to \mathbb{Z}_{\geq 0}$ with $v_i(0) = 0$ and budget $b^i \in \mathbb{Z}_{>0}$, as well as each seller j's reserve values/ ask price $r_j \in \mathbb{Z}_{\geq 0}$.

A *price vector* is a vector $p \in \mathbb{R}^{\mathcal{S}}$ with $p(0) = 0$, assigning price $p(j)$ to every good j. Bidders have quasi-linear utilities, so if bidder i receives item j under prices p, their utility is $\pi_i(j, p) = v_i(j) - p(j)$, if $p(j) \leq b^i$, and $\pi_i(j, p) = -\infty$, otherwise. An *assignment* is represented as a map $\mu : \mathcal{B} \to \mathcal{S}$ from bidders to the items they receive, where $|\mu^{-1}(\{j\})| \leq 1$ for all $j \neq 0$, so only the dummy good may be assigned to more than one bidder. An *outcome* is a pair (μ, p), where μ is an assignment and p is a price vector, such that no budget constraint is violated, i.e. $p(\mu(i)) \leq b^i$ for all $i \in \mathcal{B}$ and only sold items may have a positive price: $p(j) > 0$ implies that $|\mu^{-1}(\{j\})| = 1$. For our iterative auction in Section 4,for the sake of simplicity, we assume all reserve prices to be equal to 0. The results can be easily generalized by starting the auction at the reserve prices, and not at 0.

In neoclassic economics, a (Benthamite) social welfare function is defined as the sum of cardinally measurable values v_i of all market participants. An optimal allocation of resources is one which maximizes the social welfare in this sense:

$$\max \left\{ \sum_{i=1}^{n} v_i(\mu(i)) \ : \ \mu \text{ is an assignment} \right\}$$

This can be written in LP-form as

$$\max \sum_{i=1}^{n}\sum_{j=1}^{m} x_{ij} v_i(j) \tag{1}$$

$$\text{s.t.} \sum_{i=1}^{n} x_{ij} \leq 1 \; \forall j = 1, \ldots, m \qquad (p_j)$$

$$\sum_{j=1}^{m} x_{ij} \leq 1 \; \forall i = 1, \ldots, n \qquad (\pi_i)$$

$$x \geq 0$$

where the variables in parentheses denote the corresponding duals. This assignment problem is well-known to have an integral optimal solution and can be solved in $O(n^4)$ [Kuhn, 1955]. An integral solution x corresponds to an assignment μ via $x_{ij} = 1 \Leftrightarrow \mu(i) = j$. This notion of utilitarian welfare maximization, i.e. maximizing the sum of participants' utilities, is widely used in auction theory and competitive equilibrium theory.

New welfare economics, in the tradition of Pareto defies the idea of interpersonal utility comparisons and stipulates ordinal preferences. Pareto efficiency or Pareto optimality is the key design desideratum in this literature. A market outcome is Pareto efficient, if no market participant can be better off without making at least one other participant worse off. With cardinal utilities and interpersonal comparisons a welfare-maximizing outcome is also Pareto efficient. This is because any Pareto-improvement would increase welfare, which is not possible by definition of a welfare-maximizing allocation. It has also been shown that the converse is true [Negishi, 1960]. Another design desideratum is that of core-stability.

Definition 3.1 (Core outcome). Let (μ, p) be an outcome. A bidder-seller pair $(i, j) \in \mathcal{B} \times \mathcal{S}$ is called a *blocking pair*, if $\pi_i(j, p) > \pi_i(\mu(i), p)$ and $p(j) < b^i$. (μ, p) is a *core outcome*, if there are no blocking pairs. We also say that (μ, p) is *core-stable* in this case.

The idea of a blocking pair (i, j) is that both bidder i and seller j would strictly increase their utility, if i received item j instead of $\mu(i)$: if i pays $p(j) + \varepsilon$ for item j, then still $\pi_i(j, p) - \varepsilon > \pi_i(\mu(i), p)$, and at the same time, the profit of seller j is increased by ε.

In the literature, a core outcome is often alternatively defined in the following way: an outcome (μ, p) is in the core if there are no subsets $\mathcal{B}' \subseteq \mathcal{B}$ and $\mathcal{S}' \subseteq \mathcal{S}$ and an outcome (μ', p') on $\mathcal{B}' \times \mathcal{S}'$ such that $\pi_i(\mu'(i), p') > \pi_i(\mu(i), p)$ for all $i \in \mathcal{B}'$ and $p'(j) > p(j)$ for all $j \in \mathcal{S}'$ (see for example [Zhou, 2017]). These definitions can easily be shown to be equivalent: first suppose that such subsets $\mathcal{B}'$ and $\mathcal{S}'$ do exist. Then it is easy to see that both sets are nonempty. In particular, let $i \in \mathcal{B}'$ and $j = \mu'(i)$. Then $p'(j) > p(j)$, so $p(j) < b^i$. Furthermore, we have $\pi_i(j, p) > \pi_i(j, p') > \pi_i(\mu(i), p)$, so (i, j) is a blocking pair. On the other hand, if (i, j) is a blocking pair, then as in the above paragraph, we can set $p'(j) = p(j) + \varepsilon$ and get $\pi_i(j, p') > \pi_i(\mu(i), p)$ and $p'(j) > p(j)$. Thus we can choose $\mathcal{B}' = \{i\}$, $\mathcal{S}' = \{j\}$, $\mu'(i) = j$ and $p'(j) = p(j) + \varepsilon$ in the alternative definition.

We focus on the problem of finding welfare-maximizing core outcomes:

$$\max \left\{ \sum_{i=1}^{n} v_i(\mu(i)) \; : \; (\mu, p) \text{ is a core outcome} \right\}. \tag{2}$$

If budgets are not binding, i.e., $b^i > v_i(j)$ for all bidders i and all goods j, core-stability coincides with the definition of a competitive equilibrium. For this, let us first define the *demand set* of bidder i, which consist of the most preferred, affordable among all items at prices p:

$$D_i(p) = \left\{ j \; : \; p(j) \leq b^i \; \wedge \; \pi_i(j, p) \geq \pi_i(k, p) \, \forall k \text{ with } p(k) \leq b^i \right\}.$$

Definition 3.2 (Competitive equilibrium). An outcome (μ, p) is a *competitive equilibrium*, if $\mu(i) \in D_i(p)$ for all bidders i.

The next proposition summarizes well-known equivalences of the different notions for markets where budgets are not binding.

PROPOSITION 3.3 (BIKHCHANDANI AND MAMER [1997]). *Suppose that $b^i > v_i(j)$ for all i and j, and let (μ, p) be an outcome. Then the following statements are equivalent.*

(1) *(μ, p) is a core outcome.*
(2) *(μ, p) is a competitive equilibrium.*
(3) *The variables defined by $x_{ij} = 1 \Leftrightarrow \mu(i) = j$ solve the linear program (1) and $p_j = p(j)$ is a corresponding dual solution.*
(4) *(μ, p) is a welfare-maximizing core outcome.*

This equivalence no longer remains true if bidders have binding budgets: in general, a core outcome needs not be a competitive equilibrium, and different core outcomes might generate very different welfare.

Example 3.4. For a very simple example, consider two bidders 1, 2 and one item A. Suppose that $v_1(A) = 6$ and $v_2(A) = 10$. Both bidders have the same budget $b^1 = b^2 = 1$. It is easy to see that there are two core outcomes: either bidder 1 or 2 receives A for a price of 1, while the other bidder does not receive an item. Both core outcomes are no competitive equilibria, since the bidder i not receiving A does not receive an item in $D_i(p) = \{A\}$. This bidder thus envies the other. Moreover, one core outcome generates a welfare of 6, while the other generates a welfare of 10. Ignoring budgets, the above LP-formulation would assign item A to bidder 2 at a price $p(A) \in [6, 10]$ - no such price is feasible when considering the budget constraints.

Besides, as we show, finding a welfare-maximizing core outcome is in general NP-complete, so we cannot expect a simple LP-formulation as above to exist. Note that efficient algorithms for determining core outcomes under budget constraints have been discussed in the literature. However, desirable properties like bidder-optimality and incentive-compatibility are only guaranteed if additional assumptions on the bidders' preferences are made. Aggarwal et al. [2009] introduced the notion of *general position*, a sufficient condition for ascending auctions to indeed find the welfare-maximizing core-stable outcome. As this condition has received considerable attention in the literature, we provide a brief discussion:

Definition 3.5 (Aggarwal et al. [2009]). Consider a directed bipartite graph with edges between bidders $\mathcal{B}$ and goods $\mathcal{S}$ (including dummy good 0): For $i \in \mathcal{B}$ and $j \in \mathcal{S}$, there is a

- forward-edge from i to j with weight $-v_i(j)$
- backward-edge from j to i with weight $v_i(j)$
- maximum-price edge from i to j with weight $b^i - v_i(j)$
- terminal edge from i to the dummy good 0 with weight 0.

The auction is in *general position*, if for every bidder i, there are no two alternating walks, following alternating forward and backward edges and ending with a distinct maximum-price or terminal edge, having the same total weight.

Example 3.6. Consider an auction with two bidders 1 and 2 with $b^1 = b^2$. The number of goods and bidders' valuations may be chosen arbitrarily. Assume $j \in \mathcal{S}$ is any good. Consider the following path starting from bidder 1: $1 \longrightarrow j \longrightarrow 2 \longrightarrow j$, where the last edge is a maximum-price edge, with total weight $-v_1(j) + v_2(j) + (b^2 - v_2(j)) = b^2 - v_1(j)$. Now consider the path $1 \longrightarrow j$, where the only edge is a maximum-price edge, with weight $b^1 - v_1(j)$. Since $b^1 = b^2$, the total weight of both paths is equal, so the auction is not in general position.

As the example shows, the general position condition implies that in an ascending auction, no two bidders may reach their budget limits at the same time. Henzinger and Loitzenbauer [2015] claim the general position condition is rather restrictive, as it excludes, for instance, symmetric bidders. They additionally show that no polynomial-time algorithm can determine whether a set of valuations is in general position. The general position condition is sufficient but not necessary for the existence of a unique bidder-optimal stable matching [Aggarwal et al., 2009], which is thus also welfare-maximizing by our results below. As we will see, there are valuations not in general position, but where a welfare-maximizing core allocation can still be computed efficiently with our auction.

Let us now introduce an iterative auction that always finds a core-stable outcome in markets with budget constrained buyers and, if a simple ex-post condition is satisfied, maximizes welfare among all core outcomes.

4 AN ITERATIVE AUCTION

Our auction is based on the well-known auction by Demange et al. [1986] (denoted as DGS auction from now on), which implements the Hungarian algorithm. Contrary to Aggarwal et al. [2009], where the underlying assumption is that bidders report their valuations and budgets to the auctioneer, in our auction, they only have to report their demand sets at certain prices, similar to other ascending auctions [Mishra and Parkes, 2007]. We will provide conditions when reporting demand sets truthfully is incentive-compatible. Thus, we provide a natural generalization of the DGS auction to markets where bidders have binding budget constraints. The simple ascending nature of our auction also naturally motivates an ex-post optimality condition for the returned allocation. Without loss of generality, we will assume $r_j = 0$ in this section.

In the auction process, we may need to "forbid" some bidder to demand a certain item. We model this by introducing subsets $R_1, \ldots, R_n \subseteq \mathcal{S}$ of goods for every bidder and define the *restricted demand set* to be

$$D_i(p, R_i) = \left\{ j \in R_i \ : \ p(j) \leq b^i \ \wedge \pi_i(j, p) \geq \pi_i(k, p) \, \forall k \in R_i \text{ with } p(k) \leq b^i \right\}.$$

Note that our definition of the restricted demand set coincides with the definition of demand sets by van der Laan and Yang [2016]. The set consists of all affordable items that generate the highest utility among all items in R_i. We introduce the well-known notions of over- and underdemanded sets [Demange et al., 1986, Mishra and Talman, 2006], adjusted to our notion of restricted demand sets.

Definition 4.1. Let a price vector p and sets $R_1, \ldots, R_n \subseteq \mathcal{S}$ with $R_i \neq \emptyset \ \forall i$ be given. A set $T \subseteq \mathcal{S}$ is

- *overdemanded*, if $0 \notin T$ and $|\{i \in \mathcal{B} \ : \ D_i(p, R_i) \subseteq T\}| > |T|$, and
- *underdemanded*, if $p(j) > 0$ for all $j \in T$ and $|\{i \in \mathcal{B} \ : \ D_i(p, R_i) \cap T \neq \emptyset\}| < |T|$.

T is *minimally over-/underdemanded*, if it does not contain a proper over-/underdemanded subset.

Finally, we define the strict budget set of bidder i by $B_i(p) = \{j \in \mathcal{S} \ : \ p(j) < b^i\}$. It consists of all items with prices strictly less than the bidder's budget.

4.1 The Auction Algorithm

Algorithm 1 describes our auction. It is based on the following observation.

Lemma 4.2. *An outcome (μ, p) is in the core if and only if there are sets $R_1, \ldots, R_n \subseteq \mathcal{S}$ such that $B_i(p) \subseteq R_i$ and $\mu(i) \in D_i(p, R_i)$ for all i.*

Proof. Suppose first that (μ, p) is a core outcome. Set $R_i = \{j \in \mathcal{S} : p(j) < b^i\} \cup \{\mu(i)\}$. Then $\mu(i) \in D_i(p, R_i)$, since otherwise there would exist an item j with $p(j) < b^i$ generating a higher utility than $\mu(i)$ - this would constitute a blocking pair.

Now let's assume that there are sets R_i as described with $\mu(i) \in D_i(p, R_i)$ for all i. Suppose there is a blocking pair (i, j). Then j costs strictly less than b^i, so $j \in R_i$, and j generates a higher utility than $\mu(i)$. This would contradict $\mu(i) \in D_i(p, R_i)$. Thus, (μ, p) is a core outcome. □

Computing a core outcome can thus also be interpreted as computing a "competitive equilibrium" with respect to the restricted demand sets $D_i(p, R_i)$. This is quite similar to the definition of an *equilibrium under allotment* of van der Laan and Yang [2016]. However, they have other requirements on the sets R_i, which, in general, cause their equilibria not to lie in the core.

In view of Lemma 4.2, the goal of our auction procedure is to determine prices p together with sets R_i, such that there are neither over- nor underdemanded sets of items. As observed in Mishra and Talman [2006], this implies existence of an assignment $\mu : \mathcal{B} \to \mathcal{S}$, such that every bidder receives an item in their demand set, and every item with positive price gets assigned to some bidder. The following result is due to the aforementioned work.

Proposition 4.3. *Suppose that with respect to the $D_i(p, R_i)$, there is no over- or underdemanded set of items. Then there is an assignment $\mu : \mathcal{B} \to \mathcal{S}$ such that $\mu(i) \in D_i(p, R_i)$ for all i, and for all $j \in \mathcal{S}$ with $p(j) > 0$, there is some i with $\mu(i) = j$.*

Note that they considers markets without budgets and demand sets without restrictions. However, their proof only uses combinatorial properties of the demand sets, so it can be directly adapted to our setting. Thus, we omit a proof here.

ALGORITHM 1: Iterative Auction

1. Set $p^1 = (0, \dots, 0)$ and $R_i^1 = \mathcal{S}$ for all bidders i. Set $t = 1$, $O^0 = \emptyset$ and $I^1 = \emptyset$.
2. Request $D_i(p^t, R_i^t)$ from all bidders. If $t > 1$ and the set
$$I^t = \{i \in \mathcal{B} : D_i(p^{t-1}, R_i^{t-1}) \subseteq O^{t-1} \wedge D_i(p^{t-1}, R_i^{t-1}) \setminus D_i(p^t, R_i^t) \neq \emptyset\}$$
is nonempty, go to Step 3. Otherwise, if there is an overdemanded set, go to Step 4. Else, go to Step 5.
3. Choose a bidder $i \in I^t$ and define $J_i^t = D_i(p^{t-1}, R_i^{t-1}) \setminus D_i(p^t, R_i^t)$. Set $R_i^{t+1} = R_i^t \setminus J_i^t$, $O^t = \emptyset$ and $p^{t+1} = p^{t-1}$. For all other bidders i', the sets $R_{i'}^{t+1} = R_{i'}^t$ are unchanged. Set $t = t + 1$ and go to Step 2.
4. Choose a minimally overdemanded set O^t. For all $j \in O^t$, set $p^{t+1}(j) = p^t(j) + 1$. The prices for all other goods, as well as the sets R_i^t remain unchanged. Set $t = t + 1$ and go to Step 2.
5. Compute an assignment μ, such that $\mu(i) \in D_i(p^t, R_i^t)$ for all bidders i and $\mu(\mathcal{B}) \subseteq \{j \in \mathcal{S} : p^t(j) > 0\}$. Set $p = p^t$ and return (μ, p).

Step 3 of the auction ensures that we do not end up with underdemanded sets of items. Moreover, sets R_i^t always contain at least all items that cost strictly less than the bidder's budget b^i. Our proof of correctness is similar to the one by van der Laan and Yang [2016]: due to the budget constraints, underdemanded sets of items may appear. We show that Step 3 of the auction takes care of these sets.

Lemma 4.4. *Let O be minimally overdemanded and $T \subseteq O$ with $T \neq \emptyset$. Let prices p and sets R_i be given. Then*
$$|\{i : D_i(p, R_i) \subseteq O \wedge D_i(p, R_i) \cap T \neq \emptyset\}| > |T|.$$
In particular, T is not underdemanded.

The proof of this lemma can be found in the Appendix.

Lemma 4.5. *For all bidders $i \in \mathcal{B}$ and all iterations t of the algorithm, we have that $B_i(p) \subseteq R_i^t$. In particular, since $p^t(0) = 0$, $R_i^t \neq \emptyset$.*

Proof. Assume to the contrary that there is a minimal iteration $t + 1$, such that a bidder i^* and a good j^* exist with $p^{t+1}(j^*) < b^{i^*}$, but $j^* \notin R_{i^*}^{t+1}$. Then in iteration t, Step 3 was executed, since otherwise $p^t \leq p^{t+1}$ and $R_{i^*}^t = R_{i^*}^{t+1}$, so $t + 1$ would not be minimal. Hence, in iteration t, we have $j^* \in J_{i^*}^t$ and in particular $j^* \in D_{i^*}(p^{t-1}, R_{i^*}^{t-1})$. Because Step 3 is executed, we have $O^{t-1} \neq \emptyset$, so in iteration $t - 1$ Step 4 was executed and $p^t(j^*) = p^{t-1}(j^*) + 1 = p^{t+1}(j^*) + 1 \leq b^{i^*}$. Thus, since from iteration $t - 1$ to t, all prices for all preferred goods of bidder i^* were raised and i^* can still afford j^* at prices p^t, $j^* \in D_{i^*}(p^t, R_{i^*}^t)$, so $j^* \notin J_{i^*}^t$. This is a contradiction. □

Proposition 4.6. *For every iteration t in the auction, it holds:*

(1) *if there is a minimally underdemanded set of items T, then $T \subseteq O^{t-1}$ and Step 3 is executed*
(2) *if Step 3 is executed, there is no underdemanded set of items with respect to the $D_i(p^{t+1}, R_i^{t+1})$.*

Proof. We prove this by induction on t. For $t = 1$, there clearly is no underdemanded set of items, and Step 3 is not executed.

Suppose now that $t > 1$ and that the statement is true for all $1 \leq s < t$.

First suppose that there exists an underdemanded set of items T. Therefore, by induction, in iteration $t-1$, Step 4 must have been executed - otherwise, there would not exist an underdemanded set. But then, using the same inductive reasoning, there was no underdemanded set in iteration $t-1$. It is thus easy to see that, since in iteration $t-1$ only prices for items in O^{t-1} were raised, only the demand for those items could decrease, so T must be a subset of O^{t-1}. By Lemma 4.4, we have

$$|\{i \in \mathcal{B} \,:\, D_i(p^{t-1}, R_i^{t-1}) \subseteq O^{t-1} \,\cap\, D_i(p^{t-1}, R_i^{t-1}) \cap T \neq \emptyset\}| > |T|.$$

Thus, since $|\{i \in \mathcal{B} \,:\, D_i(p^t, R_i^t) \cap T \neq \emptyset\}| < |T|$, there must be a bidder i^* with $D_{i^*}(p^{t-1}, R_{i^*}^{t-1}) \subseteq O^{t-1}$ and $D_{i^*}(p^{t-1}, R_{i^*}^{t-1}) \cap T \neq \emptyset$, but $D_{i^*}(p^t, R_{i^*}^t) \cap T = \emptyset$. This implies that $i^* \in I^t$, so Step 3 is executed in iteration t.

Now suppose that Step 3 is executed in iteration t. Then again, in iteration $t - 1$, Step 4 was executed, since otherwise we would have $O^{t-1} = \emptyset$, which implies $I^t = \emptyset$. By induction, there was no underdemanded set of items in iteration $t - 1$. Note that $p^{t-1} = p^{t+1}$, so only the demand of a single bidder $i^* \in I^t$ chosen in Step 3 does change. Since $D_{i^*}(p^{t-1}, R_{i^*}^{t-1}) \subseteq O^{t-1}$, so $J_{i^*}^t \subseteq O^{t-1}$, only the demand for items in O^{t-1} can decrease. However, for $T \subseteq O^{t-1}$ we have again by Lemma 4.4 that

$$|\{i \in \mathcal{B} \,:\, D_i(p^{t-1}, R_i^{t-1}) \subseteq O^{t-1} \,\cap\, D_i(p^{t-1}, R_i^{t-1}) \cap T \neq \emptyset\}| > |T|,$$

and, since we only changed $R_{i^*}^{t-1}$, the demand for items in T can at most decrease by 1. Thus, T is not underdemanded in iteration $t + 1$. □

Employing the previous lemmata, we can proceed to prove correctness of our proposed auction.

Proposition 4.7. *The auction terminates after a finite number of iterations, and an assignment μ as is described in Step 5 exists whenever this Step is reached. The returned tuple (μ, p) constitutes a core outcome.*

Proof. Whenever Step 3 is executed, at least one item is removed from the set R_i^t of one bidder. Hence, Step 3 can only be called a finite number of times. Also, prices can only be increased a finite number of times in Step 4 - if prices of goods go to infinity, they are clearly not overdemanded at some point anymore. Thus, in some iteration t^*, Step 5 is executed. By Lemma 4.6, there is no underdemanded set in iteration t^*, because otherwise Step 3 would have been executed. Similarly, there is no overdemanded set. Finally, because of Lemma 4.5, no set $R_i^{t^*}$ is empty, so by Proposition 4.3, an assignment μ as required exists. By Lemma 4.2, (μ, p) is a core outcome. □

Example 4.8. Consider the example following auction with three bidders $1, 2, 3$ and two items A and B.

	$v_i(A)$	$v_i(B)$	b^i
Bidder $i = 1$	10	0	1
Bidder $i = 2$	0	10	2
Bidder $i = 3$	10	10	10

The auction proceeds as follows.

	p^t	$D_1(p^t, R_1^t)$	$D_2(p^t, R_2^t)$	$D_3(p^t, R_3^t)$	R_1^t	R_2^t	R_3^t	O^t	I^t
$t = 1$	$(0, 0)$	$\{A\}$	$\{B\}$	$\{A, B\}$	$\mathcal{S}$	$\mathcal{S}$	$\mathcal{S}$	$\{A, B\}$	$\emptyset$
$t = 2$	$(1, 1)$	$\{A\}$	$\{B\}$	$\{A, B\}$	$\mathcal{S}$	$\mathcal{S}$	$\mathcal{S}$	$\{A, B\}$	$\emptyset$
$t = 3$	$(2, 2)$	$\{0\}$	$\{B\}$	$\{A, B\}$	$\mathcal{S}$	$\mathcal{S}$	$\mathcal{S}$	$\emptyset$	$\{1\}$
$t = 4$	$(1, 1)$	$\{0\}$	$\{B\}$	$\{A, B\}$	$\{0, B\}$	$\mathcal{S}$	$\mathcal{S}$	$\emptyset$	$\emptyset$

In iterations $t = 1, 2$, there is a unique minimally overdemanded set $O^t = \{A, B\}$, and I^t is empty. Thus, Step 4 of the auction is executed and the prices for A and B are raised. In iteration $t = 3$, the set $I^t = \{1\}$ is nonempty which indicates that bidder 1's budget was tight for A at prices $(1, 1)$. Thus, we forbid 1 to receive item A and reset the prices to $(1, 1)$. Now, in iteration $t = 4$, there is no overdemanded set and I^t is empty. Thus, there exists an assignment μ with $\mu(i) \in D_i(p^4, R_i^4)$ for all $i \in \mathcal{B}$, namely $\mu(1) = 0$, $\mu(2) = B$ and $\mu(3) = A$. It is easily checked that (μ, p) is indeed a core outcome.

Example 4.9. Let us now consider an example of an auction where a non-trivial decision has to be made in Step 3.

	$v_i(A)$	$v_i(B)$	b^i
Bidder $i = 1$	10	0	3
Bidder $i = 2$	0	11	1
Bidder $i = 3$	5	3	10

It is easy to see that after 3 iterations through Step 4 of our auction, we reach prices $p^4 = (3, 1)$, where bidder 1 demands $\{A\}$, bidder 2 demands $\{B\}$ and bidder 3 demands $\{A, B\}$. Since $\{A, B\}$ is minimally overdemanded, we execute Step 4 once again to reach $p^5 = (4, 2)$, where due to the budget constraints, we have $D_1(p^5, R_1^5) = \{0\}$ and $D_2(p^5, R_2^5) = \{0\}$, while bidder 3 still demands $\{A, B\}$. Thus, Step 3 of the auction is executed with $I^5 = \{1, 2\}$, so both bidders 1 or 2 would be valid bidders to choose in Step 3. For the choice $i = 1$, we have $J_i^5 = \{A\}$, while for the choice $i = 2$, we have $J_i^5 = \{B\}$. We could thus either remove A from R_1^5, or B from R_2^5. Depending on our choice, we get two different core outcomes, both supported by the prices $p = (3, 1)$: one where bidder 1 receives nothing, bidder 2 receives B and bidder 3 receives A, and one where bidder 1 receives A, bidder 2 receives nothing and bidder 3 receives B. The total welfare of the former allocation is 16, while the one of the latter is 13.

4.2 Economic Properties

The output produced by our iterative auction is not uniquely defined - it may depend on which bidder $i \in I^t$ is chosen whenever Step 3 is executed. Indeed, we prove the following result.

Proposition 4.10. *Let (ν, q) be an arbitrary core outcome. Then bidders $i \in I^t$ in Step 3 can be chosen in such a way, that for the resulting outcome (μ, p) we have that $p \leq q$ coefficient-wise, and $\pi_i(\mu(i), p) \geq \pi_i(\nu(i), q)$ for all bidders i.*

The proof can be found in the Appendix.

We say that a core outcome is (μ, p) *Pareto optimal for the bidders*, if for every core outcome (ν, q) with $\pi_i(\nu(i), q) > \pi_i(\mu(i), p)$ for some bidder i, there is a bidder i' with $\pi_{i'}(\mu(i'), p) > \pi_{i'}(\nu(i'), q)$. Proposition 4.10 directly implies that for every core outcome which is Pareto optimal for the bidders, there is an outcome (μ, p) reachable by the auction with $\pi_i(\mu(i), p) = \pi_i(\nu(i), q)$ for all $i \in \mathcal{B}$. Aggarwal et al. [2009] prove that their algorithm for computing a core-stable outcome always finds the bidder-optimal core outcome (μ, p), whenever the auction is in general position. Here, *bidder-optimal* means that for every other core outcome (ν, q) we have that $\pi_i(\mu(i), p) \geq \pi_i(\nu(i), q)$ for all $i \in \mathcal{B}$. Bidder-optimality thus implies Pareto optimality. We show a similar result for our auction: if the bidder to choose in Step 3 of our auction is always unique, our auction also finds a bidder-optimal core outcome.

Corollary 4.11. *Suppose that whenever Step 3 is executed, $|I^t| = 1$, i.e., there is a unique bidder to choose, and let (μ, p) be the uniquely determined outcome of the auction. Then for any core outcome (ν, q) we have that $p \leq q$ and $\pi_i(\mu(i), p) \geq \pi_i(\nu(i), q)$ for all bidders i, i.e., (μ, p) is bidder-optimal.*

Proof. Since $|I^t| = 1$ in every iteration through Step 3, the outcome (μ, p) of the auction is unique. Proposition 4.10 now directly implies that for every core outcome (ν, q), we have $\pi_i(\mu(i), p) \geq \pi_i(\nu(i), q)$ and $p \leq q$. □

In particular, if the general position condition is satisfied, it can be shown that I^t never contains more than one bidder. Thus, our auction always finds a bidder-optimal outcome like the auction by Aggarwal et al. [2009] in this case.

Proposition 4.12. *Suppose the auction is in general position. Then in every iteration through Step 3 of our iterative auction, we have that $|I^t| = 1$, and for the unique $i \in I^t$, we have $|J_i^t| = 1$.*

Note that in general that our ex-post condition $|I^t| = 1$ whenever Step 3 is reached is less demanding than the general position condition, since we do not require $|J_i^t| = 1$, and it is easy to construct examples where $|J_i^t| > 1$, but the ex-post condition is fulfilled. While our condition is only ex-post, it is straight-forward to check for the auctioneer when the auction is actually performed.

Let us now consider welfare-maximization properties of our auction. We first observe that a welfare maximizing core outcome can always be found among the ones which are Pareto optimal for the bidders.

Proposition 4.13. *Let (μ, p) and (ν, q) be core outcomes. If $\pi_i(\mu(i), p) \geq \pi_i(\nu(i), q)$ for all bidders i, then*

$$\sum_{i \in \mathcal{B}} v_i(\mu(i)) \geq \sum_{i \in \mathcal{B}} v_i(\nu(i)).$$

The proofs of Propositions 4.12 and 4.13 can be found in the Appendix.

As we described above, by Proposition 4.10, we can reach any core outcome which is Pareto optimal for the bidders with our auction, and Proposition 4.13 says that one of them must be welfare-maximizing. Now if we always have $|I^t| = 1$, the outcome of our auction is unique which proves our first main result.

Theorem 4.14. *Bidders in I^t in Step 3 of the auction can be chosen such that the outcome of the auction is a welfare-maximizing core outcome.*

In particular, if $|I^t| = 1$ whenever Step 3 is reached, the unique outcome of the auction is a welfare-maximizing core outcome.

Note that knowledge of the bidders' demand sets does not suffice in order to always choose the "correct" bidders in Step 3 to reach a welfare-maximizing outcome. Our hardness result in Section 5 implies that even with perfect knowledge of the bidders' preferences, choosing the correct bidders in Step 3 is NP-hard. However, our simple ex-post condition $|I^t| = 1$ at least gives the auctioneer a simple certificate of optimality.

If the auction is in general position, then the auction is ex post incentive-compatible, which follows from the original work by Demange et al. [1986] and the paper by Aggarwal et al. [2009]. The question is if this algorithm or any other algorithm where the bidders' preferences are no further restricted can be incentive-compatible. Unfortunately, the answer is no because incentive-compatibility goes against envy-freeness and therefore the core definition as we show next.

Theorem 4.15. *In assignment markets with payoff-maximizing but budget constrained bidders there is no incentive-compatible mechanism terminating in a core-stable solution for every input.*

Proof. By the direct revelation principle, we may assume that bidders report their exact valuations, as well as their budgets to the auctioneer. Consider a market with three bidders $1, 2, 3$ and two items A, B. Let $\mathcal{M}((v_1, b^1), \ldots, (v_3, b^3)) = (\mu, p)$ denote a mechanism, mapping the bidders' reported valuations and budgets to a core-stable outcome with respect to their reports.

We consider instances of the above described market, where all bidders have the same values for both items: $v_i(A) = v_i(B) = 10$ for $i = 1, 2, 3$. Let us consider two instances, where the bidders vary their reported budget.

(1) If all bidders report $b^i = 1$ for $i = 1, 2, 3$, then obviously, since there are only two items, one bidder does not receive one: for $(\mu, p) = \mathcal{M}((v_1, 1), (v_2, 1), (v_3, 1))$, there is an i with $\mu(i) = 0$. Without loss of generality, we assume that $i = 3$. It is easy to see that for core-stability to hold, bidders 1 and 2 both receive an item, and that $p(A) = p(B) = 1$. Bidders 1 and 2 have utility 9, while bidder 3 has utility 0.
(2) If bidder 3 reports $b^3 = 2$, and the other bidders report $b^1 = b^2 = 1$, then clearly bidder 3 receives an item in any core-stable outcome, and without loss of generality $\mu(3) = B$. Also the other item A must necessarily be assigned to some bidder. Again, without loss of generality, we assume that $\mu(1) = A$ and $\mu(2) = 0$. It is easy to see that $p(A)$ must be equal to 1 in a core outcome. Additionally, we must have $p(A) = p(B)$, since otherwise bidder 3 would strictly prefer item A to item B, which would not be envy-free. Thus, $p(A) = p(B) = 1$, and bidder 3 has a utility of 9.

This already shows that $\mathcal{M}$ is not incentive-compatible: If all bidders' true budgets are equal to 1 and they report truthfully, bidder 3 has a utility of 0. However, if bidder 3 misreports $b^3 = 2$, they would receive an item and have a utility of 9. Note, that $p(B) = p(A) = 1$ in this case, so bidder 3 can still afford the received item. □

Note that Theorem 4.15 does not preclude an incentive-compatible and welfare-maximizing auction (that is not core-stable). Overall, these iterative auctions require bidders to reveal that they are indifferent to not winning the good once the price equals the valuation of a bidder. Only this allows auctioneers to differentiate between a bidder dropping out due to reaching his valuation or his budget. In practice, bidders might not always bid the null set when price reaches value even in an incentive-compatible auction, which can lead to inefficiencies in such iterative auctions.

5 A SEALED-BID AUCTION

If in some iteration of the above algorithm Step 3 is reached with $|I^t| > 1$, the ascending auction with only demand queries does not necessarily find the welfare-maximizing core-stable outcome. In such a case, a combination of value and demand queries is required in order to obtain the desired assignment. To this end, for the remaining part of this paper, we assume the auctioneer has unlimited access to all valuations for all objects and the budgets of all bidders. Unfortunately, the outcome and pricing problem does not only need a different oracle but it also becomes NP-complete as we will show next. Besides, as we mentioned above, the property of strategyproofness does not hold.

5.1 A MILP Formulation

First, we show that the problem belongs to complexity class NP by modeling it as a mixed integer linear program (MILP). Once a problem is modeled as such, there is a polynomial-time non-deterministic algorithm, where we guess the values of integer variables and solve the resulting linear program (LP) in polynomial time. Bi-linear terms present in the quadratic formulation (q-BC), namely products of continuous prices $p(j)$ and binary variables, can easily be linearized to obtain the resulting MILP.

$$
\text{(q-BC)}\quad
\begin{array}{lllr}
\text{maximize} & \sum_{i\in\mathcal{B}} \pi_i + \sum_{j\in\mathcal{S}} \pi_j & & \\
\text{subject to} & \pi_i = \sum_{j\in\mathcal{S}} (v_i(j) - p(j))m_i(j) & \forall i \in \mathcal{B} & (1)\\
& \pi_j = \sum_{i\in\mathcal{B}} (p(j) - r_j)m_i(j) & \forall j \in \mathcal{S} & (2)\\
& \sum_{j\in\mathcal{S}} m_i(j) \leq 1 & \forall i \in \mathcal{B} & (3)\\
& \sum_{i\in\mathcal{B}} m_i(j) \leq 1 & \forall j \in \mathcal{S} & (4)\\
& \pi_i \geq \big(v_i(j) - p(j)\big)\alpha_i(j) & \forall i \in \mathcal{B}, j \in \mathcal{S} & (5)\\
& \pi_j \geq min(v_i(j), b^i)(1 - y_i(j)) & \forall i \in \mathcal{B}, j \in \mathcal{S} & (6)\\
& b^i \geq p_j(1 - \beta_i(j)) & \forall i \in \mathcal{B}, j \in \mathcal{S} & (7)\\
& p_j \geq b^i \beta_i(j) & \forall i \in \mathcal{B}, j \in \mathcal{S} & (8)\\
& (1 - \alpha_i(j)) + (1 - \beta_i(j)) - 2 \leq 2(1 - y_i(j)) + \epsilon y_i(j) & \forall i \in \mathcal{B}, j \in \mathcal{S} & (9)\\
& r_j m_i(j) \leq p(j) \leq min(v_i(j), b^i)m_i(j) + M(1 - m_i(j)) & \forall i \in \mathcal{B}, j \in \mathcal{S} & (10)\\
& m_i(j) \in \{0, 1\} & \forall i \in \mathcal{B}, j \in \mathcal{S} & (11)\\
& y_i(j) \in \{0, 1\} & \forall i \in \mathcal{B}, j \in \mathcal{S} & (12)\\
& \alpha_i(j) \in \{0, 1\} & \forall i \in \mathcal{B}, j \in \mathcal{S} & (13)\\
& \beta_i(j) \in \{0, 1\} & \forall i \in \mathcal{B}, j \in \mathcal{S} & (14)\\
& p(j) \geq 0 & \forall j \in \mathcal{S} & (15)
\end{array}
$$

In this section, an assignment of buyer i to seller j is denoted as a binary variable $m_i(j)$, since solving the MILP requires variable definitions. If the resulting assignment assigns the aforementioned pair in this fashion, then $m_i(j) = 1$, and for all other buyers except i, $m_{-i}(j) = 0$. The equivalence to the previous definitions is $m_i(j) = 1 \Leftrightarrow \mu(i) = j$. In order to check for the existence of deviating coalitions, two additional binary variables are introduced. Setting $\alpha_i(j) = 0$ represents the case where bidder i has a benefit from deviating by trading with seller j, and $a_i(j) = 1$ means that the bidder is best satisfied under the current assignment. The second helper binary variable is set to $\beta_i(j) = 0$ if i possesses a sufficient amount of money to purchase j, and set to $\beta_i(j) = 1$, if the budget of bidder i is insufficient to acquire item of seller j, namely the set price of item j exceeds i's budget constraint. Variable $y_i(j) = 0$ reflects the case where bidder i prefers to trade with seller j and has sufficient budget, and the variable is set to 1 if one of the two necessary conditions does not hold.

Utilities of buyers and sellers are defined as previously argued in Section 3. With r_j, we describe the reserve value or ask price of seller j. Constraints (1) and (2) represent the utilities of buyers and sellers, respectively. Constraints (5) and (6) guarantee core-stability. We examine all possible deviating combinations of buyer-seller pairs for a given outcome. Constraint (5) examines whether the corresponding payoff π_i a buyer i receives in the selected assignment m is higher or equal to the alternative assignment (i, j) in question. In particular, this constraint checks whether an assignment (i, j) yields a higher payoff for buyer i, in which case $\alpha_i(j) = 0$. Constraint (6) tests whether a seller j's payoff π_j on the optimal matching m is higher or equal to the minimum value between any buyer i's budget constraint b^i and i's valuation for the item $v_i(j)$, which represents the maximum possible payment seller j could receive from any buyer. One or both of these conditions need to be true. Put differently, if both buyer and seller had a higher payoff under an alternative assignment (i, j), outcome m is not core-stable. In essence, core-stability can be expressed as logical *or* constraint. Constraints (7) and (8) examine whether bidder i has a sufficient budget to obtain item j under price p_j. Constraint (9) is responsible for handling the value of $y_i(j)$ in an appropriate manner, to reflect whether a deviating coalition of (i, j) is indeed profitable and budget-feasible, for any positive value $\epsilon < 1$. The value of $y_i(j)$ depends on binary values $\alpha_i(j), \beta_i(j)$, and we verify our claim by examining the inequalities formed by the different value combinations (the tuple on the left side represents values $\alpha_i(j), \beta_i(j), y_i(j)$):

$$(0,0,0) \Rightarrow 1 + 1 - 2 \leq 2 \cdot 1 - \epsilon \cdot 0 \quad (1^*)$$

$$(0,0,1) \Rightarrow 1 + 1 - 2 \leq 2 \cdot 0 - \epsilon \cdot 1 \quad (2^*)$$

$$(0,1,1) \Rightarrow 1 + 0 - 2 \leq 2 \cdot 0 - \epsilon \cdot 1 \quad (3^*)$$

$$(1,0,1) \Rightarrow 0 + 1 - 2 \leq 2 \cdot 0 - \epsilon \cdot 1 \quad (4^*)$$

$$(1,1,1) \Rightarrow 0 + 0 - 2 \leq 2 \cdot 0 - \epsilon \cdot 1 \quad (5^*)$$

In cases (1^*) and (2^*), agent i has a sufficient budget and can profit from deviating. However, inequality (2^*) is infeasible, therefore the value of $y_i(j)$ cannot be set to 1 for this combination of $\alpha_i(j), \beta_i(j)$ and is forced to 0. For the remaining cases, either i does not have sufficient budget ($\beta_i(j) = 1$), or has no profit from trading with j ($\alpha_i(j) = 1$), or both conditions hold. In all the aforementioned cases, $y_i(j) = 1$, and thus reflects the case where no deviation is preferable from the buyer's side.

Constraint (10) then makes sure that if an item j is assigned to buyer i, then the price is less than the minimum of the budget of this buyer or his value, and it is higher than the reserve price of the seller. We can conclude that the above formulation always results in a the welfare-maximizing core-stable outcome for assignment markets with budget constraints.

5.2 Complexity Analysis

We now proceed to the proof of NP-completeness. The existence of endogenous pricing variables renders the proof non-trivial: there is no standard method of encoding both continuous and discrete variables, for prices and assignment, in a problem instance. Therefore, the proof requires a more complicated structure and special assumptions regarding the handling of item prices. First, we formally define the decision version of the problem.

Maximum Welfare Budget Constrained Stable Bipartite Matching (MBSBM)

Input: Two disjoint sets $\mathcal{S}$ (sellers) and $\mathcal{B}$ (buyers) of n agents each, a budget b^i for each agent $i \in \mathcal{B}$ and a reserve value r_j for each seller $j \in \mathcal{S}$, a value $v_i(j)$ for each pair of agents $i \in \mathcal{B}$ and $j \in \mathcal{S}$, and a non-negative integer k.

Output: Boolean value

Question: Does there exist a stable outcome μ such that the total value $\sum_{i \in \mathcal{B}}(v_i(\mu(i)) - r_{\mu(i)}) \geq k$?

We have already discussed that the problem is in NP, because MBSBM can be modeled as a MILP where all the reserve values r_j of sellers $j \in \mathcal{S}$ are set to 0, and such problems are contained in NP [Del Pia et al., 2017]. We know that for the case when all budgets are too high and therefore non-binding, $b^i \geq v_i(j)$ for all $i \in \mathcal{B}$ and $j \in \mathcal{S}$, the problem is equivalent to the maximum-weight bipartite matching, which admits a polynomial time solution via the Hungarian algorithm [Kuhn, 1955]. Core-stable prices can be derived from the duals of the corresponding linear program [Shapley and Shubik, 1971]. Therefore, the case of interest that can result in increased computational cost is when budgets are binding for participating buyers. For our problem, we reduce from the Maximum Independent Set (MIS) problem, which is known to be APX-hard, thus implying NP-hardness. Chen et al. [2021] follows a similar approach for fractional matching without transferable utility. However, an environment with partially transferable utility as in our case, requires special attention and leads to significant differences.

Maximum Independent Set (MIS)

Input: A graph $G = (V, E)$, with vertices V and edges E, and a non-negative integer k.

Output: Boolean value

Question: Does there exist an Independent Set (IS) of size at least k, where as IS we define a set of vertices no two of which are adjacent?

The proof uses a specific construction in which we introduce an individual vertex and an edge gadget for each vertex and edge of the original MIS problem. In complexity theory, when performing a reduction from computational problem A to problem B, the term gadget refers to a subset of a problem instance of problem B, that simulates the behavior of certain units of problem A. Drawing from graph theory, the vertex and edge gadgets are bipartite graphs where each edge gadget is connected to two vertex gadgets, corresponding to the two endpoints of the original edge. Each vertex has a degree of three and thus each vertex gadget is connected to three edge gadgets. The edge gadget allows for two matchings between buyer and seller nodes, which all lead to the same welfare. The vertex gadgets also allow for two feasible stable matchings, where the welfare differs by one. The edges between an edge and a vertex gadget are such that it is not possible to select the welfare-maximizing matching in two consecutive vertex gadgets of two neighboring vertices, because it would generate pairs of blocking agents in the edge gadget of the connecting edge, i.e. a matching in the edge gadget would not be core-stable. Similar to the original MIS problem, where there cannot be two adjacent vertices in an IS, in our construction, there cannot be two adjacent vertex gadgets with a high welfare matching. While in the MIS problem, we need to find the IS that is maximal, in the MBSBM problem we need to determine the stable matching of the overall bipartite graph that maximizes welfare. Note that both the vertex and edge gadgets contain a pair of buyers that admit the same budget constraint, thus leading to a violation of the general position condition. Let us now discuss the construction and proof in detail.

5.2.1 Construction. Assume an instance of MIS defined on a cubic graph $G = ((V, E), k)$. A cubic graph is a graph in which all vertices have degree three. We define the transformed instance as a bipartite graph $G' = (V', E')$, with $V' = \mathcal{B} \cup \mathcal{S}$, and functions $\pi_i : \mathcal{S} \to \mathbb{R}_{\geq 0}$ for each buyer $i \in \mathcal{S}$, and $\pi_j : \mathcal{B} \to \mathbb{R}_{\geq 0}$ for each seller $j \in \mathcal{S}$ that represents agents' payoffs.

- V' represents the total set of agents
- $\mathcal{B}$ and $\mathcal{S}$ denote the sets of buyers and sellers respectively
- E' represents the potential transactions between buyers and sellers
- π_i specifies the difference between the true valuation of buyer $i \in \mathcal{B}$ for the item of seller $j \in \mathcal{S}$ and her total budget b^i, namely $v_i(j) - b^i$

- π_j specifies difference between the budget of buyer $i \in \mathcal{B}$ and the reserve value of seller $s \in \mathcal{S}$, namely $b^i - r_j$

Since we assume that all assigned buyers pay prices for items equal to their budgets, we omit prices from the payoff formulae π_i, π_j. In Lemma 5.2 we show that this restriction of the prices to be at the budgets is without loss of generality for the construction. Observe that in our construction multiple buyers share the same budget constraint, which violates the general position condition from Definition 3.5.

An *assignment* $\mu : \mathcal{B} \longrightarrow \mathcal{S}$ in G' assigns each edge $e \in E'$ according to the condition $\sum_{j \in \mathcal{S}} |\mu(i) = j| \leq 1$ for any agent $i \in \mathcal{B}$. The *utility* of agent $i \in \mathcal{B}$ under assignment μ of (G', π_i, π_j) is defined as $\pi_i(\mu) := \pi_i(\mu(i))$ and similarly for agent $j \in \mathcal{S}$ it holds that $\pi_j(\mu) := \pi_j(\mu^{-1}(j))$. Given an assignment μ, an edge $(i, j) \in E'$ is a *blocking pair/ edge* if $\pi_i(\mu(i)) < \pi_i(j)$ and $\pi_j(\mu^{-1}(j)) < \pi_j(i)$. An assignment μ is *stable* if it does not contain any blocking pair of agents.

Edge Gadget. Starting from an edge $e \in E$ of the original graph G, we construct the edge gadget $G'_e = \mathcal{B}_e \cup \mathcal{S}_e$ as a bipartite graph, with preferences for the subgraph defined as mentioned above and noted as $\pi_i : \mathcal{S}_e \longrightarrow \mathbb{R}_{\geq 0}$ and $\pi_j : \mathcal{B}_e \longrightarrow \mathbb{R}_{\geq 0}$ for each of the two sets respectively. The vertices represent participating *agents*.

For each edge $e = (u, u') \in E$, we proceed to the following construction:

- Add to $\mathcal{B}_e$ three agents $\beta_e, \gamma_e, \delta_e$
- Add to $\mathcal{S}_e$ three agents $\eta_e, \alpha_e^u, \alpha_e^{u'}$
- Add to $\mathcal{B}_e$ two extra agents ϵ_u and $\epsilon_{u'}$, if not present already. These represent the connection from edge gadget to vertex gadget.

Vertices α_e^u and $\alpha_e^{u'}$ represent the gates to the vertex gadgets, which are connected to the vertices $\epsilon_u, \epsilon_{u'}$ of the vertex gadgets of u, u' of the original graph.

For each edge $e \in E$ of the original graph G, the corresponding edge gadget consists of subgraph that agents $\{\beta_e, \gamma_e, \delta_e\} \cup \{\eta_e, \alpha_e^u, \alpha_e^{u'}\}$ induce. The detailed construction is depicted in Figure 1. Any edge not present in the figure is assigned a value of zero.

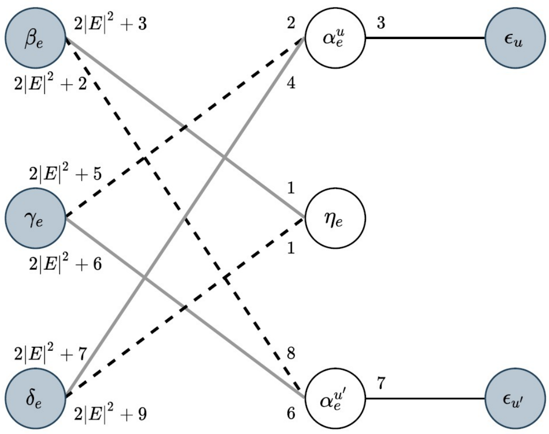

Fig. 1. The edge gadget, with corresponding buyer and seller preferences π_i, π_j. If u belongs to the independent set $IS(G)$, then we integrally match the pairs corresponding to the solid edges, otherwise we match the dashed edges.

An important observation is, that any edge gadget G'_e should contain two feasible stable assignments, with an equal total social welfare. In every stable assignment μ, at least one of $\{\alpha_e^u, \alpha_e^{u'}\}$ is

	α_e^u	η_e	$\alpha_e^{u'}$	b^i
β_e	0	$2\lvert E\rvert^2 + 12$	$2\lvert E\rvert^2 + 11$	9
γ_e	$2\lvert E\rvert^2 + 12$	0	$2\lvert E\rvert^2 + 13$	7
δ_e	$2\lvert E\rvert^2 + 16$	$2\lvert E\rvert^2 + 18$	0	9
ϵ_u	3	0	0	8
$\epsilon_{u'}$	0	0	7	8
r_j	5	8	1	

Table 1. Valuation table representing valuations $v_i(j)$ where $i \in \{\beta_e, \gamma_e, \delta_e, \epsilon_u, \epsilon_{u'}\}$ and $j \in \{\alpha_e^u, \eta_e, \alpha_e^{u'}\}$. b^i corresponds to the budget of each buyer and r_j to the reserve value of each seller.

unsatisfied under μ, namely agents α_e^u or $\alpha_e^{u'}$ have a preference towards ϵ_u or $\epsilon_{u'}$ respectively.

Vertex Gadget. In a similar manner, we construct the vertex gadget corresponding to vertex $u \in V$.

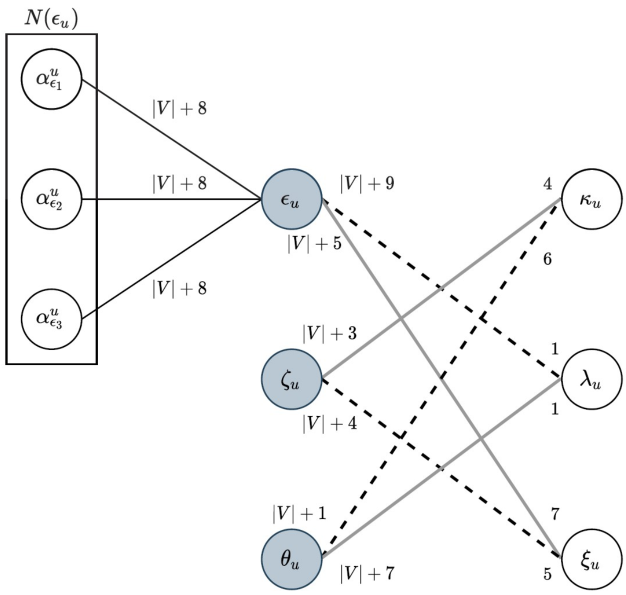

Fig. 2. The vertex gadget, with corresponding buyer and seller preferences. If u belongs to the independent set $IS(G)$, then we integrally match the pairs corresponding to the solid edges, otherwise, we match the dashed edges.

	κ_u	λ_u	ξ_u	$N(\epsilon_u)$	b^i
ϵ_u	0	$\lvert V\rvert + 17$	$\lvert V\rvert + 13$	$n + 16$	8
ζ_u	$\lvert V\rvert + 9$	0	$\lvert V\rvert + 10$	0	6
θ_u	$\lvert V\rvert + 9$	$\lvert V\rvert + 15$	0	0	8
r_j	2	7	1	5/1	

Table 2. Valuation table representing valuations $v_i(j)$ where $i \in \{\epsilon_u, \zeta_u, \theta_u\}$ and $j \in \{\kappa_u, \lambda_u, \xi_u, N(\epsilon_u)\}$. $N(\epsilon_u)$ consists of the neighboring vertices of ϵ_u in the edge gadgets of $\alpha^u_{\epsilon_1}, \alpha^u_{\epsilon_2}, \alpha^u_{\epsilon_3}$, since each vertex u has a degree of 3. b^i corresponds to the budget of each buyer and r_j to the reserve value of each seller in the vertex gadget.

Figure 2 illustrates the vertex gadget G'_u corresponding to a vertex $u \in V$, along with the preferences of buyers and sellers. Agents colored in gray belong to set $\mathcal{B}_u$, and agents in white belong to $\mathcal{S}_u$, defined in an analogous manner to $\mathcal{B}_e, \mathcal{S}_e$. Two feasible stable outcomes exist in G'_u, resulting in welfare that differs by 1: in case an assignment indicated by the solid edges is chosen, induced welfare equals $|V| + 27$, while for an assignment of the dashed edges, welfare has a value of $|V| + 26$. The set $N(\epsilon_u)$ represents the neighborhood of ϵ, and consists of three vertices from three distinct edge gadgets.

The following lemma provides a useful bound for the achieved social welfare of an optimal assignment in an edge gadget.

Lemma 5.1. *Let μ be a stable assignment for edge gadget G'_e. The total welfare SW achieved by μ is at most $|E| \cdot (6|E|^2 + 27)$. If there is an edge $(u, u') \in E$, where for the gadgets it holds that $\pi_{\epsilon_u}(\mu(\epsilon_u)) < \pi_{\epsilon_u}(\alpha_e^u)$ and $\pi_{\epsilon_{u'}}(\mu(\epsilon_{u'})) < \pi_{\epsilon_{u'}}(\alpha_e^{u'})$, then $SW < |E| \cdot (6|E|^2 + 27) - |V|$.*

This lemma implies that, in any welfare-maximizing assignment, edges between $\{\alpha_e^u, \epsilon_u\}$ and $\{\alpha_e^{u'}, \epsilon_{u'}\}$ are never chosen. Furthermore, at most one of the conditions $\pi_{\epsilon_u}(\mu(\epsilon_u)) < \pi_{\epsilon_u}(\alpha_e^u)$ and $\pi_{\epsilon_{u'}}(\mu(\epsilon_{u'})) < \pi_{\epsilon_{u'}}(\alpha_e^{u'})$ must hold, otherwise assignment μ is not maximizing welfare.

In our construction, we fix the price for each matching to the budget constraint of the buyer to whom a seller is matched. Next, we show that this restriction is without loss of generality. We show in each bipartite graph, regardless of whether it belongs to the vertex or edge gadget, the current scheme yields all possible welfare-maximizing, core-stable outcomes.

Lemma 5.2. *Assuming that all buyers pay a price equal to their budget for their assigned items does not impact generality, namely there does not exist a welfare-maximizing, core-stable outcome that is not reachable through this pricing scheme.*

Finally, we build the main complexity result for MBSBM. The proof uses no assumption on specific properties (such as general position) to hold.

Theorem 5.3. *MBSBM is NP-complete.*

The proofs of Lemmata 5.1, 5.2 and of Theorem 5.3 can be found in the Appendix.

6 CONCLUSIONS

In this work, we showed that there is no incentive-compatible mechanism that selects an outcome in the core for bidders with unit-demand valuations. However, it is possible to expand the auction algorithm by Demange et al. [1986], based on demand queries, which terminates with a core allocation if bidders truthfully reveal their demand set in each round. With an additional condition on the unit-demand valuations, this mechanism is incentive-compatible and maximizes welfare. It is typically unknown ex-ante whether this condition holds.

Overall, the analysis of markets with unit-demand bidders is important as it aids our understanding about how restrictive the assumptions need to be for a market to be efficient. In his seminal paper, Vickrey [1961] showed that markets can be designed such that it is a dominant strategy for participants to reveal their preferences truthfully. He allows for general valuations (including substitutes and complements) and only poses a seemingly innocuous assumption of payoff-maximizing bidders. Budget constraints are wide-spread and violate the assumptions of the VCG mechanism. We prove that, even in the simplest markets where bidders only have unit-demand valuations, incentive-compatibility and core-stability are conflicting. In fact, even if all unit-demand valuations and budgets were known to the auctioneer, computing a core-stable and efficient outcome is NP-complete. As a result, claims about the efficiency of simple (polynomial-time) market designs need to be considered with care in the many markets where financial constraints of bidders play a role.

REFERENCES

Gagan Aggarwal, S. Muthukrishnan, Dávid Pál, and Martin Pál. 2009. General Auction Mechanism for Search Advertising *(WWW '09)*. Association for Computing Machinery, New York, NY, USA, 241–250.

Saeed Alaei, Kamal Jain, and Azarakhsh Malekian. 2016. Competitive equilibria in two-sided matching markets with general utility functions. *Operations Research* 64, 3 (2016), 638–645.

Kenneth J Arrow and Gerard Debreu. 1954. Existence of an equilibrium for a competitive economy. *Econometrica* 22, 3 (1954), 265–290.

Lawrence Ausubel and Paul Milgrom. 2006. The Lovely but Lonely Vickrey Auction. In *Combinatorial Auctions*. MIT Press, Cambridge, MA, Chapter 1, 17–40.

Elizabeth Baldwin and Paul Klemperer. 2019. Understanding preferences: demand types, and the existence of equilibrium with indivisibilities. *Econometrica* 87, 3 (2019), 867–932.

Martin Bichler and Jacob K Goeree. 2017. *Handbook of spectrum auction design*. Cambridge University Press.

Martin Bichler and Stefan Waldherr. 2022. Core-Pricing in Combinatorial Exchanges with Financially Constrained Buyers: Computational Hardness and Algorithmic Solutions. *Operations Research* 70, 1 (2022), 241 – 264.

Sushil Bikhchandani and John W Mamer. 1997. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory* 74, 2 (1997), 385–413.

Sushil Bikhchandani and Joseph M Ostroy. 2002. The package assignment model. *Journal of Economic Theory* 107, 2 (2002), 377–406.

Yeon-Koo Che and Ian Gale. 2000. The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic Theory* 92, 2 (2000), 198–233.

Jiehua Chen, Sanjukta Roy, and Manuel Sorge. 2021. Fractional Matchings under Preferences: Stability and Optimality. In *IJCAI*.

Alberto Del Pia, Santanu S Dey, and Marco Molinaro. 2017. Mixed-integer quadratic programming is in NP. *Mathematical Programming* 162, 1 (2017), 225–240.

Gabrielle Demange, David Gale, and Marilda Sotomayor. 1986. Multi-item auctions. *Journal of Political Economy* 94, 4 (1986), 863–872.

Shahar Dobzinski, Ron Lavi, and Noam Nisan. 2008. Multi-unit auctions with budget limits. In *Foundations of Computer Science (FOCS'08)*. IEEE, 260–269.

Paul Dütting, Monika Henzinger, and Martin Starnberger. 2016. Auctions for heterogeneous items and budget limits. *ACM Transactions on Economics and Computation (TEAC)* 4, 1 (2016).

Paul Dütting, Monika Henzinger, and Ingmar Weber. 2013. Sponsored search, market equilibria, and the Hungarian Method. *Inform. Process. Lett.* 113, 3 (2013), 67–73.

Edmund Eisenberg and David Gale. 1959. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics* 30, 1 (1959), 165–168.

Satoru Fujishige and Akihisa Tamura. 2007. A Two-Sided Discrete-Concave Market with Possibly Bounded Side Payments: An Approach by Discrete Convex Analysis. *Mathematics of Operations Research* 32, 1 (2007), 136–155.

David Gale. 1984. Equilibrium in a discrete exchange economy with money. *International Journal of Game Theory* 13, 1 (1984), 61–64.

David Gale and Lloyd S Shapley. 1962. College admissions and the stability of marriage. *The American Mathematical Monthly* 69, 1 (1962), 9–15.

Faruk Gul and Ennio Stacchetti. 1999. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory* 87 (1999), 95–124.

Monika Henzinger and Veronika Loitzenbauer. 2015. Truthful unit-demand auctions with budgets revisited. *Theoretical Computer Science* 573 (2015), 1–15.

Alexander S. Kelso and Vincent P. Crawford. 1982. Job matching, coalition formation , and gross substitute. *Econometrica* 50, 6 (1982), 1483–1504.

Harold W Kuhn. 1955. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly* 2, 1-2 (1955), 83–97.

Renato Paes Leme. 2017. Gross substitutability: An algorithmic survey. *Games and Economic Behavior* 106 (2017), 294–316.

Renato Paes Leme and Sam Chiu-wai Wong. 2017. Computing Walrasian Equilibria: Fast Algorithms and Structural Properties. In *Proceedings of the28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '17)* (Barcelona, Spain). 632–651.

Debasis Mishra and David C Parkes. 2007. Ascending price Vickrey auctions for general valuations. *Journal of Economic Theory* 132, 1 (2007), 335–366.

Debasis Mishra and Adolphus Talman. 2006. Overdemand and Underdemand in Economies with Indivisible Goods and Unit Demands. *CentER Discussion Paper* No. 2006-84 (2006).

Takashi Negishi. 1960. Welfare economics and existence of an equilibrium for a competitive economy. *Metroeconomica* 12, 2-3 (1960), 92–97.
James B Orlin. 2010. Improved algorithms for computing Fisher's market clearing prices: Computing Fisher's market clearing prices. In *Proceedings of the 42th Annual ACM Symposium on Theory of Computing (STOC '10)*. 291–300.
Martine Quinzii. 1984. Core and competitive equilibria with indivisibilities. *International Journal of Game Theory* 13, 1 (1984), 41–60.
Lloyd S Shapley and Herbert Scarf. 1974. On cores and indivisibility. *Journal of mathematical economics* 1, 1 (1974), 23–37.
Lloyd S Shapley and Martin Shubik. 1971. The assignment game I: The core. *International Journal of Game Theory* 1, 1 (1971), 111–130.
Gerard van der Laan and Zaifu Yang. 2016. An ascending multi-item auction with financially constrained bidders. *Journal of Mechanism and Institution Design* 1 (2016), 109–149.
William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16, 1 (1961), 8–37.
Yu Zhou. 2017. A multi-item auction with budget-constrained bidders and price controls. *Economics Letters* 155 (2017), 76–79.

A APPENDIX

A.1 Proof of Lemma 4.4

Proof. Since $O \setminus T \subseteq O$ and O is minimally overdemanded, $O \setminus T$ is not overdemanded, so

$$|O| - |T| = |O \setminus T| \geq |\{i : D_i(p, R_i) \subseteq O \setminus T\}|.$$

Now

$$\{i : D_i(p, R_i) \subseteq O \setminus T\} = \{i : D_i(p, R_i) \subseteq O\} \setminus \{i : D_i(p, R_i) \subseteq O \wedge D_i(p, R_i) \cap T \neq \emptyset\}$$

so

$$|O| - |T| \geq |\{i : D_i(p, R_i) \subseteq O\}| - |\{i : D_i(p, R_i) \subseteq O \wedge D_i(p, R_i) \cap T \neq \emptyset\}|.$$

By rearranging terms we get

$$|\{i : D_i(p, R_i) \subseteq O \wedge D_i(p, R_i) \cap T \neq \emptyset\}| \geq |\{i : D_i(p, R_i) \subseteq O\}| - |O| + |T|$$

and since O is overdemanded, this implies

$$|\{i : D_i(p, R_i) \subseteq O \wedge D_i(p, R_i) \cap T \neq \emptyset\}| > |T|.$$

□

In the following we need a simple auxiliary lemma.

Lemma A.1. *Let (v, q) be an arbitrary outcome, and let $O \subseteq \mathcal{S}$ be a minimally overdemanded set with respect to the demand sets $D_i(q, R_i)$. Let $\emptyset \neq J \subsetneq O$. Then there exists a bidder i^* with $D_{i^*}(q, R_{i^*}) \subset O$, $D_{i^*}(q, R_{i^*}) \cap J \neq \emptyset$ and $v(i^*) \notin J$.*

A.2 Proof of Proposition 4.10

Proof. We prove the following statements by induction on the iteration t, which imply Proposition 4.10.

- If in iteration t Step 4 or 5 is executed, then $p^t \leq q$.
- If in iteration t Step 3 is executed, then $p^{t+1} \leq q$ and we can choose bidder $i \in I^t$ such that $v(i) \notin J_i^t$, and consequently, by choosing this bidder, we have $v(i) \in R_i^{t+1}$ for all i.

For iteration $t = 1$ this is obviously true, since only Step 4 or Step 5 can be executed, and $p^1 = (0, \ldots, 0) \leq q$.

Now suppose that $t > 1$. First assume that in iteration t Step 4 or 5 is executed. If in Step $t-1$ Step 3 was executed, we have that $p^t = p^{t-2}$, and since Step 3 cannot be executed twice in a row, we have by induction $p^t = p^{t-2} \leq q$. Now assume that Step 4 was executed in iteration $t-1$. Then by induction we have $p^{t-1} \leq q$. Towards a contradiction, assume that the set $J = \{j \in O^{t-1} : p^t(j) > q(j)\}$ is not empty. Note that for all $j \in J$ we must have $q(j) = p^{t-1}(j)$. By Lemma A.1 there exists a bidder i^* with $D_{i^*}(p^{t-1}, R_{i^*}^{t-1}) \subset O^{t-1}$, $D_{i^*}(p^{t-1}, R_{i^*}^{t-1}) \cap J \neq \emptyset$ and $v(i^*) \notin J$. Thus, we have $v(i^*) \notin O^{t-1}$ or $p^t(j) \leq q(j)$. Since $v(i^*) \in R_{i^*}^{t-1}$, we have for every $j \in D_{i^*}(p^{t-1}, R_{i^*}^{t-1}) \cap J$ that $\pi_{i^*}(v(i^*), q) < \pi_{i^*}(j, q)$. Consequently, since (v, q) is a core outcome, so (i^*, j) is no blocking pair, we must have that $q(j) = b^{i^*}$. But since $q(j) = p^{t-1}(j)$ and $j \in O^{t-1}$, it would follow that $j \notin D_{i^*}(p^t, R_i^t)$, so $I^t \neq \emptyset$ and Step 3 is executed in iteration t. This contradicts our assumption that Step 4 is executed.

Now consider the case where Step 3 is executed in iteration t. Since Step 3 cannot be executed twice in a row, we have by induction that $p^{t+1} = p^{t-1} \leq q$. It remains to show that there is some bidder $i \in I^t$ with $v(i) \notin J_i^t$. Again, towards a contradiction, assume that for all $i \in I^t$ we have $v(i) \in J_i^t$. Hence, $p^{t-1}(v(i)) = b^i$ for all $i \in I^t$, and since $q \geq p^{t-1}$ by induction, we have $q(v(i)) = b^i$ for all $i \in I^t$. Our argument is now very similar to the one above: Consider the set

$J = \{j \in O^{t-1} \;:\; q(j) = p^{t-1}(j)\}$. Then $\nu(I^t) \subseteq J$, and in particular J^t is nonempty. Hence there is a bidder i^* with $D_{i^*}(q, R_{i^*}) \subset O$, $D_{i^*}(q, R_{i^*}) \cap J \neq \emptyset$ and $\nu(i^*) \notin J$. Again, we have $\nu(i^*) \notin O^{t-1}$, or $q(\nu(i^*)) > p^{t-1}(\nu(i^*))$. In both cases, i^* would prefer any good j in the intersection $D_{i^*}(q, R_{i^*}) \cap J \neq \emptyset$ to $\nu(i^*)$ at prices q. But since $i^* \notin I^t$, $b^{i^*} > q(j)$, so (i^*, j) would form a blocking pair.

$$I = \{i \in \mathcal{B} \;:\; D_i(p^{t-1}, O^{t-1}) \subseteq O^{t-1} \;\wedge\; \nu(i) \in O^{t-1} \;\wedge\; p^{t-1}(\nu(i)) = q^{t-1}(\nu(i))\}.$$

Then $I^t \subseteq I$, so I is not empty. By Lemma A.1, there is a bidder $i^* \notin I$ with $D_{i^*}(p^{t-1}, R^{t-1}_{i^*}) \subseteq O^{t-1}$ and $\nu(I) \cap D_{i^*}(p^{t-1}, R^{t-1}_{i^*}) \neq \emptyset$. The $\nu(i^*) \notin O^{t-1}$, or $q(\nu(i^*)) > p^{t-1}(\nu(i^*))$. Since by induction we have $\nu(i^*) \in R^{t-1}_{i^*}$, we have that $\pi_{i^*}(\nu(i^*), q) < \pi_{i^*}(j, q)$ for all $j \in \nu(I) \cap D_{i^*}(p^{t-1}, R^{t-1}_{i^*})$. But since (ν, q) is a core outcome, for no such j, (i^*, j) can be a blocking pair - implying that $b^{i^*} = q(j) = p^{t-1}(j)$. This is a contradiction, since it would follow that $i^* \in I$.

□

A.3 Proof of Proposition 4.12

Proof. Consider the undirected bipartite graph G with vertex set $I \dot{\cup} O^{t-1}$, where $I = \{i \in \mathcal{B} \;:\; D_i(p^{t-1}, R^{t-1}_i) \subseteq O^{t-1}\}$, containing the edge $\{i, j\}$ if and only if $j \in D_i(p^{t-1}, R^{t-1}_i)$. Since O^{t-1} is minimally overdemanded, this graph is connected by Lemma A.1. For the sake of clarity, we denote the bidders in I by $1, \ldots, r$, and the goods in O^{t-1} by $1', \ldots, s'$, where $r = |I|$ and $s = |O^{t-1}|$.

Let $J = \{j' \in O^{t-1} \;:\; \exists i \in I \;:\; j' \in D_i(p^{t-1}, R^{t-1}_i) \wedge p(j') = b^i\}$ be the set of all goods j' demanded by some bidder in I with tight budget for j'. We show that $|J| = 1$. To that goal, assume that this is not the case. Then, after a possible re-indexing, there exists a path

$$j_1 \longrightarrow i_1 \longrightarrow j_2 \longrightarrow \cdots \longrightarrow i_{m-1} \longrightarrow j_m$$

alternating between goods and bids in G, where j_1 and j_m are distinct goods contained in J. Then by following the corresponding forward- and backward edges according to Definition 3.5, the total weight of this walk is

$$\sum_{k=1}^{m-1} v_{i_k}(j_k) - v_{i_k}(j_{k+1}).$$

For all $k \in \{1, \ldots, m-1\}$, we have that $\pi_{i_k}(j_k, p^{t-1}) = \pi_{i_k}(j_{k+1}, p^{t-1})$, which is equivalent to $v_{i_k}(j_k) - v_{i_k}(j_{k+1}) = p^{t-1}(j_k) - p^{t-1}(j_{k+1})$. Thus, the weight of the above walk is equal to $p^{t-1}(j_1) - p^{t-1}(j_m)$. Now let $i_0, i_m \in I$ be bidders demanding j_1, respectively j_m with $b^{i_0} = p(j_1)$ and $b^{i_m} = p(j_m)$. Let us consider two walks starting from i_0: the first one consists of a single maximum-price edge $i_0 \longrightarrow j_1$ and has weight $b^{i_0} - v_{i_0}(j_1)$. The second one is

$$i_0 \longrightarrow j_1 \longrightarrow i_1 \longrightarrow j_2 \longrightarrow \cdots \longrightarrow i_{m-1} \longrightarrow j_m \longrightarrow i_m \longrightarrow j_m$$

where the last edge is a maximum-price edge. This walk has weight

$$-v_{i_0}(j_1) + (p^{t-1}(j_1) - p^{t-1}(j_m)) + v_{i_m}(j_m) + (b^{j_m} - v_{i_m}(j_m)).$$

Since $b^{i_0} = p(j_1)$ and $b^{i_m} = p(j_m)$, this is equal to $b^{i_0} - v_{i_0}(j_1)$. Thus, we found two distinct paths with the same total weight, so the auction is not in general position. It follows that $|J| = 1$.

It remains to show that $|I^t| = 1$. To that goal, assume that $|I^t| > 1$ and let $i_1, i_2 \in I^t$ be distinct. Since $|J| = 1$, there is a good $j \in O^{t-1}$, demanded by both i_1 and i_2 with $p(j) = b^{i_1} = b^{i_2}$. The walk $i_1 \longrightarrow j$ consisting of a single maximum-price edge has weight $b^{i_1} - v_{i_1}(j)$, and the walk $i_1 \longrightarrow j \longrightarrow i_2 \longrightarrow j$ ending with a maximum-price edge, also has weight $-v_{i_1}(j) + v_{i_2}(j) + (b^{i_2} - v_{i_2}(j)) = b^{i_1} - v_{i_1}(j)$. Thus, the auction is not in general position, a contradiction. Hence, $|I^t| = 1$, and since $J = |1|$, the good $j \in D_i(p^{t-1}, R^{t-1}_i)$ with $p(j) = b^i$ for the bidder $i \in I^t$ is unique. □

A.4 Proof of Proposition 4.13

Proof. We have that

$$\sum_{i\in\mathcal{B}} v_i(\mu(i)) = \sum_{i\in\mathcal{B}} \pi_i(\mu(i),p) + \sum_{j\in\mathcal{S}} p(j) \geq \sum_{i\in\mathcal{B}} \pi_i(\mu(i),p) + p(\nu(i))$$

and

$$\sum_{i\in\mathcal{B}} v_i(\nu(i)) = \sum_{i\in\mathcal{B}} \pi_i(\nu(i),q) + \sum_{j\in\mathcal{S}} q(j) = \sum_{i\in\mathcal{B}} \pi_i(\nu(i),q) + q(\nu(i))$$

since each good j with $q(j) > 0$ is assigned to some bidder by ν. Thus, it suffices to show that

$$\sum_{i\in\mathcal{B}} \pi_i(\mu(i),p) - \pi_i(\nu(i),q) + p(\nu(i)) - q(\nu(i)) \geq 0.$$

We show that each summand is non-negative by distinguishing two cases:

Case 1: $p(\nu(i)) = b^i$. Then, since bidder i receives $\nu(i)$ in the outcome (ν, q), $q(\nu(i)) \leq b^i = p(\nu(i))$. Since by assumption $\pi_i(\mu(i),p) \geq \pi_i(\nu(i),q)$, it follows that $\pi_i(\mu(i),p) - \pi_i(\nu(i),q) + p(\nu(i)) - q(\nu(i)) \geq 0$.

Case 2: $p(\nu(i)) < b^i$. Then $\pi_i(\mu(i),p) \geq \pi_i(\nu(i),p)$ - otherwise, $(i,\nu(i))$ would be a blocking pair with respect to prices p. Since $\pi_i(\nu(i),p) = \pi_i(\nu(i),q) + q(\nu(i)) - p(\nu(i))$, we again get that $\pi_i(\mu(i),p) - \pi_i(\nu(i),q) + p(\nu(i)) - q(\nu(i)) \geq 0$. □

A.5 Proof of Lemma 5.1

Proof. We examine the edge gadget G'_e of edge $e = (u,u') \in E$, and define $\mathcal{B}_e := \{\beta_e, \gamma_e, \delta_e\}$ as the set of buyers and $\mathcal{S}_e := \{\alpha_e^u, \eta_e, \alpha_e^{u'}\}$ as the set of sellers. The total social welfare that can be achieved by agents in G'_e under assignment μ is $SW_e = \sum_{i\in\mathcal{B}_e} \pi_i(\mu(i)) + \sum_{j\in\mathcal{S}_e} \pi_j(\mu^{-1}(j))$.

For each individual edge gadget G'_e, we prove that the maximum total achieved welfare is bound by

$$SW_e \leq 3 \cdot 2|E|^2 + 27$$

We prove the claim as follows:

$$\begin{aligned}
SW_e &= \sum_{\substack{(i,j)\in\mathcal{B}_e\times\mathcal{S}_e\\(i,j)\in E(G_e)}} \Big(|\mu(i)=j|\cdot(\pi_i(j)+\pi_j(i))\Big) + \Big(|\mu(\epsilon_u)=\alpha_e^u|\cdot\pi_{\alpha_e^u}(\epsilon_u)\Big) + \Big(|\mu(\epsilon_{u'})=\alpha_e^{u'}|\cdot\pi_{\alpha_e^{u'}}(\epsilon_{u'})\Big)\\
&= \sum_{i\in\{\gamma_e,\delta_e\}} \Big(|\mu(i)=\alpha_e^u|\cdot(\pi_i(\alpha_e^u)+\pi_{\alpha_e^u}(i))\Big) + \sum_{i\in\{\beta_e,\gamma_e\}} \Big(|\mu(i)=\alpha_e^{u'}|\cdot(\pi_i(\alpha_e^{u'})+\pi_{\alpha_e^{u'}}(i))\Big)\\
&+ \sum_{i\in\{\beta_e,\delta_e\}} \Big(|\mu(i)=\eta_e|\cdot(\pi_i(\eta_e)+\pi_{\eta_e}(i))\Big) + \Big(|\mu(\epsilon_u)=\alpha_e^u|\cdot\pi_{\alpha_e^u}(\epsilon_u)\Big) + \Big(|\mu(\epsilon_{u'})=\alpha_e^{u'}|\cdot\pi_{\alpha_e^{u'}}(\epsilon_{u'})\Big)\\
&\leq 3\cdot 2|E|^2 + 27
\end{aligned}$$

We observe that each seller η_e is connected to two buyer nodes inside the edge gadget G'_e, therefore, under assignment μ, strictly one of conditions $|\mu(\beta_e) = \eta_e|$ and $|\mu(\delta_e) = \eta_e|$ returns a value of 1, while the other is set to 0. Sellers α_e^u and $\alpha_e^{u'}$ can be potentially matched to 3 different buyers: two of them within G'_e, and one belonging to the vertex gadget, u or u' respectively. We compute the upper bound on the total sum of valuations of sellers and buyers, on any possible assignment μ, maintaining that every time, one pair is matched, and the remaining conditions output 0 for all other edges connected to the pair. The resulting social welfare SW_e is equal to $3 \cdot 2|E|^2 + 27$ only for the case where assignment μ does not satisfy neither condition $|\mu(\epsilon_u) = \alpha_e^u|$ nor $|\mu(\epsilon_{u'}) = \alpha_e^{u'}|$.

We now proceed to bound the total welfare achieved under the condition that $\pi_{\epsilon_u}(\mu(\epsilon_u)) < \pi_{\epsilon_u}(\alpha_e^u)$ and $\pi_{\epsilon_{u'}}(\mu(\epsilon_{u'})) < \pi_{\epsilon_{u'}}(\alpha_e^{u'})$ hold simultaneously. Since μ is a stable outcome, from the first condition, we derive that

$$\pi_{\alpha_e^u}(\mu^{-1}(\alpha_e^u)) \geq \pi_{\alpha_e^u}(\epsilon_u) = 3 \tag{3}$$

Similarly, from the second condition, we get that

$$\pi_{\alpha_e^{u'}}(\mu^{-1}(\alpha_e^{u'})) \geq \pi_{\alpha_e^{u'}}(\epsilon_{u'}) = 7 \tag{4}$$

In order to ensure that Equation (3) holds, assignment μ needs to force node α_e^u to be matched with δ_e, as $\pi_{\alpha_e^u}(\delta_e) = 4 > 3$ in this case. If the pair $\{\gamma_e, \alpha_e^u\}$ was matched instead, it would hold that $\pi_{\alpha_e^u}(\gamma_e) = 2 < 3$, thus contradicting condition (3).

In a similar manner, for Equation (4), μ matches pair $\{\beta_e, \alpha_e^{u'}\}$, achieving utility $\pi_{\alpha_e^{u'}}(\beta_e) = 8 > 7$.

The total welfare of assignment μ is therefore:

$$\begin{aligned} SW_e &= \sum_{\substack{(i,j)\in\mathcal{B}_e\times\mathcal{S}_e \\ (i,j)\in E(G_e)}} \Big(|\mu(i)=j| \cdot (\pi_i(j) + \pi_j(i))\Big) + \Big(|\mu(\epsilon_u) = \alpha_e^u| \cdot \pi_{\alpha_e^u}(\epsilon_u)\Big) + \Big(|\mu(\epsilon_{u'}) = \alpha_e^{u'}| \cdot \pi_{\alpha_e^{u'}}(\epsilon_{u'})\Big) \\ &= |\mu(\delta_e) = \alpha_e^u| \cdot (\pi_{\delta_e}(\alpha_e^u) + \pi_{\alpha_e^u}(\delta_e)) + |\mu(\beta_e) = \alpha_e^{u'}| \cdot (\pi_{\beta_e}(\alpha_e^{u'}) + \pi_{\alpha_e^{u'}}(\beta_e)) \\ &\leq 2|E|^2 + 11 + 2|E|^2 + 10 \\ &\leq 4|E|^2 + 22 \\ &< 6|E|^2 + 27 - |V| \end{aligned}$$

since, without loss of generality, we can assume that $2|E|^2 - |V| + 5 > 0$, as we are referring to cubic graphs, where $3|V| = 2|E|$ and as a result the claim trivially holds. □

In order to prove NP-hardness for MBSBM, we need to provide an intermediate result concerning the existence of blocking pairs under certain assignments.

Lemma A.2. *Consider edge $e = (u, u')$ of G, and let μ be an assignment for edge gadget G'_e. The following two statements hold:*

(1) *If assignment μ satisfies $\mu(\beta_e) = \eta_e$, $\mu(\gamma_e) = \alpha_e^{u'}$ and $\mu(\delta_e) = \alpha_e^u$ (solid edges in Figure 1), then no blocking pair of μ involves any agent from $\{\beta_e, \gamma_e, \delta_e, \alpha_e^u, \eta_e\}$.*
(2) *If assignment μ satisfies $\mu(\beta_e) = \alpha_e^{u'}$, $\mu(\gamma_e) = \alpha_e^u$ and $\mu(\delta_e) = \eta_e$ (dashed edges in Figure 1), then no blocking pair of μ involves any agent from $\{\beta_e, \gamma_e, \delta_e, \alpha_e^{u'}, \eta_e\}$.*

Proof. For statement (1), we assume outcome μ sets values as suggested. Computing the utilities under μ for each agent, we get:

$$\pi_{\beta_e}(\eta_e) = 2|E|^2 + 3,\ \pi_{\gamma_e}(\alpha_e^{u'}) = 2|E|^2 + 6,\ \pi_{\delta_e}(\alpha_e^u) = 2|E|^2 + 7$$

$$\pi_{\alpha_e^u}(\delta_e) = 4,\ \pi_{\eta_e}(\beta_e) = 1$$

One can easily verify that agents $\beta_e, \gamma_e, \alpha_e^u$ are maximizing their utility under μ, and since agent η_e is indifferent between agents δ_e and β_e, they have no incentive to deviate by forming a blocking pair with agent δ_e. Therefore, there does not exist a blocking pair that includes agents $\{\beta_e, \gamma_e, \delta_e, \alpha_e^u, \eta_e\}$. For statement (2), we assume outcome μ sets values as suggested. Computing the utilities under μ for each agent, we get:

$$\pi_{\beta_e}(\alpha_e^{u'}) = 2|E|^2 + 2,\ \pi_{\gamma_e}(\alpha_e^u) = 2|E|^2 + 5,\ \pi_{\delta_e}(\eta_e) = 2|E|^2 + 9$$

$$\pi_{\alpha_e^{u'}}(\beta_e) = 8,\ \pi_{\eta_e}(\delta_e) = 1$$

One can easily verify that agents $\delta_e, \alpha_e^{u'}$ are maximizing their utility under μ, and since agent η_e is indifferent between agents δ_e and β_e, they have no incentive to deviate by forming a blocking pair with agent β_e. Given that agent $\alpha_e^{u'}$ is assigned to agent β_e, and is maximizing their utility, agent γ_e is assigned to agent α_e^u, as this constitutes their best available choice. Therefore, there does not exist a blocking pair that includes agents $\{\beta_e, \gamma_e, \delta_e, \alpha_e^{u'}, \eta_e\}$. □

A.6 Proof of Lemma 5.2

Proof. We examine the edge and vertex gadgets separately, and argue that, in both cases, setting prices equal to the winning bidders' budgets yields all feasible welfare-maximizing, core-stable outcomes. Formally, the set of all welfare-maximizing, core-stable assignments coincides with the set of optimal assignments when prices are set at the budget limit.

In a two-sided matching the welfare is defined as the gains from trade, the value of the buyers minus that of the sellers. This means, for each match between buyer i and seller j, the corresponding welfare is computed as the sum of buyer $\pi_i(j) = v_i(j) - p_j$ and seller payoff $\pi_j(i) = p_j - r_j$, therefore prices are not included in the final sum, which is the result of the difference between assigned items' valuations and seller reserve values.

We begin by analyzing the edge gadget, as seen in Figure 1. The values depicted in Table 1 represent the true valuation of each buyer among $\{\beta_e, \gamma_e, \delta_e\}$ for the item of each seller among $\{\alpha_e^u, \eta_e, \alpha_e^{u'}\}$. One can trivially observe that there exist 6 feasible assignments between buyers and sellers in the bipartite graph. As stated above, prices do not participate in the welfare computation, and thus we can calculate the welfare-maximizing assignment based on the buyer valuation table. Since valuations $v_{\beta_e}(\alpha_e^u) = v_{\gamma_e}(\eta_e) = v_{\delta_e}(\alpha_e^{u'}) = 0$, only 2 among 6 assignments are stable, and simultaneously welfare-maximizing. The two assignments μ_1, μ_2 are $\mu_1(\beta_e) = \eta_e, \mu_1(\gamma_e) = \alpha_e^{u'}, \mu_1(\delta_e) = \alpha_e^u$ and $\mu_2(\beta_e) = \alpha_e^{u'}, \mu_2(\gamma_e) = e_j^i, \mu_2(\delta_e) = \eta_e$. The total welfare admitted by assignments μ_1, μ_2 is equal to $6|E|^2 + 27$, while the remaining feasible assignments yield a strictly lower welfare. Additionally, according to Lemma A.2, assignments μ_1, μ_2 (grey, dotted) with seller payoffs set as $\pi_{\alpha_e^u}(\delta_e) = b_{\delta_e} - r_{\alpha_e^u} = 9 - 5 = 4$, $\pi_{\eta_e}(\beta_e) = b_{\beta_e} - r_{\eta_e} = 9 - 8 = 1$, $\pi_{\alpha_e^{u'}}(\gamma_e) = b_{\gamma_e} - r_{\alpha_e^{u'}} = 7 - 1 = 6$ under assignment μ_1, and $\pi_{\alpha_e^u}(\gamma_e) = b_{\gamma_e} - r_{\alpha_e^u} = 7 - 5 = 2$, $\pi_{\eta_e}(\delta_e) = b_{\delta_e} - r_{\eta_e} = 9 - 8 = 1$, $\pi_{\alpha_e^{u'}}(\beta_e) = b_{\beta_e} - r_{\alpha_e^{u'}} = 9 - 1 = 8$ under assignment μ_2, do not generate any blocking pairs. This argument leads to the conclusion that, for the edge gadget and corresponding valuation table, setting prices for winning buyers equal to their budgets produces every feasible core-stable, welfare-maximizing assignment, namely setting prices to lower values cannot yield different stable assignments of higher or equal welfare.

In a similar analysis, we observe that, for the vertex gadget of Figure 2, the welfare-maximizing assignment μ_3, based on valuations described on Table 2, is $\mu_3(\epsilon_u) = N(\epsilon_u), \mu_3(\zeta_u) = \xi_u, \mu_3(\theta_u) = \lambda_u$. However, as suggested in Lemma 5.1, assignment $\mu_0(\epsilon_u)$ results in sub-optimal welfare for the overall subgraph including vertex and edge gadget. Therefore, the welfare-maximizing assignment μ_4 is a result of an assignment between the vertices within the gadget. The aforementioned assignment is $\mu_4(\epsilon_u) = \xi_u, \mu_3(\zeta_u) = \kappa_u, \mu_3(\theta_u) = \lambda_u$ (solid assignment in Figure 2), admitting welfare equal to $3|V| + 27$. In Theorem 5.3, we have shown that assignment μ_4, for winning buyer prices equal to their budgets is core-stable. Thus, the initial claim is true for each vertex gadget, concluding the proof of the lemma. □

A.7 Proof of Theorem 5.3

Proof. Let $G = (V, E)$ be a cubic graph, with sizes of vertex and edge sets defined as $|V|$ and $|E|$ respectively. Since G is assumed to be cubic, it does not possess any isolated vertices. An instance of MIS is defined by G and an integer k.

A key property of cubic graphs is that all nodes must have a degree of 3. According to the *Handshaking lemma*, it holds that $\sum_{u \in V} deg(u) = 2|E|$, namely the sum of degrees of all vertices of the graph is twice as large as than the size of the edge set.

We construct an instance $< G', \pi_i, \pi_j, SW >$ of MBSBM, where $G' = (V', E')$ is a bipartite graph, with $SW = |V| \cdot (3|V| + 26) + k + |E| \cdot (6|E|^2 + 27)$. For each edge $e \in E$, buyer and seller sets of the corresponding edge gadget are defined as $\mathcal{B}_e, \mathcal{S}_e$. For each vertex $u \in V$, the respective sets for the corresponding vertex gadget are defined as $\mathcal{B}_u, \mathcal{S}_u$. Unifying for all edge gadgets in G', we define $\mathcal{B}_E = \bigcup_{e \in E} \mathcal{B}_e$, $\mathcal{S}_E = \bigcup_{e \in E} \mathcal{S}_e$, and for all vertex gadgets, sets $\mathcal{S}_V = \bigcup_{u \in V} \mathcal{S}_u$ and $\mathcal{B}_V = \bigcup_{u \in V} \mathcal{B}_u$. Vertex set V' consists of the union over all vertex and edge gadgets, and therefore $\mathcal{B} = \mathcal{B}_E \cup \mathcal{B}_V$ and $\mathcal{S} = \mathcal{S}_E \cup \mathcal{S}_V$ represent the total number of buyers and sellers respectively. Each edge gadget corresponding to edge $e \in E$ consists of vertices $\mathcal{B}_e = \{\beta_e, \gamma_e, \delta_e\}$ and $\mathcal{S}_e = \{\alpha_e^u, \eta_e, \alpha_e^{u'}\}$. Similarly, the vertex gadget corresponding to vertex $u \in V$ consists of two disjoint sets of vertices defined as $\mathcal{B}_u = \{\epsilon_u, \zeta_u, \theta_u\}$ and $\mathcal{S}_u = \{\kappa_u, \lambda_u, \xi_u\}$. The vertex and edge gadgets are defined for each vertex $u \in V$ and each edge $e \in E$ of the original graph G. The total size of each set is $|\mathcal{B}| = |\mathcal{S}| = 6|E| + 6|V|$.

Formally, the following claim should be proven: G has an independent set $IS(G)$ of size at least k **if and only if** G' admits a stable outcome with welfare at least SW.

The transformation from an instance of MIS to an instance of MBSBM is performed in polynomial time. An important aspect of our construction is the assumption that, both in edge and vertex gadgets, there exist pairs of buyers with equal budgets. From Tables 1 and 2, one can observe that pairs of agents $\{\beta_e, \delta_e\}$ and $\{\epsilon_u, \theta_u\}$ verify the claim. Thus, the instance is not in general position, and therefore the hardness proof holds for cases where the property is violated.

We now proceed to the forward part of the proof, namely prove that, given an independent set $IS(G)$ of size at least k, we construct an assignment μ as follows:

- For each $u \in V$, if $u \in IS(G)$, set $\mu(\epsilon_u) = \xi_u, \mu(\zeta_u) = \kappa_u, \mu(\theta_u) = \lambda_u$.
- For each $u \in |V|$, if $u \notin IS(G)$, set $\mu(\epsilon_u) = \lambda_u, \mu(\zeta_u) = \xi_u, \mu(\theta_u) = \kappa_u$.
- For each edge $e = \{u, u'\} \in E$ do:
  - If $u \in IS(G)$, set $\mu(\beta_e) = \eta_e, \mu(\gamma_e) = \alpha_e^{u'}, \mu(\delta_e) = \alpha_e^u$, which matches condition (1) of Lemma A.2.
  - If $u \notin IS(G)$, set $\mu(\beta_e) = \alpha_e^{u'}, \mu(\gamma_e) = \alpha_e^u, \mu(\delta_e) = \eta_e$, which matches condition (2) of Lemma A.2.
- For any pair of agents $(i, j) \in G'$ not assigned above, set $\mu(i) = 0, \mu^{-1}(j) = 0$.

Since the size of the independent set is at least k, we can trivially verify that, according to the aforementioned rules, the welfare attained by the agents of all vertex gadgets is equal to $|V| \cdot (3|V| + 26) + k$ and of all edge gadgets is $|E| \cdot (6|E|^2 + 27)$, thus achieving a total welfare of SW. In what follows, we prove that outcome μ is indeed stable.

Edge Gadget. Examining the edge gadget of edge $e = \{u, u'\} \in E$, we need to prove that there does not exist any blocking pair of agents. There are two distinct cases, as mentioned previously.

(1) If $u \in IS(G)$, then μ corresponds to condition (1) of Lemma A.2. Therefore, no blocking pair involving agents $\{\beta_e, \gamma_e, \delta_e, \alpha_e^u, \eta_e\}$ exists. In this case, assigning $\alpha_e^{u'}$ to $\epsilon_{u'}$ would yield a higher utility for agent $\alpha_e^{u'}$. However, since $u' \notin IS(G)$, by definition $\mu(\epsilon_{u'}) = \lambda_u$ and $\epsilon_{u'}$ is assigned to her most preferred agent, with no incentive to deviate. We conclude that no blocking pair involving agents $\{\alpha_e^{u'}, \epsilon u'\}$, guaranteeing stability for μ.

(2) If $u \notin IS(G)$, then μ corresponds to condition (2) of Lemma A.2. Therefore, no blocking pair involving agents $\{\beta_e, \gamma_e, \delta_e, \alpha_e^{u'}, \eta_e\}$ exists. In this case, assigning α_e^u to ϵ_u would yield a

higher utility for agent α_e^u. However, since $u \notin IS(G)$, by definition $\mu(\epsilon_u) = \lambda_u$ and ϵ_u is integrally matched to her most preferred agent, with no incentive to deviate. We conclude that no blocking pair involving agents $\{\alpha_e^u, \epsilon_u\}$, guaranteeing stability for μ.

We prove a similar result for the vertex gadgets of G'.

Vertex Gadget. Examining the vertex gadget of vertex $u \in V$, we again detect two cases.

(1) If $u \in IS(G)$, then no blocking pair of μ involves agents $\{\xi_u, \theta_u, \lambda_u\}$, as μ assigns them to their most preferred agents. Therefore, no blocking pair involves agents $\{\zeta_u, \kappa_u\}$, since no possible combination can yield improved payoff for all participants. As previously argued, no agent from the edge gadget participates in a blocking pair, and thus ϵ_u is also not involved in a blocking pair.
(2) If $u \notin IS(G)$, then no blocking pair of μ involves agents $\{\kappa_u, \zeta_u, \lambda_u\}$, as μ assigns them to their most preferred agents. The same holds for agent ϵ_u. Therefore, no blocking pair involves agents $\{\theta_u, \xi_u\}$, since no possible combination can yield improved payoff for all participants.

We conclude that neither the edge nor vertex gadget contain agents involved in blocking pairs under μ. Thus, μ is stable and achieves a welfare of SW.

For the reverse direction, assume μ is a stable outcome for $< G', \pi_i, \pi_j >$ with welfare $welfare(\mu) \geq SW$. Defining a subset of the vertices as $IS(G) = \{u \in V \mid \mu(\epsilon_u) = \xi_u\}$, we prove that $IS(G)$ is an independent set of G of size at least k.

Firstly, we prove the desired lower bound on $|IS(G)|$. We define $\mu(\epsilon_u) = N(\epsilon_u)$ as the assignment between vertex ϵ_u and any of the neighboring vertices from the edge gadgets. Summing up for all $|V|$ vertices of G, we get:

$$\sum_{u \in V} \Big(\sum_{i \in \{\epsilon_u, \zeta_u, \theta_u\}} \pi_i(\mu(i)) + \sum_{j \in \{\kappa_u, \lambda_u, \xi_u\}} \pi_j(\mu^{-1}(j)) \Big)$$

$$= \sum_{u \in V} (|V| + 10) \cdot |\mu(\epsilon_u) = \lambda_u| + (|V| + 12) \cdot |\mu(\epsilon_u) = \xi_u| + (|V| + 8) \cdot |\mu(\epsilon_u) = N(\epsilon_u)| +$$

$$(|V| + 7) \cdot |\mu(\zeta_u) = \kappa_u| + (|V| + 9) \cdot |\mu(\zeta_u) = \xi_u| + (|V| + 7) \cdot |\mu(\theta_u) = \kappa_u| + (|V| + 8) \cdot |\mu(\theta_u) = \lambda_u|$$

$$\leq \sum_{u \in V} (|V| + 10)|\mu(\epsilon_u) = \lambda_u| + |\mu(\epsilon_u) = \xi_u| + |\mu(\epsilon_u) = N(\epsilon_u)| + (|V| + 8) \cdot |\mu(\zeta_u) = \kappa_u| + |\mu(\zeta_u) = \xi_u| +$$

$$(|V| + 8) \cdot |\mu(\theta_u) = \kappa_u| + |\mu(\theta_u) = \lambda_u| + \sum_{u \in V} |\mu(\epsilon_u) = \lambda_u|$$

$$\leq |V| \cdot (3|V| + 26) + \sum_{u \in V} |\mu(\epsilon_u) = \lambda_u|$$

Since μ corresponds to a binary assignment, only one condition holds true for each of $\epsilon_u, \zeta_u, \theta_u$, and examining all possible outcomes, we use the mean of values to provide an upper bound on utilities. The sum of values for each vertex $\kappa_u, \lambda_u, \xi_u$ is at most 1, which is also taken into account when computing the upper bound. We can trivially verify that, the total welfare of μ is maximized if and only if pair ϵ_u, ξ_u is integrally matched.

We now need to show that $|IS(G)| \geq k$. Set $IS(G)$ has been defined as the set of vertices ϵ_u for $u \in V$, that are integrally matched to ξ_u. This can be expressed as $|IS(G)| = \sum_{u \in V} |\mu(\epsilon_u) = \xi_u|$. It then suffices to prove inequality $\sum_{u \in V} |\mu(\epsilon_u) = \xi_u| \geq k$, to prove the lower bound on the size of the independent set. Since the induced welfare from all $|E|$ edge gadgets is at most $welfare(\mu) \leq |E| \cdot (6|E|^2 + 27)$. At least $|V| \cdot (3|V| + 26) + \sum_{u \in V} |\mu(\epsilon_u) = \xi_u|$ must be derived from the vertex gadgets.

Using the upper bound provided above, we conclude that $\sum_{u \in V} |\mu(\epsilon_u) = \xi_u| \geq k$, as required, thus proving the lower bound on the size of $IS(G)$.

Finally, we need to prove that $IS(G)$ is in fact an independent set of G. We prove the claim by contradiction. Suppose there is an edge $e = \{u, u'\} \in V$ and $\{u, u'\} \subseteq IS(G)$. Then, for both nodes u, u' under assignment μ, it must hold that $\pi_{\epsilon_u}(\mu(\epsilon_u)) < |V| + 8 = \pi_{\epsilon_u}(\alpha_e^u)$ and $\pi_{\epsilon_{u'}}(\mu(\epsilon_{u'})) < |V| + 8 = \pi_{\epsilon_{u'}}(\alpha_e^{u'})$, as $IS(G)$ is defined as the set of nodes u that are integrally matched to ξ_u. From Lemma 5.1, the induced welfare from agents of the edge gadgets is at most $6|E|^2 + 27 - |V|$, and using the previously shown bound, the welfare induced from agents of the vertex gadgets is at most $3|V| + 26 + |V|$. Therefore, the total welfare of G' under μ is lower than the original assumption, which is a contradiction. Note that the restriction of our prices to be at the budget constraint b^i of the buyers $i \in \mathcal{B}$ in the edge and vertex gadgets is without loss of generality as shown in Lemma 5.2. Concluding, we have proven that indeed $IS(G)$ is an independent set of G.

□

4 Contribution 2: Budget-Feasible Market Design for Biodiversity Conservation: Considering Incentives and Spatial Coordination

Peer-Reviewed Conference Paper

Title: Budget-Feasible Market Design for Biodiversity Conservation: Considering Incentives and Spatial Coordination

Authors: Eleni Batziou, Martin Bichler

In: 18th International Conference on Wirtschaftsinformatik (WI ’23)

Abstract: How to best incentivize farmers to conserve biodiversity on private land is an important policy question. Conservation auctions provide a mechanism to elicit farmers’ opportunity costs, but their design is challenging and often suffer from low participation due to strategic complexity. Conservation auctions should ideally be incentive-compatible, address spatial synergies that maximize biodiversity gains, and respect the predefined budget of the government. Recent advances in mechanism design suggest budget-feasible auctions, but little is known about average-case efficiency. Based on this line of research, we introduce an incentivecompatible conservation auction mechanism that considers the bid taker’s spatial synergies and respects budget. The results are compared against the celebrated Vickrey-Clarke-Groves mechanism. Our numerical results estimate the efficiency loss that can be expected for different assumptions on the synergistic values of the government. We provide evidence that budget-feasible mechanisms provide a new tool for policymakers in this domain.

Contribution of dissertation author: Methodology, software, experimental design, investigation, visualization, joint paper management.

Citation: Batziou and Bichler (2023)

Budget-Feasible Market Design for Biodiversity Conservation: Considering Incentives and Spatial Coordination

Research Paper

Eleni Batziou[1] and Martin Bichler[1]

School of Computation, Information and Technology, Technical University of Munich, Germany
{batziou,bichler}@cit.tum.de

Abstract. How to best incentivize farmers to conserve biodiversity on private land is an important policy question. Conservation auctions provide a mechanism to elicit farmers' opportunity costs, but their design is challenging and often suffer from low participation due to strategic complexity. Conservation auctions should ideally be incentive-compatible, address spatial synergies that maximize biodiversity gains, and respect the predefined budget of the government. Recent advances in mechanism design suggest budget-feasible auctions, but little is known about average-case efficiency. Based on this line of research, we introduce an incentive-compatible conservation auction mechanism that considers the bid taker's spatial synergies and respects budget. The results are compared against the celebrated Vickrey-Clarke-Groves mechanism. Our numerical results estimate the efficiency loss that can be expected for different assumptions on the synergistic values of the government. They provide evidence that budget-feasible mechanisms provide a new tool for policymakers in this domain.
Keywords: Conservation auction, incentive compatibility, budget feasibility, mechanism design

1 Introduction

Wildlife populations have declined by more than two-thirds in less than 50 years according to the World Wildlife Foundation.[1] As in many studies, it is shown that biodiversity is being destroyed at a rate unprecedented in history. Given the prominent role of private land use in achieving improvements to biodiversity, changing the behavior of those who manage private land in a manner that benefits biodiversity is a key objective (Armsworth et al. 2012). Incentive-based agri-environmental schemes (AES) that encourage landowners to undertake costly ecosystem services or conservation activities for biodiversity have been globally growing in popularity. AES combine information systems and market design to coordinate relevant stakeholders (Gholami et al. 2016). So far, incentives have largely been offered as simple posted-price schemes. In most of these schemes, landholders are paid a fixed price for a service, such as price per hectare for a land piece converted from agricultural use to grassland by a farmer. Prices can be too high

[1] `https://www.bbc.com/news/science-environment-54091048`

leading to over-provisioning or too low such that they do not provide enough incentive to participate. Information asymmetries between farmers render the implementation of biodiversity goals in an efficient way challenging (Gómez-Limón et al. 2019), or create an adverse selection problem (Lundberg et al. 2018).

Conservation auctions have been promoted as an alternative to such posted-price schemes. These auctions are a specific type of payment for ecosystem service mechanisms where a single buyer (government or regulatory body) elicits bids from private producers to enter into contracts regulating how land is to be managed. Through competition between farmers, the government hopes to reduce the information rents which farmers earn, and thus improve the cost-effectiveness of the budget spent on biodiversity measures. Conservation Reserve Program in the US is arguably the world's largest payments-for-ecosystem services program. It has been in operation since 1985 and uses pay-as-bid auctions (Cramton et al. 2021). The benefits of a switch to conservation auctions have been estimated within range of 24% to 33% for Australian conservation programs (Windle & Rolfe 2008). Rousseau & Moons (2007) argue that the adoption of an auction mechanism can increase social welfare by 22% compared to the fixed-price scheme being employed under the afforestation program in Belgium.

Running an auction is not common in the EU even though it is permitted (Article 49.3 of Regulation (EU) No 1305/2013) (Grohe, 2009; Allen et al. (2014); Latacz-Lohmann & Breustedt (2019)). One reason is the difficulty in designing appropriate market mechanisms (Schilizzi 2017). Some questions are similar to those in other market design problems. For example, one needs to decide the product being auctioned, i.e., whether payments are based on the performance or actions of bidders. The payment rule matters, and so do the bid selection criteria, and the bidding format (single-round or multiple-round). However, a few aspects of conservation auctions are specific:

(1) To be environmentally effective for biodiversity conservation, there is a need for an auction design that *encourages spatial coordination* of conservation actions, since the successful realization of many biodiversity objectives depends on conservation actions occurring on neighboring sites or along wildlife corridors. A wildlife corridor is a habitat area connecting wildlife populations separated by human activities or structures. The absence of spatial coordination of conservation efforts leads to far from optimal results. This problem constitutes a major source of AES ineffectiveness (Nguyen et al. 2022). In the past, conservation auctions did not promote landscape-level coordination of conservation efforts across landholders but focused on outcomes on a farm level (Nguyen et al. (2022), Hanley et al. (2022), Iftekhar & Latacz-Lohmann (2017*a*), Lamb et al. (2016)). The conservation of several patches of land that are adjacent might lead to a wildlife corridor and be preferable to the conservation of two non-adjacent patches of the same size. The switch from a farm-scale and fragmented conservation approach to a landscape-scale approach has been argued to be an important change for AES to meet their environmental goals more effectively (Westerink et al. 2017). In summary, spatial and temporal dependencies among ecosystem services need to be considered (Wätzold et al. 2016), requiring non-standard auction formats to get to optimal land-use patterns.

(2) *Participation* in conservation auctions has long been a concern (Rolfe et al. 2018). Participants in such auctions are laymen and simplicity and transparency are key. The complexity of auctions was raised as a central reason for low participation in

conservation auctions (Palm-Forster et al. 2015). To address this concern, a version of the Vickrey-Clarke-Groves (VCG) mechanism was proposed to implement a welfare-maximizing outcome (Polasky et al. 2014). Maximizing welfare implies minimizing costs for governments in conservation auctions. The VCG mechanism is the unique strategyproof mechanism (i.e., dominant strategy incentive-compatible) for general valuations that achieves this goal, a significant advantage over alternative auctions. However, the underlying assumption is that buyers and sellers have no financial constraints.

(3) Instead of cost-minimization governments typically want to maximize the level of conservation of endangered species and habitats at a given budget. In other words, the government wants to minimize cost subject to a hard budget constraint. Incentive compatibility and *budget feasibility* are incompatible in general, which follows from the fact that the VCG mechanism is unique. As a result, the VCG mechanism does not provide a viable option to governments.

Conservation auctions used today violate one or more of these properties (Nguyen et al. 2022). While incentive compatibility and budget feasibility clash for general valuations, one can hope for auction designs that address these properties for restricted valuations. Indeed, progress is made on this front in the mechanism design literature.

This work adds to the literature on procurement auctions. While industrial procurement auctions focus on cost minimization (Dasgupta & Spulber 1989, Bichler et al. 2006, Anton & Yao 1992, Kokott et al. 2019), the main difference of conservation auctions is the presence of budget constraints. A small theoretical literature exists on such budget-feasible auctions, but applications in the field have not yet been explored.

A number of authors proposed strategyproof budget-feasible procurement auctions for single-minded bidders, i.e. bidders who are only interested in selling a particular set of goods or services and not multiple combinations. In conservation auctions, this is usually satisfied. Farmers own land parcels that they can either develop and profit or conserve for agri-environmental purposes such as wildlife protection. However, until recently the proposed mechanisms were randomized sealed-bid mechanisms (Singer 2010). Balkanski et al. (2022) proposed a deterministic clock auction that respects budgets. Budget-feasible auction mechanisms have not yet been explored for conservation auctions.

Collusion can be a particular challenge in mechanism design, since it can lead to lower income for the principal. The VCG algorithm is prone to collusion, which may lead to arbitrarily low auctioneer revenue, as argued in Polasky et al. (2014). On the contrary, BFA satisfies *weak group-strategyproofness*, namely no coalition of sellers can misreport such that they all profit. This property guarantees collusion-resistance. This argument strengthens our claim that the BFA poses as an attractive alternative for policymakers, as it eliminates the risk of spatial synergies.

The price to be paid for such strong incentive properties and budget feasibility is efficiency. The clock auction by Balkanski et al. (2022) but also the sealed-bid predecessors are approximation mechanisms and do not implement the welfare-maximizing outcome. For monotone submodular valuations, the worst-case approximation ratio of BFA is 4.75. Thus, welfare could be almost five times worse than the maximum welfare solution. For non-monotone submodular valuations, the worst-case ratio is 64. For submodular and subadditive valuations, the auction matches the bounds of the best randomized budget-feasible auctions. While such bounds are interesting from an algorithmic perspective, the worst-case might be too pessimistic. A recent stream of work in theoretical computer science seeks insights beyond a purely worst-case analysis (Roughgarden 2019).

In this paper, we study the average-case efficiency of incentive-compatible and budget-feasible conservation auctions, in particular, the clock auction by Balkanski et al. (2022), which provides the most recent realization and is deterministic and group-strategyproof. For these reasons, we focus on this auction among the set of budget-feasible auctions. The strong incentive properties of this auction allow us to focus on numerical experiments as we can expect bidders to follow their dominant strategy.

In standard combinatorial auctions, synergies are on the bidders' side and are typically unknown to the auctioneer. In conservation auctions, the synergistic valuations are on the bid taker's side and thus known in advance to the auctioneer. Based on our analysis and knowledge about complementarities, a government can decide whether the expected efficiency loss due to budget-feasible auctions is acceptable for a particular application. This sheds light on the question of whether budget-feasible auctions can provide a practical policy tool to implement agri-environmental services.

We show that efficiency loss due to BFA in the case of submodular values is at around 25% on average. For additive or superadditive valuations, the efficiency loss is between 36% and 51% with very high levels of superadditivity. Domain-specific value models for wildlife corridors achieve similar relative welfare gains, but the budget violation with VCG is even worse. We argue that with or without hard budget constraints, such variations in the payment will be unacceptable for most policymakers.

Auction mechanisms used in some countries are neither incentive-compatible nor consider landscape-level complementarities for the government. The complexity of conservation auction mechanisms is the key reason behind low participation rates of farmers. Indeed, a major shortcoming of sealed-bid auctions is the lack of *transparency*, as participants must trust the auctioneer to correctly implement the auction algorithm and not mishandle their private information. By design, clock auctions, and as such the BFA, are *simple*, meaning that participants only see the price offered to them and need not understand the intrinsic details of the mechanism implemented. Against this background, we argue that BFA provides a powerful new tool for policymakers in this field.

2 Related Work

There is extensive literature on conservation auctions to incentivize biodiversity measures (Nguyen et al. 2022). One line of research concerns combinatorial auctions. Such auctions play a role when farmers have significant synergies across services ((Nemes et al. 2008, Saïd & Thoyer 2007, Iftekhar et al. 2009)). We focus on farmland conservation, where individual landowners can decide to develop or conserve a patch of land. Synergies may arise on the side of the government and on a landscape level across multiple land patches. The goals of such auctions include enhancing the population of farmland birds, protecting native vegetation, or restoring wetlands on farms. On a landscape level, it can be useful to conserve adjacent patches of land (e.g., wildlife corridors) or other combinations across the landscape, which all promise different biodiversity gains to the government. Polasky et al. (2014) focused on this widespread scenario and proposed a VCG mechanism. The VCG mechanism is strategyproof if the regulator aims for a particular conservation target and does not have budget constraints. We refer to such auctions as target-constrained auctions.

Typically, governments do not decide on a biodiversity target, but on a budget devoted to biodiversity (Hellerstein 2017). In the US Conservation Reserve Program, each potential supplier can make a bid detailing how much they will accept for agreeing to the contract terms. The buyer then orders the bids in terms of either bid price alone, or bid price weighted by some environmental metric, and selects the most cost-effective bids until budget is exhausted, some quantitative program target is achieved (e.g., cumulative hectares enrolled), or a reserve price is reached for bids (Hellerstein 2017). However, such auction formats neither consider spatial coordination nor are incentive-compatible.

Nguyen et al. (2022) provide a thorough review of economic mechanisms built with the purpose of generating incentives for spatial coordination. In traditional AES, landowners usually operate in isolation, which leads to suboptimal outcomes when it comes to ecosystem preservation, given the scattered nature of preserved parcels, or even failure to meet preservation targets. Hence, spatial coordination arises as a significant property, exceeding the capabilities of conventional design models. The existing auction formats to address the problem determine a bonus added on top of individual payments in case of coordination.

Budget-feasible but incentive-compatible mechanisms have been developed in the literature on algorithmic mechanism design but received little attention outside. In particular, they have not been explored in the context of conservation auctions, even though they address central design desiderata as discussed earlier. A few sealed-bid (randomized), incentive-compatible and budget-constrained auctions were proposed for restricted types (e.g., additive or submodular) of buyer preferences (e.g., Singer (2010) or Gravin et al. (2020)). However, randomized mechanisms, while elegant algorithmically, are rarely used in practice. In recent work, Balkanski et al. (2022) introduced deterministic auctions that are also budget-feasible and incentive-compatible and match or even improve on the best-known randomized approximation ratio. These auctions are run as clock auctions and match the best-known worst-case approximation bounds by any other polynomial-time strategyproof auctions for restricted types of valuations such as submodular and subadditive. These worst-case approximation ratios might not be appealing to decision-makers, but also not appropriate for decision-making. In this article, we aim to understand the average-case efficiency of budget-feasible auctions and whether they provide a viable option for policymakers.

3 Preliminaries

3.1 Notation

Our analysis is inspired by Polasky et al. (2014), who consider a grid of land parcels each belonging to one farmer $N = \{1, 2, \dots, n\}$. An auctioneer (*buyer* or government) seeks to acquire land parcels, in the form of bundles of parcels, and we thus define a set of combinations or packages of parcels as $S \subseteq N$. The landowner of each parcel $i \in N$ (*seller* or farmer) has two possible courses of action: *conserve*, translated to no income for the owner but positive effect on biodiversity, or *develop*, leading to little or no ecosystem value but profit for the individual. The action is represented as a binary variable y_i, where $y_i = 0$ denotes the case where the items in bundle S are developed (and thus bring no benefit to the auctioneer), and $y_i = 1$ in case of preservation.

For each individual land parcel i, the corresponding owner has an opportunity cost (c_i) for developing the parcel, which remains private information. She would only accept to preserve parcel i if price $p_i \geq c_i$. The price or payment vector $p = (p_1, p_2, \ldots, p_n)$ is determined by an auction mechanism. The auctioneer is subject to a hard budget limit of B, and thus the sum of individual payments should not exceed this constraint. In addition, the auctioneer has a publicly known value function $v : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ for each combination of parcels. We examine three different families of valuation functions: additive, submodular, and superadditive to cover a broad range of possible valuations. These are the types of set functions that are widely analyzed in the algorithmic literature:

– **Additive:** $\forall S, T \subseteq \{N\} : \;\; v(S) + v(T) = v(S \cup T) - v(S \cap T)$
– **Superadditive:** $\forall S, T \subseteq \{N\} : \;\; v(S) + v(T) \leq v(S \cup T)$
– **Submodular:** $\forall S, T \subseteq \{N\} : \;\; v(S) + v(T) \geq v(S \cup T) + v(S \cap T)$

3.2 Optimal Allocation

We introduce a model to compute the optimal allocation (OPT), which serves as a baseline to compare against incentive-compatible auction mechanisms. This integer program does not generate prices but guarantees that at least the total cost of the farmers is within budget. This does not mean that the result of an auction mechanism where prices might be higher than costs will still be within budget. However, this constrained welfare maximization provides a useful baseline to compare outcomes of auction mechanisms.

$$\begin{array}{lll} \max \sum_{S \subseteq N} v(S)x(S) - \sum_{i \in N} c_i y_i & & \\ \text{s.t.} \quad \sum_{S \subseteq N} x(S) \leq 1 & & \\ \quad \sum_{S: i \in S} x(S) \leq y_i & \forall i \in N & \\ \quad \sum_{i \in N} c_i y_i \leq B & & \textbf{(OPT)} \\ \quad x(S) \in \{0, 1\} & \forall S \subseteq N & \\ \quad y_i \in \{0, 1\} & \forall i \in N & \end{array}$$

The objective function maximizes gains from trade. The first constraint makes sure that only one package of the government is selected, and the second enforces that a package that is bought also has farmers selling their corresponding pieces of land. The third constraint guarantees that the final allocation respects budget. The last constraints enforce the integrality of the solution. $V^B(X)$ is defined as the objective function value of OPT for a feasible allocation $X = \{x(S)\}_{S \subseteq N}$, and $V(X)$ the objective function value of OPT for a feasible allocation X without the budget constraint. The welfare maximizing allocation, ignoring the budget constraint, is then given by $X^* = argmax_X[V(X)]$.

3.3 The VCG Mechanism

Based on the above notation, it is straightforward to define the VCG mechanism. In this sealed-bid mechanism, sellers i report their costs c_i to the auctioneer. The auctioneer then computes the maximum welfare $V(X^*)$. In order to determine the payments for each winning seller i, the auctioneer computes the optimal allocation X^*_{-i} without the winning seller i. The difference $\Delta V_i = V_i(X^*_i) - V_i(X^*_{-i})$ defines the VCG payment p_i for a winning farmer i. In single-sided (conservation) auctions, these payments are such that the auctioneer does not make a loss. However, if the auctioneer has an additional exogenous and binding budget B, they might make a loss with respect to this budget.

4 Budget-Feasible Auctions

The core of the algorithm by Balkanski et al. (2022) is a backward greedy technique designed for submodular function maximization. In this auction, the auctioneer with budget B proposes a price $p_{i,t}$ to each seller (or landowner) i in phase t, that is non-increasing such that $p_{i,t} \leq p_{i,t-1}$, and is computed using the available public information. The sellers have the option to either reject the price and permanently exit the auction, or continue by accepting this lower price. In the latter case, they are included in the set A_t of *active* sellers at round t. Since in every iteration, we expect more sellers to reject, naturally it holds that $A_t \subseteq A_{t-1} \subseteq \cdots \subseteq A_1 \subseteq N$. At phase $\hat{t}$, the auction terminates, a subset of $A_{\hat{t}}$ is chosen as the winning set W, and each seller receives a price equal to the last accepted offer. The auction is *budget-feasible* if $\sum_{i \in W} p_{i,\hat{t}} \leq B$.

The auction outcome is compared against the resulting welfare of OPT, which produces the optimal allocation for the case when all private costs are known, with a price $p_i = c_i$ paid to each seller. The approximation ratio is measured as the factor $\rho = \frac{OPT}{BFA}$, where BFA describes the welfare achieved with the BFA algorithm. The theoretical guarantee provided in this work is that, for monotone submodular value functions, the approximation ratio for BFA is 4.75.

The clock auction is described as a two-stage process. During the first stage, for each phase t, an estimate of the value of the optimal solution, $\tilde{OPT}$, is updated. $\tilde{OPT}$ is initiated with a pessimistic estimation, equal to the maximum value single item. In each phase, we seek to determine a set of sellers S_t, such that the total value of their union exceeds $\tilde{OPT}$. The seller with the highest marginal contribution $v(\{i\}|S_t)$ to set S_t is considered in each iteration and is offered a price equal to $p_{i,t} = min\{p_{i,t}, v(\{i\}|S_t) \cdot \frac{B}{\tilde{OPT}}\}$. This price $p_{i,t}$ corresponds to the marginal contribution of i to S_t scaled to reach $\tilde{OPT}$ within budget B. If the price is accepted by i, the auctioneer adds i to set S_t, or i is removed from A otherwise. When $v(S_t) \geq \tilde{OPT}$ or there are no active sellers left to be offered a price, phase t comes to an end.

As a new phase $t > 1$ begins, the estimate is updated to $\tilde{OPT}_t = 2 \cdot \tilde{OPT}_{t-1}$, and set S_{t-1} represents the selected sellers of the previous phase. If the union $S_{t-1} \cup S_t$ contains all active sellers at the end of phase t, the mechanism terminates.

At the second stage of the process, once the clock auction has terminated, sets S_{t-1} and S_t are fed to a routine for submodular function maximization: a bidder is removed from the set of phase $t-1$ and added to phase t, to maintain budget feasibility. Finally, the maximum value budget-feasible set is selected, as a combination of the elements belonging to sets S_{t-1} and S_t. A detailed pseudo-code is provided in the appendix.

5 Research Design

In our numerical experiments, we compare OPT with the VCG and BFA mechanisms. Farmers grow a range of crops and employ different production technology, hence we assume a standard private-value model, where valuations are drawn from a certain distribution. The treatment variables of our numerical experiments include the value model, budget constraint, and grid size. The opportunity costs c_i of the landowners are drawn i.i.d from a uniform distribution in the range $[0, 50]$. Results are reported on a 3x3 grid. We conduct experiments with larger grid sizes, up to sizes 7x7, but the impact of grid size on the results was negligible.

The auctioneer's value function, $v(S)$, for packages of parcels $S \subseteq N$ is divided into several classes of set functions: additive, submodular, superadditive valuations, and domain-specific functions. This is a central treatment variable in our experiments. The first three classes are typical in the auction design literature, while the latter is motivated by the domain of conservation auctions. Note that additive $\subset$ submodular functions.

First, we report purely *additive functions*, adopting Polasky et al. (2014). Each item is assigned a value v_i drawn i.i.d. from the uniform distribution within range $[0, 100]$. Here, the value of a set of parcels is the sum of their individual values to the auctioneer.

Superadditive value functions capture general complementarities that result from aggregating pieces of land. If the auctioneer can gather multiple parcels, the added value is more beneficial to ecosystem preservation. For the case of superadditive functions, for each package, the sum of values is incremented by a multiplier in the range $(1, 2]$.

We generate all packages or combinations of items $S \subseteq N$, regardless of their position on the grid, and assign a value for each subset by adding the individual values v_i in a bottom-up manner and multiplying with the multiplier.

Submodular functions are generated following the rule: $v(S \cup T) \leq v(S) + v(T) - v(S \cap T)$. For the generation of submodular valuations, we define the maximum value resulting from adding any two subsets S, T and subtracting the value of their intersection. This value provides an upper bound on the value that $S \cup T$ is allowed to take. The minimum allowed value for $S \cup T$ is the maximum value of S, T. Additive, submodular, and superadditive set functions are widely used in discrete mathematics and algorithm design, but do not always adequately capture the specifics of conservation auctions.

Iftekhar & Tisdell (2014) emphasize wildlife corridors as a landscape pattern with significant biodiversity benefits. Such corridors can be seen as rows on a grid (see Figure 1). Once a wildlife corridors is achieved, the value of conserving parcels in this row becomes superadditive. However, the value of preserving multiple corridors satisfies submodularity, thus although beneficial to conserve a corridor, as more are added, the marginal contribution of each new one to the conservation pattern diminishes. All other combinations of parcels have additive values. Landowners need to coordinate their bids to form valid corridors and compete with other valid corridors to be successful (Iftekhar & Latacz-Lohmann 2017*a*). Other conservation targets aim at reducing soil erosion or water pollution, where the auctioneer has preferences for specific patterns as well. We focus on the goal of wildlife corridors as a particularly illustrative and widespread conservation target that has received significant attention (Iftekhar & Tisdell

2014, Iftekhar & Latacz-Lohmann 2017*b*, Dijk et al. 2017). An interesting application is the auction program to create habitat corridors in Australia (Nguyen et al. 2022).

The choice of landscape corridors as a specific type of value model is due to its relevance to biodiversity and conservation auctions. The literature on biodiversity-based value models on a landscape scale is scarce and differs from application to application (e.g., protecting turtle doves is different from landscape level patterns that help bees). Knowledge about beneficial patterns is currently emerging.

A motivating argument for studying submodular functions is the fact that BFA achieves the highest efficiency for this family of functions, yielding a close approximation of the optimal solution. Bordewich & Semple (2011) model phylogenetic diversity, a measure of biodiversity of a species collection, under the assumption of submodularity. Similar to additive and superadditive, which can be intuitively perceived as natural candidate classes, submodularity is an interesting property to examine.

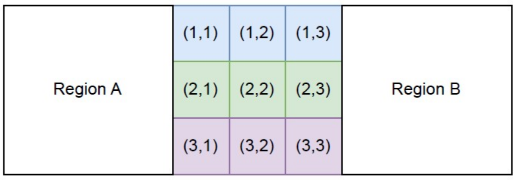

Figure 1. Corridors on the landscape grid allowing for wildlife movement between regions A and B highlighted in different colors.

Budgets can be binding or non-binding. The case of non-binding budgets is trivial: the OPT and VCG outcomes are equivalent, and all items can be purchased for a certain price. For the more interesting case of binding budgets, we draw values at random in the range from which costs are drawn, so that the auctioneer can only afford a strict subset of items. In particular, since costs are drawn i.i.d. from a uniform distribution in [0,50], we draw budget values i.i.d. from a uniform distribution in [50, 150]: in this way, the auctioneer has the power to purchase at least one single land item.

To the best of our knowledge, conservation auction data is not made publicly available, as principal values and land owner opportunity costs are documented by government bodies. We thus resort to synthetic data generation to support and confirm our claims.

6 Results

6.1 Relative Efficiency

We first compute the welfare of BFA relative to OPT (BFA/OPT) and then relative to VCG (BFA/VCG). The welfare of the latter does not consider the budget constraint and as such describes unconstrained (and higher) welfare. We report average values and standard deviations for 50 auction instances with randomly generated valuations. In Tables 1-6 we include the statistical measure of p-values. We define symbolic notations that represent the significance level of each reported result. [2]

[2] P-value ranges symbolically: $[10^{-60}, 10^{-40}]$:***, $[10^{-40}, 10^{-20}]$:**, $[10^{-20}, 10^{-01}]$:*

Result 1 *The average efficiency of BFA relative to OPT is 0.64 or 0.76 respectively (see table 1). With general superadditive valuations, the average efficiency (BFA/OPT) can be as low as 0.491 for high levels of superadditivity (see table 2).*

The worst-case theoretical approximation ratio for submodular valuations is 4.75, which translates to a welfare of 0.21% in the maximization problem. Taking OPT as a baseline, the group-strategyproof BFA achieves high levels of efficiency of 0.64 for additive and 0.77 for submodular valuations (see table 1).

Table 1. Additive and submodular valuations: mean and standard deviation of relative efficiency.

	additive	submodular
BFA/VCG $(mean)$	0.454**	0.757*
BFA/ VCG (std)	0.106	0.149
BFA/OPT $(mean)$	0.640*	0.763*
BFA/ OPT (std)	0.129	0.130

We report superadditive valuations in detail since average efficiency depends heavily on the level of superadditivity. Table 2 shows the average efficiency values for different multipliers. A multiplier of 2 (rightmost column) means that the value of the package is twice the value of the individual parcels, and indicates a high level of superadditivity.

Table 2. Superadditive valuations: mean and standard deviation of relative efficiency for varying value of superadditivity multiplier.

	1.1	1.2	1.3	1.4	1.5	1.7	2.0
BFA/VCG $(mean)$	0.459**	0.430**	0.423***	0.394***	0.408***	0.376***	0.332***
BFA/VCG (std)	0.123	0.140	0.136	0.130	0.118	0.136	0.132
BFA/OPT $(mean)$	0.638*	0.610*	0.599*	0.569*	0.599*	0.542*	0.491*
BFA/OPT (std)	0.131	0.141	0.146	0.156	0.126	0.162	0.164

Result 2 *The value model with superadditive corridors combines superadditive and submodular valuations. The level of superadditivity within corridors determines the overall auction efficiency. The values for BFA/OPT do not significantly differ from those for superadditive valuations, while for BFA/VCG are significantly lower (see table 3).*

The slightly lower efficiency values for BFA/VCG deserve some discussion. The reason behind this discrepancy lies in the nature of the BFA algorithm. Starting from the maximum single item value, the algorithm adds to a current set S_t the item with the largest marginal contribution. This item can be anywhere on the grid, as the greedy algorithm has no spatial understanding or knowledge that completing a corridor increases value, while all other packages are additive. If the auctioneer has acquired one parcel in every row, the algorithm simply selects the item with highest marginal contribution to the current set welfare, which might lie in another corridor. With binding budgets, no corridor might be established in cases where it would have been possible.

Table 3. Submodular corridors: mean and standard deviation of relative efficiency for varying values of superadditivity within a corridor.

	1.1	1.2	1.3	1.4	1.5	1.7	2.0
BFA/VCG $(mean)$	0.384***	0.342***	0.301***	0.275***	0.252***	0.225***	0.183***
BFA/VCG (std)	0.108	0.097	0.073	0.073	0.076	0.075	0.074
BFA/OPT $(mean)$	0.615*	0.63*	0.601*	0.588*	0.573*	0.545*	0.479*
BFA/OPT (std)	0.129	0.152	0.126	0.128	0.168	0.191	0.178

6.2 Payments

Next, we study how much the VCG mechanism exceeds the budget on average, and how much budget is left over with BFA.

Result 3 *The VCG payments for submodular valuations do not significantly exceed the budgets constraint. However, with superadditive valuations, the payments are even 12.9 times higher than the budget with a superadditivity multiplier of 2 (see tables 4 and 5). For submodular corridors, this ratio goes up to 68 for a multiplier of 2 (see table 6). In contrast, the payments with the BFA algorithm are always significantly below budget. For superadditive valuations, the payments are always more than 20% below budget.*

Table 4. Relative payments for additive and submodular valuations compared to given budget constraint.

	additive	submodular
VCG/ Budget $(mean)$	6.6***	1.7*
VCG / Budget (std)	2.0	0.6
BFA/ Budget $(mean)$	0.8*	0.8*
BFA/ Budget (std)	0.1	0.1

Table 5. Relative payments for superadditive valuations compared to given budget constraints, for varying value of superadditivity multiplier.

	1.1	1.2	1.3	1.4	1.5	1.7	2.0
VCG/ Budget $(mean)$	7.2***	7.9***	8.5***	9.2***	9.9***	11.1***	12.9***
VCG/ Budget (std)	2.2	2.4	2.8	3.0	3.3	3.5	3.9
BFA/ Budget $(mean)$	0.8*	0.8*	0.8*	0.8*	0.8*	0.8*	0.8*
BFA/Budget (std)	0.1	0.1	0.1	0.2	0.2	0.2	0.2

Budget violation in submodular corridors is significantly higher than in superadditive valuations (see table 6). This is due to the structure of the VCG payment rule and the value model specifics. Once a corridor is achieved, the marginal welfare contribution of a single item is very high, since this addition leads to a high value corridor in the welfare. As the multiplier value increases, so do the gains from preserving a corridor. This high marginal contribution of a single seller is reflected in the VCG payment. In the case of superadditive valuations, many packages with high superadditive valuations exist and the effect is much reduced. BFA payments are slightly higher but largely unaffected.

Table 6. Relative payments for submodular corridor dataset compared to given budget constraints, for varying value of superadditivity multiplier.

	1.1	1.2	1.3	1.4	1.5	1.7	2.0
VCG/ Budget $(mean)$	29.7***	36.0***	41.8***	44.3***	49.2***	55.9***	68.4***
VCG/ Budget (std)	10.0	12.7	13.3	14.3	15.1	19.5	23.0
BFA/ Budget $(mean)$	0.8*	0.8*	0.8*	0.8*	0.8*	0.8*	0.8*
BFA/Budget (std)	0.1	0.1	0.1	0.1	0.1	0.1	0.1

7 Conclusions

Conservation auctions have received significant attention worldwide. Recent insights from natural sciences show that agri-environmental services need to be coordinated on a landscape level to maximize biodiversity gains. Conservation of wildlife corridors in a landscape is a standard example with many applications. Policymakers aim at designing auctions that are incentive-compatible, respect governments' budget constraints, and account for complementarities of agri-environmental services. Recent results in algorithmic mechanism design propose auction formats that satisfy these goals at the expense of efficiency. The contribution by Balkanski et al. (2022) is remarkable as it is a transparent clock auction that is robust against collusion and runs in polynomial time.

Worst-case bounds for specific and stylized value functions might not provide sufficient guidance for policymakers when selecting a mechanism. We analyze standard value models motivated by the application domain and compute the average-case efficiency loss a regulator can expect. Our experiments confirm that, primarily for the case of submodular, but, to some degree, for additive functions, the BFA achieves high efficiency, verifying that, in practice, the approximation observed is much closer than the theoretical worst-case bound. We show that, with modest levels of superadditivity in conservation auctions, the welfare loss is still reasonable. At the same time, the auction guarantees budget feasibility, accommodates the value function of the government, and is strategyproof even with respect to coalitions of farmers. In contrast, payments by the government render the VCG mechanism unsuitable for most applications in this domain where the government has budget restrictions. Therefore, even though the BFA is designed as a submodular maximization algorithm, one could argue that, since it maintains important properties and stays within budget, it can be considered as a promising alternative for policymakers, when budget and truthfulness are strict constraints. In future research, laboratory experiments would be useful to understand the difference to non-truthful auction formats.

References

Allen, B., Hart, K., Radley, G., Tucker, G., Keenleyside, C., Oppermann, R., Underwood, E., Menadue, H., Poux, X., Beaufoy, G., Herzon, I., Povellato, A., Vanni, F., Pražan, J., Hudson, T. & Yellachich, N. (2014), *Biodiversity protection through results based remuneration of ecological achievement: Report Prepared for the European Commission, DG Environment, Contract No ENV.B.2/ETU/2013/0046, Institute for European Environmental Policy, London*, Institute for European Environmental Policy, United Kingdom.

Anton, J. J. & Yao, D. A. (1992), 'Coordination in split award auctions', *The Quarterly Journal of Economics* **107**(2), 681–707.

Armsworth, P. R., Acs, S., Dallimer, M., Gaston, K. J., Hanley, N. & Wilson, P. (2012), 'The cost of policy simplification in conservation incentive programs', *Ecology Letters* **15**(5), 406–414.

Balkanski, E., Garimidi, P., Gkatzelis, V., Schoepflin, D. & Tan, X. (2022), *Deterministic Budget-Feasible Clock Auctions*, pp. 2940–2963.

Bichler, M., Davenport, A., Hohner, G. & Kalagnanam, J. (2006), 'Industrial procurement auctions', *Combinatorial Auctions* pp. 593–612.

Bordewich, M. & Semple, C. (2011), 'Budgeted nature reserve selection with diversity feature loss and arbitrary split systems', *Journal of mathematical biology* **64**, 69–85.

Cramton, P., Hellerstein, D., Higgins, N., Iovanna, R., López-Vargas, K. & Wallander, S. (2021), 'Improving the cost-effectiveness of the Conservation Reserve Program: A laboratory study', *Journal of Environmental Economics and Management* **108**(C).

Dasgupta, S. & Spulber, D. F. (1989), 'Managing procurement auctions', *Information Economics and Policy* **4**(1), 5–29.

Dijk, J., Ansink, E. & van Soest, D. (2017), 'Buyouts and agglomeration bonuses in wildlife corridor auctions'.

Gholami, R., Watson, R. T., Hasan, H., Molla, A. & Bjorn-Andersen, N. (2016), 'Information systems solutions for environmental sustainability: How can we do more?', *Journal of the Association for Information Systems* **17**(8), 2.

Gravin, N., Jin, Y., Lu, P. & Zhang, C. (2020), 'Optimal budget-feasible mechanisms for additive valuations', *ACM Trans. Econ. Comput.* **8**(4).

Gómez-Limón, J. A., Gutiérrez-Martín, C. & Villanueva, A. J. (2019), 'Optimal design of agri-environmental schemes under asymmetric information for improving farmland biodiversity', *Journal of Agricultural Economics* **70**(1), 153–177.

Hanley, N., Banerjee, S. & Simpson, K. (2022), Auctions for biodiversity conservation: what works best? and would a biodiversity offset market work better?, Technical report, University of Glasgow.

Hellerstein, D. M. (2017), 'The us conservation reserve program: The evolution of an enrollment mechanism', *Land Use Policy* **63**, 601–610.

Iftekhar, M. S. & Latacz-Lohmann, U. (2017*a*), 'How well do conservation auctions perform in achieving landscape-level outcomes? a comparison of auction formats and bid selection criteria', *Australian Journal of Agricultural and Resource Economics* **61**(4), 557–575.

Iftekhar, M. S. & Latacz-Lohmann, U. (2017*b*), 'How well do conservation auctions perform in achieving landscape-level outcomes? a comparison of auction formats and bid selection criteria', *Australian Journal of Agricultural and Resource Economics* **61**(4), 557–575.

Iftekhar, M. & Tisdell, J. (2014), 'Wildlife corridor market design: An experimental analysis of the impact of project selection criteria and bidding flexibility', *Ecological Economics* **104**, 50–60.

Iftekhar, S., Hailu, A. & Lindner, R. (2009), *Auction designs for native biodiversity conservation on private lands*, Nova Science, pp. 1– 45.

Kokott, G.-M., Bichler, M. & Paulsen, P. (2019), 'The beauty of dutch: Ex-post split-award auctions in procurement markets with diseconomies of scale', *European Journal of Operational Research* **278**(1), 202–210.

Lamb, A., Balmford, A., Green, R. E. & Phalan, B. (2016), 'To what extent could edge effects and habitat fragmentation diminish the potential benefits of land sparing?', *Biological Conservation* **195**, 264–271.

Latacz-Lohmann, U. & Breustedt, G. (2019), 'Using choice experiments to improve the design of agri-environmental schemes', *European Review of Agricultural Economics* **46**(3), 495–528.

Lundberg, L., Persson, U. M., Alpizar, F. & Lindgren, K. (2018), 'Context matters: exploring the cost-effectiveness of fixed payments and procurement auctions for pes', *Ecological Economics* **146**, 347–358.

Nemes, V., Plott, C. & Stoneham, G. (2008), 'Electronic bushbroker exchange: Designing a combinatorial double auction for native vegetation offsets', *SSRN Archive* .

Nguyen, C., Latacz-Lohmann, U., Hanley, N., Schilizzi, S. & Iftekhar, S. (2022), 'Spatial coordination incentives for landscape-scale environmental management: A systematic review', *Land Use Policy* **114**, 105936.

Palm-Forster, L. H., Swinton, S. M., Lupi, F. & Shupp, R. S. (2015), Too Burdensome to Bid: Transaction Costs and Pay-for-Performance Conservation, 2016 Allied Social Sciences Association (ASSA) Annual Meeting, January 3-5, 2016, San Francisco, California 212816, Agricultural and Applied Economics Association.

Polasky, S., Lewis, D. J., Plantinga, A. J. & Nelson, E. (2014), 'Implementing the optimal provision of ecosystem services', *Proceedings of the National Academy of Sciences* **111**(17), 6248–6253.

Rolfe, J., Schilizzi, S., Boxall, P., Latacz-Lohmann, U., Iftekhar, S., Star, M. & O'Connor, P. (2018), 'Identifying the causes of low participation rates in conservation tenders', *International Review of Environmental and Resource Economics* **12**(1), 1–45.

Roughgarden, T. (2019), 'Beyond worst-case analysis', *Communications of the ACM* **62**(3), 88–96.

Rousseau, S. & Moons, E. (2007), 'The potential of auctioning contracts for conservation policy', *European Journal of Forest Research* **127**, 183–194.

Saïd, S. & Thoyer, S. (2007), Agri-environmental auctions with synergies, Working papers, LAMETA, Universtiy of Montpellier.

Schilizzi, S. G. (2017), 'An overview of laboratory research on conservation auctions', *Land Use Policy* **63**, 572–583.

Singer, Y. (2010), Budget feasible mechanisms, *in* '2010 IEEE 51st Annual Symposium on Foundations of Computer Science', pp. 765–774.

Westerink, J., Jongeneel, R., Polman, N., Prager, K., Franks, J., Dupraz, P. & Mettepenningen, E. (2017), 'Collaborative governance arrangements to deliver spatially coordinated agri-environmental management', *Land Use Policy* **69**, 176–192.

Windle, J. & Rolfe, J. (2008), 'Exploring the efficiencies of using competitive tenders over fixed price grants to protect biodiversity in australian rangelands', *Land Use Policy* **25**(3), 388–398.

Wätzold, F., Drechsler, M., Johst, K., Mewes, M. & Sturm, A. (2016), 'A novel, spatiotemporally explicit ecological-economic modeling procedure for the design of cost-effective agri-environment schemes to conserve biodiversity', *American Journal of Agricultural Economics* **98**(2), 489–512.

A Appendix

A.1 BFA Algorithm Details

We provide the detailed pseudo-code of the two-stage BFA algorithm, adopting the notation of Balkanski et al. (2022). Starting from Mechanism 1, the set of active sellers A initially contains the entire set N, and whoever rejects the offered price is removed in each iteration, leading to a gradually shrinking active set. The OPT estimate is initiated with the value maximizing item to the auctioneer and is refined as the auction proceeds. Set S_t of currently considered sellers in phase t initially contains the index of the value maximizing seller. Given S_t, we update an index i that contains the seller whose value, when aggregated with the set S_t, offers the maximum marginal contribution. To this seller i, a price offering equal to the minimum between the previous price and marginal contribution scaled by a factor $\frac{B}{\tilde{OPT}}$ that guarantees budget feasibility is made. If this offer p_i is accepted by the seller, set S_t is updated to include i, otherwise i is removed from A and permanently exits the auction. While the value of set S_t is smaller than the current estimate $\tilde{OPT}$ and the set of remaining sellers that can be considered is not empty, the algorithm proceeds in considering the next best addition to S_t, until the condition on the value of S_t is met. If the value of S_t exceeds this estimate, a new phase begins: $\tilde{OPT}$ is doubled, and S_t is initialized as an empty set for the new phase. When the set of remaining sellers, that are active but do not belong to either S_t or S_{t-1} - the two sets of accepted sellers of the last two phases, is empty, the auction terminates. We examine sets S_{t-1} and S_t, now denoted as W_1 and $\bar{W}_2$. Starting from W_1, if the sum of prices of all sellers belonging to W_1 exceeds the given budget B, then the last added seller is removed from W_1, and is instead offered a price equal to the marginal contribution to the set S_t. In this way, the algorithm aims at moving the seller from one set, where the budget is exceeded for the current price, to another set, with the hope of achieving a budget-feasible payment scheme. If the updated price is accepted by the seller, the set S_t, denoted as $\bar{W}_2$ is updated to include this seller. If the sum of prices of sellers in W_1 is within budget, the swap of the last added seller does not take place and the algorithm proceeds to the final stage: invoking a routine for value maximization subject to a knapsack constraint (the budget).

Algorithm 2 examines the two sets, W_1 and $\bar{W}_2$. As a first step, the longest budget-feasible sequence of agents is picked from $\bar{W}_2$, since the sets are built in a manner that supports retrieval of the order of insertion, and is denoted as W_2. The two sets W_1 and W_2 are now both within budget: a final set W_3 is defined as the union between W_2 and the longest sequence of sellers from W_1 such that budget-feasibility of the union is maintained. The algorithm chooses the set among W_1 and W_3 that yields the highest value, which corresponds to the final set of sellers that are chosen in the allocation.

MECHANISM 1: ITERATIVE-PRUNING, a deterministic budget-feasible clock auction for monotone submodular valuation functions

Input : Budget B, valuation function $v : 2^N \to \mathbb{R}$

initialize $A \leftarrow \mathcal{N}$, $S_0 \leftarrow \emptyset, S_1 \leftarrow \left\{\text{argmax}_{i \in \mathcal{N}}\, v(\{i\})\right\}$, $\tilde{\texttt{OPT}} \leftarrow v(S_1)$, $t \leftarrow 1$, $p_i \leftarrow B$ for all $i \in \mathcal{N}$

while $A \setminus (S_{t-1} \cup S_t) \neq \emptyset$ **do**

    update $t \leftarrow t+1$, $\tilde{\texttt{OPT}} \leftarrow 2\tilde{\texttt{OPT}}$ and initialize $S_t \leftarrow \emptyset$; `// start a new phase`

    while $v(S_t) < \tilde{OPT}$ *and* $A \setminus (S_{t-1} \cup S_t) \neq \emptyset$ **do**

        let $i \leftarrow \text{argmax}_{i \in A \setminus (S_{t-1} \cup S_t)} v(\{i\} \mid S_t)$;

        update $p_i \leftarrow \min\left\{p_i,\ v(\{i\} \mid S_t) \cdot \frac{B}{\tilde{\texttt{OPT}}}\right\}$

        if *seller i accepts price p_i* **then**

            update $S_t \leftarrow S_t \cup \{i\}$; `// add seller` i `to current solution`

        else

            update $A \leftarrow A \setminus \{i\}$; `// permanently eliminate seller` i

Let $W_1 \leftarrow S_{t-1}$ and $\overline{W}_2 \leftarrow S_t$

if $\sum_{i \in W_1} p_i > B$ **then** `// enforce budget feasibility of` W_1

    let $j^* \leftarrow$ the last seller added to S_{t-1}

    update $p_{j^*} \leftarrow \min\{p_{j^*},\ v(\{j^*\} \mid S_t) \cdot \frac{B}{\tilde{\texttt{OPT}}_t}\}$

    update $W_1 \leftarrow W_1 \setminus \{j^*\}$

    if *seller j^* accepts price p_{j^*}* **then**

        update $\overline{W}_2 \leftarrow \overline{W}_2 \cup \{j^*\}$; `// move the last seller` j^* `to` $\overline{W}_2$

return MAXIMIZE-VALUE($W_1, \overline{W}_2, p$)

MECHANISM 2: MAXIMIZE-VALUE, an algorithm for maximizing value subject to knapsack constraint

Input : W_1, $\overline{W}_2$ and the prices p_i for all $i \in W_1 \cup \overline{W}_2$

let $W_2 \leftarrow$ the longest budget-feasible prefix of $\overline{W}_2$

let $W_3 \leftarrow W_2 \cup T$, where T is the longest prefix of W_1 such that $W_2 \cup T$ is budget-feasible

Let $W \in \{W_1, W_3\}$ be the set with the largest value $v(W)$

return *W and the corresponding prices*

5 Conclusion

The study of budget constraints in auctions is an important and practically motivated problem, focusing on the design of market mechanisms that achieve social objectives while respecting financial limits. Budgets serve as a foundational tool in managing resources, controlling costs, and mitigating risks while providing a financial framework for decision-making. Domains such as display advertising and spectrum auctions are common examples where such constraints play a prominent role. The existence of budget constraints poses a challenge in the design of auction mechanisms: crucial properties such as incentive compatibility, core stability, and welfare maximization clash for general valuations. Our work addresses the problems posed by such constraints and concentrates on the design and analysis of algorithmic schemes that overcome these limitations.

Competitive equilibria constitute an important tool in market analysis due to their property of matching the supply and demand of participants. The existence of competitive equilibria serves as a guarantee of stability in a market, as no participant has an incentive to deviate from the outcome while also being welfare-maximizing. Another important solution concept describing stability is the core, which consists of a set of outcomes from which no participant has an incentive to deviate and collude with others in the market to achieve favorable alternative outcomes. In the presence of budget constraints, a core outcome is not guaranteed to maximize welfare.

The fulfillment of conditions necessary to ensure the existence of a competitive equilibrium is frequently unmet in real-world markets. Therefore, it becomes imperative to devise more generalized notions of stability to address such market scenarios. This challenge is linked to the development of novel iterative auctions designed to compute outcomes in these situations.

The quasi-linear utility model stands as a robust standard assumption in competitive equilibrium theory but is often not realized in practical scenarios. Under quasi-linear utilities, bidders are assumed to face no financial constraints; irrespective of the price, they can always afford their preferred bundle. However, deviating from this model and considering scenarios where bidders encounter strict upper limits on spending can lead to instances where a competitive equilibrium does not exist. In this context, utility is no longer transferable in the form of payments: an agent can only transfer up to the spending limit and fails to express preferences from that point on.

Recent findings indicate that computing outcomes satisfying a subset of competitive equilibria properties, specifically core stability and welfare maximization, is computationally intractable. In the more general setting of combinatorial exchanges, Bichler and Waldherr (2022) demonstrate that, when bidders are allowed arbitrary valuation functions, the problem becomes Σ_2^p-hard.

In the first publication included in this dissertation (Batziou et al., 2022a), we address this computational challenge. Restricting preferences to the simplest market model, the assignment market with unit-demand and budget-constrained bidders, we prove that finding core-stable welfare-maximizing outcomes is hard for the complexity class NP, even with full access to participants' private information.

In an assignment market, each bidder aims to acquire, at most, a single item. Employing only demand queries, we introduce an iterative auction algorithm, an extension of the DGS auction (Demange et al., 1986), that establishes core-stable outcomes without knowledge of bidders' true valuations, accounting for budget limits. Our auction process terminates in a core allocation if bidders truthfully reveal their demand set in each round. This mechanism, when subject to an additional condition on unit-demand valuations, proves to be incentive-compatible and welfare-maximizing, though the ex-ante determination of this condition remains uncertain. In an additional interesting finding of this work, we demonstrate that there is no incentive-compatible mechanism selecting a core outcome for bidders with unit-demand valuations.

While Vickrey's groundbreaking study (Vickrey, 1961) illustrated that markets could be structured to induce truthful preferences as a dominant strategy, the widespread prevalence of budget constraints poses a challenge to these assumptions, undermining the conditions of the VCG mechanism. We demonstrate that, even in the most basic markets with unit-demand valuations, achieving both incentive compatibility and core stability becomes a conflicting objective.

We conclude that assertions regarding the efficiency of polynomial-time market designs should be approached cautiously in markets where the financial constraints of bidders come into play. An interesting research direction arising from this work revolves around analyzing the computational complexity of the multi-unit-demand setting. In this context, bidders remain constrained by a budget but seek to acquire multiple items. Lying in the realm between unit-demand and combinatorial exchanges, the problem of determining welfare-maximizing core outcomes in such markets might belong to a complexity class in between NP and Σ_2^p. While one may show that the problem is intuitively harder than NP, the exact class inclusion remains an open question.

While the aforementioned work involves scenarios where buyers are subject to financial limitations, the second contribution included in this dissertation (Batziou and Bichler, 2023) explores settings where the constraint is imposed on the auctioneer.

Before the influential paper of Singer (2010), research on budget-feasible mechanisms was limited and mainly operated under the assumption of access to infinite resources for incentive-compatibility. In mechanism design, incentive compatibility is a crucial property, ensuring agents do not misreport private values to the auctioneer for higher utility.

Budget-feasible procurement, common in government auctions, aims to maximize social value within budget constraints, resembling the knapsack problem. Singer's work sparked research on budget-feasible auctions, with the assumption of submodularity and monotonicity yielding improved approximations.

The majority of work on budget-feasible mechanism design involves sealed-bid auctions, where participants often misreport in practice, as shown in (Kagel et al., 1987). Milgrom and Segal (2020) introduced clock auctions, progressing through rounds with decreasing prices until the budget limit, offering benefits over sealed-bid auctions. Unlike previous work, Balkanski et al. (2022) proposed a clock auction respecting budgets and achieving a constant factor of 4.75 approximation for monotone submodular valuations. This algorithm, based on a backward greedy method, ensures properties like resistance to collusion, incentive compatibility, and transparency with deterministic polynomial runtime.

In the second work of this dissertation (Batziou and Bichler, 2023), we approach biodiversity conservation from a mechanism design standpoint, addressing an urgent environmental concern. With a notable decline in wildlife populations, motivating private landowners to engage in conservation becomes imperative. Appropriate incentives, typically in the form of monetary payments, should devised to dissuade agricultural use. The governing bodies overseeing such markets encounter financial constraints, reserving only a fraction of funds for environmental purposes, which are then utilized for payments.

In this study, we empirically assess the clock auction algorithm introduced by Balkanski et al. (2022). Our experimental findings reveal that this approximation algorithm achieves substantial welfare compared to the optimal solution and is incentive-compatible and budget-respecting, thus making it a promising alternative for policymakers. Notably, our results caution against the use of the widely employed VCG auction algorithm (Vickrey, 1961; Clarke, 1971; Groves, 1973), considered the state-of-the-art method in conservation auctions (Polasky et al., 2014). This traditional approach may lead to payments that arbitrarily exceed the budget, rendering it undesirable in practical conservation scenarios.

Recent discoveries in natural sciences underscore the imperative need for coordinated agri-environmental services at a landscape level to optimize biodiversity gains. Exploring algorithmic frameworks derived from mechanism design and auction literature represents a main avenue for future research, enhancing current methods to advance environmental conservation efforts. Another objective involves developing appropriate scoring rules aligned with practical considerations and developed in collaboration with environmental scientists to accurately assess the overall value of auction outcomes. This approach aims to amplify biodiversity gains and bolster conservation impact.

An extension of the research presented in this dissertation seeks to narrow the divide between experimental findings and practical implementation. This involves formulating realistic scoring rules to evaluate proposed auction designs. The collaboration between environmental scientists and economists focusing in market design emerges as a critical factor, holding the potential for research that significantly influences ecosystem preservation.

6 Licenses and Copyright Information

6.1 Part I: Core-Stability in Assignment Markets with Financially Constrained Buyers

Title: Core-Stability in Assignment Markets with Financially Constrained Buyers.

Authors: Eleni Batziou, Martin Bichler, Maximilian Fichtl.

In: Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22).

Relevant copyright information is provided as follows:

- The article was accepted for presentation at the 23rd ACM Conference on Economics and Computation (EC '22).
- In the conference proceedings, only a one-page abstract was published: https://doi.org/10.1145/3490486.3538262. The copyrights for this one-page abstract remain with ACM, see the copyright note here: https://dl.acm.org/action/showFmPdf?doi=10.1145%2F3490486.
- For the original full-version of the article, all copyright remains with the authors. This is explicitly granted, see here: https://ec22.sigecom.org/call-for-contributions-acm/papers/index.html:
To accommodate the publishing traditions of different fields, authors of accepted papers can ask that only a one-page abstract of the paper appear in the proceedings, along with a URL pointing to the full paper. Authors should guarantee the link to be reliable for at least two years. This option is available to accommodate subsequent publication in journals that would not consider results that have been published in preliminary form in conference proceedings.
- The full version of the paper – as included in this dissertation – is published on arxiv. The published one-page abstract also points to the arxiv article. The arxiv article can be accessed here: https://arxiv.org/abs/2205.06132.
- The arxiv article is published under the CC-BY license. This license allows for re-publishing in the dissertation:
Core-Stability in Assignment Markets with Financially Constrained Buyers © 2022 by Eleni Batziou, Martin Bichler, Maximilian Fichtl is licensed under CC BY 4.0.
- The ACM copyright note (https://authors.acm.org/author-resources/author-rights) and

ACM Author Rights

ACM exists to support the needs of the computing community. For over sixty years ACM has developed publications and publication policies to maximize the visibility, impact, and reach of the research it publishes to a global community of researchers, educators, students, and practitioners. ACM has achieved its high impact, high quality, widely-read portfolio of publications with:

- Affordably priced publications
- Liberal Author rights policies
- Wide-spread, perpetual access to ACM publications via a leading-edge technology platform
- Sustainability of the good work of ACM that benefits the profession

Choose

ACM gives authors the opportunity to choose between two levels of rights management for their work. Note that both options obligate ACM to defend the work against improper use by third parties:

- **Exclusive Licensing Agreement:** Authors choosing this option will retain copyright of their work while providing ACM with exclusive publishing rights.
- **Non-exclusive Permission Release:** Authors who wish to retain all rights to their work must choose ACM's author-pays option, which allows for perpetual open access to their work through ACM's digital library. Choosing this option enables authors to display a Creative Commons License on their works.

Post

Otherwise known as "Self-Archiving" or "Posting Rights", all ACM published authors of magazine articles, journal articles, and conference papers retain the right to post the pre-submitted (also known as "pre-prints"), submitted, accepted, and peer-reviewed versions of their work in any and all of the following sites:

CC BY 4.0

ATTRIBUTION 4.0 INTERNATIONAL

Deed

Canonical URL : https://creativecommons.org/licenses/by/4.0/

See the legal code

You are free to:

Share — copy and redistribute the material in any medium or format for any purpose, even commercially.

Adapt — remix, transform, and build upon the material for any purpose, even commercially.

The licensor cannot revoke these freedoms as long as you follow the license terms.

Under the following terms:

Attribution — You must give appropriate credit , provide a link to the license, and indicate if changes were made . You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

No additional restrictions — You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Notices:

You do not have to comply with the license for elements of the material in the public domain or where your use is permitted by an applicable exception or limitation .

No warranties are given. The license may not give you all of the permissions necessary for your intended use. For example, other rights such as publicity, privacy, or moral rights may limit how you use the material.

Notice

This deed highlights only some of the key features and terms of the actual license. It is not a license and has no legal value. You should carefully review all of the terms and conditions of the actual license before using the licensed material.

Creative Commons is not a law firm and does not provide legal services. Distributing, displaying, or linking to this deed or the license that it summarizes does not create a lawyer-client or any other relationship.

Creative Commons is the nonprofit behind the open licenses and other legal tools that allow creators to share their work. Our legal tools are free to use.

6.2 Part II: Budget-Feasible Market Design for Biodiversity Conservation: Considering Incentives and Spatial Coordination

Title: Budget-Feasible Market Design for Biodiversity Conservation: Considering Incentives and Spatial Coordination.

Authors: Eleni Batziou, Martin Bichler.

In: 18th International Conference on Wirtschaftsinformatik (WI '23).

Relevant copyright information is provided as follows:

- The article was accepted for presentation at the 18th International Conference on Wirtschaftsinformatik (WI '23).
- The conference proceedings are published in the AIS eLibrary and the article is made available under the following link: https://aisel.aisnet.org/wi2023/44/.
- The copyrights for articles published in the AIS eLibrary remain solely with the authors, see here: https://aisnet.org/page/PoliciesGuidelines: *As a result, the AIS Policy on Conference and Journal Publication encourages conference papers to be published in journals. This policy applies to AIS conferences (so there are no copyright issues to prevent your AIS conference paper from being published in any journal (AIS or not)), and to AIS journals (so there are no restrictions preventing you from submitting your paper from any conference (AIS or not) to an AIS journal).*
- In addition, the AIS eLibrary operates as part of the Digital Commons platform, see bottom of page here: https://aisel.aisnet.org/
- As explicitly mentioned in the website https://network.bepress.com/about/: *The Digital Commons Network provides free access to full-text scholarly articles and other research from hundreds of universities and colleges worldwide.*
- All copyrights remain with the authors allowing for republishing in the dissertation.

Policies and Guidelines

PUBLISHING CONFERENCE PAPERS IN AIS JOURNALS

AIS believes that it is the best interests of the field for research to be first presented at a conference and then published in a more refined format in a journal. The conference review process and the presentation of the study to colleagues helps improve the research, so the research has the opportunity to be much more fully developed for review by journals. As a result, the AIS Policy on Conference and Journal Publication encourages conference papers to be published in journals. This policy applies to AIS conferences (so there are no copyright issues to prevent your AIS conference paper from being published in any journal (AIS or not)), and to AIS journals (so there are no restrictions preventing you from submitting your paper from any conference (AIS or not) to an AIS journal). Many non-AIS journals have policies similar to the AIS policy, so you can submit your AIS conference paper to them; this includes all the AIS Senior Scholars' Basket Journals (European Journal of Information Systems, Information Systems Journal, Information Systems Research, Journal of the AIS, Journal of Information Technology, Journal of MIS, Journal of Strategic Information Systems, and MIS Quarterly), as well as Decision Support Systems, Information and Management, Information and People, MIS Quarterly Executive, Systemes d'Information et Management, Business & Information Systems Engineering, International Journal of Knowledge Management, The Data Base for Advances in Information Systems, The Australasian Journal of Information Systems, The Scandinavian Journal of Information Systems, Journal of Information Technology Management, Journal of Information Systems Education, International Journal of Public Administration in the Digital Age, International Journal of Information Systems for Crisis Response and Management, and Information Systems Frontiers. (If your journal has a similar policy and you would like to include it here, please contact Alan Dennis at ardennis@indiana.edu (mailto:ardennis@indiana.edu)).

GREEN OPEN ACCESS POLICY

AIS has a green Open Access Policy. Authors may immediately self-archive their articles without an embargo period when publishing in these AIS journals:

- *Journal of the Association for Information Systems (JAIS) (https://aisel.aisnet.org/jais/)*
- *Communications of the Association for Information Systems (CAIS) (https://aisel.aisnet.org/cais)*
- *Pacific Asia Journal of the Association for Information Systems (PAJAIS) (https://aisel.aisnet.org/pajais)*
- *Transactions on Human-Computer Interaction (THCI) (https://aisel.aisnet.org/thci)*
- *Transactions on Replication Research (TRR) (https://aisel.aisnet.org/trr)*
- *MIS Quarterly Executive (MISQE) (https://aisel.aisnet.org/misqe/)*

AIS's self-archiving option pertains to the version of the article that received the official acceptance letter from the editor after successfully passing peer review (i.e., post-print or Author's Accepted Manuscript) or the final, official published version after the publisher has copy-edited, formatted and distributed it (i.e., version of record). Authors may self-archive the version of their research article that they choose on their website, or on a publicly-accessible institutional or subject-based repository

AIS RESEARCH CODE OF CONDUCT

AIS members are expected to adhere to the AIS Research Code of Conduct (/AdmBullCResearchCond), which has been developed by the Research Conduct Committee (RCC), whose members are listed below, and endorsed by AIS Council.

The RCC has developed guidelines for journal editors and conference program chairs (https://cdn.ymaws.com/sites/ais.site-ym.com/resource/resmgr/Docs/AIS_RCC_Guidelines_for_Edit.pdf) and guidelines for victims of plagiarism (https://aisnet.org/?GuidelinesforVictim). These guidelines are designed to help individuals deal with situations where inappropriate scholarly conduct may have occurred.

All AIS journal editors and conference program chairs should review the Guidelines for Editors for advice on how to deal with suspected or potential scholar misconduct issues, including plagiarism.

If you think that someone may have plagiarized your work, please look at the Guidelines for Victims. These guidelines have been written with victims in mind and are informed by previous plagiarism cases. They provide actionable advice.

Finally, to understand how cases of scholarly misconduct are handled by AIS, please refer to AIS Policy, 5.4.3 Scholarly Misconduct (/resource/resmgr/Docs/Excerpt_policy_5.4.3_-_Schol.pdf). This policy explains how cases of scholarly misconduct are reported and handled, the ways in which the AIS investigates suspected cases of academic misconduct, the role of the RCC in this process, and the array of penalties that exist for scholarly misconduct.

We recommend that doctoral students and all new AIS members review all four of these documents.

Reports related to potential violations of the AIS Research Code of Conduct may be submitted confidentially and anonymously through the third-party NAVEX EthicsPoint (https://secure.ethicspoint.com/domain/media/en/gui/74273/index.html) online reporting platform. Please visit conduct.aisnet.org (http://conduct.aisnet.org/) for additional information about the complaint resolution process and to file a report. If you have questions, comments or suggestions, please contact the chair of the RCC (mailto:rcc@aisnet.org). or the president of AIS (mailto:president@aisnet.org).

The Research Conduct Committee:

- **Maung Sein, University of Agder (Committee Chair)**
- Guillermo Rodríguez-Abitia, William & Mary (VP of Meetings and Conferences)
- Chee-Wee Tan, The Hong Kong Polytechnic University (VP of Publications)

Quick Links:

- AIS Research Code of Conduct (/AdmBullCResearchCond)
- Guidelines for Journal Editors and Conference Program Chairs (https://cdn.ymaws.com/sites/ais.site-ym.com/resource/resmgr/Docs/AIS_RCC_Guidelines_for_Edit.pdf)
- Guidelines for Victims of Plagiarism (https://aisnet.org/?GuidelinesforVictim)
- AIS Policy, 5.4.3 Scholarly Misconduct (/resource/resmgr/Docs/Excerpt_policy_5.4.3_-_Schol.pdf)
- Did I Plagiarize? Check out this infographic! (http://thevisualcommunicationguy.com/wp-content/uploads/2014/09/Infographic_Did-I-Plagiarize1.jpg)
- (http://thevisualcommunicationguy.com/wp-content/uploads/2014/09/Infographic_Did-I-Plagiarize1.jpg)AIS Policy on Conference and Journal Publication (/resource/resmgr/Docs/Policy_on_Conference_and_Jou.pdf)

Bibliography

G. Amanatidis, P. Kleer, and G. Schäfer. Budget-feasible mechanism design for non-monotone submodular objectives: Offline and online. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 901–919, 2019.

N. Anari, G. Goel, and A. Nikzad. Mechanism design for crowdsourcing: An optimal 1-1/e competitive budget-feasible mechanism for large markets. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 266–275. IEEE, 2014.

K. J. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.

L. M. Ausubel. An efficient dynamic auction for heterogeneous commodities. *American Economic Review*, 96(3):602–629, 2006.

L. M. Ausubel and P. Milgrom. Ascending auctions with package bidding. *The B.E. Journal of Theoretical Economics*, 1(1):1–44, 2002.

L. M. Ausubel, P. Crampton, and P. Milgrom. The clock-proxy auction: A practical combinatorial auction design. Discussion Papers 03-034, Stanford Institute for Economic Policy Research, 2004.

A. Badanidiyuru, S. Dobzinski, and S. Oren. Optimization with demand oracles. In *Proceedings of the 13th ACM conference on electronic commerce*, pages 110–127, 2012.

E. Baldwin and P. Klemperer. Implementing walrasian equilibrium - the language of product-mix auctions. Technical report, Oxford University, 2019a.

E. Baldwin and P. Klemperer. Understanding preferences: demand types, and the existence of equilibrium with indivisibilities. *Econometrica*, 87(3):867–932, 2019b.

E. Balkanski and J. D. Hartline. Bayesian budget feasibility with posted pricing. In *Proceedings of the 25th International Conference on World Wide Web*, pages 189–203, 2016.

E. Balkanski, P. Garimidi, V. Gkatzelis, D. Schoepflin, and X. Tan. *Deterministic Budget-Feasible Clock Auctions*, pages 2940–2963. 2022. doi: 10.1137/1.9781611977073.114.

E. Batziou and M. Bichler. Budget-feasible market design for biodiversity conservation: Considering incentives and spatial coordination. In *Wirtschaftsinformatik (WI) 2023 Proceedings*, 2023.

E. Batziou, K. A. Hansen, and K. Høgh. Strong approximate consensus halving and the borsuk-ulam theorem. In *International Colloquium on Automata, Languages and Programming*, 2021.

E. Batziou, M. Bichler, and M. Fichtl. Core-stability in assignment markets with financially constrained buyers. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, EC '22, page 473–474, New York, NY, USA, 2022a. Association for Computing Machinery.

E. Batziou, G. Bouritsas, and A. Tsigonias-Dimitriadis. Iterative combinatorial auctions: An improved approach using reinforcement learning. In *Working Paper*, 2022b.

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.

I. Bello, H. Pham, Q. V. Le, M. Norouzi, and S. Bengio. Neural combinatorial optimization with reinforcement learning. *CoRR*, abs/1611.09940, 2016. URL http://arxiv.org/abs/1611.09940.

J.-P. Benoît and V. Krishna. Multiple-Object Auctions with Budget Constrained Bidders. *The Review of Economic Studies*, 68(1):155–179, 01 2001. doi: 10.1111/1467-937X.00164.

M. Beyeler, G. Brero, B. Lubin, and S. Seuken. iMLCA: Machine learning-powered iterative combinatorial auctions with interval bidding. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. ACM, 2021. doi: 10.1145/3465456.3467535.

M. Bichler and J. K. Goeree. *Handbook of spectrum auction design*. Cambridge University Press, 2017.

M. Bichler and S. Waldherr. Core-pricing in combinatorial exchanges with financially constrained buyers: Computational hardness and algorithmic solutions. *Operations Research*, 70(1):241 – 264, 2022.

M. Bichler, P. Shabalin, and J. Wolf. Do core-selecting combinatorial clock auctions always lead to high efficiency? an experimental analysis of spectrum auction designs. *Experimental Economics*, 16(4):511–545, 2013.

S. Bikhchandani and J. W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74(2):385–413, 1997.

S. Bikhchandani and J. M. Ostroy. The package assignment model. *Journal of Economic Theory*, 107(2):377–406, 2002.

L. Blumrosen and N. Nisan. *Combinatorial Auctions*, page 267–300. Cambridge University Press, 2007. doi: 10.1017/CBO9780511800481.013.

C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget-constrained bidders. EC '05, page 44–51, New York, NY, USA, 2005. Association for Computing Machinery. doi: 10.1145/1064009.1064014.

C. Borgs, J. Chayes, N. Immorlica, K. Jain, O. Etesami, and M. Mahdian. Dynamics of bid optimization in online advertisement auctions. WWW '07, page 531–540, New York, NY, USA, 2007. Association for Computing Machinery. doi: 10.1145/1242572.1242644.

G. Brero, B. Lubin, and S. Seuken. Combinatorial auctions via machine learning-based preference elicitation. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 128–136, 7 2018.

G. Brero, B. Lubin, and S. Seuken. Machine learning-powered iterative combinatorial auctions. *CoRR*, abs/1911.08042, 2019. URL http://arxiv.org/abs/1911.08042.

Y.-K. Che and I. Gale. The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic Theory*, 92(2):198–233, 2000.

N. Chen, N. Gravin, and P. Lu. On the approximability of budget feasible mechanisms. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 685–699. SIAM, 2011.

E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971. URL http://www.jstor.org/stable/30022651.

R. J. Cole, N. R. Devanur, V. Gkatzelis, K. K. Jain, T. Mai, V. V. Vazirani, and S. Yazdanbod. Convex program duality, fisher markets, and nash social welfare. *Proceedings of the 2017 ACM Conference on Economics and Computation*, 2016.

R. Colini-Baldeschi, M. Henzinger, S. Leonardi, and M. Starnberger. On multiple round sponsored search auctions with budgets. *ArXiv*, abs/1112.6361, 2011. URL https://api.semanticscholar.org/CorpusID:106172.

V. Conitzer, C. Kroer, D. Panigrahi, O. Schrijvers, E. Sodomka, N. Stier-Moses, and C. Wilkens. Pacing equilibrium in first-price auction markets. pages 587–587, 06 2019. doi: 10.1145/3328526.3329600.

V. Conitzer, C. Kroer, E. Sodomka, and N. E. Stier-Moses. Multiplicative pacing equilibria in auction markets. *Oper. Res.*, 70(2):963–989, mar 2022. doi: 10.1287/opre.2021.2167.

S. A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing*, STOC '71, page 151–158, New York, NY, USA, 1971. Association for Computing Machinery. doi: 10.1145/800157.805047.

R. Day and P. Milgrom. Core-selecting package auctions. *international Journal of game Theory*, 36(3-4):393–407, 2008.

R. Day and S. Raghavan. Fair payments for efficient allocations in public sector combinatorial auctions. *Management Science*, 53:1389–1406, 2007.

G. Demange, D. Gale, and M. Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–872, 1986.

M. Deudon, P. Cournut, A. Lacoste, Y. Adulyasak, and L.-M. Rousseau. Learning heuristics for the TSP by policy gradient. In *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pages 170–181. Springer International Publishing, 2018. doi: 10.1007/978-3-319-93031-2_12.

S. Dobzinski, R. Lavi, and N. Nisan. Multi-unit auctions with budget limits. In *Foundations of Computer Science (FOCS'08)*, pages 260–269. IEEE, 2008.

S. Dobzinski, C. H. Papadimitriou, and Y. Singer. Mechanisms for complement-free procurement. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 273–282, 2011.

S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. *Journal of Computer and System Sciences*, 78(1):15–25, 2012. doi: https://doi.org/10.1016/j.jcss.2011.02.010.

P. Dütting, M. Henzinger, and M. Starnberger. Auctions for heterogeneous items and budget limits. *ACM Transactions on Economics and Computation (TEAC)*, 4(1), 2016.

P. Dütting, V. Gkatzelis, and T. Roughgarden. The Performance of Deferred-Acceptance Auctions. *Mathematics of Operations Research*, 42(4):897–914, November 2017. doi: 10.1287/moor.2016.0835.

S. Fujishige and Z. Yang. A note on kelso and crawford's gross substitutes condition. *Mathematics of Operations Research*, 28(3):463–469, 2003.

D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973. URL http://www.jstor.org/stable/1914083.

G. Goel, A. Nikzad, and A. Singla. Mechanism design for crowdsourcing markets with heterogeneous tasks. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 2, 2014.

N. Gravin, Y. Jin, P. Lu, and C. Zhang. Optimal budget-feasible mechanisms for additive valuations. *ACM Trans. Economics and Comput.*, 8(4):21:1–21:15, 2020. doi: 10.1145/3417746.

J. Green and J.-J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45(2):427–438, 1977. URL http://www.jstor.org/stable/1911219.

T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973. URL http://www.jstor.org/stable/1914085.

F. Gul, E. Stacchetti, et al. The english auction with differentiated commodities. *Journal of Economic Theory*, 92(1):66–95, 2000.

J. Hartmanis and R. E. Stearns. On the computational complexity of algorithms. *Transactions of the American Mathematical Society*, 117:285–306, 1965. URL https://api.semanticscholar.org/CorpusID:40185800.

B. Holmström. Groves' scheme on restricted domains. *Econometrica*, pages 1137–1144, 1979.

R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47(1):104–123, 2004. doi: https://doi.org/10.1016/S0899-8256(03)00184-2.

T. Horel, S. Ioannidis, and S. Muthukrishnan. Budget feasible mechanisms for experimental design. In *Latin American Symposium on Theoretical Informatics*, pages 719–730. Springer, 2014.

L. Hurwicz, R. Radner, and S. Reiter. A stochastic decentralized resource allocation process: Part i. *Econometrica*, 43(2):187–221, 1975. URL http://www.jstor.org/stable/1913581.

P. Jalaly and É. Tardos. Simple and efficient budget feasible mechanisms for monotone submodular valuations. *ACM Trans. Economics and Comput.*, 9(1):4:1–4:20, 2021.

J. H. Kagel, R. M. Harstad, and D. Levin. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica: Journal of the Econometric Society*, pages 1275–1304, 1987.

A. S. Kelso and V. P. Crawford. Job matching, coalition formation , and gross substitute. *Econometrica*, 50(6):1483–1504, 1982.

W. Kool, H. van Hoof, and M. Welling. Attention, learn to solve routing problems! In *International Conference on Learning Representations*, 2019.

H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955.

R. Lavi and C. Swamy. Truthful and near-optimal mechanism design via linear programming. *J. ACM*, 58(6), dec 2011. doi: 10.1145/2049697.2049699.

R. Lavi, A. Mu'Alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS '03)*, pages 574–583. IEEE, 2003.

D. Lehmann, L. I. Oćallaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *J. ACM*, 49(5):577–602, sep 2002. doi: 10.1145/585265.585266.

D. Lehmann, R. Müller, and T. Sandholm. The winner determination problem. In *Combinatorial Auctions*, pages 297–318. The MIT Press, 2005. doi: 10.7551/mitpress/9780262033428.003.0013.

R. P. Leme. Gross substitutability: An algorithmic survey. *Games and Economic Behavior*, 106:294–316, 2017.

H. B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–79, 1983.

L. W. McKenzie. On the existence of general equilibrium for a competitive market. *Econometrica*, pages 54–71, 1959.

P. Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108(2):245–272, 2000. doi: 10.1086/262118.

P. Milgrom and I. Segal. Clock auctions and radio spectrum reallocation. *Journal of Political Economy*, 128(1):1–31, 2020.

R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981. URL http://www.jstor.org/stable/3689266.

R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265 – 281, 1983. doi: http://dx.doi.org/10.1016/0022-0531(83)90048-0. URL http://www.sciencedirect.com/science/article/pii/0022053183900480.

G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical Programming*, 14:265–294, 12 1978. doi: 10.1007/BF01588971.

N. Nisan. Introduction to mechanism design (for computer scientists). In *Algorithmic Game Theory*. 2007.

N. Nisan and A. Ronen. Computationally feasible vcg mechanisms. *J. Artif. Int. Res.*, 29 (1):19–47, may 2007.

C. Papadimitriou, M. Schapira, and Y. Singer. On the hardness of being truthful. In *Foundations of Computer Science (FOCS'08)*, pages 250–259. IEEE, 2008.

D. C. Parkes. Iterative combinatorial auctions. In *Combinatorial Auctions*, pages 41–78. The MIT Press, 2005. doi: 10.7551/mitpress/9780262033428.003.0003.

S. Polasky, D. J. Lewis, A. J. Plantinga, and E. Nelson. Implementing the optimal provision of ecosystem services. *Proceedings of the National Academy of Sciences*, 111(17):6248–6253, 2014. doi: 10.1073/pnas.1404484111.

A. Roth and G. Schoenebeck. Conducting truthful surveys, cheaply. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 826–843, 2012.

A. E. Roth. The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy*, 92(6):991–1016, 1984. doi: 10.1086/261272.

A. E. Roth and M. Sotomayor. The college admissions problem revisited. *Econometrica*, 57(3):559–570, 1989. URL http://www.jstor.org/stable/1911052.

A. E. Roth, T. Sönmez, and M. U. Ünver. Kidney exchange. *The Quarterly Journal of Economics*, 119(2):457–488, 2004. URL http://www.jstor.org/stable/25098691.

T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1):1–54, 2002. doi: https://doi.org/10.1016/S0004-3702(01)00159-X.

M. A. Satterthwaite. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975. doi: https://doi.org/10.1016/0022-0531(75)90050-2.

T. Scheffel, G. Ziegler, M. Bichler, and R. Jacob. On the impact of cognitive limits in combinatorial auctions: An experimental study in the context of spectrum auction design. *Experimental Economics*, 15, 12 2012. doi: 10.1007/s10683-012-9321-0.

L. S. Shapley and M. Shubik. The assignment game I: The core. *International Journal of Game Theory*, 1(1):111–130, 1971.

A. Shioura and A. Tamura. Gross substitutes condition and discrete concavity for multi-unit valuations: a survey. *Journal of the Operations Research Society of Japan*, 58(1):61–103, 2015.

Y. Singer. Budget feasible mechanisms. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 765–774. IEEE, 2010.

Y. Singer. How to win friends and influence people, truthfully: influence maximization mechanisms for social networks. In *Proceedings of the fifth ACM international conference on Web search and data mining*, pages 733–742, 2012.

Y. Singer and M. Mittal. Pricing mechanisms for crowdsourcing markets. In *Proceedings of the 22nd international conference on World Wide Web*, pages 1157–1166, 2013.

D. Talman and Z. Yang. An efficient multi-item dynamic auction with budget constrained bidders. *International Journal of Game Theory*, 44:769 – 784, 2011.

G. van der Laan and Z. Yang. An ascending multi-item auction with financially constrained bidders. *Journal of Mechanism and Institution Design*, 1:109–149, 2016.

V. V. Vazirani. *Combinatorial Algorithms for Market Equilibria*, page 103–134. Cambridge University Press, 2007. doi: 10.1017/CBO9780511800481.007.

V. V. Vazirani and M. Yannakakis. Market equilibrium under separable, piecewise-linear, concave utilities. *J. ACM*, 58(3), jun 2011. doi: 10.1145/1970392.1970394.

W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

L. Walras. 1954. *Elements of pure economics*, 1874.

M. Weiss, B. Lubin, and S. Seuken. Sats: A universal spectrum auction test suite. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '17, page 51–59, 2017.

J. Weissteiner and S. Seuken. Deep learning-powered iterative combinatorial auctions. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI-20).*, 2020.

J. Weissteiner, J. Heiss, J. Siems, and S. Seuken. Monotone-value neural networks: Exploiting preference monotonicity in combinatorial assignment. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, 2022a. doi: 10.24963/ijcai.2022/77.

J. Weissteiner, C. Wendler, S. Seuken, B. Lubin, and M. Püschel. Fourier analysis-based iterative combinatorial auctions. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, 2022b. doi: 10.24963/ijcai.2022/78.

J. Weissteiner, J. Heiss, J. Siems, and S. Seuken. Bayesian optimization-based combinatorial assignment. *CoRR*, abs/2208.14698, 2023. doi: https://doi.org/10.48550/arXiv.2208.14698.

R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3-4):229–256, 1992. doi: 10.1007/bf00992696.