

Cap-and-Penalize: Competitive Mechanisms for Multi-Phase Regularized Online Allocation

Seyedehkimia Alaviyar, Faraz Zargari, John Tyler, Yunwei Ryan Li, Xiaoqi Tan

University of Alberta, Edmonton, AB, Canada

{alaviyar, fzargari, jetyler, yunwei.li, xiaoqi.tan}@ualberta.ca

Abstract

This paper introduces a novel mechanism for online allocation with multi-phase, non-separable regularizers, termed Cap-and-Penalize (CNP), inspired by real-world applications such as cap-and-tax policies in carbon pricing. The CNP regularizer models a multi-phase cost structure, imposing a monotone convex penalty when total allocation exceeds a pre-defined level (soft cap) and enforcing a strict limit (hard cap) beyond which allocation is prohibited. Our contributions are twofold: (1) we propose an online mechanism for CNP-regularized allocation without per-step resource constraints, which operates as a simple and intuitive posted-price mechanism, but achieves the best-possible guarantee among all possible online algorithms; (2) we tackle the more complex setting with per-step resource constraints by decomposing the regularizer into local components, yielding a similar mechanism with time-dependent marginal pricing functions. To establish the tightness of our results in both settings, we introduce a representative function-based approach that transforms the lower-bound proof into the problem of solving an ordinary differential equation with boundary conditions. We believe that this technique has the potential to be applied to other similar online optimization problems.

1 Introduction

Online allocation problems address the challenge of allocating limited resources to an online sequence of requests while maximizing social welfare or revenue and adhering to supply constraints. Compared to resource allocation in offline settings, online allocation problems introduce an additional layer of complexity, as decisions must be made without complete knowledge of future requests. Motivated by numerous high-stakes applications, online allocation problems have been extensively studied in computer science and operations research, with applications spanning Internet advertising [Mehta, 2013; Buchbinder *et al.*, 2007], ride-sharing [Dickerson *et al.*, 2021], electric vehicle charging [Guo *et al.*, 2017; Gerding *et al.*, 2011; Sun *et al.*, 2018], and revenue

management [Ball and Queyranne, 2009; Ma and Simchi-Levi, 2020].

Despite extensive studies, the literature on online allocation problems predominantly focuses on optimizing objectives that are additive and separable. For instance, in the online knapsack problem [Chakrabarty *et al.*, 2008], items arrive sequentially, each associated with a value and weight. The objective is to maximize the total value of packed items, which is separable among items and across different steps, while adhering to the knapsack’s capacity constraint. However, in many real-world scenarios, decision-makers must address non-separable objectives due to coupling effects, such as production costs [Tan *et al.*, 2020c; Huang and Kim, 2015; Azar *et al.*, 2016; Tan *et al.*, 2025; Jazi *et al.*, 2025], fairness among participants [Zargari *et al.*, 2025; Balseiro *et al.*, 2021], or regulatory constraints, such as load balancing [Balseiro *et al.*, 2021], that inherently couple all participants throughout the allocation process. For example, [Tan *et al.*, 2025] investigates an online allocation problem with convex production costs, where the objective is inherently non-separable. Similarly, [Balseiro *et al.*, 2021] shows that fairness among participants can be achieved by incorporating a regularizer that is not separable over time into the objective function.

Motivated by practical applications such as the cap-and-tax carbon pricing policy implemented in countries like Canada [Government of Canada, 2025], this paper focuses on a novel regularized online allocation problem with a multi-phase, non-separable regularizer termed Cap-and-Penalize (CNP). The CNP regularizer incorporates a multi-phase cost model for resource provision and imposes a penalty when the total allocation exceeds a predetermined production level, referred to as the *soft cap*—the maximum production level allowed without incurring additional charges (i.e., the penalty-free region). Additionally, the CNP regularizer enforces a strict upper limit on total resource production, referred to as the *hard cap*, beyond which no further allocation is permitted. Our objective is to maximize social welfare, defined as the difference between the aggregate individual utilities (which are separable across agents) and the cost imposed by the non-separable CNP regularizer. Given the inherent trade-off between these two factors in the objective, we aim to identify the optimal balance between aggregate utility and regulatory costs.

Papers	Value functions	Resource types	Resource constraints	Regularizer	Main Results
[Sun <i>et al.</i> , 2021]	Concave	Divisible	✓	None	Near-optimal CRs
[Azar <i>et al.</i> , 2016]	Linear	Divisible	✗	Non-Separable	Tight CRs
[Balseiro <i>et al.</i> , 2021]	General	Divisible	✓	Non-Separable	Sub-linear regret; Asymptotic CR
[Ma <i>et al.</i> , 2024]	General	Divisible	✓	Non-Separable	Optimal regret
[Tan <i>et al.</i> , 2020c]	Linear	Indivisible	✓	Separable	Optimal CRs
[Tan <i>et al.</i> , 2020a]	Linear	Indivisible	✓	Separable	Optimal CRs
This paper	Concave	Divisible	✗	Multi-Phase; Non-Separable	Optimal CRs
			✓	Multi-Phase; Non-Separable	Best-known CRs

Table 1: Comparison of our work with the existing literature on online allocation problems. The term “CRs” refers to competitive ratios, formally defined in Section 2.3. In [Balseiro *et al.*, 2021], “asymptotic CRs” are defined under the assumption that the time horizon approaches infinity (i.e., $T \rightarrow \infty$), whereas all other competitive ratios, including the one in this paper, are defined without this assumption.

1.1 Main Results, Techniques, and Insights

Our main results can be summarized as follows:

- We initiate the study of CNP-regularized online allocation with a multi-phase, non-separable regularizer. Unlike prior studies that focus on developing online mechanisms with sublinear regret over the time horizon, our work emphasizes competitive analysis, deriving the optimal online mechanism with the best-possible competitive ratio in the setting without resource constraints. Our upper-bound results build upon recent advancements in threshold-based algorithms and online primal-dual analysis. To establish the optimality of the proposed mechanism, we propose a novel representative function-based approach, transforming the proof of the lower bound into the search for key functions (we call them utilization functions) represented by an ordinary differential equation with two boundary conditions.
- We extend our optimal design to the more challenging setting with per-step resource constraints. The core of this extension lies in decoupling the non-separable, non-smooth regularizer across time slots, resulting in the design of a series of time-dependent marginal pricing functions. We theoretically prove that this online mechanism achieves a bounded competitive ratio in the general setting when the CNP regularizer is convex and further demonstrate that it provides a near-optimal guarantee in the linear case.

Our mechanisms belong to a family of online mechanisms known as posted price mechanisms (PPMs) [Einav *et al.*, 2018; Parkes, 2007], where the supplier posts a price signal to sequentially arriving agents, allowing them to solve their own local utility-maximization problems. PPMs are incentive compatible [Nisan *et al.*, 2007] and, more importantly, are typically easy to implement, logistically simple, and intuitive. As a result, they are widely used in practice for welfare and revenue maximization [Ferreira *et al.*, 2021; Feldman *et al.*, 2015; Tan *et al.*, 2020b; Etzion *et al.*, 2006; Tan *et al.*, 2020c]. An interesting insight from our results is the close connection between the proposed multi-phase regularizer and the resulting optimal multi-phase pricing scheme. Specifically, we show that the penalty-free region of the regularizer plays a key role in shaping the multi-phase pricing function design. While expected in theory, this insight offers

valuable guidance for designing pricing models in real-world applications like carbon pricing.

1.2 Related Work

This work builds on extensive research in resource allocation problems across computer science, economics, and operations research. Below, we review them based on the presence of (non-separable) regularizers.

Online allocation without regularizers. Online allocation problems without regularizers have been extensively studied in various forms, including online knapsack [Chakrabarty *et al.*, 2008; Sun *et al.*, 2021], bin packing [Lee and Lee, 1985; Angelopoulos *et al.*, 2023], one-way trading [El-Yaniv *et al.*, 2001; Lorenz *et al.*, 2009; Lin *et al.*, 2019], online matching [Mehta *et al.*, 2007; Devanur and Jain, 2012; Ma and Simchi-Levi, 2020] (see [Mehta, 2013] and [Huang *et al.*, 2024] for a comprehensive survey), online routing [Awerbuch *et al.*, 2005; Buchbinder and Naor, 2006], and online auctions [Blum *et al.*, 2003]. For example, [Ma and Simchi-Levi, 2020] examines an extension of online matching where resources are offered at multiple prices, achieving the best competitive ratio of $1 - \frac{1}{e}$ for deterministic algorithms. A particularly challenging variant is the generalized one-way trading problem with inventory constraints, optimally addressed in [Lin *et al.*, 2019]. Our problem, involving multiple time slots, is closely related to [Sun *et al.*, 2021], which studies an online fractional multiple knapsack problem with assignment restrictions, where each item can only be assigned to a subset of knapsacks [Kellerer *et al.*, 2004]. They propose an algorithm with a competitive ratio between $1 + \ln(U/L)$ and $2 + \ln(U/L)$, where U/L is the ratio of maximum to minimum marginal utilities. However, their work focuses solely on maximizing aggregate utilities without regularizers.

Online allocation with regularizers. A closely related line of research examines allocation problems with regularizers, modeled as negative rewards or penalties that are non-separable over time [Tan *et al.*, 2020c; Tan *et al.*, 2025; Ma *et al.*, 2024; Balseiro *et al.*, 2021; Huang and Kim, 2015; Azar *et al.*, 2016]. [Tan *et al.*, 2020c] studies online mechanism design for a single knapsack, proposing an optimal posted price mechanism for convex and linear cost functions,

but their work is limited to single-period, single-resource settings. Extensions to multi-period or multi-resource allocation appear in studies like [Tan *et al.*, 2020a; Balseiro *et al.*, 2021]. For example, [Tan *et al.*, 2020a] develops an optimal algorithm for online combinatorial auctions with convex separable cost functions, while [Balseiro *et al.*, 2021] addresses non-separable regularized problems, achieving sublinear regret and bounded asymptotic competitive ratios. In adversarial environments, the optimal competitive ratio for problems with non-separable regularizers and multiple allocation options remains open, with no optimal design yet available. [Azar *et al.*, 2016] explores online convex packing problems, proposing a competitive algorithm under specific structural assumptions but without capacity constraints. In contrast, our work focuses on decoupling multi-phase, convex regularizers while incorporating capacity constraints, avoiding additional restrictive assumptions.

Table 1 summarizes key related works in the online allocation domain and compares them with our problem.

2 Problem Statement

We first formally define the resource allocation model and then discuss the required background and assumptions. Throughout the paper, we denote $[N] = \{1, 2, \dots, N\}$ for any positive integer N .

2.1 CNP-ROA: The Model

We consider a generic online allocation problem with divisible resource and a finite time horizon T , and assume the following agent and supply models.

Agents and Valuations. A sequence of N agents arrives sequentially, each characterized by a type vector $\theta_n = (v_n, \mathcal{X}_n, \mathcal{T}_n)$, where $v_n : \mathcal{X}_n \rightarrow \mathbb{R}_+$ represents a private valuation function, \mathcal{X}_n denotes the feasible set of allocation actions, and $\mathcal{T}_n = \{t_n^{\text{on}}, t_n^{\text{on}} + 1, \dots, t_n^{\text{off}}\}$ specifies the availability interval during which the agent can receive resources (i.e., agent n arrives at time t_n^{on} and departs at time t_n^{off}). Upon the arrival of agent $n \in [N]$ at time t_n^{on} , an immediate and irrevocable decision $\mathbf{x}_n := \{x_n^t\}_{t \in \mathcal{T}_n} \in \mathcal{X}_n$ is made, and the agent will receive a valuation of $v_n(\mathbf{x}_n)$. In particular, we assume that $\mathcal{X}_n = \{\mathbf{x}_n \mid \sum_{t \in \mathcal{T}_n} x_n^t \leq D_n, \text{ and } x_n^t \in [0, d_n], \forall t \in \mathcal{T}_n\}$. Here, d_n represents the rate limit for the resource consumption per step, and D_n denotes the total demand of resources requested by agent n .

Multi-Phase Regularized Supply. We consider a regularized supply model in which resource production incurs multi-phase regulatory costs. Specifically, given the allocation profile $\{\mathbf{x}_n\}_{n \in [N]}$ for all the agents, let y be defined as

$$y = \sum_{t=1}^T \sum_{n=1}^N w_t x_n^t,$$

which represents the total resource production weighted by the coefficient $w_t \geq 0, \forall t \in [T]$. For any given $y \geq 0$, let $R(y)$ be the function defined as

$$R(y) = \begin{cases} 0 & \text{if } y \in [0, C_{\text{soft}}], \\ f(y - C_{\text{soft}}) & \text{if } y \in (C_{\text{soft}}, C_{\text{hard}}], \\ \infty & \text{if } y \in (C_{\text{hard}}, \infty), \end{cases} \quad (1)$$

where C_{soft} represents the penalty-free cap, or the *soft cap*, which allows the supplier to produce up to C_{soft} without incurring penalties. The parameter C_{hard} denotes the *hard cap*, imposing a strict upper limit on the total amount of resources being produced. Hereinafter, we refer to R as the Cap-and-Penalize (CNP) regularizer and the above CNP-regularized resource allocation problem as CNP-ROA.

Social Welfare and Mechanisms. We focus on designing an online mechanism comprising an allocation rule and a payment rule. For each agent $n \in [N]$, let the allocation and payment be denoted by \mathbf{x}_n and π_n , respectively. The utility of agent n is then given by $v_n(\mathbf{x}_n) - \pi_n$. The objective of the mechanism is to maximize social welfare:

$$\sum_{n=1}^N v_n(\mathbf{x}_n) - R(y),$$

where the payment terms cancel out. We consider a strategic setting where agents may act strategically regarding their valuation v_n to improve their outcomes. If all agents are assumed to be non-strategic, the CNP-ROA problem reduces to a standard online optimization problem, and the online mechanism simplifies to an online algorithm with only an allocation rule.

2.2 Illustrating Examples

The CNP-ROA problem can capture many interesting problems. Below we give two illustrative examples.

- **Carbon-regularized online job scheduling.** The primary application motivating CNP-ROA is the carbon-regularized job scheduling problem in data centers [Acun *et al.*, 2023; Anderson *et al.*, 2023] and electric vehicle charging [Cheng *et al.*, 2022; Lechowicz *et al.*, 2023]. For instance, agents can represent electric vehicles that arrive sequentially in an online manner to connect to a charger during their parking periods \mathcal{T}_n . The weight coefficient w_t represents the carbon intensity of energy generation at time t , while y denotes the total carbon emissions over the entire charging process. The supplier is allowed to emit up to C_{soft} units of carbon without incurring a penalty (i.e., zero cost if $y \leq C_{\text{soft}}$). The hard cap, C_{hard} , represents the maximum allowable carbon emissions, enforcing the constraint $y \leq C_{\text{hard}}$ at all times.
- **Online resource allocation with production costs.** If $w_t = 1$ for all $t \in [T]$, y represents the total amount of resources produced. Under this assumption, the CNP regularizer generalizes the production cost model employed in prior online allocation studies. Specifically, when $C_{\text{soft}} = 0$, CNP-ROA recovers the allocation problems examined in [Tan *et al.*, 2020c]. Furthermore, when both $C_{\text{soft}} = 0$ and $C_{\text{hard}} = \infty$ hold, CNP-ROA reduces to the online convex packing problem studied in [Huang and Kim, 2015; Azar *et al.*, 2016].

In addition, if we set $f(y - C_{\text{soft}}) = 0$ for all $y \in (C_{\text{soft}}, C_{\text{hard}}]$, the CNP-ROA problem simplifies to a typical online allocation problem without a regularizer or cost, a setting extensively studied in various contexts. For instance, if we assume $T = 1$ (i.e., a single-period allocation problem)

and further assume that the valuation function is linear (e.g., $v_n(x_n) = \nu_n x_n$) and $d_n = D_n = 1$, then CNP-ROA reduces to the well-known online (fractional) knapsack problem [Chakrabarty *et al.*, 2008; Sun *et al.*, 2021]. In this case, C_{hard} can be interpreted as the capacity of the knapsack.

2.3 Competitive Ratios

To specify all the information available to the mechanism designer in the online setting, we define $\mathcal{I}_{\text{known}}$ as:

$$\mathcal{I}_{\text{known}} = \{R, T, \mathbf{w}\}.$$

Given $\mathcal{I}_{\text{known}}$, the mechanism has complete knowledge of the regularizer R (including C_{soft} , C_{hard} , and f), the time horizon T , and the weight vector $\mathbf{w} = \{w_t\}_{t \in [T]}$. However, all information about the agents' types $\{\theta_n\}_{n \in [N]}$, including the total number of agents N , remains unknown a priori.

To quantify the performance of an online mechanism, we follow the celebrated competitive analysis framework [Borodin and El-Yaniv, 1998] to define the competitive ratio (CR) of an online mechanism as follows:

$$\alpha = \max_I \frac{\text{OPT}(I)}{\text{ALG}(I)},$$

where $I = \{\theta_1, \dots, \theta_N\}$ denotes an arrival instance. Given an arrival instance I , $\text{ALG}(I)$ is the welfare achieved by our online mechanism (only knows $\mathcal{I}_{\text{known}}$), and $\text{OPT}(I)$ is the optimal social welfare achievable in an offline setting assuming the type information of all arrival agents in I (in addition to $\mathcal{I}_{\text{known}}$) is given in advance. We say an online mechanism is α -competitive if it achieves a competitive ratio of α . Intuitively, $\alpha \geq 1$ holds all the time and a smaller α indicates a better performance.

2.4 Assumptions

Throughout the paper we make the following assumptions.

Assumption 1 (VALUATION FUNCTIONS). *For any agent n , the value function v_n satisfies the following conditions*

- (i) v_n is concave, monotone, and differentiable.
- (ii) $v_n(\mathbf{0}) = 0$, i.e., agents have zero utility if they are not receiving any resource.
- (iii) The partial derivative of $v_n(\mathbf{x}_n)$ is bounded, i.e., $\frac{\partial v_n}{\partial x_n^t} / w_t \in [L, U]$ for all $t \in \mathcal{T}_n$ and $n \in [N]$, where $\{w_t\}_{t \in [T]}$ are coefficients given from the regularizer. We also assume that L and U are given a priori.

We also define $\Delta = U/L$, which represents the fluctuation ratio of agents' valuations. For instance, $\Delta = 1$ indicates that all agents are homogeneous, with identical linear valuation functions, whereas a larger Δ signifies greater variation and heterogeneity among agents.

Assumption 2 (COST FUNCTION). *The cost function f is convex, monotonically non-decreasing, and differentiable. In addition, $f(0) = 0$, i.e., the regularizer $R(y)$ is a continuous function for $y \in [0, C_{\text{hard}}]$.*

In practice, the cost function f may be piecewise linear, such as the taxation rate implemented under the cap-and-tax policy in countries like Canada [Government of Canada,

2025]. We argue that assuming f to be convex and differentiable provides a reasonable approximation of the piecewise linear, convex tax function while being technically more convenient.

Assumption 3 (TRUTHFUL IN DURATIONS). *Agents do not lie about their availability durations, i.e., each agent $n \in [N]$ is truthful in following their true \mathcal{T}_n (but may be strategic about the valuation v_n).*

Assumption 3 is common in the online mechanism design literature, also known as *limited misreport* [Parkes, 2007], indicating that the strategic behaviors may exist for the valuations, but not with respect to the durations.

3 Optimal Online Mechanism for CNP-ROA

In this section, we present our main results regarding the design of an optimal online mechanism for solving the CNP-ROA problem.

3.1 PPM- Φ : How It Works

We propose a posted-price mechanism in Algorithm 1, termed PPM- Φ , where Φ denotes the *marginal pricing* function, a key function that guides the execution of the online mechanism and also its competitive analysis.

Algorithm 1 Online Mechanism for CNP-ROA (PPM- Φ)

- 1: **Inputs:** $\mathcal{I}_{\text{known}}$ and Φ
- 2: **Initialization:** Capacity utilization $u^{(0)} = 0$
- 3: **while** a new agent n arrives **do**
- 4: Post pricing function $p^{(n)}$ for agent n :

$$p^{(n)}(z) = \int_{u^{(n-1)}}^{u^{(n-1)}+z} \Phi(s) ds. \quad (2)$$

- 5: Agent n decides its own profile $\hat{\mathbf{x}}_n$:

$$\hat{\mathbf{x}}_n = \arg \max_{\mathbf{x}_n \in \mathcal{X}_n} v_n(\mathbf{x}_n) - p^{(n)}\left(\sum_{t \in \mathcal{T}_n} w_t x_n^t\right). \quad (3)$$

- 6: Agent n makes the payment π_n :

$$\pi_n = p^{(n)}\left(\sum_{t \in \mathcal{T}_n} w_t \hat{x}_n^t\right). \quad (4)$$

- 7: Update the utilization: $u^{(n)} = u^{(n-1)} + \sum_{t \in \mathcal{T}_n} w_t \hat{x}_n^t$.
 - 8: **end while**
-

When an agent arrives, PPM- Φ posts the pricing function based on Eq. (2), and the agent determines their allocation profile by solving the utility maximization problem in Eq. (3). The marginal pricing function Φ must strike a balance between accepting agents too readily and delaying in anticipation of future better arrivals. For instance, if Φ is set too low (e.g., $\Phi(y) = f'(y)$), most agents will find the price inexpensive and consume as many resources as possible according to Eq. (3). Conversely, if Φ is set too high (e.g., $\Phi(y) = Uy$), most agents will find the price unaffordable and walk away

with zero consumption (i.e., $\hat{\mathbf{x}}_n = \mathbf{0}$). Thus, ensuring the effective performance of PPM- Φ hinges on designing an appropriate marginal pricing function Φ . In the next section, we show that there exists an optimal marginal pricing function Φ so that PPM- Φ achieves the best-possible competitive ratio among all online mechanisms.

3.2 Optimal Design of Φ

We begin by giving a lemma characterizing the major parameters in our design.

Lemma 1. Let $F(u|\alpha)$ be defined as:

$$F(u|\alpha) = \int_{s=C_{\text{soft}}}^u \frac{-\alpha}{C_{\text{hard}}} \exp\left(-\frac{\alpha s}{C_{\text{hard}}}\right) f'(s) ds. \quad (5)$$

There exists a unique pair of $\alpha \geq 1$ and $\beta \in [C_{\text{soft}}, C_{\text{hard}}]$ which satisfy the following system of equations:

$$\alpha = \ln(U) - \ln\left(F(C_{\text{hard}}|\alpha) - F(\max(\beta, C_{\text{soft}})|\alpha) + Le^{-\frac{\alpha\beta}{C_{\text{hard}}}}\right), \quad (6a)$$

$$L\beta - R(\beta) = \frac{1}{\alpha}(LC_{\text{hard}} - R(C_{\text{hard}})). \quad (6b)$$

Based on Lemma 1, we introduce our design of the marginal pricing function Φ in the theorem below.

Theorem 1 (CNP-ROA: CONVEX). Given $\mathcal{I}_{\text{known}}$ and under Assumptions 1-3, PPM- Φ is α -competitive when Φ is designed as follows:

- **(Low-Soft-Cap)** If $C_{\text{soft}} \leq \beta$:

$$\Phi(u) = \begin{cases} L & u \in [0, \beta], \\ \exp\left(\frac{\alpha u}{C_{\text{hard}}}\right) \cdot (F(u|\alpha) + \Lambda) & u \in [\beta, C_{\text{hard}}]. \end{cases}$$

- **(High-Soft-Cap)** If $C_{\text{soft}} > \beta$:

$$\Phi(u) = \begin{cases} L & u \in [0, \beta], \\ L \exp\left(\frac{\alpha}{C_{\text{hard}}}(u - \beta)\right) & u \in [\beta, C_{\text{soft}}], \\ \exp\left(\frac{\alpha u}{C_{\text{hard}}}\right) \cdot (F(u|\alpha) + \Lambda) & u \in [C_{\text{soft}}, C_{\text{hard}}]. \end{cases}$$

For both cases, $\Lambda = Le^{-\frac{\alpha\beta}{C_{\text{hard}}}} - F(\max(\beta, C_{\text{soft}})|\alpha)$; α and β are the solution to the system of equations in Eq. (6).

Here, we outline some key intuitions. Figure 1 illustrates the design of Φ for both the Low-Soft-Cap and High-Soft-Cap cases. In both scenarios, the marginal pricing function Φ begins with a flat phase, where $\Phi(u) = L$ for all $u \in [0, \beta]$, with L representing the minimum marginal valuation, as specified in Assumption 1. Beyond the threshold β , Φ transitions to an increasing phase for $u \in [\beta, C_{\text{hard}}]$. In the High-Soft-Cap case, where $C_{\text{soft}} \geq \beta$, the marginal pricing function continues to increase in a third phase for $u \in [C_{\text{soft}}, C_{\text{hard}}]$.

- **About the flat phase during $[0, \beta]$.** Lemma 1 demonstrates that for any given $\alpha \geq 1$, there exists a unique $\beta \in [C_{\text{soft}}, C_{\text{hard}}]$ satisfying Eq. (6b), which determines

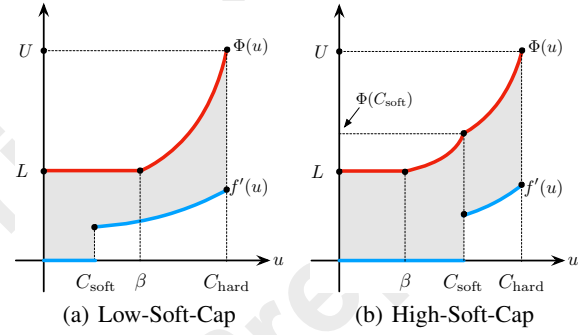


Figure 1: Illustration of the marginal pricing function Φ designed based on Theorem 1. Red curves denote the marginal pricing function $\Phi(u)$ and blue curves denote the marginal cost function $f'(u)$. The gray regions illustrate the gap between Φ and f' .

the threshold β such that the marginal price remains constant for all $u \in [0, \beta]$ (i.e., $\Phi(u) = L$), as established in Theorem 1. The rationale behind Eq. (6b) is to ensure that the first β -portion of the resource is always allocated, even to agents with the minimum marginal valuation L . This design guarantees that PPM- Φ remains α -competitive in the worst-case scenario, where all arriving agents have the lowest marginal valuation.

- **About the increasing phase during $[\beta, C_{\text{hard}}]$.** For both cases in Theorem 1, the marginal pricing function Φ enters the increasing phase after exceeding the threshold β , reflecting the fundamental principle that “scarcity of resources leads to higher prices.” Notably, the soft cap C_{soft} plays a critical role in shaping the design of the increasing phase. A lower C_{soft} implies that a single increasing phase is sufficient, whereas a higher C_{soft} necessitates a two-segment design for the increasing phase. This distinction highlights the key difference between our results and existing studies with single-phase regularizers, including the special case without a regularizer.

Theorem 2 (LOWER BOUND). Under Assumptions 1-3, PPM- Φ with the marginal pricing function Φ specified in Theorem 1 achieves the optimal competitive ratio.

Theorem 2 demonstrates that PPM- Φ is not only an optimal posted price mechanism but also optimal among all online algorithms, including those with randomization. This result indicates that the existence of strategic behaviors in CNP-ROA does not introduce additional difficulty in the online setting. In other words, focusing on posted price mechanisms does not result in any loss of generality when aiming to achieve the best possible competitive ratio among all online algorithms.

We provide the following remark to illustrate how Theorem 1 and Theorem 2 generalize existing state-of-the-art results for specific choices of the soft cap C_{soft} .

Remark 1 (CONNECTION TO EXISTING RESULTS). If $C_{\text{soft}} = 0$, the multi-phase regularizer simplifies to a single-phase penalty (i.e., $R(y) = f(y)$ for all $y \in [0, C_{\text{hard}}]$). In this case, Theorem 1 reproduces the optimal design proposed in [Tan et al., 2020c]. Conversely, if $C_{\text{soft}} = C_{\text{hard}}$, then

$R(y) = 0$ for all $y \in [0, C_{\text{hard}}]$. In this scenario, PPM- Φ , with the pricing function design specified in Theorem 1, reproduces the design proposed in [Sun et al., 2021], achieving the optimal competitive ratio:

$$\alpha = 1 + \ln \Delta,$$

where $\Delta = U/L$ denotes the fluctuation ratio of agents' marginal valuations.

3.3 Case Study: Linear Cost ($f(y) = \sigma y$)

To illustrate the design of the marginal pricing function Φ more clearly, we consider the special case where $f(y) = \sigma y$. This allows us to refine the results of Theorem 1 by providing a more explicit characterization of Φ and calculating key parameters, such as α and β .

Proposition 1 (CNP-ROA: LINEAR). Given $\mathcal{I}_{\text{known}}$ and under Assumptions 1-3, PPM- Φ is α -competitive when Φ is designed as follows:

- **(Low-Soft-Cap)** If $C_{\text{soft}} \leq \beta$:

$$\Phi(u) = \begin{cases} L & u \in [0, \beta), \\ (L - \sigma)e^{\frac{\alpha}{C_{\text{hard}}}(u-\beta)} + \sigma & u \in [\beta, C_{\text{hard}}], \end{cases}$$

where the competitive ratio α and β are given by

$$\alpha = \ln \left(\frac{U}{Le^{\left(\frac{\sigma(C_{\text{hard}} - C_{\text{soft}})}{LC_{\text{hard}}} - 1\right)} - \sigma(e^{-\frac{C_{\text{soft}}\alpha}{C_{\text{hard}}}} - e^{-\alpha})} \right),$$

$$\beta = \frac{C_{\text{hard}}}{\alpha} - \frac{\sigma C_{\text{soft}}(\alpha - 1)}{\alpha(L - \sigma)}.$$

- **(High-Soft-Cap)** If $C_{\text{soft}} > \beta$:

$$\Phi(u) = \begin{cases} L & u \in [0, \beta), \\ Le^{\frac{\alpha}{C_{\text{hard}}}(u-\beta)} & u \in [\beta, C_{\text{soft}}), \\ Ae^{\frac{\alpha}{C_{\text{hard}}}(u-C_{\text{soft}})} + \sigma & u \in [C_{\text{soft}}, C_{\text{hard}}], \end{cases}$$

where α , β , and A are given as follows:

$$\alpha = 1 + \frac{1}{1 + \frac{C_{\text{soft}}}{C_{\text{hard}}} \cdot \frac{\sigma}{L-\sigma}} \cdot \ln \left(\frac{U - \sigma}{L - \sigma} \right),$$

$$\beta = \frac{C_{\text{hard}}}{\alpha} - \frac{\sigma(C_{\text{hard}} - C_{\text{soft}})}{\alpha L},$$

$$A = L \exp \left(\frac{\alpha}{C_{\text{hard}}}(C_{\text{soft}} - \beta) \right) - \sigma.$$

Proposition 1 directly follows from Theorem 1 by substituting $f'(y) = \sigma$ into Eq. (6a), resulting in an analytical expression for $F(u|\alpha)$. According to Proposition 1, if $C_{\text{soft}} = C_{\text{hard}}$ or $\sigma = 0$, PPM- Φ is $(1 + \ln \Delta)$ -competitive, which is known to be optimal based on prior studies [Chakrabarty et al., 2008; Tan et al., 2020c; Sun et al., 2021].

4 Extensions to CNP-ROA with Resource Constraints

The mechanism PPM- Φ , designed in Section 3, solves CNP-ROA optimally in the absence of resource constraints for each time slot. In this section, we demonstrate that the design can be extended to address the more complex scenario where resource supply is limited at each step.

4.1 PPM- ϕ : A Modified PPM- Φ

Let c_t denote the resource constraint at time $t \in [T]$. Thus, the total allocation must satisfy the following constraint:

$$\sum_{n \in [N]} x_n^t \leq c_t \quad \forall t \in [T].$$

To extend PPM- Φ to account for the above resource constraints, a key step is to change the design of a single marginal pricing function Φ to a set of time-dependent marginal pricing functions, denoted as $\phi = \{\phi_t\}_{t \in [T]}$. We can then modify PPM- Φ based on ϕ as follows:

- The supplier posts the pricing functions $\{p_t^{(n)}\}_{t \in \mathcal{T}_n}$ to agent n upon its arrival:

$$p_t^{(n)}(z) = \int_{u_t^{(n-1)}}^{u_t^{(n-1)} + z} \phi_t(s) ds \quad \forall t \in \mathcal{T}_n. \quad (7)$$

- Agent n decides its consumption profile by solving the following utility-maximization problem

$$\hat{\mathbf{x}}_n = \arg \max_{\mathbf{x}_n \in \mathcal{X}_n} v_n(\mathbf{x}_n) - \sum_{t \in \mathcal{T}_n} p_t^{(n)}(x_n^t), \quad (8)$$

and makes the payment $\pi_n = \sum_{t \in \mathcal{T}_n} p_t^{(n)}(\hat{x}_n^t)$.

- The supplier updates the total utilization for each time slot $u_t^{(n)} = u_t^{(n-1)} + \hat{x}_n^t, \forall t \in \mathcal{T}_n$.

Hereinafter, we refer to PPM- Φ as PPM- ϕ to highlight the role of time-dependent marginal pricing functions ϕ . While the execution of PPM- ϕ is similar to PPM- Φ , its analysis requires a different approach due to the non-separability of the regularizer across time slots. To address this, we decompose the regularizer into multiple local regularizers, each tied to a specific time slot. Details of this decomposition are provided in the next subsection.

4.2 Core Idea: Decoupling the Regularizer

The CNP framework considers an aggregate soft cap for the entire allocation process across all time slots. To facilitate the design of a marginal pricing function for each time slot, Definition 1 introduces a method to decompose this global soft cap and assign an appropriate share, b_t , to each time slot, referred to as the local soft cap.

Definition 1 (LOCAL SOFT-CAP). Given C_{soft} as the global soft-cap, let b_t denote the local soft-cap at time t , defined as

$$b_t = \frac{c_t}{\sum_{t=1}^T w_t c_t} \cdot C_{\text{soft}}, \quad \forall t \in [T].$$

Intuitively, $\sum_t w_t b_t = C_{\text{soft}}$, indicating that b_t represents the normalized soft cap for each time slot, weighted by the consumption rate w_t . Using the local soft cap b_t , we define the local regularizer r_t to decouple the regularizer cost across time slots.

Definition 2 (LOCAL REGULARIZER). Let r_t denote the local regularizer at time t , defined as

$$r_t(y) = \begin{cases} 0 & \text{if } y \in [0, b_t), \\ (y - b_t) \cdot \frac{\partial}{\partial y} R\left(w_t y + \sum_{\tau \neq t} w_\tau c_\tau\right) & \text{if } y \in [b_t, c_t]. \end{cases}$$

For the local regularizer r_t defined in Definition 2, we compute the partial derivative of R with respect to the utilization level of the t -th time slot, while keeping the utilization levels of all other time slots fixed at their respective capacities to eliminate interdependencies. Note that the partial derivative always exists because $w_t y + \sum_{\tau \neq t} w_\tau c_\tau > C_{\text{soft}}$ always holds.

Based on Definition 2, we show that a similar group of design parameters, α and $\{\beta_t\}_{\forall t}$, still exist and the details are given in the lemma below.

Lemma 2. Let $F_t(u|\alpha)$ be defined as:

$$F_t(u|\alpha) = \int_{s=b_t}^u \frac{-\alpha}{c_t} \exp\left(-\frac{\alpha s}{c_t}\right) r'_t(s) ds. \quad (9)$$

There exists a unique set of $\alpha \geq 1$ and $\{\beta_t\}_{\forall t}$ that satisfy the following system of equations:

$$\alpha = \ln(U) - \min_{t \in [T]} \ln\left(F_t(c_t|\alpha) - F_t(\max(\beta_t, b_t)|\alpha) + Le^{-\frac{\alpha \beta_t}{c_t}}\right), \quad (10a)$$

$$L\beta_t - r_t(\beta_t) = \frac{1}{\alpha}(Lc_t + L\beta_t), \quad \forall t \in [T]. \quad (10b)$$

Lemma 2 extends Lemma 1 by introducing a time-dependent $\beta_t \in [0, c_t]$ for each slot $t \in [T]$. The next subsection demonstrates that the design of β_t is crucial for developing a time-dependent marginal pricing function ϕ_t for each time slot $t \in [T]$.

4.3 Design of $\{\phi_t\}_{\forall t}$

Building on Lemma 2, we extend Theorem 1 to design $\phi = \{\phi_t\}_{\forall t \in T}$ by modifying (F, β, Λ) to $(F_t, \beta_t, \Lambda_t)$. To better illustrate this approach, we consider the specific case where the cost function is linear, e.g., $f(y) = \sigma y$, and extend Proposition 1 to solve CNP-ROA under resource constraints. The full details are presented in Proposition 2 below.

Proposition 2 (CNP-ROA WITH RESOURCE CONSTRAINTS: LINEAR COSTS). Given $\mathcal{I}_{\text{known}}$ with resource capacity $\mathbf{c} = \{c_t\}_{\forall t}$ and under Assumptions 1-3, PPM- ϕ is α -competitive when ϕ_t is designed as follows:

- **(Low-Soft-Cap)** If $C_{\text{soft}} \leq \frac{1}{\alpha-1} \cdot C_{\text{hard}}$:

$$\phi_t(u) = \begin{cases} L & u \in [0, \beta_t], \\ (L - q_t)e^{\frac{\alpha}{c_t}(u-\beta_t)} + q_t & u \in [\beta_t, c_t], \end{cases}$$

where the competitive ratio α and β_t are given by

$$\alpha - \alpha \cdot \frac{L - \alpha \cdot q_{\max} \cdot \frac{C_{\text{soft}}}{C_{\text{hard}}}}{(\alpha - 1)L - \alpha q_{\max}} = \ln\left(\frac{U - q_{\max}}{L - q_{\max}}\right),$$

$$\beta_t = \frac{c_t L - \alpha q_t b_t}{(\alpha - 1)L - \alpha q_t}, \quad \forall t \in [T].$$

- **(High-Soft-Cap)** If $C_{\text{soft}} > \frac{1}{\alpha-1} \cdot C_{\text{hard}}$:

$$\phi_t(u) = \begin{cases} L & u \in [0, \beta_t], \\ Le^{\frac{\alpha}{c_t}(u-\beta_t)} & u \in [\beta_t, b_t], \\ A_t e^{\frac{\alpha}{c_t}(u-b_t)} + q_t & u \in [b_t, c_t], \end{cases}$$

where α , β_t , and A_t are given as follows:

$$\exp\left(\alpha - \frac{\alpha C_{\text{soft}}}{C_{\text{hard}}}\right) = \frac{U - q_{\max}}{L \exp\left(\frac{\alpha C_{\text{soft}}}{C_{\text{hard}}} - \frac{\alpha}{\alpha-1}\right) - q_{\max}},$$

$$\beta_t = \frac{c_t}{\alpha - 1}, \quad \forall t \in [T],$$

$$A_t = L \exp\left(\frac{\alpha}{c_t}(b_t - \beta_t)\right) - q_t, \quad \forall t \in [T].$$

In both cases, $q_t = \sigma w_t$ and $q_{\max} = \max_{t \in [T]} \sigma w_t$.

The following corollary shows that PPM- ϕ achieves a near-optimal competitive ratio in the High-Soft-Cap case.

Corollary 1 (LOWER BOUND). For the High-Soft-Cap case in Proposition 2, the competitive ratio of PPM- ϕ is $O(2 + \ln \hat{\Delta})$, where $\hat{\Delta} = \frac{U - q_{\max}}{L - q_{\max}}$, and the optimal competitive ratio of this problem is $\Omega(1 + \ln \hat{\Delta})$.

Corollary 1 demonstrates that when the soft-cap is sufficiently large, such that the forced flat segment concludes before the penalty cost impacts the pricing functions, the influence of the regularizer on the competitive ratio remains bounded.

Corollary 2 (CONNECTION TO EXISTING RESULTS). When $C_{\text{soft}} = C_{\text{hard}}$, namely, $R(y) = 0$ holds for all $y \in [0, C_{\text{hard}}]$, PPM- ϕ with the time-dependent marginal pricing functions specified by Proposition 2 is α -competitive with α given by

$$\alpha - 1 - \frac{1}{\alpha - 1} = \ln \Delta, \quad (11)$$

where $\Delta = U/L$ denotes the fluctuation ratio of agents' marginal valuations.

Notably, [Sun *et al.*, 2021] examines a similar online allocation problem without a regularizer, achieving the same competitive ratio as in Eq. (11). Here, we demonstrate that our design generalizes the results of [Sun *et al.*, 2021] as a special case when $C_{\text{soft}} = C_{\text{hard}}$.

5 Conclusion and Future Work

Motivated by real-world applications such as *cap-and-tax* carbon pricing, this paper introduces a novel online optimization problem, *cap-and-penalize* (CNP)-regularized online allocation (CNP-ROA). We derived an optimal online mechanism for CNP-ROA without per-step resource constraints and presented the best-known results for the more challenging setting with such constraints. To prove optimality, we introduced a novel representative function-based approach, reducing the lower-bound proof to solving an ordinary differential equation with boundary conditions.

This work opens several promising avenues for future research. Two key directions include addressing the gap in the CNP-ROA problem under resource constraints, which involves solving a multi-knapsack problem with multidimensional weights, and tackling the issue of temporal fairness. The latter challenge stems from the increasing marginal pricing function, which disproportionately penalizes later arrivals. Incorporating fairness constraints into this framework represents a significant theoretical and practical challenge.

References

- [Acun *et al.*, 2023] Bilge Acun, Benjamin Lee, Fiodar Kazhamiaka, Kiwan Maeng, Udit Gupta, Manoj Chakkaravarthy, David Brooks, and Carole-Jean Wu. Carbon explorer: A holistic framework for designing carbon aware datacenters. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2023.
- [Anderson *et al.*, 2023] Thomas Anderson, Adam Belay, Mosharaf Chowdhury, Asaf Cidon, and Irene Zhang. Treehouse: A case for carbon-aware datacenter software. *SIGENERGY Energy Inform. Rev.*, 2023.
- [Angelopoulos *et al.*, 2023] Spyros Angelopoulos, Shahin Kamali, and Kimia Shadkani. Online bin packing with predictions. *Journal of Artificial Intelligence Research*, 78:1111–1141, December 2023.
- [Awerbuch *et al.*, 2005] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. *STOC*, 2005.
- [Azar *et al.*, 2016] Yossi Azar, Niv Buchbinder, TH Hubert Chan, Shahar Chen, Ilan Reuven Cohen, Anupam Gupta, Zhiyi Huang, Ning Kang, Viswanath Nagarajan, Joseph Naor, et al. Online algorithms for covering and packing problems with convex objectives. *FOCS*, 2016.
- [Ball and Queyranne, 2009] Michael O Ball and Maurice Queyranne. Toward robust revenue management: Competitive analysis of online booking. *Operations Research*, 57(4):950–963, 2009.
- [Balseiro *et al.*, 2021] Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. Regularized online allocation problems: Fairness and beyond. *ICML*, 2021.
- [Blum *et al.*, 2003] Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. Online learning in online auctions. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '03, page 202–204, USA, 2003. Society for Industrial and Applied Mathematics.
- [Borodin and El-Yaniv, 1998] Allan Borodin and Ran El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.
- [Buchbinder and Naor, 2006] Niv Buchbinder and Joseph Naor. Improved bounds for online routing and packing via a primal-dual approach. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS '06)*, October 2006.
- [Buchbinder *et al.*, 2007] Niv Buchbinder, Kamal Jain, and Joseph (Seffi) Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. Springer Berlin Heidelberg, 2007.
- [Chakrabarty *et al.*, 2008] Deeparnab Chakrabarty, Yunhong Zhou, and Rajan Lukose. Online knapsack problems. In *WINE*, 2008.
- [Cheng *et al.*, 2022] Kai-Wen Cheng, Yuexin Bian, Yuanyuan Shi, and Yize Chen. Carbon-aware ev charging. In *IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, 2022.
- [Devanur and Jain, 2012] Nikhil R. Devanur and Kamal Jain. Online matching with concave returns. In *Proceedings of the 44th symposium on Theory of Computing - STOC '12*, page 137, New York, New York, USA, 2012. ACM Press.
- [Dickerson *et al.*, 2021] John P. Dickerson, Karthik A. Sankararaman, Aravind Srinivasan, and Pan Xu. Allocation problems in ride-sharing platforms: Online matching with offline reusable resources. *ACM Trans. Econ. Comput.*, 9(3), June 2021.
- [Einav *et al.*, 2018] Liran Einav, Chiara Farronato, Jonathan Levin, and Neel Sundaresan. Auctions versus posted prices in online markets. *Journal of Political Economy*, 2018.
- [El-Yaniv *et al.*, 2001] R. El-Yaniv, A. Fiat, R. M. Karp, and G. Turpin. Optimal search and one-way trading online algorithms. *Algorithmica*, 30(1):101–139, 2001.
- [Etzion *et al.*, 2006] Hila Etzion, Edieal Pinker, and Abraham Seidmann. Analyzing the simultaneous use of auctions and posted prices for online selling. *Manufacturing & Service Operations Management*, 8(1):68–91, 2006.
- [Feldman *et al.*, 2015] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, page 123–135, USA, 2015. Society for Industrial and Applied Mathematics.
- [Ferreira *et al.*, 2021] Matheus VX Ferreira, Daniel J Moroz, David C Parkes, and Mitchell Stern. Dynamic posted-price mechanisms for the blockchain transaction-fee market. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, 2021.
- [Gerding *et al.*, 2011] Enrico H. Gerding, Valentin Robu, Sebastian Stein, David C. Parkes, Alex Rogers, and Nicholas R. Jennings. Online mechanism design for electric vehicle charging. In *AAMAS*, 2011.
- [Government of Canada, 2025] Government of Canada. Putting a price on carbon pollution: Industry. <https://www.canada.ca/en/environment-climate-change/services/climate-change/pricing-pollution-how-it-will-work/putting-price-on-carbon-pollution/industry.html>, 2025. Accessed: 2025-01-21.
- [Guo *et al.*, 2017] Linqi Guo, Karl F Erliksson, and Steven H Low. Optimal online adaptive electric vehicle charging. In *IEEE Power & Energy Society General Meeting*, 2017.
- [Huang and Kim, 2015] Zhiyi Huang and Anthony Kim. Welfare maximization with production costs: A primal dual approach. Society for Industrial and Applied Mathematics, 2015.
- [Huang *et al.*, 2024] Zhiyi Huang, Zhihao Gavin Tang, and David Wajc. Online matching: A brief survey. *SIGecom Exch.*, 22(1):135–158, October 2024.

- [Jazi *et al.*, 2025] Hossein Nekouyan Jazi, Bo Sun, Raouf Boutaba, and Xiaoqi Tan. Posted price mechanisms for online allocation with diseconomies of scale. *WWW*, 2025.
- [Kellerer *et al.*, 2004] Hans Kellerer, Ulrich Pferschy, and David Pisinger. *Multiple Knapsack Problems*. Springer Berlin Heidelberg, 2004.
- [Lechowicz *et al.*, 2023] Adam Lechowicz, Nicolas Christianson, Jinhang Zuo, Noman Bashir, Mohammad Hajiesmaili, Adam Wierman, and Prashant Shenoy. The online pause and resume problem: Optimal algorithms and an application to carbon-aware load shifting. *ACM SIGMETRICS*, 2023.
- [Lee and Lee, 1985] C. C. Lee and D. T. Lee. A simple on-line bin-packing algorithm. *J. ACM*, 1985.
- [Lin *et al.*, 2019] Qiulin Lin, Hanling Yi, John Pang, Minghua Chen, Adam Wierman, Michael Honig, and Yuanzhang Xiao. Competitive online optimization under inventory constraints. *ACM SIGMETRICS*, 2019.
- [Lorenz *et al.*, 2009] Julian Lorenz, Konstantinos Panagiotou, and Angelika Steger. Optimal algorithms for k -search with application in option pricing. *Algorithmica*, 2009.
- [Ma and Simchi-Levi, 2020] Will Ma and David Simchi-Levi. Algorithms for online matching, assortment, and pricing with tight weight-dependent competitive ratios. *Operations Research*, 2020.
- [Ma *et al.*, 2024] Wanteng Ma, Ying Cao, Danny HK Tsang, and Dong Xia. Optimal regularized online allocation by adaptive re-solving. *Operations Research*, 2024.
- [Mehta *et al.*, 2007] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *Journal of the ACM (JACM)*, 2007.
- [Mehta, 2013] Aranyak Mehta. Online matching and ad allocation. *Found. Trends Theor. Comput. Sci.*, 8(4):265–368, October 2013.
- [Nisan *et al.*, 2007] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [Parkes, 2007] David C Parkes. Online mechanisms. *Cambridge University Press*, 2007.
- [Sun *et al.*, 2018] Bo Sun, Xiaoqi Tan, and Danny HK Tsang. Eliciting multi-dimensional flexibilities from electric vehicles: A mechanism design approach. *IEEE Transactions on Power Systems*, 2018.
- [Sun *et al.*, 2021] Bo Sun, Ali Zeynali, Tongxin Li, Mohammad Hajiesmaili, Adam Wierman, and Danny H.K. Tsang. Competitive algorithms for the online multiple knapsack problem with application to electric vehicle charging. *ACM SIGMETRICS*, 2021.
- [Tan *et al.*, 2020a] Xiaoqi Tan, Alberto Leon-Garcia, Yuan Wu, and Danny H.K. Tsang. Online combinatorial auctions for resource allocation with supply costs and capacity limits. *IEEE Journal on Selected Areas in Communications*, 2020.
- [Tan *et al.*, 2020b] Xiaoqi Tan, Alberto Leon-Garcia, Yuan Wu, and Danny H.K. Tsang. Posted-price retailing of transactive energy: An optimal online mechanism without prediction. *IEEE Journal on Selected Areas in Communications*, 38(1):5–16, 2020.
- [Tan *et al.*, 2020c] Xiaoqi Tan, Bo Sun, Alberto Leon-Garcia, Yuan Wu, and D.H.K Tsang. Mechanism design for online resource allocation: A unified approach. *ACM SIGMETRICS*, 2020.
- [Tan *et al.*, 2025] Xiaoqi Tan, Siyuan Yu, Raouf Boutaba, and Alberto Leon-Garcia. Threshold policies with tight guarantees for online selection with convex costs. *ACM Trans. Econ. Comput.*, 13(2), April 2025.
- [Zargari *et al.*, 2025] Faraz Zargari, Hossein Nekouyan, Bo Sun, and Xiaoqi Tan. Online allocation with multi-class arrivals: Group fairness vs individual welfare. *ACM SIGMETRICS*, 2025.