Breaking Barriers, Finding Boundaries: Not Obviously Manipulable Budget-Feasible Mechanism Design

Bart de Keijzer[*1], **Guido Schäfer**[*2, 3], **Artem Tsikiridis**[*4], **Carmine Ventre**[*1]

[1]Department of Informatics, King's College London, UK
[2]Centrum Wiskunde & Informatica (CWI), The Netherlands
[3]Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands
[4]Department of Computer Science, Technical University of Munich, Germany
bart.de_keijzer@kcl.ac.uk, g.schaefer@cwi.nl, artem.tsikiridis@tum.de, carmine.ventre@kcl.ac.uk

Abstract

Strategyproofness has been the holy grail in mechanism design for decades, providing strong incentive compatibility guarantees under the assumption of perfectly rational agents. However, this assumption is questionable when agents exhibit bounded rationality. Moreover, strategyproofness often imposes strong impossibility results that prevent mechanisms from surpassing certain approximation barriers. We study this tension in budget-feasible mechanism design, where a designer wants to procure services of maximum value from agents subject to a budget constraint. Here, strategyproofness imposes approximation barriers of 2.41 and 2 for deterministic and randomized mechanisms, respectively. We investigate how much we can potentially gain under bounded rationality. We adopt the weaker notion of not obviously manipulable (NOM), which only prevents "obvious" strategic deviations. We fully resolve the achievable approximation guarantees under NOM: We derive a deterministic 2-approximate NOM mechanism under the general class of monotone subadditive valuations. We also show that this bound is tight (even for additive valuations). Additionally, we provide a simple randomized NOM mechanism that is approximately optimal. These results demonstrate a clear separation between strategyproof and NOM mechanisms. Our mechanisms use Golden Tickets and Wooden Spoons as natural design primitives, arising from our characterization of NOM mechanisms.

Extended version — https://arxiv.org/abs/2502.12306

1 Introduction

Incentive compatibility is a central goal in mechanism design, ensuring that agents cannot benefit from misrepresenting their private information. Among the various notions of incentive compatibility, *strategyproofness (SP)* has been a fundamental objective in mechanism design for decades. The reason for its importance lies in the strong guarantee it provides: agents have provably no incentive to misreport their true types (private information), regardless of the strategic choices of the other agents. However, this notion relies on the assumption that agents are perfectly rational, meaning that they always make fully informed and optimal decisions. This assumption becomes questionable when agents exhibit bounded rationality as observed in real-world settings. For example, there is empirical evidence that agents may fail to identify profitable deviations; see (Troyan and Morrill 2020) for a more detailed discussion. Moreover, because strategyproofness is such a demanding notion, it often imposes strong impossibility results that prevent mechanisms from surpassing certain approximation barriers in terms of economic efficiency.

These limitations have recently led to the study of weaker incentive compatibility notions that potentially offer improved performance while still providing meaningful strategic guarantees. One such notion is *not obviously manipulable (NOM)*, introduced by Troyan and Morrill (2020). NOM guarantees incentive compatibility only with respect to deviations that are "obviously beneficial". Basically, NOM ensures that each agent has no incentive to misreport when the other agents choose a best or worst opposing strategy profile. NOM is a weaker notion than SP and thus enlarges the class of incentive compatible mechanisms. Also, NOM is backed by empirical evidence of behavioral biases; see (Troyan and Morrill 2020) for more details.

The tension between strategyproofness and approximate economic efficiency is particularly pronounced in the context of *budget-feasible mechanism design*, a fundamental mechanism design problem introduced by Singer (2010). This setting models a procurement auction in which a principal has a fixed budget B and seeks to procure services from a set of agents $N = \{1, \dots, n\}$. The principal's preferences are described by a valuation function V that assigns a value to each subset of agents. Each agent i incurs a private cost for delivering their service, which they may strategically misreport. The principal can hire any subset of agents but must pay at least their declared costs, without exceeding the total budget B. The objective is to choose a subset $X \subseteq N$ that maximizes $V(X)$ while satisfying the budget constraint.

The design of strategyproof budget-feasible mechanisms has been extensively studied and is by now relatively well

[*]These authors contributed equally.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

understood. In the basic setting with additive valuations, Gravin et al. (2020) present the best known deterministic SP mechanism achieving a 3-approximation. Chen, Gravin, and Lu (2011) establish a lower bound of $1+\sqrt{2} \approx 2.41$ on the approximation ratio achievable by any deterministic mechanism. Randomization allows for strictly better performance: Gravin et al. (2020) give a randomized SP mechanism achieving a 2-approximation in expectation, which they also show is optimal among randomized SP mechanisms. Importantly, these lower bounds hold independently of any computational constraints—the approximation barriers stem solely from the requirement of strategyproofness.

In this paper, we initiate the study of budget-feasible mechanism design under the weaker incentive compatibility notion of not obviously manipulable. The main question that we address is: *How much can NOM budget-feasible mechanisms improve upon the approximation barriers imposed by strategyproofness?* The primary motivation behind our investigations is to understand the potential improvements in economic efficiency by adopting NOM[1]. Given that NOM is a weaker notion, we expect the possible improvements with respect to strategyproofness to be significant.

Our Contributions. We present a general template for designing budget-feasible mechanisms satisfying NOM and demonstrate its effectiveness by deriving mechanisms with optimal approximation guarantees, even for the broad class of monotone subadditive valuations.

Full understanding of approximation landscape: We resolve the achievable approximation guarantees of budget-feasible mechanisms under NOM, for both deterministic and randomized mechanisms. We derive a deterministic 2-approximate NOM mechanism for the class of monotone subadditive valuations and show that this is best possible (even for additive valuations). Further, we design a simple randomized NOM mechanism achieving a 1-approximation in expectation. These bounds assume access to an optimal algorithm for the underlying NP-hard knapsack problem. If instead a poly-time γ-approximate algorithm is used with $\gamma \geq 1$, the approximation ratios degrade to $\max\{2, \gamma\}$ and $\max\{1, \gamma\}$, respectively.

Complete characterization of NOM mechanisms: We characterize the class of budget-feasible NOM mechanisms. More specifically, NOM requires that for each agent, the maximum and minimum utility achievable through truthful reporting (over all reports by others) are no worse than the corresponding utility extremes achievable through misreporting; these conditions are termed *best-case NOM (BNOM)* and *worst-case NOM (WNOM)*, respectively. We derive a characterization for BNOM and WNOM independently, refining the characterization given in (Archbold, de Keijzer, and Ventre 2023a). In contrast to their characterization, we obtain an explicit description of the allocation and payment functions for any mechanism satisfying BNOM or WNOM in our budget-feasible setting.

[1]In particular, we take no position on whether NOM represents the 'right' notion of incentive compatibility for boundedly rational agents. We refer the reader to (Troyan and Morrill 2020) for its rationale and motivational applications.

Optimal mechanisms for BNOM and WNOM: We derive impossibility results for deterministic budget-feasible mechanisms satisfying BNOM or WNOM only. We prove lower bounds of 2 and $\varphi = (1+\sqrt{5})/2$ (golden ratio) respectively, on the achievable approximation ratios. Our lower bounds hold for additive valuations. As a consequence, our deterministic 2-approximate (B)NOM mechanism for subadditive valuations is best possible. We also obtain a WNOM φ-approximation mechanism for subadditive valuations, which is thus optimal.

Significance and Technical Merits. Our results establish a clear separation in approximation guarantees between strategyproof and NOM mechanisms. We break the known approximation barriers of $1+\sqrt{2}$ and 2 for deterministic and randomized strategyproof mechanisms by obtaining approximation guarantees of 2 and 1, respectively. When computational efficiency matters, these approximation ratios degrade to $\max\{2, \gamma\}$ and $\max\{1, \gamma\}$, where γ is the best known approximation ratio for the underlying knapsack problem.

Archbold, de Keijzer, and Ventre (2024) introduced the class of *Willy Wonka mechanisms* to achieve NOM. These mechanisms ensure BNOM and WNOM by issuing *golden tickets* and *wooden spoons*. In our context, golden tickets mean that for each agent i and cost declaration c_i, there is a cost profile of the other agents for which i is hired and paid the whole budget B. Similarly, wooden spoons guarantee that there is a cost profile of the other agents for which i is not hired. We adapt this framework to design deterministic and randomized NOM budget-feasible mechanisms with strong approximation guarantees. Notably, our mechanisms extend without loss in approximation guarantee to monotone subadditive valuations, which are significantly more expressive than additive or submodular functions.

Our framework yields a deterministic NOM budget-feasible mechanism that achieves an approximation ratio of $\max\{2, \gamma\}$. For the special case of monotone submodular valuations, the best known SP mechanism achieves an approximation ratio of 4.45 (Han, Zhang, and Cui 2025). The algorithmic strength of NOM becomes even more pronounced with randomization. By selecting golden tickets and wooden spoons at random, we obtain a simple yet versatile randomized mechanism that is universally NOM and achieves approximation ratio $\max\{1, \gamma\}$. For the special case of monotone submodular valuations, the best known randomized SP mechanism achieves an approximation ratio of 4.08 (Han, Zhang, and Cui 2025). This substantial improvement highlights the flexibility of our framework. Moreover, our randomized scheme is more broadly applicable and extends beyond the binary allocation setting here.

Our characterization reveals how NOM relaxes SP through two threshold-based conditions: BNOM requires a threshold b_i for each agent i such that i is never selected when their declared cost $c_i \geq b_i$, but is selected and paid b_i in some profile when $c_i \leq b_i$. WNOM instead uses a dual threshold w_i: i is always selected and paid w_i when $c_i \leq w_i$, but is not selected in some profile when $c_i \geq b_i$. Using this characterization, we derive lower bounds of 2 and φ on the approximation ratios achievable by BNOM and

WNOM mechanisms, respectively, which hold even for additive valuations. This implies that our 2-approximate mechanism is optimal with respect to (B)NOM for all classes of monotone valuation functions between additive and subadditive. For WNOM there is some leeway and we derive a non-trivial budget-feasible mechanism achieving a φ-approximation for subadditive valuations, which is optimal.

Overall, our results confirm that NOM budget-feasible mechanisms can achieve improved approximation guarantees. Also, when randomization is allowed, these improvements are optimal in the sense that our mechanisms match the best-known guarantees of the underlying knapsack problem. However, non-trivial boundaries remain in the deterministic setting: Surprisingly, our lower bound reveals an inherent limitation of deterministic mechanisms: even when imposing only mild incentive compatibility requirements (i.e., BNOM), an approximation barrier of 2 is unavoidable.

Related Work. The NOM solution concept was introduced by Troyan and Morrill (2020) to model scenarios where agents, due to limited contingent reasoning abilities, might not fully exploit strategic behavior to their advantage. In their work, Troyan and Morrill (2020) characterize NOM mechanisms in both settings without monetary transfers (such as bipartite matching) and with monetary transfers (such as auctions and bilateral trade). Interestingly, for the bilateral trade setting, Archbold, de Keijzer, and Ventre (2023b) showed in a follow-up work a separation between BNOM and WNOM, which, in some sense, aligns with the results of our work. Specifically, they show that while there exist bilateral trade mechanisms that are individually rational, efficient, weakly budget-balanced, and WNOM, the same cannot be achieved when substituting WNOM with BNOM under any relaxation of weak budget-balance. NOM has also been studied in other domains, including revenue maximization (Archbold, de Keijzer, and Ventre 2024), voting rules (Aziz and Lam 2021; Elkind, Neoh, and Teh 2024), fair division (Psomas and Verma 2022; Ortega and Segal-Halevi 2022), assignment mechanisms (Troyan 2024), hedonic games (Flammini, Fomenko, and Varricchio 2025; Ferraioli and Varricchio 2025), and settings beyond direct-revelation mechanisms (Archbold, de Keijzer, and Ventre 2023a). Finally, we note that there is an alternative incentive compatibility notion that strengthens strategyproofness called *obvious strategyproofness (OSP)* due to Li (2017), see also (Ferraioli and Ventre 2023) and references within.

The study of strategyproof budget-feasible procurement auctions was initiated by Singer (2010). Beyond the canonical setting with additive valuations Singer (2010) (see also (Gravin et al. 2020; Chen, Gravin, and Lu 2011)), several other budget-feasible mechanism design settings have been explored in the literature. These include settings with more general valuation functions such as monotone submodular valuation functions (Chen, Gravin, and Lu 2011; Jalaly and Tardos 2021; Balkanski et al. 2025; Han, Zhang, and Cui 2025), non-monotone submodular valuation functions (Amanatidis, Kleer, and Schäfer 2022; Balkanski et al. 2025; Huang et al. 2023; Han, Zhang, and Cui 2025), and XOS and subadditive valuation functions (Bei et al. 2017; Amanatidis, Birmpas, and Markakis 2017; Dobzinski, Papadimitriou, and Singer 2011; Balkanski et al. 2025; Neogi, Pashkovich, and Swamy 2024, 2025). Additionally, different feasibility restrictions have been considered (Amanatidis, Birmpas, and Markakis 2016; Huang et al. 2023; Leonardi et al. 2021). Finally, the model has been studied under a large-market assumption (Anari, Goel, and Nikzad 2018) and for divisible agents (Klumper and Schäfer 2022; Amanatidis et al. 2025). For further details, we refer the reader to the survey of Liu et al. (2024).

2 Model and Preliminaries

We consider a procurement auction involving a set of agents $N = \{1, \dots, n\}$ and an auctioneer who has some budget $B \in \mathbb{R}_{>0}$ available. Each agent $i \in N$ offers a service and has a private cost parameter $t_i \in \mathbb{R}_{\geq 0}$, representing their true cost for providing this service. The auctioneer can allocate services from arbitrary subsets $X \subseteq N$ and has access to a valuation set function $V : 2^N \mapsto \mathbb{R}_{\geq 0}$, where $V(X)$ defines the value of choosing agent subset $X \subseteq N$. Throughout, we assume without loss of generality that agents in N are ordered so that $V(\{1\}) \geq \cdots \geq V(\{n\})$. Moreover, we assume that V is *monotone*, i.e., for every $S \subseteq T \subseteq N$, we have $V(S) \leq V(T)$, and *normalized*, i.e., $V(\emptyset) = 0$. We introduce some specific classes of valuation functions below.

Mechanisms. A *deterministic mechanism* $\mathcal{M}$ consists of an *allocation rule* $\mathbf{x} : \mathbb{R}^n_{\geq 0} \to \{0, 1\}^n$ and a *payment rule* $\mathbf{p} : \mathbb{R}^n_{\geq 0} \to \mathbb{R}^n_{\geq 0}$. The auctioneer collects a cost profile $\mathbf{c} = (c_i)_{i \in N} \in \mathbb{R}^n_{\geq 0}$, where $c_i \in \mathbb{R}_{\geq 0}$ is the cost *declared* by agent i (which may differ from their true cost (type) t_i). Given $\mathbf{c}$, the auctioneer determines an allocation $\mathbf{x}(\mathbf{c}) = (x_i(\mathbf{c}))_{i \in N}$, where $x_i(\mathbf{c}) \in \{0, 1\}$ indicates whether agent i is chosen or not. Given an allocation $\mathbf{x}(\mathbf{c})$, we use $X(\mathbf{c}) = \{i \in N \mid x_i(\mathbf{c}) = 1\}$ to refer to the set of agents who are selected under $\mathbf{x}(\mathbf{c})$; we use $\mathbf{x}(\mathbf{c})$ and $X(\mathbf{c})$ interchangeably. The auctioneer also determines a vector of payments $\mathbf{p}(\mathbf{c}) = (p_i(\mathbf{c}))_{i \in N}$, where $p_i(\mathbf{c}) \in \mathbb{R}_{\geq 0}$ is the payment that agent i receives for their service. We also consider *randomized mechanisms*. A randomized mechanism $\mathcal{M}$ is a probability distribution over deterministic ones.

We denote an instance of our procurement auction by the tuple $I = (N, \mathbf{c}, V, B)$; we refer to an instance simply by its declared cost profile $\mathbf{c}$, when the rest of the input is clear from context. For any agent $i \in N$, we use $\mathbf{c}_{-i}$ to denote the cost profile restricted to all agents except i.

Not Obviously Manipulable. We consider the setting where each agent $i \in N$ wants to maximize their *quasi-linear utility function* defined as $u_i^{t_i}(\mathbf{c}) = p_i(\mathbf{c}) - t_i \cdot x_i(\mathbf{c})$. Mechanisms that are individually rational and strategyproof for single parameter domains (as considered here) are generally well understood. Myerson (1981) gives a complete characterization of the properties that an allocation rule must satisfy and provides a formula to derive the corresponding payments. In this work, we focus on a less stringent notion of incentive compatibility, which is more suitable for imperfectly rational agents (see (Troyan and Morrill 2020)).

Definition 1 (Troyan and Morrill (2020))**.** *A deterministic mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ is not obviously manipulable (NOM) if it satisfies the following two properties:* Best-

Case Not Obviously Manipulable (BNOM): *for every agent $i \in N$, and all $t_i \in [0, B]$ it holds that*

$$\sup_{\mathbf{c}_{-i}} u_i^{t_i}(t_i, \mathbf{c}_{-i}) \geq \sup_{\mathbf{c}_{-i}} u_i^{t_i}(c_i, \mathbf{c}_{-i}) \qquad \forall c_i \geq 0. \tag{1}$$

Worst-Case Not Obviously Manipulable (WNOM): *for every agent $i \in N$, and all $t_i \in [0, B]$ it holds that*

$$\inf_{\mathbf{c}_{-i}} u_i^{t_i}(t_i, \mathbf{c}_{-i}) \geq \inf_{\mathbf{c}_{-i}} u_i^{t_i}(c_i, \mathbf{c}_{-i}) \qquad \forall c_i \geq 0. \tag{2}$$

A randomized algorithm is said to be *universally not obviously manipulable* if it is a probability distribution over deterministic NOM mechanisms.

Additional Design Objectives. We are interested in designing a NOM mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ that additionally satisfies the following properties for every cost profile $\mathbf{c}$: (i.) *Individual Rationality (IR):* A selected agent receives at least their cost as payment, and the payments are non-negative always, i.e., for every $i \in N$, $p_i(\mathbf{c}) \geq c_i x_i(\mathbf{c})$; (ii.) *Budget-Feasibility (BF):* The auctioneer's total amount of payments is at most the budget, i.e., $\sum_{i \in N} p_i(\mathbf{c}) \leq B$. (iii.) *Normalized Payments (NP):* A non-selected agent receives no payment, i.e., for every $i \notin X(\mathbf{c})$, $p_i(\mathbf{c}) = 0$.

Note that a mechanism $\mathcal{M}$ satisfying both IR and BF cannot select any agent $i \in N$ whose declared cost exceeds the budget, i.e., $c_i > B$. Throughout this paper, we therefore assume that $c_i, t_i \in [0, B]$ for all agents $i \in N$.

Approximation Guarantees. In budget-feasible mechanism design, the quality of an allocation computed by a mechanism is evaluated against the optimal solution to the underlying packing problem (see Singer (2010)). Given an instance $I = (N, \mathbf{c}, V, B)$, we measure the performance of a mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ by comparing $V(X(\mathbf{c}))$ with the optimal solution of the following *packing problem*:

$$\max \quad V(X) \quad \text{s.t.} \quad \sum_{i \in X} c_i \leq B, \quad X \subseteq N. \tag{3}$$

We use $X^*(\mathbf{c})$ to refer to an optimal agent-set solution maximizing (3); similarly, we use $\mathbf{x}^*(\mathbf{c})$ to denote the respective binary allocation. In case of ties, the optimal set $X^*(\mathbf{c})$ is chosen with respect to any fixed strict total order over 2^N that prefers higher-cardinality sets over lower-cardinality sets. We say that a deterministic mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ is *α-approximate* with $\alpha \geq 1$ if for every cost profile $\mathbf{c}$ it holds that $\alpha V(X(\mathbf{c})) \geq V(X^*(\mathbf{c}))$.

Our mechanism makes use of an approximation algorithm for the packing problem in (3). We use APX to refer to any such (deterministic) algorithm. Given an instance $I = (N, \mathbf{c}, V, B)$, we use $X_{\text{APX}}(\mathbf{c})$ and $\mathbf{x}_{\text{APX}}(\mathbf{c})$ to refer to the agent-set solution and binary allocation, respectively, computed by APX. We assume that APX has an approximation ratio of $\gamma \geq 1$, i.e., $\gamma V(X_{\text{APX}}(\mathbf{c})) \geq V(X^*(\mathbf{c}))$.

Classes of Valuation Functions. Generally, we assume that the valuation function V is monotone and normalized (i.e., without any further restrictions). However, sometimes we consider more specific valuation functions: *Additive Valuation Functions* V_{ADD}: For all $S, T \subseteq N$ with $S \cap T = \emptyset$, we have $V(S \cup T) = V(S) + V(T)$. Equivalently, V can be defined as $V(S) = \sum_{i \in S} v_i$ for every $S \subseteq N$, where $\mathbf{v} = (v_i)_{i \in N} \in \mathbb{R}^n_{\geq 0}$ is some value profile. *Subadditive Valuation Functions* V_{SUB}: For all $S, T \subseteq N$, we have $V(S \cup T) \leq V(S) + V(T)$. Note that an additive valuation function V is subadditive; recall that we assume that V is non-negative always.

Knapsack Approximation Guarantees. The approximation guarantee of our mechanisms depends on that of APX for the packing problem in (3), which depends on the class of valuation functions. Clearly, for $V = V_{\text{ADD}}$, the packing problem in (3) is the classical *knapsack problem*. This problem is known to be NP-hard but admits a fully polynomial-time approximation scheme (FPTAS) (see, e.g., Vazirani (2001)).[2] When $V = V_{\text{SUB}}$, the packing problem in (3) is notoriously hard in the general case. In particular, Singer (2010) has shown that obtaining a $\gamma = o(n)$ requires an exponential number of queries to V_{SUB}.[3] However, with access to the stronger demand-query oracle, Badanidiyuru, Dobzinski, and Oren (2019) present a γ-approximation with $\gamma = 2 + \varepsilon$ for any $\varepsilon > 0$, using a polynomial number of oracle calls. Finally, for the subclass of *submodular valuations*, i.e., $V(S \cup T) \leq V(S) + V(T) - V(S \cap T)$ for all $S, T \subseteq N$, there is a γ-approximation with $\gamma = e/(e-1) \approx 1.58$, even for the standard value-query model (see e.g., Khuller, Moss, and Naor (1999) and Sviridenko (2004)).

Due to space restrictions, some proofs are deferred to the extended version.

3 Design Framework for NOM Mechanisms

The main result of this section is a general-purpose deterministic mechanism that is individually rational (IR), budget-feasible (BF) and not obviously manipulable (NOM), named WILLYWONKA (Mechanism 1). Even though this mechanism achieves all our mechanism design objectives, as a stand-alone mechanism it does not offer tangible approximation guarantees. However, below we use this mechanism in certain types of *compositions* of mechanisms (both deterministic and randomized) to derive optimal approximation guarantees. More specifically, by combining our WILLYWONKA mechanism with another simple mechanism, we derive a deterministic mechanism that is IR, BF and NOM and achieves an approximation guarantee of 2 for the general class of monotone subadditive valuation functions. The approximation guarantee of our mechanism is best possible: As we show in Section 4.1, even for the more restrictive class of additive valuation functions, no deterministic mechanism satisfying IR, BF and BNOM (only) can achieve an approximation guarantee better than 2.

The section is structured as follows. In Section 3.1, we present our WILLYWONKA mechanism and show that it satisfied IR, BF and NOM. Then, in Section 3.2, we combine our WILLYWONKA mechanism with another simple mechanism (trivially satisfying IR, BF and NOM) and prove that

[2]Recall that a *fully polynomial-time approximation scheme (FPTAS)* computes a $(1 + \varepsilon)$-approximate solution, for any given $\varepsilon > 0$, in time polynomial in the input size and $1/\varepsilon$.

[3]This is true even for the special case of monotone XOS functions, see e.g., (Lehmann, Lehmann, and Nisan 2001).

agent i	$(c_i, \mathbf{c}^{\text{GT}}_{-i})$	$(c_i, \mathbf{c}^{\text{WS}}_{-i})$
1	(c_1, B, B)	$(c_1, 0, 0)$
2	$(0, c_2, B)$	$(0, c_2, 0)$
3	$(0, 0, c_3)$	(B, B, c_3)

Table 1: Golden tickets and wooden spoons for $n = 3$.

the resulting mechanism achieves an approximation ratio of 2 for monotone subadditive valuation functions.

3.1 WILLYWONKA Mechanism

We describe our new WILLYWONKA mechanism in more detail. The mechanism achieves NOM by using an approach similar to the one by Archbold, de Keijzer, and Ventre (2024), implementing *golden tickets* and *wooden spoons*. In Section 4, we introduce more refined notions of golden tickets and wooden spoons for the budget-feasible setting to obtain our characterization results for BNOM and WNOM. By exploiting this characterization, we can derive a simple mechanism using carefully designed golden tickets and wooden spoons, as described below.

High-Level Idea. The high-level idea behind this approach is to define for each agent two special cost profiles of the opposing agents, called *golden ticket* and *wooden spoon*, that trigger particular outcomes when they occur. The golden tickets realize BNOM by ensuring that for each agent $i \in N$ the left-hand side of (1) attains maximum utility. Similarly, the wooden spoons implement WNOM by ensuring that for each agent $i \in N$ the right-hand side of (2) is non-positive.

The golden tickets and wooden spoons used by our WILLYWONKA mechanism are defined as follows. For each agent $i \in N$ with $c_i \in [0, B)$, we define:

$$\mathbf{c}^{\text{GT}}_{-i} = (\underset{1}{0}, \underset{\dots}{\dots}, \underset{i-1}{0}, \underset{i+1}{B}, \underset{\dots}{\dots}, \underset{n}{B}) \tag{4}$$

$$\mathbf{c}^{\text{WS}}_{-i} = \begin{cases} (0, \dots, 0) & \text{if } i < n \\ (B, \dots, B) & \text{if } i = n. \end{cases} \tag{5}$$

Note that that these golden tickets and wooden spoons do not 'interfere', i.e., each agent i can effectuate their golden ticket and wooden spoon. In fact, for this it is crucial that they are defined for $c_i \neq B$ only; see Table 1 for an example with $n = 3$. The profile $\mathbf{c} = (0, \dots, 0)$ is the golden ticket of agent n and the wooden spoons of agents 1 to $n-1$ (which can be effectuated simultaneously). It will become clear below that the case $c_i = B$ is handled automatically (due to IR and BF of the mechanism).

Mechanism. A detailed description of our WILLYWONKA mechanism is given in Mechanism 1. The mechanism takes an instance $I = (N, \mathbf{c}, V, B)$ as input. Recall that, by assumption, the agents are ordered such that $V(\{1\}) \geq \cdots \geq V(\{n\})$. Then, the mechanism checks whether any agent admits their golden ticket with respect to the given cost profile. Note that the all-zero cost profile $\mathbf{c} = (0, \dots, 0)$ is handled through the golden ticket $\mathbf{c}^{\text{GT}}_{-n}$ of agent n with $c_n = 0$; as we will show below, this also implements the wooden spoon of each agent $i \in \{1, \dots, n-1\}$ with $c_i = 0$ correctly. After that, the mechanisms verifies if any agent admits their

---

Mechanism 1: WILLYWONKA(I)

Require: instance $I = (N, \mathbf{c}, V, B)$
1: **if** $\exists i \in N$ with $c_i \in [0, B)$ and $\mathbf{c}_{-i} = \mathbf{c}^{\text{GT}}_{-i}$ **then**
2: set $x_i = 1, p_i = B$
3: set $x_j = 1, p_j = 0$ for all $j = 1, \dots, i-1$
4: set $x_j = 0, p_j = 0$ for all $j = i+1, \dots, n$
5: **else if** $\exists i \in N$ with $c_i \in [0, B)$ and $\mathbf{c}_{-i} = \mathbf{c}^{\text{WS}}_{-i}$ **then**
6: set $x_i = 0, p_i = 0$
7: **if** $i = n$ **then**
8: set $x_1 = 1, p_1 = B$
9: set $x_j = 0, p_j = 0$ for $j \in N \setminus \{1, n\}$
10: **else**
11: set $x_j = 1, p_j = 0$ for all $j \in N \setminus \{i\}$
12: **end if**
13: **else**
14: $\mathbf{x} = \mathbf{x}_{\text{APX}}(N, \mathbf{c}, V, B)$ // *γ-approximation to* (3)
15: $\mathbf{p} = \mathbf{c} \cdot \mathbf{x}$
16: **end if**
17: **return** $(\mathbf{x}, \mathbf{p})$

---

wooden spoon. Note that for $i = n$ the wooden spoon $\mathbf{c}^{\text{WS}}_{-n} = (B, \dots, B)$ (which is structurally different to avoid interference) is implemented differently: agent 1 is allocated and paid B, while all other agents are not allocated. Finally, if no golden tickets or wooden spoons apply, the mechanism uses a γ-approximate algorithm APX to compute an approximate solution $\mathbf{x}_{\text{APX}}$ of the packing problem. In this case, the allocation is determined by $\mathbf{x}_{\text{APX}}$ and each allocated agent is paid their cost. We use $(\mathbf{x}(\mathbf{c}), \mathbf{p}(\mathbf{c})) = \text{WILLYWONKA}(I)$ to refer to the allocation and payments output by the mechanism for a given instance I.

Lemma 1. WILLYWONKA *satisfies IR, BF, NP and NOM.*

Proof. We show below that WILLYWONKA satisfies NOM; the proof for IR, BF and NP can be found in the extended version.

Consider an arbitrary agent $i \in N$. It suffices to prove that for all $t_i \in [0, B]$ the BNOM condition in (1) and the WNOM condition in (2) are satisfied.

First consider the case $t_i = B$. Then (1) and (2) follow from IR and BF: If i declares cost $t_i = B$ then their utility is $u_i^{t_i}(t_i, \mathbf{c}_{-i}) = 0$ (either i is selected and paid B by IR, or i is not selected). Alternatively, if i declares cost $c_i \in [0, B]$, we have $u_i^{t_i}(c_i, \mathbf{c}_{-i}) \leq 0$ (either i is selected and paid at most B by BF, or i is not selected). Thus, (1) and (2) hold.

Let $t_i \in [0, B)$ and consider BNOM. For every $t_i \in [0, B)$, there is a unique golden ticket $\mathbf{c}^{\text{GT}}_{-i}$ for i (as defined in (4)) that is implemented in Lines 2–4: i is selected and paid B and we have $u_i^{t_i}(t_i, \mathbf{c}^{\text{GT}}_{-i}) = B - t_i$. Since $u_i^{t_i}$ nowhere exceeds $B - t_i$, we have $\sup_{\mathbf{c}_{-i}} u_i^{t_i}(t_i, \mathbf{c}_{-i}) = B - t_i \geq \sup_{\mathbf{c}_{-i}} u_i^{t_i}(c_i, \mathbf{c}_{-i})$ for all $c_i \in [0, B]$. Thus (1) holds.

Let $t_i \in [0, B)$ and consider WNOM. Observe that for all $c_i \in (0, B)$, there is a unique wooden spoon $\mathbf{c}^{\text{WS}}_{-i}$ for i (as defined in (5)) that is implemented in Lines 6–11: i is not selected and paid 0. The same holds if $c_i = 0$ and $i = n$. Thus, we have $u_i^{t_i}(c_i, \mathbf{c}^{\text{WS}}_{-i}) = 0$ in both cases.

Mechanism 2: MaxOrWillyWonka(I)

Require: instance $I = (N, \mathbf{c}, V_{\text{SUB}}, B)$

1: $i^\star = \arg\max_{i \in N} \frac{V(\{i\})}{V(N \setminus \{i\})}$
2: **if** $V(\{i^\star\}) \geq V(N \setminus \{i^\star\})$ **then**
3:     set $x_{i^\star} = 1, p_{i^\star} = B$
4:     set $x_i = 0, p_i = 0$ for all $i \in N \setminus \{i^*\}$
5: **else**
6:     $(\mathbf{x}, \mathbf{p}) =$ WillyWonka(I)
7: **end if**
8: **return** $(\mathbf{x}, \mathbf{p})$

If $c_i = 0$ and $i \neq n$, $\mathbf{c} = (c_i, \mathbf{c}_{-i}^{\text{WS}})$ is the all-zero profile which coincides with the golden ticket of agent n with $c_n = 0$, which is implemented instead: agent $i \neq n$ is selected in this case but paid 0 (Line 3). Thus, we have $u_i^{t_i}(c_i, \mathbf{c}_{-i}^{\text{WS}}) \leq 0$ in this case. Since $u_i^{t_i}(t_i, \cdot)$ is non-negative always, we have $\inf_{\mathbf{c}_{-i}} u_i^{t_i}(t_i, \mathbf{c}_{-i}) \geq 0 \geq u_i^{t_i}(c_i, \mathbf{c}_{-i}^{\text{WS}}) \geq \inf_{\mathbf{c}_{-i}} u_i^{t_i}(c_i, \mathbf{c}_{-i})$ for all $c_i \in [0, B]$. Thus (2) holds. □

3.2 Mechanism for Subadditive Valuations

We derive a mechanism that is IR, BF, NP and NOM and achieves an approximation guarantee of $\max\{2, \gamma\}$ for the general class of subadditive valuations.

Composed Mechanism. The core idea underling our mechanism is as follows: The mechanism first checks whether there is an agent $i^\star$ who is "valuable enough" to be selected on their own, roughly compared to the optimal total value that all other agents can generate. If this is the case, the mechanism selects agent $i^\star$ and pays the entire budget B (regardless of the declared cost $c_{i^\star} \in [0, B]$). Otherwise, it calls WillyWonka (Mechanism 1) above. As it turns out, this composition allows us to prove tight approximation guarantees.[4] The resulting mechanism is referred to as MaxOrWillyWonka and given in Mechanism 2.

Theorem 2. *MaxOrWillyWonka is IR, BF, NP, NOM and $\max\{2, \gamma\}$-approximate for subadditive valuations.*

Proof. Note that whether the mechanisms executes the mechanism in the **if**-part (Lines 3 & 4) or WillyWonka(I) does not depend on the declared costs. In the former case, IR, BF and NP hold by construction. Also, the utility of each agent is constant in this case and hence the mechanism is NOM. In the latter case, IR, BF, NP and NOM are inherited from the WillyWonka mechanism (Lemma 1). It remains to prove that the approximation ratio is $\max\{2, \gamma\}$.

First, consider the case that $V(\{i^\star\}) \geq V(N \setminus \{i^\star\})$. Then $X(\mathbf{c}) = \{i^\star\}$. By the monotonicity and subadditivity of V, we have that $V(X^*(\mathbf{c})) \leq V(N) \leq V(N \setminus \{i^\star\}) + V(\{i^\star\}) \leq 2V(\{i^\star\}) = 2V(X(\mathbf{c}))$. Thus, the mechanism achieves a 2-approximation in this case.

Consider the case $V(\{i^\star\}) < V(N \setminus \{i^\star\})$. Then WillyWonka is run on I. Let $(\mathbf{x}(\mathbf{c}), \mathbf{p}(\mathbf{c}))$ be the output computed by WillyWonka(I). Below, all line numbers refer to WillyWonka. We distinguish three cases based on the profile $\mathbf{c}$.

Case 1. There is an agent $i \in N$ such that $c_i \in [0, B)$ and $\mathbf{c}_{-i} = \mathbf{c}_{-i}^{\text{GT}}$. The outcome of the mechanism is determined by Lines 2–4, i.e., $X(\mathbf{c}) = \{1, \ldots, i\}$. Note that, by the definition of $\mathbf{c}_{-i}^{\text{GT}}$, the optimal allocation $X^*(\mathbf{c})$ contains at most one agent $j \in \{i+1, \ldots, n\}$, i.e., $X^*(\mathbf{c}) \subseteq \{1, \ldots, i\} \cup \{j\}$. We thus obtain $V(X^*(\mathbf{c})) \leq V(\{1, \ldots, i\} \cup \{j\}) \leq V(\{1, \ldots, i\}) + V(\{j\}) \leq V(\{1, \ldots, i\}) + V(\{i\}) \leq 2V(\{1, \ldots, i\}) = 2V(X(\mathbf{c}))$. The first and last inequality hold by the monotonicity of V, the second inequality by the subadditivity of V, and the third inequality holds since, by assumption, $V(\{1\}) \geq \cdots \geq V(\{n\})$.

Case 2. There is an agent $i \in N$ such that $c_i \in [0, B)$ and $\mathbf{c}_{-i} = \mathbf{c}_{-i}^{\text{WS}}$. If $c_i = 0$ and $i \neq n$, then $\mathbf{c} = (c_i, \mathbf{c}_{-i}^{\text{WS}})$ is the all-zero profile which coincides with the golden ticket of agent n with $c_n = 0$. In this case, the output is determined by the golden ticket of agent n, analyzed in Case 1 above. Otherwise, the outcome of the mechanism is determined by handling the wooden spoon of i in Lines 6–11.

If $i = n$, by the definition of $\mathbf{c}_{-n}^{\text{WS}}$, the optimal allocation $X^*(\mathbf{c})$ contains at most one agent $j \in \{1, \ldots, n-1\}$. In addition, when $c_n = 0$ it also contains agent n due to the monotonicity of V. On the other hand, $X(\mathbf{c}) = \{1\}$. Therefore, $V(X^*(\mathbf{c})) \leq V(\{j, n\}) \leq V(\{j\}) + V(\{n\}) \leq 2V(\{1\}) = 2V(X(\mathbf{c}))$. The second inequality follows from the subadditivity of V and the third inequality holds since, by assumption, $V(\{1\}) \geq \cdots \geq V(\{n\})$.

If $i \neq n$, then $X(\mathbf{c}) = N \setminus \{i\}$ and $X^*(\mathbf{c}) = N$. Clearly, $V(X^*(\mathbf{c})) = V(N) \leq V(N \setminus \{i\}) + V(\{i\}) = V(N \setminus \{i\}) \cdot [1 + V(\{i\})/V(N \setminus \{i\})] \leq V(N \setminus \{i\}) \cdot [1 + V(\{i^\star\})/V(N \setminus \{i^\star\})] \leq 2V(N \setminus \{i\}) = 2V(X(\mathbf{c}))$. The first inequality follows from the subadditivity of V, the second inequality from the definition of $i^\star$ in Line 1 of MaxOrWillyWonka, and the last inequality holds because $V(\{i^\star\}) < V(N \setminus \{i^\star\})$ by assumption.

Case 3. If none of the above cases hold, the outcome of the mechanism is determined by Lines 14–15, and thus $V(X^*(\mathbf{c})) \leq \gamma V(X_{\text{APX}}(\mathbf{c})) = \gamma V(X(\mathbf{c}))$. Thus, the resulting approximation ratio is $\max\{2, \gamma\}$. □

Computational Constraints We discuss the trade-off implied by Mechanism 2 between achievable approximation ratios and computational efficiency. Without computational constraints (i.e., $\gamma = 1$), MaxOrWillyWonka achieves a 2-approximation for monotone subadditive functions, demonstrating a strict separation from deterministic, budget-feasible, individually rational, and strategyproof mechanisms, which have a known lower bound of $1 + \sqrt{2}$.

A negative result by Singer (2010) implies $\gamma = \Omega(n)$ for monotone (fractionally) subadditive valuations unless exponentially many queries to V are allowed. However, under a stronger oracle (demand queries), Badanidiyuru, Dobzinski, and Oren (2019) derive a $(2+\varepsilon)$-approximate algorithm for subadditive valuations, giving MaxOrWillyWonka the same approximation and a polynomial running time. For

[4]Note that running either one of these two mechanisms alone does not provide any non-trivial approximation guarantee.

monotone submodular valuations, a polynomial-time algorithm achieving a $\gamma = e/(e-1) \approx 1.58$ approximation due to Khuller, Moss, and Naor (1999) implies an approximation of $\max\{2, \gamma\} = 2$ for our mechanism.

4 Characterization and Impossibility Results

We derive our characterizations of budget-feasible mechanisms satisfying BNOM and WNOM. Using our BNOM characterization, we derive a lower bound of 2 on the approximation ratio of BNOM for additive valuations. In particular, this shows that the approximation ratio of MAXORWILLYWONKA (Mechanism 2) is optimal. Similarly, we establish a lower bound of $\varphi = (1+\sqrt{5})/2$ on the approximation ratio for additive valuations. We also derive a mechanism satisfying WNOM that matches this lower bound for subadditive valuations (see the extended version).

4.1 BNOM Characterization and Lower Bound

We provide a sufficient and necessary condition for mechanisms satisfying BNOM, provided that they satisfy the NP and IR properties. In words, the property states that for each agent $i \in N$ there is a particular threshold $b_i \in [0, B]$ such that: (1) declaring a type strictly above b_i guarantees i to not get selected, (2) when declaring any type strictly below b_i, there exists a bid profile of the remaining agents under which i gets selected and receives a payment of b_i, (3) when declaring type b_i, either of the above two cases apply. We formalize the above as follows.

Definition 2. *A mechanism $\mathcal{M}$ uses* threshold golden tickets *if for every $i \in N$, there exists a bid $b_i \in [0, B]$ that satisfies:*

1. $\sup_{\mathbf{c}_{-i}} x_i(c_i, \mathbf{c}_{-i}) = 0$*, for all* $c_i \in (b_i, B]$.
2. $\forall c_i \in [0, b_i) : \sup_{\mathbf{c}_{-i}} p_i(c_i, \mathbf{c}_{-i}) = b_i$.
3. $\sup_{\mathbf{c}_{-i}} p_i(b_i, \mathbf{c}_{-i}) = b_i$ *or* $\sup_{\mathbf{c}_{-i}} x_i(b_i, \mathbf{c}_{-i}) = 0$.

Proposition 3. *A mechanism $\mathcal{M}$ satisfying NP and IR is BNOM if and only if $\mathcal{M}$ uses threshold golden tickets.*

It can be seen that MAXORWILLYWONKA (Mechanism 2), which is NOM, and hence BNOM, uses threshold golden tickets with $b_i = B$ for all $i \in N$: If i declares anything less than B, then i receives a payment of B when the profile of the other agents is $\mathbf{c}_{-i}^{\mathrm{GT}}$. Furthermore, if i declares B, then by IR the payment to i is at least B, in case i gets selected. The following theorem shows that the 2-approximation achieved by MAXORWILLYWONKA is optimal among all BNOM mechanisms; in fact, this holds for additive, monotone submodular and monotone subadditive valuations.

Theorem 4. *Let $\mathcal{M}$ be any deterministic IR, BF, NP and BNOM mechanism. Then $\mathcal{M}$ is not $(2-\varepsilon)$-approximate for any $\varepsilon > 0$, for additive valuation functions.*

4.2 WNOM Characterization and Tight Bound

For WNOM, we provide a characterization in terms of thresholds that is similar to our BNOM characterization given in Section 4.1. In words, our characterization property for WNOM states for each agent $i \in N$ there is a particular threshold $w_i \in [0, B]$ such that: (1) when declaring a type strictly above w_i, there exists a bid profile of the remaining agents under which i is not selected by the mechanism, (2) declaring a type strictly below w_i guarantees i to get selected, and the minimum payment received by i is w_i, (3) when declaring type w_i, either of the above two cases apply.

Definition 3. *A mechanism $\mathcal{M}$ uses* threshold wooden spoons *if for every $i \in N$, there exists a bid $w_i \in [0, B]$ that satisfies:*

1. $\inf_{\mathbf{c}_{-i}} x_i(c_i, \mathbf{c}_{-i}) = 0$*, for all* $c_i \in (w_i, B]$.
2. $\inf_{\mathbf{c}_{-i}} p_i(c_i, \mathbf{c}_{-i}) = w_i$*, for all* $c_i \in [0, w_i)$.
3. $\inf_{\mathbf{c}_{-i}} p_i(w_i, \mathbf{c}_{-i}) = w_i$ *or* $\inf_{\mathbf{c}_{-i}} x_i(w_i, \mathbf{c}_{-i}) = 0$.

Proposition 5. *A mechanism $\mathcal{M}$ that satisfies NP and IR is WNOM if and only if $\mathcal{M}$ uses threshold wooden spoons.*

We can use the above characterization to establish a tight bound on the approximation ratio of WNOM mechanisms.

Theorem 6. *There is a deterministic IR, BF, NP, WNOM and φ-approximate mechanism for subadditive valuations. Further, the approximation ratio of φ is best possible even for additive valuations.*

5 Optimal Randomized Mechanisms

In this section, we show that for randomized mechanisms, the approximation barrier of 2 (which we established aboved for deterministic mechanisms) breaks. The idea behind our randomized mechanism $\mathcal{R}$ is very simple: $\mathcal{R}$ draws for each $i \in N$ a golden ticket $\mathbf{c}_{-i}^{\mathrm{GT}} \in [0, B]^{n-1}$ and a wooden spoon $\mathbf{c}_{-i}^{\mathrm{WS}} \in [0, B]^{n-1}$ independently and uniformly at random. The mechanism then checks whether any of the chosen golden tickets or wooden spoons applies: If for some $i \in N$ it holds that $\mathbf{c}_{-i} = \mathbf{c}_{-i}^{\mathrm{GT}}$, then only agent i is selected. If for some $i \in N$ it holds that $\mathbf{c}_{-i} = \mathbf{c}_{-i}^{\mathrm{WS}}$, no agent is selected. Otherwise, the mechanism computes an allocation $\mathbf{x} = \mathbf{x}_{\text{APX}}(\mathbf{c})$, using a γ-approximation APX for the packing problem in (3), and returns it together with the first-price payments $\mathbf{p} = \mathbf{c} \cdot \mathbf{x}$.

Theorem 7. *The randomized mechanism $\mathcal{R}$ is IR, BF, NP, universally NOM and γ-approximate in expectation.*

Note that without any computational constraints we can choose $\gamma = 1$ and obtain a 1-approximate (i.e., optimal) mechanism through randomization. $\mathcal{R}$ randomizes over infinitely many deterministic mechanisms. We can also limit the support over which the mechanism can randomize at the expense of a weaker approximation guarantee (depending on the size of the support); we defer more details to the extended version.

Note that our randomization technique is not specific to our budget-feasible setting. In fact, it applies to any mechanism design scenario where agents have sufficiently large strategy sets. This distinguishes NOM from dominant-strategy incentive compatibility: while good approximation guarantees are generally hard to obtain for the latter (see e.g., Chen, Gravin, and Lu (2011)), under NOM this solely depends on the best possible approximation guarantee γ (which is 1 without any computational constraints).

Acknowledgments

BdK was partially supported by EPSRC grant EP/X021696/1. AT, who was at CWI during the time of this work, was partially supported by the Gravitation Project NETWORKS, grant no. 024.002.003, and the EU Horizon 2020 Research and Innovation Program under the Marie Skłodowska-Curie Grant Agreement, grant no. 101034253.

References

Amanatidis, G.; Birmpas, G.; and Markakis, E. 2016. Coverage, matching, and beyond: new results on budgeted mechanism design. In *Web and Internet Economics: 12th International Conference, WINE 2016, Proceedings*, 414–428.

Amanatidis, G.; Birmpas, G.; and Markakis, E. 2017. On budget-feasible mechanism design for symmetric submodular objectives. In *Web and Internet Economics: 13th International Conference, WINE 2017, Proceedings*, 1–15.

Amanatidis, G.; Kleer, P.; and Schäfer, G. 2022. Budget-Feasible Mechanism Design for Non-monotone Submodular Objectives: Offline and Online. *Mathematics of Operations Research*, 47(3): 2286–2309.

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

Anari, N.; Goel, G.; and Nikzad, A. 2018. Budget feasible procurement auctions. *Operations Research*, 66(3): 637–652.

Archbold, T.; de Keijzer, B.; and Ventre, C. 2023a. Non-Obvious Manipulability for Single-Parameter Agents and Bilateral Trade. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023*, 2107–2115.

Archbold, T.; de Keijzer, B.; and Ventre, C. 2023b. Non-Obvious Manipulability in Extensive-Form Mechanisms: The Revelation Principle for Single-Parameter Agents. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023*, 2503–2510.

Archbold, T.; de Keijzer, B.; and Ventre, C. 2024. Willy Wonka Mechanisms. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2024*, 78–86.

Aziz, H.; and Lam, A. 2021. Obvious Manipulability of Voting Rules. In *Algorithmic Decision Theory - 7th International Conference, ADT 2021, Proceedings*, volume 13023 of *Lecture Notes in Computer Science*, 179–193.

Badanidiyuru, A.; Dobzinski, S.; and Oren, S. 2019. Optimization with demand oracles. *Algorithmica*, 81: 2244–2269.

Balkanski, E.; Garimidi, P.; Gkatzelis, V.; Schoepflin, D.; and Tan, X. 2025. Deterministic Budget-Feasible Clock Auctions. *Operations Research*.

Bei, X.; Chen, N.; Gravin, N.; and Lu, P. 2017. Worst-case mechanism design via bayesian analysis. *SIAM Journal on Computing*, 46(4): 1428–1448.

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

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

Elkind, E.; Neoh, T. Y.; and Teh, N. 2024. Temporal Elections: Welfare, Strategyproofness, and Proportionality. In *Proceedings of the 27th European Conference on Artificial Intelligence, ECAI 2024*.

Ferraioli, D.; and Varricchio, G. 2025. Non-Obvious Manipulability in Additively Separable and Fractional Hedonic Games. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, 3848–3856.

Ferraioli, D.; and Ventre, C. 2023. On the Connection between Greedy Algorithms and Imperfect Rationality. In *Proceedings of the 24th ACM Conference on Economics and Computation, EC 2023*.

Flammini, M.; Fomenko, M.; and Varricchio, G. 2025. Non-obvious Manipulability in Hedonic Games with Friends Appreciation Preferences. In *Proceedings of the 24th Conference on Autonomous Agents and Multi-Agent Systems, AAMAS 2025*.

Gravin, N.; Jin, Y.; Lu, P.; and Zhang, C. 2020. Optimal budget-feasible mechanisms for additive valuations. *ACM Transactions on Economics and Computation (TEAC)*, 8(4): 1–15.

Han, K.; Zhang, H.; and Cui, S. 2025. Efficient and effective budget-feasible mechanisms for submodular valuations. *Artificial Intelligence*, 104348.

Huang, H.; Han, K.; Cui, S.; and Tang, J. 2023. Randomized pricing with deferred acceptance for revenue maximization with submodular objectives. In *Proceedings of the ACM Web Conference 2023*, 3530–3540.

Jalaly, P.; and Tardos, É. 2021. Simple and efficient budget feasible mechanisms for monotone submodular valuations. *ACM Transactions on Economics and Computation (TEAC)*, 9(1): 1–20.

Khuller, S.; Moss, A.; and Naor, J. S. 1999. The budgeted maximum coverage problem. *Information processing letters*, 70(1): 39–45.

Klumper, S.; and Schäfer, G. 2022. Budget feasible mechanisms for procurement auctions with divisible agents. In *Algorithmic Game Theory: 15th International Symposium, SAGT 2022, Proceedings*, 78–93.

Lehmann, B.; Lehmann, D.; and Nisan, N. 2001. Combinatorial auctions with decreasing marginal utilities. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, 18–28.

Leonardi, S.; Monaco, G.; Sankowski, P.; and Zhang, Q. 2021. Budget feasible mechanisms on matroids. *Algorithmica*, 83(5): 1222–1237.

Li, S. 2017. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11): 3257–3287.

Liu, X.; Chan, H.; Li, M.; and Wu, W. 2024. Budget Feasible Mechanisms: A Survey. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, 8132–8141.

Myerson, R. 1981. Optimal auction design. *Mathematics of Operations Research*, 6(1): 58–73.

Neogi, R.; Pashkovich, K.; and Swamy, C. 2024. Budget-Feasible Mechanism Design: Simpler, Better Mechanisms and General Payment Constraints. In *15th Innovations in Theoretical Computer Science Conference, ITCS 2024*.

Neogi, R.; Pashkovich, K.; and Swamy, C. 2025. An O (log log n)-approximate budget feasible mechanism for subadditive valuations. In *Proceedings of the 26th ACM Conference on Economics and Computation, EC 2025*, 599–599.

Ortega, J.; and Segal-Halevi, E. 2022. Obvious manipulations in cake-cutting. *Social Choice and Welfare*, 59(4): 969–988.

Psomas, A.; and Verma, P. 2022. Fair and Efficient Allocations Without Obvious Manipulations. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022*.

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

Sviridenko, M. 2004. A note on maximizing a submodular set function subject to a knapsack constraint. *Operations Research Letters*, 32(1): 41–43.

Troyan, P. 2024. (Non-) obvious manipulability of rank-minimizing mechanisms. *Journal of Mathematical Economics*, 113: 103015.

Troyan, P.; and Morrill, T. 2020. Obvious manipulations. *Journal of Economic Theory*, 185: 104970.

Vazirani, V. V. 2001. *Approximation Algorithms*. Berlin, Heidelberg: Springer.