Mechanism Design for Online Resource Allocation: A Unified Approach

XIAOQI TAN, University of Toronto, Canada
BO SUN, HKUST, Hong Kong, China
ALBERTO LEON-GARCIA, University of Toronto, Canada
YUAN WU, University of Macau, Macau, China
DANNY H.K. TSANG, HKUST, Hong Kong, China

This paper concerns the mechanism design for online resource allocation in a strategic setting. In this setting, a single supplier allocates capacity-limited resources to requests that arrive in a sequential and arbitrary manner. Each request is associated with an agent who may act selfishly to misreport the requirement and valuation of her request. The supplier charges payment from agents whose requests are satisfied, but incurs a load-dependent supply cost. The goal is to design an incentive compatible online mechanism, which determines not only the resource allocation of each request, but also the payment of each agent, so as to (approximately) maximize the social welfare (i.e., aggregate valuations minus supply cost). We study this problem under the framework of competitive analysis. The major contribution of this paper is the development of a unified approach that achieves the best-possible competitive ratios for setups with different supply costs. Specifically, we show that when there is no supply cost or the supply cost function is linear, our model is essentially a standard 0-1 knapsack problem, for which our approach achieves logarithmic competitive ratios that match the state-of-the-art (which is optimal). For the more challenging setup when the supply cost is strictly-convex, we provide online mechanisms, for the first time, that lead to the optimal competitive ratios as well. To the best of our knowledge, this is the first approach that unifies the characterization of optimal competitive ratios in online resource allocation for different setups including zero, linear and strictly-convex supply costs.

CCS Concepts: • **Theory of computation → Algorithmic mechanism design**; **Online algorithms**; • **Networks → Network economics**.

Additional Key Words and Phrases: Mechanism Design, Online Algorithms, Pricing, Resource Allocation

1 INTRODUCTION

We study the mechanism design for online resource allocation problems. A single supplier who allocates capacity-limited resources (e.g., computing cycles, network bandwidth, energy, etc.) to

Authors' addresses: Xiaoqi Tan, University of Toronto, 27 King's College Cir, Toronto, ON, Canada, xiaoqi.tan@utoronto.ca; Bo Sun, HKUST, Clear Water Bay, Kowloon, Hong Kong, China, bsunaa@connect.ust.hk; Alberto Leon-Garcia, University of Toronto, 27 King's College Cir, Toronto, ON, Canada, alberto.leongarcia@utoronto.ca; Yuan Wu, University of Macau, Avenida da Universidade, Taipa, Macau, China; Danny H.K. Tsang, HKUST, Clear Water Bay, Kowloon, Hong Kong, China, eetsang@ust.hk.

Proc. ACM Meas. Anal. Comput. Syst., Vol. 4, No. 2, Article 24. Publication date: June 2020.

24

requests that arrive in a sequential and arbitrary manner. We consider a strategic setting where each request is owned by a self-interested agent who may deliberately misreport the resource requirement and value of the request. A request is satisfied if the required resource is allocated to the corresponding agent. The supplier charges payment from agents whose requests are satisfied, and affords a supply cost which is a function of the total resource allocated. The goal is to design not only the resource allocation of each request, but also the payment of each agent, so that agents are well-incentivized to follow their true preferences (i.e., incentive compatible [34]) and meanwhile, the social welfare (i.e., the aggregate value minus the supply cost) can be approximately maximized.

A prominent application of this model is the market-based resource allocation in cloud computing [26]. Here, the resource may represent computing cycles that can run at different speeds. Cloud service providers charge money from customers who purchase their services[1], but must pay a considerable amount of power and cooling cost to maintain their data centers. Moreover, such operational costs are often load-dependent and is usually an increasing function of the total resource allocated (e.g., CPU). Another application of the investigated model may arise in the context of network routing with congestion cost [4, 15, 16]. Each incoming user wants to own some pair of connections with a valuation if any feasible connection is established. Congestion cost occurs when a link is occupied by multiple users, and the cost can be modelled as a traffic-dependent increasing function [2, 15]. The target is to maximize the total valuation of routed connections minus the congestion cost. In reality, such congestion costs can either reflect the total energy needed to support the network, or simply a virtual cost to penalize the degradation of the quality-of-service (e.g., latency [4, 16]). The authors of [2] indicated that the power-rate curve for networking or computing devices (e.g., CPU, communication links, and edge routers, etc) exhibit *diseconomy-of-scale* properties (i.e., a convex curve), and thus polynomial functions are commonly used to model the power requirements of the network [15].

Motivated by the above applications, this paper primarily focuses on setups when the supply cost function is strictly-convex and differentiable, i.e., the marginal cost is strictly increasing. A basic challenge for online resource allocation in such setups is as follows: if the limited resource is allocated too aggressively, then an excessive portion of the resource may be allocated to earlier requests with low values. This will increase the total cost for the supplier and thus increase the payment for future agents, which will consequently prevent the later agents from purchasing the resources even if their valuations are higher than the earlier ones. On the other hand, if the payment is set too high, a majority of the requests may be declined, leading to a poor performance as well. This paper aims to tackle this challenge by designing an online mechanism that leads to the best-possible performance in social welfare. More importantly, our optimal results for strictly-convex costs extend to setups with no supply cost (i.e., zero cost) or linear supply cost, leading to a unified approach for online resource allocation with zero, linear or strictly-convex supply costs.

1.1 Related Results

Different variants of online resource allocation problems have been studied in both the non-strategic setting, e.g., [5, 13, 19, 27, 38, 42], and the strategic setting, e.g., [7, 8, 11, 25, 39].

In the non-strategic setting, the main focus is to design online algorithms with a competitive fraction of optimum. Classic problems in this stream of studies include online bipartite matching [27], Adwords problems [19, 31, 32], online covering and packing problems [12, 13], online knapsack

---

[1]In practice, major cloud service providers such as Amazon Web Services, Microsoft Azure, and Google Cloud often use fixed pricing schemes. However, dynamic pricing is also adopted under some circumstances. For example, customers of Amazon EC2 Spot Instances (https://aws.amazon.com/ec2/spot/) are charged based on spot prices that are adjusted in real-time. Typically, the spot prices are determined based on multiple factors such as long-term trend in demand and supply of spot instance capacity.

problems [41, 42], and one-way trading problems [23]. For example, the authors of [32] introduced the Adwords problem and characterized the optimal competitive ratio by assuming each bid is much smaller than the total budget (i.e., the infinitesimal assumption). Similarly, the authors of [42] proved that, when the weight of each item is much smaller than the capacity of the knapsack, and the value-to-weight ratio of each item is bounded within $[L, U]$, no online algorithm can achieve a competitive ratio tighter than $1 + \ln(U/L)$. In addition to the above classic problems, new variants of online resource allocation problems have also been reported, e.g., [5, 21, 28, 29]. For example, a generalization of the online matching and Adwords problem was studied in [21], where each incoming agent has a concave function representing the return or utility of this agent, in contrast to the budgeted linear utility function studied by almost all previous online matching problems. In particular, the authors of [21] derived a differential equation based on the primal-dual analysis, and then characterized the optimal competitive ratio by analyzing the differential equation via variations of calculus. Lin et al. [28] recently proposed an optimal online algorithm for a generalized one-way trading problem, where the online decisions are made in each round to maximize the per-round concave revenue function under inventory (budget) constraints. Recently, an extended online matching problem was presented by [29], in which the items can be sold at multiple feasible prices instead of a single price. Based on this setup, Ma et al. [29] designed an online algorithm with the best-possible competitive ratios that depend on the sets of multiple prices given in advance.

In the strategic setting, agents are self-interested and thus a careful payment rule must be designed along with the online allocation algorithms so as to guarantee the incentive compatibility. A classic problem in this setting is the online combinatorial auction (CA) problem, e.g., [6–8, 11, 25, 34]. Existing studies studying the online CAs focus on two cases: one is with stringent supply constraints (i.e., limited supply) and the other is unlimited supply, namely, the supplier can produce as many copies of items as possible with no cost. For example, the authors of [7] studied an online CA problem and proposed an $O(\log(v_{max}/v_{min}))$-competitive online algorithm when there are $\Omega(\log(v_{max}/v_{min}))$ copies of each item, where the customers' valuation is assumed to be in the range of $[v_{min}, v_{max}]$. Considering the shortcomings of both the limited-supply case and the unlimited-supply case in modelling real-world problems, Blum et al. [8] studied online CAs in a more general case with increasing production cost. In this case, the supplier can produce additional number of items while paying an increasing marginal cost per copy[2]. Blum et al. proposed a pricing scheme called *twice-the-index*, and gave constant competitive ratios for simple marginal cost functions such as linear and lower-degree polynomial. Huang et al. [25] later studied a similar problem and achieved a tighter competitive ratio. In particular, for polynomial cost functions $f(y) = y^s$, Huang et al. [25] proved that an optimal $s^{\frac{s}{s-1}}$-competitive online mechanism can be designed when each agent buys infinitesimally small units of items. However, the optimal competitive ratio in [25] was obtained without considering the capacity limits, and thus the results cannot capture how the limited supply affects the optimal online mechanisms. Tan et al. [39] later studied this problem with both polynomial supply costs $f(y) = y^s$ and stringent supply constraints. When the valuation-to-demand ratio of agents (i.e., valuation divided by resource demand) are upper bounded by $\bar{p}$, Tan et al. [39] proved that the optimal competitive ratio $s^{\frac{s}{s-1}}$ which is firstly characterized by [25] becomes non-achievable if $\bar{p}$ exceeds a certain threshold. Moreover, this threshold depends on the supply capacity. In this regard, the results in [25, 39] show that both the supply costs and the supply constraints can influence the optimal competitive ratio.

---

[2]The limited-supply case can be considered a special case of the model in [8] with a 0-to-∞ step function representing the marginal supply cost.

1.2 Our Contributions and Techniques

In this paper, we consider a *significant* generalization of the online resource allocation problem by allowing *arbitrary* strictly-convex supply costs and stringent supply constraints. In this setting, additional units of resources can be supplied at strictly-increasing marginal costs, and the total resource allocated must stay within the capacity limit. To the best of our knowledge, characterizing the optimal competitive ratios in such settings is an open question. This paper fills this gap and proposes an optimal online mechanism with the best-possible competitive ratios. Meanwhile, for some special cases when the supply costs are not strictly-convex, e.g., zero and linear supply cost, we prove that our design still holds (with some minor modifications) and achieves logarithmic competitive ratios which have been proven as optimal by existing literature. Therefore, our proposed approach unifies the characterization of optimal competitive ratios in online resource allocation for different setups including zero, linear and strictly-convex supply costs.

Close to our work is the design of online primal-dual algorithms, e.g., [5, 12, 13, 20, 21, 24, 25, 29] (in particular, see [14] for a comprehensive survey). The key is to construct a dual feasible objective that is close to the primal feasible objective at each round when there is a new arrival of agent, and then use weak duality to bound the performance of the online algorithm. In this paper, the construction of feasible primal and dual solutions is based on a universal monotone "pricing function", denoted by ϕ, which is designed by solving an ordinary differential equation (ODE) [3]. Unlike techniques adopted by [21, 25, 29], our derived ODE contains two boundary conditions imposed by the supply constraints, and thus is essentially a two-point boundary value problem (BVP) [1, 3, 35]. We prove that the existence of any α-competitive algorithm is equivalent to the existence of a strictly-increasing solution to the BVP, for which the uniqueness and existence conditions are rigorously proved. To the best of our knowledge, this is the first time that the design and analysis of online algorithms are based on the existence and uniqueness of solutions to BVPs parameterized by the competitive ratios. We believe that our proposed design principle sheds light on new directions for solving other online optimization problems.

2 PROBLEM SETUP

In this section, we present the problem setup with the resource allocation model, and then introduce the assumptions and definitions that will be used in our online mechanism design.

2.1 The Model

We consider a single supplier who allocates a single type of resource to a set of requests $\mathcal{N} = \{1, \cdots, N\}$. Each request $n \in \mathcal{N}$ is owned by an agent (which is also indexed by n for simplicity), and can be represented by a private type $\theta_n = (r_n, v_n)$, where r_n represents the resource requirement and v_n denotes the valuation of agent n if the request is satisfied. In practice, we can interpret v_n as the *willingness-to-pay* of agent n for obtaining the required resource.

Let us define a binary variable $x_n = \{0, 1\}$ for each request $n \in \mathcal{N}$, and assume $x_n = 1$ if request n is satisfied and $x_n = 0$ otherwise. We assume that agents have quasi-linear utilities [34], i.e., the utility of agent n is denoted by $U_n = v_n x_n - \pi_n$, where π_n denotes the payment made by agent n. Intuitively, agents whose requests are not satisfied make zero payment, i.e., $\pi_n = 0$ if $x_n = 0$. The supplier collects payment π_n from agent n, and pays a total supply cost of $f(\sum_n r_n x_n)$, where $\sum_n r_n x_n$ denotes the total resource allocated and f represents the supply cost function. Therefore, the utility of the supplier (i.e., the profit) can be denoted by $U_s = \sum_{n \in \mathcal{N}} \pi_n - f(\sum_n r_n x_n)$. We consider limited supply with a stringent capacity limit. Without loss of generality, we normalize the capacity limit to be 1, and thus $\sum_n r_n x_n \leq 1$. Meanwhile, the requirements $\{r_n\}_{\forall n}$ denote the

proportions of the normalized capacity limit accordingly. Our mechanism design relies on some properties of the cost function f, which is discussed later in Section 2.2.

We focus on social welfare maximization in an online setting, where agents arrive one-by-one in a sequential and arbitrary manner, breaking ties arbitrarily. We denote the sequence of agent arrivals by $\mathcal{A} \triangleq (\theta_1, \theta_2, \cdots, \theta_N)$, where we assume without loss of generality that the arrival times are in ascending order. In the following we refer to $\mathcal{A}$ as the arrival instance. If we assume a complete knowledge of $\mathcal{A}$, then the social welfare maximization in the offline setting can be written as follows:

$$\underset{\{x_n\}_{\forall n}}{\text{maximize}} \qquad \sum_{n \in \mathcal{N}} v_n x_n - f\left(\sum_{n \in \mathcal{N}} r_n x_n\right) \tag{1a}$$

$$\text{subject to} \qquad \sum_{n \in \mathcal{N}} r_n x_n \leq 1, \tag{1b}$$

$$x_n = \{0, 1\}, \forall n \in \mathcal{N}. \tag{1c}$$

In Problem (1), the objective is the aggregate utilities of all the agents and the supplier, where the payment terms cancel out.

REMARK 1. *If there is no supply cost, i.e., $f = 0$, then Problem (1) is a standard 0-1 knapsack problem [42]. Meanwhile, if we consider allocation of bundles of resources with a supply cost function for each type of resource, Problem (1) can be extended to model the standard online combinatorial auction problem [8, 34]. Therefore, the studied model is a generalization of a variety of classic online resource allocation problems. In this paper, we mainly focus on the basic model presented by Problem (1). For discussions of extending our model to consider multiple types of resources, as well as generalizing Problem (1) to model online resource allocation with multiple time slots (e.g., job scheduling in cloud computing), please refer to Section 5.3.2.*

2.2 Assumptions

We make the following assumptions throughout the paper.

ASSUMPTION 1. *For each request $n \in \mathcal{N}$, the valuation density, defined as v_n/r_n, is lower and upper bounded as follows:*

$$\underline{p} \leq \frac{v_n}{r_n} \leq \bar{p}, \forall n \in \mathcal{N}. \tag{2}$$

Assumption 1 is similar to the one in [7, 29, 42]. Specifically, one can interpret $\underline{p}$ as the lowest selling price per unit of resource, which is set by the supplier in advance. Thus, the request of any agent with a valuation density lower than $\underline{p}$ will automatically be rejected. In comparison, the upper bound $\bar{p}$ can be considered an inherent attribute of the arrival instance with rational agents, and in reality it exists naturally.

ASSUMPTION 2. *The resource requirement of each request is very small compared to the total capacity limit, i.e., $r_n \ll 1, \forall n \in \mathcal{N}$.*

Assumption 2 is common in designing online algorithms, e.g., [21, 25, 32, 42], since it allows us to focus on the infinitesimal nature of our problem with mathematical convenience. Meanwhile, Assumption 2 is reasonable in many real-world large-scale systems.

DEFINITION 1 (INFORMATION SETUP). *We define all the information known to the supplier a priori as a setup $\mathcal{S}$, denoted by*

$$\mathcal{S} \triangleq \{f, \underline{p}, \bar{p}\}. \tag{3}$$

In the following, we say $\mathcal{S}$ is a *nice setup* if f is monotonically non-decreasing with zero startup cost (i.e, $f(0) = 0$) and satisfies $f'(0) < \underline{p} \leq \bar{p}$. Meanwhile, we refer to a nice setup $\mathcal{S}$ as a *convex setup* if f is also *strictly-convex and differentiable*. Given a nice setup $\mathcal{S}$, the supplier only knows the information given by $\mathcal{S}$, and does not know anything about the arrival instance $\mathcal{A}$ such as the total number of requests (i.e., N).

2.3 Definitions

Below we give some definitions regarding incentive compatibility, online mechanisms, and competitive ratios.

(**Incentive Compatibility**) We consider the design of direct revelation mechanisms under a strategic setting, where each agent $n \in \mathcal{N}$ participates by declaring the type $\hat{\theta}_n = (\hat{r}_n, \hat{v}_n)$ at time $\hat{a}_n$. Since agents are self-interested, the reported preference $\hat{\theta}_n$ of agent $n \in \mathcal{N}$ may or may not be her true preference θ_n. A mechanism is *incentive compatible (IC)* if all agents achieve the best outcomes (i.e., the maximum utilities) by acting according to their true preferences.

(**Online Mechanism**) Given a convex setup $\mathcal{S}$, we aim to design an IC online mechanism to (approximately) maximize the social welfare. In the online setting, agents arrive one-by-one in a sequential order (where ties are broken arbitrarily). At each round when there is a new arrival of agent $n \in \mathcal{N}$, the supplier needs to make two types of irrevocable decisions as follows: i) the *allocation rule*, namely the decisions about whether request n should be satisfied or not (i.e., x_n); and ii) the *payment rule*, namely, how much money should be charged from agent n if the request is satisfied. The resource allocation rule and the payment rule constitute an *online mechanism*. In the strategic setting, agents may deliberately misreport their type information to better off, and thus the supplier needs to carefully design the payment rule such that agents are well-incentivized to follow their true preferences[3].

(**Competitive Ratio**) The performance of an online mechanism can be quantified by the standard competitive analysis framework [9]. Given an arrival instance $\mathcal{A}$, let us denote the optimal social welfare by $S_{\text{offline}}(\mathcal{A})$, which is the optimal objective value of Problem (1) in the offline setting. Let $S_{\text{online}}(\mathcal{A})$ denote the social welfare achieved by an online mechanism based on the knowledge of $\mathcal{S}$ only. The competitive ratio of an online mechanism is defined as

$$\alpha \triangleq \max_{\text{all possible } \mathcal{A}} \frac{S_{\text{offline}}(\mathcal{A})}{S_{\text{online}}(\mathcal{A})}, \tag{4}$$

where $\alpha \geq 1$ and the closer to 1 the better. Note that in the online setting the supplier does not known any future information except $\mathcal{S}$, and thus the competitive ratio α should depend on $\mathcal{S}$ only.

Our target is to design an IC online mechanism such that $S_{\text{online}}(\mathcal{A})$ is as close to $S_{\text{offline}}(\mathcal{A})$ as possible for all possible $\mathcal{A}$'s, namely, an online mechanism with an competitive ratio that is as small as possible. To define how small the competitive ratio could possible be, we give the following definition regarding the optimal competitive ratio for a given nice setup $\mathcal{S}$.

DEFINITION 2 (OPTIMAL COMPETITIVE RATIO). *A competitive ratio is optimal if no other online algorithms can achieve a smaller one under Assumption 1 and Assumption 2. In particular, given a nice setup $\mathcal{S}$, the optimal competitive ratio $\alpha_*(\mathcal{S})$ is defined as*

$$\alpha_*(\mathcal{S}) \triangleq \inf \max_{\text{all possible } \mathcal{A}} \frac{S_{\text{offline}}(\mathcal{A})}{S_{\text{online}}(\mathcal{A})}, \tag{5}$$

where the inf *operator is taken w.r.t. all possible online algorithms.*

---

[3]In a non-strategic environment, only an allocation rule is needed, and thus in this case an online mechanism is simply an online algorithm.

3 ONLINE MECHANISM DESIGN

This section presents our proposed online mechanism. We start by giving some preliminaries and notations regarding the primal and dual of Problem (1), and then describe the design principle and design procedures of our online mechanism.

Table 1. Notations

Symbol	Description
v_n	value of agent n
x_n	binary decision variable of agent n
r_n	resource requirement of agent n
f	supply cost function
$\underline{p}$ ($\bar{p}$)	lower (upper) bound of v_n/r_n
$\underline{c}$ ($\bar{c}$)	minimum (maximum) marginal cost
$\underline{\rho}$ ($\bar{\rho}$)	maximum resource utilization when $\frac{v_n}{r_n} = \underline{p}$ ($\bar{p}$), $\forall n$
$\phi(y)$	pricing function with three segments
$\varphi(y)$	increasing segment of $\phi(y)$ when $y \in [\omega, \bar{\rho}]$
ω	resource utilization threshold so that $\phi(\omega) = \underline{c}$
u	resource utilization threshold so that $\phi(u) = \bar{c}$
$\hat{\pi}_n$	final payment made by agent n
$\hat{p}_n$	posted price for agent $n + 1$
$\hat{y}_n$	total resource utilization after agent n

3.1 Preliminaries

We first introduce some notations to help derive the dual of Problem (1). Given a convex setup $\mathcal{S}$, we define $\underline{c}$ and $\bar{c}$ as follows:

$$\underline{c} \triangleq f'(0), \bar{c} \triangleq f'(1), \tag{6}$$

where $\underline{c}$ and $\bar{c}$ denote the minimum and maximum marginal cost, respectively. Based on the capacity limit, we define the *extended cost function* $\bar{f}(y)$ as follows:

$$\bar{f}(y) = \begin{cases} f(y) & \text{if } y \in [0, 1], \\ +\infty & \text{if } y \in (1, +\infty), \end{cases} \tag{7}$$

which extends the domain of $f(y)$ to all $y \in [0, +\infty)$. Based on the extended cost function $\bar{f}$, we define $F_p(y)$ as follows:

$$F_p(y) \triangleq py - \bar{f}(y), y \in [0, +\infty), \tag{8}$$

which represents the profit of selling y units of resources at price p, where py represents the revenue and $\bar{f}(y)$ represents the supply cost. In the following, $F_p(y)$ is referred to as the *profit function*.

(**Relaxed Primal Problem**) Our online mechanism is designed based on a principled primal-dual analysis of Problem (1). We consider the following relaxed welfare maximization problem

$$\underset{\boldsymbol{x}, y}{\text{maximize}} \qquad \sum_{n \in \mathcal{N}} v_n x_n - \bar{f}(y) \tag{9a}$$

$$\text{subject to} \qquad \sum_{n \in \mathcal{N}} r_n x_n \leq y, \qquad\qquad (p) \tag{9b}$$

$$x_n \leq 1, \forall n \in \mathcal{N}, \qquad (\gamma_n) \tag{9c}$$

$$\boldsymbol{x} \geq \boldsymbol{0}, y \geq 0, \tag{9d}$$

where $\boldsymbol{x} = \{x_n\}_{\forall n}$. In Problem (9), (p, γ_n) are the Lagrange multipliers associated with the corresponding constraint.

When f is non-decreasing, Problem (9) is equivalent to Problem (1) except the relaxation of $\{x_n\}_{\forall n}$. The reason is as follows. First, compared to the original cost function f with the capacity limit constraint (1b), the introduction of $\bar{f}$ in Problem (9) is an equivalent transformation. Second, the relaxation of the equality constraint (1b) to an inequality one in Eq. (9b) is lossless since the inequality will always be binding when f is non-decreasing.

(**Dual Problem**) Problem (9) is a continuous optimization problem, whose dual can be expressed as follows:

$$\underset{p, \boldsymbol{\gamma}}{\text{minimize}} \qquad \sum_{n \in \mathcal{N}} \gamma_n + h(p) \tag{10a}$$

$$\text{subject to} \qquad \gamma_n \geq v_n - pr_n, \forall n \in \mathcal{N}, \tag{10b}$$

$$p \geq 0, \boldsymbol{\gamma} \geq \boldsymbol{0}, \tag{10c}$$

where $\boldsymbol{\gamma} = \{\gamma_n\}_{\forall n}$. We interpret p as the price per unit of resource, and γ_n as the utility of agent n. In Eq. (10a), $h(p)$ is given by

$$h(p) = \max_{y \geq 0} \ py - \bar{f}(y) = \max_{y \geq 0} \ F_p(y), \tag{11}$$

which can be written as follows:

$$h(p) = \begin{cases} F_p\left(f'^{-1}(p)\right) & \text{if } p \in [\underline{c}, \bar{c}], \\ F_p(1) & \text{if } p \in (\bar{c}, +\infty). \end{cases} \tag{12}$$

The above function $h(p)$ is known as the convex conjugate of $\bar{f}$ [10]. An economic interpretation of Eq. (11) is that $h(p)$ represents the optimal profit when the selling price is p, where py represents the revenue and $\bar{f}(y)$ represents the supply cost. In this regard, the dual objective in Eq. (10a) is the aggregate utilities of all the agents plus the utility of the supplier (i.e., the profit). In comparison, the primal objective in Eq. (9) is the aggregate values of all the satisfied requests minus the supply cost. Both of them represent the social welfare of the system.

Given an arrival instance $\mathcal{A}$, let us denote the optimal primal objective of Problem (9) by $P_*(\mathcal{A})$. Similarly, denote by $D_*(\mathcal{A})$ the optimal dual objective of Problem (10). Then, we have

$$S_{\text{offline}}(\mathcal{A}) \leq P_*(\mathcal{A}) \leq D_*(\mathcal{A}), \tag{13}$$

where the first inequality is due to the relaxation of $\{x_n\}_{\forall n}$ and the second one is because of weak duality[4].

---

[4]Note that when f is convex, strong duality holds here, i.e., $P_*(\mathcal{A}) = D_*(\mathcal{A})$. However, throughout the paper we do not need the strong duality.

3.2 Design Principles

In the online setting, we need to process agents in a sequential manner. For a given arrival instance $\mathcal{A}$, after processing the n-th request, we denote the objective values of Problem (9) and Problem (10) by $P_n(\mathcal{A})$ and $D_n(\mathcal{A})$, respectively. For simplicity, we drop the parenthesis and simply write P_n and D_n hereinafter. Meanwhile, we say that P_n (D_n) is feasible if P_n (D_n) is a feasible objective value of Problem (1) (Problem (10)) after processing the n-th request. Proposition 3.1 below shows that if the sequences of $\{P_n\}_{\forall n}$ and $\{D_n\}_{\forall n}$ are feasible and satisfy a group of inequalities parameterized by α, then the corresponding online algorithm is guaranteed to be α-competitive.

PROPOSITION 3.1. *An online algorithm is α-competitive if the following two conditions are satisfied, namely, i) the sequences of $\{P_n\}_{\forall n}$ and $\{D_n\}_{\forall n}$ are feasible, and ii) there exists an index $k \in \mathcal{N}$ such that the following inequalities hold*

$$\begin{cases} P_k \geq \frac{1}{\alpha} D_k, \\ P_n - P_{n-1} \geq \frac{1}{\alpha} \left(D_n - D_{n-1} \right), \forall n \in \{k+1, \cdots, N\}. \end{cases} \tag{14}$$

PROOF. The feasibility of $\{P_n\}_{\forall n}$ is trivial since any online algorithm must first produce a feasible solution to the original problem in Eq. (1). We next show how the feasibility of $\{D_n\}_{\forall n}$ and the inequalities in Eq. (14) lead to an α-competitive online algorithm. Note that it suffices to prove $P_N \geq \frac{1}{\alpha} D_N$ since

$$S_{\text{online}} = P_N \geq \frac{1}{\alpha} D_N \overset{(i)}{\geq} \frac{1}{\alpha} D_* \overset{(ii)}{\geq} \frac{1}{\alpha} P_* \overset{(iii)}{\geq} \frac{1}{\alpha} S_{\text{offline}}, \tag{15}$$

where inequalities (ii) and (iii) directly follow Eq. (13), and inequality (i) holds since D_N is a feasible objective value of Problem (10), while D_* is the optimum (minimum).

Suppose there exists an index $k \in \mathcal{N}$ such that the second inequality in Eq. (14) hold for all $n \in \{k+1, \cdots, N\}$, then

$$P_N - P_k = \sum_{n=k+1}^{N} (P_n - P_{n-1}) \geq \frac{1}{\alpha} \sum_{n=k+1}^{N} (D_n - D_{n-1}) = \frac{1}{\alpha}(D_N - D_k),$$

which leads to $P_N \geq \frac{1}{\alpha} D_N$ after substituting $P_k \geq \frac{1}{\alpha} D_k$ into the above equation. We thus complete the proof. □

We refer to the first inequality in Eq. (14) as the *initial inequality*, and the second one in Eq. (14) as the *incremental inequality*. Note that when $P_0 = D_0 = 0$ and $k = 1$, the two types of inequalities can be combined. Proposition 3.1 in this case simply follows the standard online primal-dual approach [14]. In this regard, Proposition 3.1 is a more general principle for designing online algorithms.

3.3 Posted Price Mechanism

Our online mechanism is designed based on Proposition 3.1. The key idea is to construct a set of feasible primal and dual solutions at each round when there is a new arrival of request $n \in \mathcal{N}$, and then guarantee that the resulting sequences of $\{P_n\}_{\forall n}$ and $\{D_n\}_{\forall n}$ satisfy the inequalities in Eq. (14).

(**A Two-Step Design**) Note that the dual variables $\{\gamma_n\}_{\forall n}$ in Problem (10) are associated with each individual agent, while p is a global one couples all the agents. Once the final price p is known to the supplier, we can easily decouple different agents and design feasible primal and dual solutions for each individual request $n \in \mathcal{N}$. However, when there is no future information, it is impossible to know the exact value of p a priori. Our idea is to adopt a two-step design procedures as follows:

- **Step-1**: design a trajectory of $\{\hat{p}_n\}_{\forall n}$, where $\hat{p}_n$ denotes the supplier's prediction of the final price p after processing agent n.
- **Step-2**: when there is a new arrival of request $n \in \mathcal{N}$, based on $\hat{p}_{n-1}$, perform the following decision-making:
  - Set the dual variable $\hat{\gamma}_n$ by

$$\hat{\gamma}_n = \max\left\{v_n - \hat{p}_{n-1}r_n, 0\right\}, \forall n \in \mathcal{N}. \tag{16}$$

  - Set the primal variable $\hat{x}_n$ by:

$$\hat{x}_n = \begin{cases} 0 & \text{if } v_n - \hat{p}_{n-1}r_n < 0, \\ 1 & \text{if } v_n - \hat{p}_{n-1}r_n \geq 0 \text{ and } \hat{y}_{n-1} + r_n \leq 1, \end{cases}$$

    where $\hat{y}_{n-1}$ denotes the total resource utilization after processing agent $n - 1$. Intuitively, we have $\hat{y}_0 = 0$.
  - Set the final payment $\hat{\pi}_n$ for agent n by

$$\hat{\pi}_n = \hat{p}_{n-1}r_n\hat{x}_n, \tag{17}$$

  - Update the total resource utilization by

$$\hat{y}_n = \hat{y}_{n-1} + r_n\hat{x}_n. \tag{18}$$

Based on the above two-step design procedure, the terminal value of the total resource utilization is $\hat{y}_N$, and the terminal value of the predicted final price is $\hat{p}_N$. Together with $\{\hat{x}_n\}_{\forall n}$ and $\{\hat{\gamma}_n\}_{\forall n}$, these variables constitute a complete set of online primal and dual solutions, which are denoted by $\mathcal{V}_P$ and $\mathcal{V}_D$ as follows:

$$\mathcal{V}_P \triangleq \left(\{\hat{x}_n\}_{\forall n}, \hat{y}_N\right), \mathcal{V}_D \triangleq \left(\hat{p}_N, \{\hat{\gamma}_n\}_{\forall n}\right).$$

Note that to differentiate between offline and online settings, we place a hat on top of variables that denote the decisions made online.

(**Pricing Function in Step-1**) To enable an online implementation, the price predictions in **Step-1** must be performed based on causal information only. One natural way of designing such price predictions (without future information) is to relate $\hat{p}_n$ to the current total resource utilization as follows:

$$\hat{p}_n = \phi(\hat{y}_n), \forall n \in \mathcal{N}, \tag{19}$$

where $\hat{y}_n$ denotes the total resource utilization *after* allocating the required resources to request n. Based on our interpretation of $\hat{p}_n$, ϕ is referred to as the *pricing function* hereinafter. Eq. (19) indicates that our prediction of the final price will be updated whenever the total resource utilization changes.

(**Posted Price Mechanism:** PM_ϕ) The above two-step design indicates that the pricing function ϕ plays an important role in influencing $\mathcal{V}_P$ and $\mathcal{V}_D$, as well as the sequences of $\{P_n\}_{\forall n}$ and $\{D_n\}_{\forall n}$. Therefore, the design of ϕ is directly related to the inequalities in Proposition 3.1, and thus determines the competitive ratio of the online mechanism described above. The techniques of how to design ϕ in **Step-1** constitute the major results of this paper, and the details are deferred to Section 4. In the following, we temporarily assume that the pricing function ϕ is given and summarize our proposed online mechanism in Algorithm 1, dubbed PM_ϕ.

An interesting observation about PM_ϕ is that it can be implemented in a posted price manner. Unlike auctions [18, 34, 37], the supplier running posted price [22, 40] simply publishes the selling price (i.e., line 4) and does not collect any information from the agents. The decisions are made by each individual agent in the manner of *take-it-or-leave-it* (i.e., line 5-line 14). By virtue of posted-price mechanisms [17, 22, 40], PM_ϕ is IC, privacy-preserving and computationally-efficient.

Therefore, we argue that this is an extra advantage of our design, although we do not commit to posted price mechanisms a priori.

Algorithm 1: Posted Price Mechanism (PM_ϕ)

1: **Inputs:** A given setup $\mathcal{S} = \{f, \underline{p}, \bar{p}\}$ and ϕ.
2: **Initialize**: $\hat{y}_0 = 0$ and $\hat{p}_0 = \phi(\hat{y}_0)$.
3: **while** a new agent n arrives **do**
4:     Supplier publishes the price $\hat{p}_{n-1}$.
5:     **if** $v_n - \hat{p}_{n-1}r_n < 0$ **then**
6:         Agent n leaves (i.e., set $\hat{x}_n = 0$)
7:     **else if** $\hat{y}_{n-1} + r_n > 1$ **then**
8:         Request n is rejected (i.e., set $\hat{x}_n = 0$)
9:     **else**
10:        Request n is satisfied (i.e., set $\hat{x}_n = 1$)
11:        Collect the payment $\hat{\pi}_n$ by Eq. (17).
12:        Update the total resource utilization by Eq. (18).
13:        Update the price by $\hat{p}_n = \phi(\hat{y}_n)$.
14:     **end if**
15: **end while**

(Feasibility and Rationality of PM_ϕ**)** As shown by Proposition 3.1, the feasibility of $\mathcal{V}_P$ and $\mathcal{V}_D$ is crucial to prove the competitive ratio of the designed online mechanism. In Proposition 3.2 below, we show that our above design of $\mathcal{V}_P$ and $\mathcal{V}_D$ is feasible as long as ϕ is monotone.

PROPOSITION 3.2. *The primal solutions in $\mathcal{V}_P$ are always feasible to Problem* (1) *and Problem* (9). *The dual solutions in $\mathcal{V}_D$ are feasible to Problem* (10) *as long as ϕ is monotonically non-decreasing.*

PROOF. It is obvious that the design of $(\{\hat{x}_n\}_{\forall n}, \hat{y}_N)$ in $\mathcal{V}_P$ are feasible to Problem (1). For each agent $n \in \mathcal{N}$, our design of $\hat{\gamma}_n$ in Eq. (16) indicates that $\hat{\gamma}_n \geq v_n - \hat{p}_{n-1}r_n$. If ϕ is non-decreasing, we have $\hat{p}_{n-1} \leq \hat{p}_n \leq \hat{p}_N, \forall n \in \mathcal{N}$. Therefore, we have $\hat{\gamma}_n \geq v_n - \hat{p}_N r_n$ holds for all $n \in \mathcal{N}$. We thus complete the proof. □

Before leaving this section, it is worth mentioning that in economics and game theory [30, 34], our design of a dual feasible $\hat{\gamma}_n$ in Eq. (16) is known to guarantee the *individual rationality* of the mechanism, namely, no agent suffers from negative utility by participating in the mechanism (i.e., line 6 in Algorithm 1). Note that the individual rationality is also equivalent to saying that no agent is forced to participate in the mechanism.

4 MAIN RESULTS AND TECHNIQUES

Proposition 3.2 shows that our previous design of $\mathcal{V}_P$ and $\mathcal{V}_D$ is feasible if the pricing function ϕ is monotone. In this section, we present our major results regarding **Step-1**, namely, the design of a monotone pricing function ϕ so that PM_ϕ achieves a competitive performance in social welfare.

4.1 Sufficient Conditions

To aid our following presentation, let us define $\underline{\rho}$ and $\bar{\rho}$ as follows:

$$\underline{\rho} \triangleq \arg\max_{y \geq 0} \underline{p}y - \bar{f}(y), \quad \bar{\rho} \triangleq \arg\max_{y \geq 0} \bar{p}y - \bar{f}(y). \tag{20}$$

For a convex setup when $\bar{f}(y)$ is strictly-convex and differentiable in $y \in [0, 1]$, $\underline{\rho}$ and $\bar{\rho}$ can be respectively written as:

$$\underline{\rho} = \begin{cases} f'^{-1}(\underline{p}) & \text{if } \underline{p} \in (\underline{c}, \bar{c}), \\ 1 & \text{if } \underline{p} \geq \bar{c}, \end{cases} \; ; \; \bar{\rho} = \begin{cases} f'^{-1}(\bar{p}) & \text{if } \bar{p} \in (\underline{c}, \bar{c}), \\ 1 & \text{if } \bar{p} \geq \bar{c}. \end{cases}$$

The definitions of $\underline{\rho}$ and $\bar{\rho}$ can be interpreted as follows. Suppose there are infinitely-many identical agents whose valuation densities are all $\underline{p}$ ($\bar{p}$), then the maximum (optimal) resource utilization level is $\underline{\rho}$ ($\bar{\rho}$). Both $\underline{\rho}$ and $\bar{\rho}$ are capped by 1 due to the capacity limit.

Based on the above definitions of $\underline{\rho}$ and $\bar{\rho}$, we next give Theorem 4.1 which shows the sufficient conditions for ϕ so that PM_ϕ is IC and achieves a bounded competitive ratio.

THEOREM 4.1. *Given a convex setup $\mathcal{S}$, PM_ϕ is IC and α-competitive if ϕ is given by*

$$\phi(y) = \begin{cases} \underline{p} & \text{if } y \in [0, \omega), \\ \varphi(y) & \text{if } y \in [\omega, \bar{\rho}], \\ +\infty & \text{if } y \in (\bar{\rho}, +\infty), \end{cases} \tag{21}$$

where ω is a resource utilization threshold that satisfies

$$F_{\underline{p}}(\omega) \geq \frac{1}{\alpha} h(\underline{p}) \text{ and } 0 \leq \omega \leq \underline{\rho}, \tag{22}$$

and $\varphi(y)$ is an increasing function that satisfies

$$\begin{cases} \varphi'(y) \leq \alpha \cdot \frac{\varphi(y) - f'(y)}{h'(\varphi(y))}, y \in (\omega, \bar{\rho}); \\ \varphi(\omega) = \underline{p}, \varphi(\bar{\rho}) \geq \bar{p}. \end{cases} \tag{23}$$

In Eq. (22), $F_{\underline{p}}$ is the profit function defined in Eq. (8) and h is given by Eq. (11). In Eq. (23), h' represents the derivative of h.

Theorem 4.1 shows that PM_ϕ with any pricing function given by Eq. (21) is IC and α-competitive, provided that ω and φ satisfy certain conditions. The sufficient conditions in Theorem 4.1 are derived based on Proposition 3.1. In particular, Eq. (22) and Eq. (23) correspond to the initial inequality and the incremental inequality in Eq. (14), respectively. A rigorous proof is given in Appendix A, while discussions of intuitions are given in Section 5.

We note that a pricing function ϕ given by Eq. (21) consists of three segments, namely, the *flat-segment* $[0, \omega]$, the *increasing-segment* $[\omega, \bar{\rho}]$, and the *infinite-segment* $(\bar{\rho}, +\infty)$. Since ω is a resource utilization threshold that separates the first two segments, and plays a critical role in shaping the curvature of ϕ, we refer to ω as the *critical threshold* hereinafter. Recall that the valuation densities of all the agents are lower bounded by $\underline{p}$, and thus the incoming requests will always be satisfied when the total resource utilization is below ω, regardless of their valuations.

4.2 Necessary Conditions

An interesting result proved by this paper is as follows: existence of a pricing function given by Eq. (21) is not only sufficient to guarantee a bounded competitive ratio for PM_ϕ, but also necessary to the existence of any α-competitive online algorithm. The result is given by the following Theorem 4.2.

THEOREM 4.2. *Given a convex setup $\mathcal{S}$, if there exists an α-competitive online algorithm, then there must exist a critical threshold ω which satisfies Eq. (22) so that the following claims hold simultaneously:*

- *There exists a case when the total resource utilization is ω and all the accepted agents have the same valuation density $\underline{p}$.*

- *There exists a strictly-increasing function $\psi(p)$ that satisfies*

$$\begin{cases} \psi'(p) = \frac{1}{\alpha} \cdot \frac{h'(p)}{p - f'(\psi(p))}, p \in (\underline{p}, \bar{p}), \\ \psi(\underline{p}) = \omega, \psi(\bar{p}) \leq \bar{\rho}. \end{cases} \tag{24}$$

- *There exists a strictly-increasing function $\varphi(y)$ that satisfies*

$$\mathrm{BVP}(\omega, \alpha) \begin{cases} \varphi'(y) = \alpha \cdot \frac{\varphi(y) - f'(y)}{h'(\varphi(y))}, y \in (\omega, \bar{\rho}), \\ \varphi(\omega) = \underline{p}, \varphi(\bar{\rho}) \geq \bar{p}. \end{cases} \tag{25}$$

- ψ *and* φ *are inverse to each other, i.e.,* $\psi = \varphi^{-1}$ *or* $\varphi = \psi^{-1}$.

PROOF. The proof of this theorem is based on constructing a resource utilization level ω and a strictly-increasing function ψ for any α-competitive online algorithm under two special arrival instances. The complete proof is deferred to Appendix B. □

The first necessary condition in Theorem 4.2 argues that the existence of a critical threshold ω is directly related to the existence of any α-competitive algorithm. Therefore, selling the resource at the lowest price $\underline{p}$ during $[0, \omega]$ is necessary for PM_ϕ to achieve a bounded competitive ratio. Meanwhile, the value of the critical threshold ω must stay within a certain range defined by Eq. (22).

The second and the third conditions in Theorem 4.2 are related to ODEs with two boundary conditions, which are often termed as first-order two-point BVPs in mathematics [3, 36]. For a given setup, Eq. (25) is related to the critical threshold ω and the competitive ratio α, and thus we refer to Eq. (25) by $\mathrm{BVP}(\omega, \alpha)$ hereinafter. Note that $\mathrm{BVP}(\omega, \alpha)$ is equivalent to Eq. (23) after enforcing the equality of the differential inequality. Meanwhile, the sufficient conditions in Theorem 4.1 only require an increasing φ that satisfies Eq. (23), while a strictly-increasing φ is needed in Theorem 4.2 to guarantee the existence of an arbitrary α-competitive online algorithm.

(**Principles of Optimal Design**) Based on Theorem 4.2, if we can find a competitive ratio parameter α such that

- for some ω that satisfies Eq. (22), we can find a strictly-increasing solution to $\mathrm{BVP}(\omega, \alpha)$;
- for all ω that satisfies Eq. (22), there exists no strictly-increasing solution to $\mathrm{BVP}(\omega, \alpha - \epsilon), \forall \epsilon > 0$,

then this α must be the optimal (minimum) competitive ratio achievable by all online algorithms[5]. In the next two subsections, we use this principle to characterize the optimal competitive ratios for different setups.

4.3 Optimal Designs for Convex Setups

Theorem 4.3 below characterizes the design of a pricing function to achieve the optimal competitive ratio for a given convex setup.

THEOREM 4.3. *Given a convex setup* S, *there exists a unique optimal critical threshold* ω_* *so that the following claims hold:*

- *The optimal competitive ratio* $\alpha_*(S)$ *is given by*

$$\alpha_*(S) = \frac{h(\underline{p})}{F_{\underline{p}}(\omega_*)} = \begin{cases} \frac{\underline{p} f'^{-1}(\underline{p}) - f(f'^{-1}(\underline{p}))}{\underline{p}\omega_* - f(\omega_*)} & \text{if } \underline{p} \in (\underline{c}, \bar{c}), \\ \frac{\underline{p} - f(1)}{\underline{p}\omega_* - f(\omega_*)} & \text{if } \underline{p} \in [\bar{c}, +\infty). \end{cases} \tag{26}$$

---

[5]Note that the same design principle can also be applied for the BVP in Eq. (24). Here, we choose to deal with $\mathrm{BVP}(\omega, \alpha)$ because in the following we need to compute φ as part of our pricing function ϕ.

- *There exists a unique optimal pricing function ϕ_* given by*

$$\phi_*(y) = \begin{cases} \underline{p} & \text{if } y \in [0, \omega_*), \\ \varphi_*(y) & \text{if } y \in [\omega_*, \bar{\rho}], \\ +\infty & \text{if } y \in (\bar{\rho}, +\infty), \end{cases} \tag{27}$$

  so that PM_{ϕ_*} is IC and $\alpha_*(\mathcal{S})$-competitive. In Eq. (27), φ_* is the unique strictly-increasing solution to $BVP(\omega_*, \alpha_*(\mathcal{S}))$.

Theorem 4.3 illustrates our major results regarding the convex setup, namely, the existence and uniqueness of the optimal pricing function ϕ_* so that PM_{ϕ_*} achieves the optimal competitive ratio $\alpha_*(\mathcal{S})$. The proof of Theorem 4.3 is based on analyzing the solution structures of $BVP(\omega, \alpha)$, and the details are given in Section 6. We emphasize that Theorem 4.3 shows the existence and uniqueness of the optimal critical threshold ω_* without discussing how to quantify it. For the details about how to calculate ω_*, please refer to Theorem 6.3 in Section 6 as well.

COROLLARY 4.4. *Given a convex setup $\mathcal{S}$, we have:*
- *$\alpha_*(\mathcal{S})$ is strictly decreasing in $\underline{p} \in (\underline{c}, \bar{p}]$ for a given $\bar{p} \in (\underline{c}, +\infty)$.*
- *$\alpha_*(\mathcal{S})$ is strictly increasing in $\bar{p} \in [\underline{p}, +\infty)$ for a given $\underline{p} \in (\underline{c}, +\infty)$.*
- *$\alpha_*(\mathcal{S}) = 1$ and $\omega_* = \rho$ when $\underline{p} = \bar{p} \in (\underline{c}, +\infty)$.*

In Eq. (26), $\alpha_*(\mathcal{S})$ explicitly depends on f and $\underline{p}$, and implicitly depends on $\bar{p}$ through the optimal critical threshold ω_*. This leads to the monotonicity of $\alpha_*(\mathcal{S})$ w.r.t. to $\underline{p}$ and $\bar{p}$ in Corollary 4.4. In particular, the third bullet in Corollary 4.4 can be interpreted as follows: when $\underline{p} = \bar{p} \in (\underline{c}, +\infty)$, the agents are identical in terms of their valuation densities, and thus it makes no difference to know all the future arrival information, i.e., $\alpha_*(\mathcal{S}) = 1$. The proof of Corollary 4.4 is deferred to Appendix H.

Before leaving this subsection, we give the following remark regarding the uniqueness of ϕ_*.

REMARK 2 (UNIQUENESS). *We emphasize that the property of uniqueness in Theorem 4.3 does not mean that there exists only one $\alpha_*(\mathcal{S})$-competitive online mechanism/algorithm. Instead, Theorem 4.3 only argues that PM_{ϕ_*} can achieve the optimal competitive ratio $\alpha_*(\mathcal{S})$ with a unique pricing function ϕ_*.*

4.4 Optimal Designs for Linear Supply Costs

For a given convex setup $\mathcal{S}$, Theorem 4.3 shows that the optimal competitive ratio $\alpha_*(\mathcal{S})$ and the optimal pricing function ϕ_* are directly related to the optimal critical threshold ω_*. It is worth emphasizing that ω_* is a design parameter which cannot be given in analytical forms, so does ϕ_* (in fact, f itself is arbitrary, and thus this is not surprising). Nevertheless, for some special cost functions, Corollary 4.5 below shows that logarithmic competitive ratios can be obtained via analytical designs of ϕ_* and ω_*.

COROLLARY 4.5. *Given a nice setup $\mathcal{S} = \{f, \underline{p}, \bar{p}\}$, if the cost function $f(y) = qy$, where $q \geq 0$, then there exists a unique pricing function ϕ_* given by*

$$\phi_*(y) = \begin{cases} \underline{p} & \text{if } y \in [0, \omega_*), \\ (\underline{p} - q) \cdot \exp\left(y/\omega_* - 1\right) + q & \text{if } y \in [\omega_*, 1], \\ +\infty & \text{if } y \in (1, +\infty), \end{cases} \tag{28}$$

such that PM_{ϕ_*} is IC and $\alpha_*(\mathcal{S})$-competitive, where $\alpha_*(\mathcal{S})$ is given by

$$\alpha_*(\mathcal{S}) = 1 + \ln\left(\frac{\bar{p} - q}{\underline{p} - q}\right).$$

In Eq. (28), the optimal critical threshold $\omega_* = \frac{1}{\alpha_*(\mathcal{S})}$.

We note that the logarithmic competitive ratios in Corollary 4.5 are not new and have been discussed in the literature, e.g., [42]. Based on [42], such logarithmic competitive ratios are optimal, unless extra assumptions are introduced. Here, based on our above sufficient and necessary conditions for convex setups, we provide new proofs, which are simple and intuitive, for the results in Corollary 4.5. The details about the proof are deferred to Appendix C.

Summarizing our results in Theorem 4.3 and Corollary 4.5, we argue that PM_ϕ is a unified mechanism for online resource allocation with or without supply costs. In particular, we obtain optimal competitive ratios for different setups including zero cost (i.e, without supply cost, which corresponds to $q = 0$ in Corollary 4.5), linear cost, and strictly-convex cost.

5 INTERPRETATION, INTUITION AND GENERALIZATION OF THEOREM 4.1

In this section, we give a geometric interpretation of Theorem 4.1 and discuss the intuitions of our geometric interpretation via worst-case analysis. We also show that the sufficient conditions in Theorem 4.1 can be generalized in multiple directions.

5.1 A Geometric Interpretation of Theorem 4.1

Since $\phi(\bar\rho) = \varphi(\bar\rho) \geq \bar p$, the highest-possible resource utilization level under PM_ϕ is $\phi^{-1}(\bar p)$. Let us denote the final resource utilization level under PM_ϕ by $\rho \in [0, \phi^{-1}(\bar p)]$. Intuitively, if $\rho \in [0, \omega]$, then the decisions made by PM_ϕ and its offline counterpart are the same as long as ϕ has the flat-segment $[0, \omega]$, namely, both are to satisfy all the requests. Hence, if $\rho \in [0, \omega]$, the competitive ratio of PM_ϕ is 1. We next focus on the more general case when $\rho \in [\omega, \phi^{-1}(\bar p)]$.

For any $\rho \in [\omega, \phi^{-1}(\bar p)]$, let us denote the final price by $p = \phi(\rho)$. The pricing function $\phi(y)$ satisfying Eq. (23) indicates that for any given $\rho \in [\omega, \phi^{-1}(\bar p)]$, we have

$$\int_\omega^\rho \left(\phi(y) - f'(y) \right) dy \geq \int_{\underline p}^p \frac{1}{\alpha} h'\left(\phi(y) \right) d\phi(y). \tag{29}$$

Meanwhile, based on Eq. (22), we have

$$F_{\underline p}(\omega) = \underline p \omega - f(\omega) \geq \frac{1}{\alpha} h(\underline p), \tag{30}$$

where $F_{\underline p}(\omega)$ is the profit function defined by Eq. (8). The combination of Eq. (29) and Eq. (30) leads to the following condition

$$\underline p \omega + \int_\omega^\rho \phi(y) dy - f(\rho) \geq \frac{1}{\alpha} h(\phi(\rho)). \tag{31}$$

Eq. (31) must hold for all possible ρ's, i.e., all $\rho \in [\omega, \phi^{-1}(\bar p)]$, and thus we have the following expression of α:

$$\alpha = \max_{\rho \in [\omega, \phi^{-1}(\bar p)]} \frac{h(\phi(\rho))}{\underline p \omega + \int_\omega^\rho \phi(y) dy - f(\rho)}. \tag{32}$$

To illustrate the geometric meaning of Eq. (32), let us assume $\bar p \geq \bar c$ for simplicity. We illustrate the geometric interpretation of Eq. (32) in Fig. 1. Specifically, Fig. 1(a) shows the case when $p = \phi(\rho) \leq \bar c$, i.e., $\rho \in [\omega, u]$, where u is a resource utilization level such that $\phi(u) = \bar c$ or equivalently, $u = \phi^{-1}(\bar c)$; Fig. 1(b) shows the case when $p = \phi(\rho) > \bar c$, i.e., $\rho \in [u, 1]$. Based on the expression of $h(p)$ in Eq. (12), for both subfigures in Fig. 1, the area of the grey region and the two blue regions (which is referred to as '**grey+blue**') represents the numerator of the fraction in Eq. (32), while the denominator is the area of the grey region only (which is referred to as '**grey**'). Therefore, based on Eq. (32), an
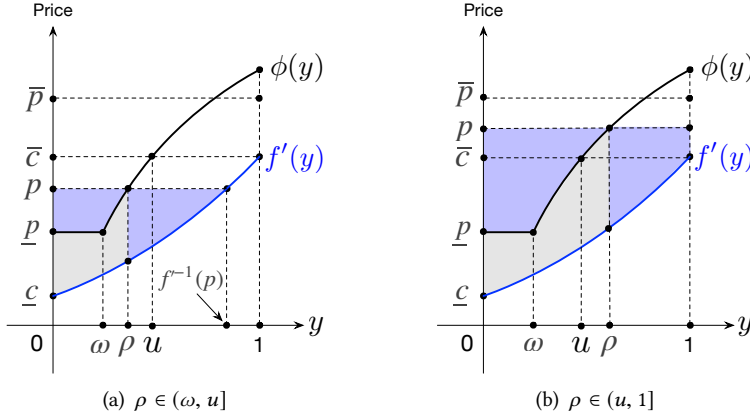
Fig. 1. Illustration of a feasible pricing function ϕ. The parameter u is a resource utilization level such that $\phi(u) = \bar{c}$.

α-competitive PM_ϕ is equivalent to a pricing function ϕ so that the ratio between '**grey+blue**' and '**grey**' less than or equal to α for all possible ρ's in $[\omega, \phi^{-1}(\bar{p})]$.

In summary, given a pricing function ϕ that satisfies the sufficient conditions in Theorem 4.1, when the final resource utilization level $\rho \in [0, \omega]$, the competitive ratio of PM_ϕ is 1; when $\rho \in [\omega, \phi^{-1}(\bar{p})]$, the competitive ratio of PM_ϕ equals the maximum ratio between '**grey+blue**' and '**grey**' for all $\rho \in [\omega, \phi^{-1}(\bar{p})]$. Since there is no prior information of ρ, the competitive ratio of PM_ϕ is dominated by the latter case, leading to the expression of α in Eq. (32).

5.2 Intuitions via Worst-Case Analysis

Our above analysis only visualizes the geometric meaning of Eq. (32), but reveals little intuition and rationality about why the maximum ratio between '**grey+blue**' and '**grey**' determines the competitive ratio of PM_ϕ. Below we show that, Eq. (32) can be traced back to the original definition of α in Eq. (4), based on which the rationality of Theorem 4.1 can be demonstrated by the performance of PM_ϕ under a special arrival instance in the worst-case scenario.

To be more specific, let us again focus on the case when the final resource utilization level under PM_ϕ is $\rho \in [\omega, \bar{\rho}]$. For any given $\rho \in [\omega, \bar{\rho}]$, let us assume the interval $[0, \rho]$ is discretized into $K + B$ blocks with each block Δ long, i.e., $\Delta = \rho/(K + B)$, where we assume Δ is infinitesimally small. Let us consider a special arrival instance, denoted by $\mathcal{A}_\rho$, which consists of three groups of agents as follows: *the first group of agents are similar to our construction of $\mathcal{A}_p$ in Section 4.2, where we have K identical agents with valuation density $\underline{p}$ and requirement Δ. Here, we assume $K\Delta = \omega$. The second group of agents are indexed by $b \in \{1, 2, \cdots, B\}$, and the third group of agents are also identical and are indexed by i, where $i = \{1, 2, \cdots, I\}$ and $I \geq 1/\Delta$. For each request $b \in \{1, 2, \cdots, B\}$ in the second group, we assume the valuation of agent b is given by $v_b = \phi(\omega + b\Delta) \cdot \Delta$, where $\phi(\omega + b\Delta)$ and Δ denote the valuation density of agent b and the resource requirement of her request, respectively. Similarly, for each request $i \in \{1, 2, \cdots, I\}$ in the third group, we assume the valuation of agent i is given by $v_i = \phi(\rho)\Delta$, where $\phi(\rho)$ and Δ denote the valuation density of agent i and the requirement of her request, respectively.*

Given an arrival instance $\mathcal{A}_\rho$ with $\rho \in [\omega, \bar{\rho}]$, let us denote $p = \phi(\rho)$. The optimal social welfare in hindsight, denoted by $S_{\text{offline}}(\mathcal{A}_\rho)$, is to reject all the requests in the first two groups but satisfy all

the requests in the third group until reaching the resource utilization level y, at which the marginal cost equals either p (i.e., $y = f'^{-1}(p)$) or $\underline{c}$ (i.e., $y = 1$). Thus, $S_{\text{offline}}(\mathcal{A}_\rho)$ is given by

$$S_{\text{offline}}(\mathcal{A}_\rho) = \begin{cases} pf'^{-1}(p) - f\left(f'^{-1}(p)\right) & \text{if } p \in (\underline{c}, \bar{c}], \\ p - f(1) & \text{if } p \in (\bar{c}, +\infty), \end{cases} \tag{33}$$

which equals the numerator of the fraction in Eq. (32) based on the expression of $h(p)$ in Eq. (11). Note that $I \geq 1/\Delta$ guarantees that the final resource utilization level $y = f'^{-1}(p)$ or $y = 1$ can be reached by the offline optimal strategy.

In comparison, given the arrival instance $\mathcal{A}_\rho$, our online mechanism PM_ϕ will satisfy all the requests in the first two groups and reject all in the third group. Thus, the social welfare is given by

$$S_{\text{online}}(\mathcal{A}_\rho) = \underline{p}K\Delta + \sum_{b=1}^{B} \phi(\omega + b\Delta) \cdot \Delta - f(\rho), \tag{34}$$

which equals the denominator of the fraction in Eq. (32) since $K\Delta = \omega$ and Δ is infinitesimally small.

Based on (32)-(34), the competitive ratio of PM_ϕ is given by

$$\alpha = \max_{\rho \in [\omega, \bar{\rho}]} \frac{S_{\text{offline}}(\mathcal{A}_\rho)}{S_{\text{online}}(\mathcal{A}_\rho)} = \max_{\text{all possible } \mathcal{A}} \frac{S_{\text{offline}}(\mathcal{A})}{S_{\text{online}}(\mathcal{A})},$$

where the second equality is from the original definition of α in Eq. (4). It becomes clear now that the sufficient conditions in Theorem 4.1 lead to the expression of α in Eq. (32), which in principle captures the worst-case performance ratio between the optimal strategy in hindsight and our proposed online mechanism PM_ϕ under a continuum of special arrival instances $\{\mathcal{A}_\rho\}_{\{\forall \rho \in [\omega, \bar{\rho}]\}}$. In this regard, the arrival instance $\mathcal{A}_\rho$ is not constructed randomly; it indeed captures the worst-case scenario in the context.

5.3 Generalizations of Theorem 4.1

5.3.1 *General Increasing Pricing Functions*. Based on Eq. (32), we can generalize Theorem 4.1 to the following Corollary 5.1 with a broader class of competitive pricing functions.

COROLLARY 5.1. *Given a convex setup $\mathcal{S}$, if ϕ is given by*

$$\phi(y) = \begin{cases} \underline{p} & \text{if } y \in [0, \omega), \\ \varphi(y) & \text{if } y \in [\omega, 1], \\ +\infty & \text{if } y \in (1, +\infty), \end{cases} \tag{35}$$

where $\omega \in (0, \underline{\rho}]$, and φ is increasing in $[\omega, 1]$ with $\varphi(\omega) = \underline{p}$ and $\varphi(1) \geq \bar{c}$, then PM_ϕ is IC and $\alpha(\omega, \varphi)$-competitive, where $\alpha(\omega, \varphi)$ is

$$\alpha(\omega, \varphi) = \max \left\{ \frac{h(\underline{p})}{F_{\underline{p}}(\omega)}, \frac{h(\bar{p})}{\underline{p}\omega + \int_\omega^{\bar{\rho}_\varphi} \varphi(y)dy - f(\bar{\rho}_\varphi)}, \max_{\rho \in [\omega, \bar{\rho}_\varphi]} \frac{h(\varphi(\rho))}{\underline{p}\omega + \int_\omega^\rho \varphi(y)dy - f(\rho)} \right\}. \tag{36}$$

In Eq. (36), $\bar{\rho}_\varphi$ is the maximum resource utilization level defined as:

$$\bar{\rho}_\varphi \triangleq \begin{cases} \varphi^{-1}(\bar{p}) & \text{if } \bar{p} \leq \varphi(1), \\ 1 & \text{if } \bar{p} > \varphi(1). \end{cases}$$

Corollary 5.1 argues that any continuous pricing function with a flat-segment and an increasing-segment can lead to a bounded competitive ratio for PM_ϕ. Recall that when $\omega = 0$, our previous necessary conditions in Theorem 4.2 argue that it is impossible to achieve a bounded competitive

ratio. This can be illustrated by Eq. (36), where the first term in the bracket is unbounded since $F_p(0) = 0$. Note that the third term in Eq. (36) is derived from Eq. (32). Meanwhile, when $\varphi(1) \geq \bar{p}$, the second term in the bracket of Eq. (36) is not needed as it is contained by the third term. The proof of Corollary 5.1 is given in Appendix D.

REMARK 3. *Based on Theorem 4.3 and Corollary 5.1, we argue that $\alpha_*(\mathcal{S})$ is the optimal solution to the following optimization problem:*

$$\alpha_*(\mathcal{S}) = \underset{\omega, \varphi}{\text{minimize}} \qquad \alpha(\omega, \varphi)$$
$$\text{subject to} \qquad \omega \in [0, \underline{\rho}],$$
$$\varphi'(y) \geq 0, \forall y \in [\omega, 1],$$
$$\varphi(\omega) = \underline{p}, \varphi(1) \geq \bar{c}.$$

In this regard, our previous design of ω_* and φ_* in Theorem 4.3 essentially provides a method of solving the above optimization problem in functional spaces.

5.3.2 *Pricing Functions for Multiple Time Slots.*

Our mechanism also extends to more general settings when there exists multiple time slots. Specifically, let us consider the following model:

$$\underset{\boldsymbol{x},\boldsymbol{y}}{\text{maximize}} \qquad \sum_{n \in \mathcal{N}} v_n x_n - \sum_{t \in \mathcal{T}} \bar{f}_t(y_t) \tag{37a}$$

$$\text{subject to} \qquad \sum_{n \in \mathcal{N}} r_n^t x_n = y_t, \forall t \in \mathcal{T}, \tag{37b}$$

$$x_n \in \{0, 1\}, \forall n \in \mathcal{N}. \tag{37c}$$

In Problem (37), we consider a discrete time system and index different time slots by $t \in \mathcal{T} \triangleq \{1, 2, \cdots, T\}$. In Eq. (37a), the supply cost at $t \in \mathcal{T}$ is directly written as an extended cost function $\bar{f}_t$, i.e., $\bar{f}_t(y) = f_t(y)$ if $y \in [0, 1]$, and $\bar{f}_t(y) = +\infty$ if $y \in (1, +\infty)$. In Eq. (37b), r_n^t denotes the requirement of agent n at $t \in \mathcal{T}_n$, where $\mathcal{T}_n$ represents the time duration of request n. We assume $r_n^t = 0$ if $t \notin \mathcal{T}_n$ and $r_n^t > 0$ if $t \in \mathcal{T}_n$ so that we can simply denote the resource requirement of request n by $\{r_n^t\}_{\forall t \in \mathcal{T}}$. Meanwhile, similar to our previous basic resource allocation model, each agent $n \in \mathcal{N}$ can be represented by $\boldsymbol{\theta}_n = (\{r_n^t\}_{\forall t \in \mathcal{T}_n}, v_n)$, where v_n denotes the valuation of agent n if all the requirements $\{r_n^t\}_{\forall t \in \mathcal{T}_n}$ are satisfied.

While most of our previous definitions and notations can be reused by simply adding a time index $t \in \mathcal{T}$, e.g., $h_t(p)$, $\underline{c}_t$, and $\bar{c}_t$, we redefine the lower and upper bounds of the valuation density by

$$\min_{n \in \mathcal{N}: r_n^t \neq 0} \left\{ \frac{v_n}{r_n^t} \right\} \geq \underline{p}_t, \max_{n \in \mathcal{N}: r_n^t \neq 0} \left\{ \frac{v_n}{r_n^t} \right\} \leq \bar{p}_t, \forall t \in \mathcal{T},$$

where $\underline{p}_t$ and $\bar{p}_t$ correspond to $\underline{p}$ and $\bar{p}$ in Assumption 1, respectively. Based on $\underline{p}_t$ and $\bar{p}_t$, we can define $\underline{\rho}_t$ and $\bar{\rho}_t$ in the similar way as $\underline{\rho}$ and $\bar{\rho}$ in Eq. (20). Here we omit the details for brevity.

Following the principle discussed in Section 3.3, we can design a pricing function $\hat{p}_t^{(n)} = \phi_t(y_t^{(n)})$ for each time slot $t \in \mathcal{T}$, where $y_t^{(n)}$ denotes the total resource utilization after processing agent n, similar to our definition of $\hat{y}_n$ in Eq. (18). Based on the pricing functions $\{\phi_t\}_{\forall t}$, we can set the utility of agent $n \in \mathcal{N}$ by

$$\hat{\gamma}_n = \max \left\{ v_n - \sum_{t \in \mathcal{T}_n} r_n^t \cdot \phi_t \left(y_t^{(n-1)} \right), 0 \right\}, \forall n \in \mathcal{N},$$

where $\sum_{t \in \mathcal{T}_n} r_n^t \cdot \phi_t(y_t^{(n-1)})$ denotes the payment of agent n if $\hat{\gamma}_n > 0$. The mechanism PM_ϕ can thus be generalized to a posted price mechanism PM_ϕ with a vector of pricing functions $\boldsymbol{\phi} \triangleq \{\phi_t\}_{\forall t}$.

Based on the above discussions, we give a generalized version of Theorem 4.1 in the following Theorem 5.2.

THEOREM 5.2. *Given a convex setup $\mathcal{S} = \{f_t, \underline{p}_t, \bar{p}_t\}_{\forall t}$, $PM_{\boldsymbol{\phi}}$ is IC and $\max_t\{\alpha_t\}$-competitive if for each $t \in \mathcal{T}$, ϕ_t is given by*

$$\phi_t(y) = \begin{cases} \underline{p}_t & \text{if } y \in [0, \omega_t), \\ \varphi_t(y) & \text{if } y \in [\omega_t, 1], \\ +\infty & \text{if } y \in (1, +\infty), \end{cases} \tag{38}$$

where ω_t is the critical threshold that satisfies

$$F_{\underline{p}_t}(\omega_t) \geq \frac{1}{\alpha_t} h_t(\underline{p}_t) \text{ and } 0 \leq \omega_t \leq \rho_t, \tag{39}$$

and $\varphi_t(y)$ is an increasing function that satisfies

$$\begin{cases} \varphi_t'(y) \leq \alpha_t \cdot \frac{\varphi_t(y) - f_t'(y)}{h'(\varphi_t(y))}, y \in (\omega_t, 1) \\ \varphi_t(\omega_t) = \underline{p}_t, \varphi_t(1) \geq \bar{p}_t + \sum_{\hat{t} \in \mathcal{T} \setminus \{t\}} h_{\hat{t}}(\bar{p}_{\hat{t}}). \end{cases} \tag{40}$$

Note that Eq. (40) is the same as Eq. (23) except the second boundary condition, which depends on parameters related to all the time slots $\mathcal{T}$. Our previous analysis regarding the basic resource allocation model shows that the optimal competitive ratio $\alpha_*(\mathcal{S})$ depends on the boundary conditions in Eq. (25). Therefore, in the case with multiple time slots, the final competitive ratio $\alpha = \max_t\{\alpha_t\}$ depends on $|\mathcal{T}|$ as well. For the proof of Theorem 5.2, as well as discussions of the properties of $\{\phi_t\}_{\forall t}$ based on Theorem 5.2, please refer to Appendix I.

REMARK 4 (APPLICATIONS IN CLOUD COMPUTING). *Problem (37) can be used to model the online resource allocation in cloud computing [26, 41]. A cloud service provider allocates a single type of resources (e.g., CPU) to a set of jobs $\mathcal{N} = \{1, \cdots, N\}$ that arrive in a sequential order, and each job $n \in \mathcal{N}$ is active in duration $\mathcal{T}_n \subset \mathcal{T}$. Based on this interpretation, Problem (37) can be regarded as maximizing the social welfare of online resource allocation in cloud computing with server costs (e.g., energy costs). In particular, the arrival and departure times of job n can be taken into account by $\mathcal{T}_n$, and agents are flexible to set their resource requirements in each time slot based on the length of their active durations.*

We end this section with the following remark regarding the setting with multiple types of resources.

REMARK 5 (MULTIPLE TYPES OF RESOURCES). *Problem (37) can also be interpreted in the context of resource allocation with multiple types of resources. Specifically, we can consider that in each time slot we have a different type of resource so that $\mathcal{T}$ denotes the set of resource types. Based on this notational system, r_n^t denotes the resource requirement of agent n for resource type t, and $\mathcal{T}_n$ represents the set of resource types required by agent n. For this reason, it is mathematically equivalent to consider multiple types of resources and multiple time slots. In addition, we can also consider the setup with multiple resource types and multiple time slots simultaneously. We argue that our previous design principle still applies to such more general and complex setups. The difference is that we need to design a pricing function for each type of resource at each time slot[6].*

---

[6]The rationality is that in the worst case agents may only require a single type of resource at a single time slot.

(a) **Category of the three cases**

(b) **Case-1**: $\underline{c} < \underline{p} < \bar{c} < \bar{p}$

(c) **Case-2**: $\underline{c} < \bar{c} \leq \underline{p} \leq \bar{p}$

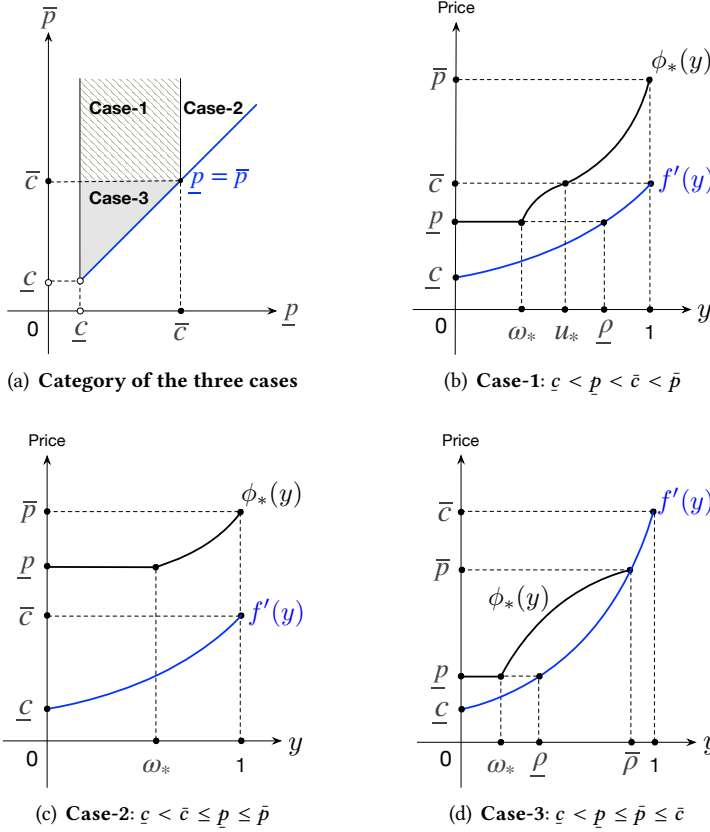(d) **Case-3**: $\underline{c} < \underline{p} \leq \bar{p} \leq \bar{c}$

Fig. 2. Category of the three cases for a given setup and illustrations of the optimal pricing functions in the three cases.

6 PROOF OF THEOREM 4.3

In this section, we provide the proof of Theorem 4.3. Our proof is organized in three cases based on the relationship between $\underline{c}, \bar{c}, \underline{p}$, and $\bar{p}$. After the proof, we give an example (quadratic supply cost $f(y) = \frac{1}{2}y^2$) to show the calculation of $\alpha_*(\mathcal{S})$, the optimal critical threshold ω_*, and the optimal pricing function ϕ_*.

6.1 Overview of Our Three-Case Proof

Our proof heavily follows the sufficient conditions in Theorem 4.1 and the necessary conditions in Theorem 4.2. For the sake of better reference, here we revisit BVP(ω, α) as follows:

$$\text{BVP}(\omega, \alpha) \begin{cases} \varphi'(y) = \alpha \cdot \frac{\varphi(y) - f'(y)}{h'(\varphi(y))}, y \in (\omega, \bar{\rho}), \\ \varphi(\omega) = \underline{p}, \varphi(\bar{\rho}) \geq \bar{p}. \end{cases} \tag{41}$$

Based on the expression of h in Eq. (11), the denominator of the ODE in Eq. (41) can be written as

$$h'(\varphi(y)) = \begin{cases} f'^{-1}(\varphi(y)) & \text{if } \varphi(y) \in [\underline{c}, \bar{c}], \\ 1 & \text{if } \varphi(y) \in (\bar{c}, +\infty), \end{cases} \tag{42}$$

Eq. (42) indicates that the ODE in Eq. (41) can be equivalently transformed into two ODEs based on whether $\varphi(y)$ is larger than $\bar{c}$ or not. Moreover, notice that the second boundary condition of BVP(ω, α) is $\varphi(\bar{\rho}) \geq \bar{p}$, where $\bar{p}$ can be either less than or larger than $\bar{c}$. For this reason, based on the values of $\underline{p}$ and $\bar{p}$, as well as the two boundary conditions in Eq. (41), we can categorize the whole possibilities into three cases as follows.

- **Case-1**: $\underline{c} < \underline{p} < \bar{c} < \bar{p}$. This case captures the setup when the minimum valuation density is small but the maximum valuation density is large.
- **Case-2**: $\underline{c} < \bar{c} \leq \underline{p} \leq \bar{p}$. This case captures the setup when the minimum valuation density is large.
- **Case-3**: $\underline{c} < \underline{p} \leq \bar{p} \leq \bar{c}$. This case captures the setup when the maximum valuation density is small.

We illustrate the above three cases in Fig. 2(a). A nice setup $\mathcal{S}$ must have $\bar{p} \geq \underline{p}$, and thus we only focus on the upper-left triangular part. We next present our proof of Theorem 4.3 in **Case-1**. The proofs of Theorem 4.3 in **Case-2** and **Case-3** are similar and thus are deferred to Appendix G.

6.2 Proof of Theorem 4.3 in Case-1

In this case, $\underline{p} < \bar{c} < \bar{p}$ indicates that $\rho < 1 = \bar{\rho}$. Based on the two boundary conditions in Eq. (41), there must exist a resource utilization threshold $u \in (\omega, 1)$ so that $\phi(u) = \bar{c}$, as illustrated in Fig. 1(b). Based on Theorem 4.1 and Theorem 4.2, we give the following Corollary 6.1 which summarizes the sufficient and necessary conditions in **Case-1**.

COROLLARY 6.1. *Given a convex setup $\mathcal{S}$ in **Case-1**, PM_ϕ is IC and α-competitive if there exists a pair of resource utilization thresholds $(\omega, u) \in [F_{\underline{p}}^{-1}\left(\frac{1}{\alpha}h(\underline{p})\right), \rho] \times (\omega, 1)$ such that ϕ is given by*

$$\phi(y) = \begin{cases} \underline{p} & \text{if } y \in [0, \omega], \\ \varphi_1(y) & \text{if } y \in [\omega, u], \\ \varphi_2(y) & \text{if } y \in [u, 1], \\ +\infty & \text{if } y \in (1, +\infty), \end{cases} \tag{43}$$

where $\varphi_1(y)$ is a solution to the following BVP:

$$\text{BVP}_1(\omega, u, \alpha) \begin{cases} \varphi_1'(y) = \alpha \cdot \frac{\varphi_1(y) - f'(y)}{f'^{-1}(\varphi_1(y))}, y \in (\omega, u); \\ \varphi_1(\omega) = \underline{p}, \varphi_1(u) = \bar{c}, \end{cases} \tag{44}$$

and $\varphi_2(y)$ is a solution to the following BVP:

$$\text{BVP}_2(u, \alpha) \begin{cases} \varphi_2'(y) = \alpha \cdot (\varphi_2(y) - f'(y)), y \in (u, 1); \\ \varphi_2(u) = \bar{c}, \varphi_2(1) \geq \bar{p}. \end{cases} \tag{45}$$

On the other hand, if there exists an α-competitive online algorithm, then there must exist some pair of $(\omega, u) \in [F_{\underline{p}}^{-1}\left(\frac{1}{\alpha}h(\underline{p})\right), \rho] \times (\omega, 1)$ such that $\text{BVP}_1(\omega, u, \alpha)$ and $\text{BVP}_2(u, \alpha)$ are well-defined, and each of them has a strictly-increasing solution.

In Corollary 6.1, $F_{\underline{p}}^{-1}$ represents the inverse[7] of the profit function $F_{\underline{p}}$ defined in Eq. (8), and $F_{\underline{p}}^{-1}\left(\frac{1}{\alpha}h(\underline{p})\right)$ is derived from Eq. (22). The flat-segment $[0, \omega]$ of ϕ directly follows Theorem 4.1. The two BVPs in Eq. (44) and Eq. (45), as well as their corresponding solutions $\varphi_1(y)$ and $\varphi_2(y)$, follow

---

[7]Based on Eq. (8), $F_{\underline{p}}(y)$ is strictly increasing in $y \in [0, \rho]$ since $f'(y) \leq \underline{p}$ holds for all $y \in [0, \rho]$. Thus, the inverse of $F_{\underline{p}}$, denoted by $F_{\underline{p}}^{-1}$, is well-defined.

Theorem 4.2 after substituting $h'(\varphi(y))$ from Eq. (42) into $\text{BVP}(\omega, \alpha)$. The sufficiency and necessity of Corollary 6.1 thus follow.

Corollary 6.1 not only shows how to justify whether a given pricing function is competitive or not, but also argues that, to find the optimal competitive ratio, we simply need to find the minimum α so that both $\text{BVP}_1(\omega, u, \alpha)$ and $\text{BVP}_2(u, \alpha)$ have strictly-increasing solutions. Below we give Proposition 6.2 which shows the existence conditions of solutions to both $\text{BVP}_1(\omega, u, \alpha)$ and $\text{BVP}_2(u, \alpha)$.

PROPOSITION 6.2. *Given a convex setup $\mathcal{S}$ in **Case-1**, the following claims regarding $\text{BVP}_1(\omega, u, \alpha)$ and $\text{BVP}_2(u, \alpha)$ are true:*

- *For each given pair of $(\omega, u) \in [F_{\underline{p}}^{-1}\left(\frac{1}{\alpha}h(\underline{p})\right), \underline{\rho}] \times (\omega, 1)$, there exists a well-defined function $\Gamma_1(\omega, u)$ so that $\text{BVP}_1(\omega, u, \alpha)$ has a unique strictly-increasing solution if and only if $\alpha = \Gamma_1(\omega, u)$.*
- *For each given $u \in (0, 1)$, $\text{BVP}_2(u, \alpha)$ has a unique strictly-increasing solution if and only if $\alpha \geq \Gamma_2(u)$, where $\Gamma_2(u)$ is the unique root to the following equation in variable Γ_2:*

$$\int_u^1 \frac{\Gamma_2 f'(y)}{\exp(y\Gamma_2)} dy = \frac{\bar{p}}{\exp(\Gamma_2)} - \frac{\bar{c}}{\exp(u\Gamma_2)}. \tag{46}$$

Moreover, for each given $u \in (0, 1)$, when $\alpha = \Gamma_2(u)$, the unique solution to $\text{BVP}_2(u, \alpha)$ satisfies $\varphi_2(1) = \bar{p}$.

The proof of Proposition 6.2 is based on the two ODEs given in Eq. (44) and Eq. (45), and the details are deferred to Appendix E. Based on the above Proposition 6.2, below we give the optimal competitive ratio in **Case-1**.

THEOREM 6.3. *Given a convex setup $\mathcal{S}$ in **Case-1**, the optimal competitive ratio achievable by all online algorithms is given by:*

$$\alpha_*(\mathcal{S}) = \frac{h(\underline{p})}{F_{\underline{p}}(\omega_*)} = \Gamma_1(\omega_*, u_*) = \Gamma_2(u_*), \tag{47}$$

where $u_* \in (0, 1)$ is the unique root that satisfies

$$\Gamma_1\left(F_{\underline{p}}^{-1}\left(\frac{h(\underline{p})}{\Gamma_2(u_*)}\right), u_*\right) = \Gamma_2(u_*), \tag{48}$$

and $\omega_* \in [0, \underline{\rho}]$ is given by

$$\omega_* = F_{\underline{p}}^{-1}\left(\frac{h(\underline{p})}{\Gamma_2(u_*)}\right).$$

Meanwhile, PM_{ϕ_*} is $\alpha_*(\mathcal{S})$-competitive if and only if ϕ_* is given by Eq. (43) with $(\omega, u, \alpha) = (\omega_*, u_*, \alpha_*(\mathcal{S}))$. Moreover, we have $\phi_*(1) = \bar{p}$.

PROOF. This theorem shows the unique existence of ω_*, u_*, and $\alpha_*(\mathcal{S})$. Based on these parameters, we can calculate the unique optimal pricing function ϕ_* based on Eq. (43). The proof of the unique existence of ω_* and u_* is based on Proposition 6.2. For the detailed proof, please refer to Appendix F. □

Based on Theorem 6.3, Theorem 4.3 directly follows in **Case-1**. Due to space limitations, we defer the proof of Theorem 4.3 in **Case-2** and **Case-3** to Appendix G.

Fig. 2(b) illustrates the optimal pricing function ϕ_* in **Case-1**. We can see that in addition to the flat-segment $[0, \omega_*]$, the increasing-segment is further divided into two parts, namely, $[\omega_*, u_*]$ and $[u_*, 1]$. Meanwhile, Fig. 2(c) and Fig. 2(d) illustrate the optimal pricing functions in **Case-2** and **Case-3**, respectively. Note that by PM_{ϕ_*}, the highest-possible resource utilization levels in **Case-1**

and **Case-2** are 1, while in **Case-3** is $\bar{\rho}$ and $\bar{\rho} < 1$. This is consistent with our definition of $\bar{\rho}$ in Eq. (20).

For all the three cases, the optimal pricing function ϕ_* cannot be given in analytical forms. We argue that the computation of ω_*, u_*, and $\alpha_*(\mathcal{S})$ is light-weight and can be performed efficiently via various numerical methods such as bisection searching. More importantly, ω_* and u_* can be computed offline before the start of PM_ϕ. For more detailed discussions of how to quantify ω_*, u_*, and $\alpha_*(\mathcal{S})$, please refer to Appendix J.

6.3 An Example: Quadratic Supply Costs

To better show how to calculate $\omega_*, u_*, \alpha_*(\mathcal{S})$, and ϕ_* in Theorem 6.3, in this subsection we perform a case study based on quadratic supply costs. Let us assume $f(y) = \frac{1}{2}y^2$ (i.e., $f'(y) = y$), and thus $\underline{c} = f'(0) = 0$ and $\bar{c} = f'(1) = 1$. Based on the definitions of h and F_p in Section 3.1, we have

$$h(\underline{p}) = \begin{cases} \frac{1}{2}\underline{p}^2 & \text{if } p \in [\underline{c}, \bar{c}] \\ \underline{p} - \frac{1}{2} & \text{if } p \in (\bar{c}, +\infty) \end{cases}, \quad F_{\underline{p}}(\omega) = \underline{p}\omega - \frac{1}{2}\omega^2.$$

Therefore, the lower bound of the critical threshold ω can be written as a function of α as $F_{\underline{p}}^{-1}\left(\frac{1}{\alpha}h(\underline{p})\right) = \underline{p}\left(1 - \sqrt{1 - 1/\alpha}\right)$.

We next show how to compute $\omega_*, u_*, \alpha_*(\mathcal{S})$ in **Case-1**. When $f'(y) = y$, $\text{BVP}_1(\omega, u, \alpha)$ in Eq. (44) can be written as

$$\begin{cases} \varphi_1'(y) = \alpha \cdot \frac{\varphi_1(y)-y}{\varphi_1(y)}, y \in (\omega, u); \\ \varphi_1(\omega) = \underline{p}, \varphi_1(u) = 1. \end{cases} \tag{49}$$

The ODE in Eq. (49) can be solved by separation of variables [36]. Substituting the two boundary conditions into the solution of the ODE in Eq. (49) leads to the following equation of α, ω, and u:

$$\int_{\underline{p}/\omega}^{1/u} \frac{-\eta}{\eta^2 - \alpha\eta + \alpha} d\eta = \ln\left(\frac{u}{\omega}\right). \tag{50}$$

Based on Proposition 6.2, Eq. (50) has a unique root in $\alpha \geq 1$ for each given pair of $(\omega, u) \in [F_p^{-1}\left(\frac{1}{\alpha}h(\underline{p})\right), \underline{\rho}] \times (\omega, 1)$. Therefore, $\alpha = \Gamma_1(\omega, u)$ is well-defined through Eq. (50).

Similarly, $\text{BVP}_2(u, \alpha)$ in Eq. (45) can be written as

$$\begin{cases} \varphi_2'(y) = \alpha \cdot (\varphi_2(y) - y), y \in (u, 1); \\ \varphi_2(u) = \bar{c}, \varphi_2(1) = \bar{p}, \end{cases} \tag{51}$$

where the second boundary condition is written as an equality since the optimal design should have $\phi_*(1) = \bar{p}$, as indicated by Theorem 6.3. Solving the ODE in Eq. (51) is elementary. Substituting the two boundary conditions into the analytical solution of the ODE in Eq. (51) leads to the following equation of α and u:

$$\exp\left(\alpha(1-u)\right) = \frac{\alpha(1+\bar{p})+1}{\alpha(u+\bar{c})+1}. \tag{52}$$

Based on Proposition 6.2, Eq. (52) has a unique root in α for each given $u \in (0, 1)$, and thus $\alpha = \Gamma_2(u)$ is well defined through Eq. (52). Note that Eq. (52) is a simplification of Eq. (46) when $f'(y) = y$.

Combining the above results, when $f'(y) = y$, the calculation of ω_*, u_* and $\alpha_*(\mathcal{S})$ (which is denoted as α_* below for simplicity) in Eq. (47) can be written as the following system of three

(a) $\underline{p} = 0.3$
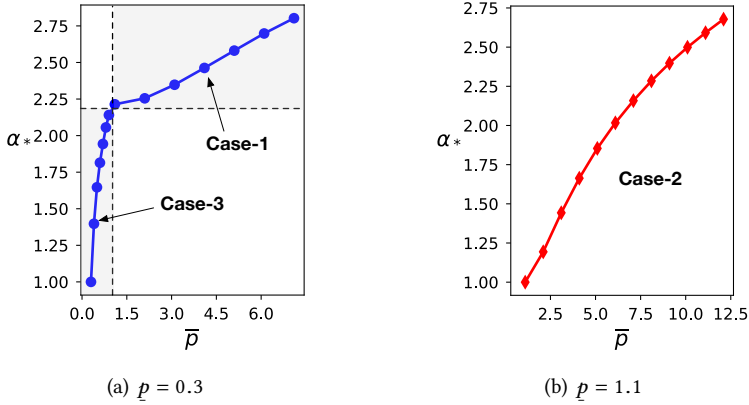
(b) $\underline{p} = 1.1$

Fig. 3. Relationship between $\alpha_*(\mathcal{S})$ and $\bar{p}$ when $f(y) = \frac{1}{2}y^2$.

equations:

$$\omega_* = \underline{p}\left(1 - \sqrt{1 - 1/\alpha_*}\right), \tag{53a}$$

$$\int_{\underline{p}/\omega_*}^{1/u_*} \frac{-\eta}{\eta^2 - \alpha_*\eta + \alpha_*}\,d\eta = \ln\left(\frac{u_*}{\omega_*}\right), \tag{53b}$$

$$\exp\left(\alpha_*(1 - u_*)\right) = \frac{\alpha_*(1 + \bar{p}) + 1}{\alpha_*(u_* + 1) + 1}. \tag{53c}$$

Based on Theorem 6.3, the above system of equations has a unique pair of solutions $(\omega_*, u_*, \alpha_*)$. In particular, the optimal competitive ratio for the quadratic cost setup can be written as

$$\alpha_*(\mathcal{S}) = \frac{h(\underline{p})}{F_{\underline{p}}(\omega_*)} = \begin{cases} \frac{\underline{p}^2}{2\underline{p}\omega_* - \omega_*^2} & \text{if } \underline{p} \in [\underline{c}, \bar{c}], \\ \frac{2\underline{p} - 1}{2\underline{p}\omega_* - \omega_*^2} & \text{if } \underline{p} \in (\bar{c}, +\infty). \end{cases}$$

Recall that for the linear cost setup in Corollary 4.5, $\alpha_*(\mathcal{S}) = \frac{1}{\omega_*}$.

When calculating the unique solution to Eq. (53), we can substitute ω_* given by Eq. (53a) into Eq. (53b), and then leverage the monotonicity of both sides of Eq. (53b) and Eq. (53c) to locate the unique values of α_* and u_*. Based on α_* and u_*, the optimal critical threshold ω_* can be easily calculated based on Eq. (53a). Theorem 6.3 shows that the optimal pricing function ϕ_* can be obtained as long as ω_*, u_* and α_* are given. For example, in **Case-1**, we can solve $\mathrm{BVP}_1(\omega_*, u_*, \alpha_*)$ in Eq. (49) to get the optimal pricing function φ_1^* in the interval $[\omega_*, u_*]$, and solve $\mathrm{BVP}_2(u_*, \alpha_*)$ in Eq. (51) to get the optimal pricing function φ_2^* in the interval $[u_*, 1]$. The complete optimal pricing function ϕ_* (in **Case-1**) can thus be determined.

We numerically solve the system of equations in Eq. (53) and illustrate the relationship between α_* and $\bar{p}$ in Fig. 3. Specifically, we fix $\underline{p} = 0.3$ and vary $\bar{p}$ from 0.3 to 7.1 in Fig. 3(a) so that our parameter setting covers both **Case-1** (i.e., $\bar{p} > 1$) and **Case-3** (i.e., $\bar{p} \le 1$). Meanwhile, we fix $\underline{p} = 1.1$ and vary $\bar{p}$ from 1.1 to 12.1 in Fig. 3(b) to simulate **Case-2** (i.e., $\underline{p} \ge 1$). Fig. 3 shows that $\alpha_*(\mathcal{S}) = 1$ when $\bar{p} = \underline{p}$, and is strictly increasing in $\bar{p} \in [\underline{p}, +\infty)$. This validates our results in Corollary 4.4. Note that the systems of equations for the calculation of α_* in **Case-2** and **Case-3** are different from Eq. (53). For the details, please refer to Appendix G.

7 CONCLUSION AND FUTURE WORK

We studied the mechanism design for general online resource allocation problems in the presence of supply costs and capacity limits. In this setting, the supplier can produce additional units of resources at increasing marginal cost, but with a stringent capacity limit; agents are strategic and may deliberately misreport their preference to be better off. The major contribution of the paper is the development of a unified incentive compatible online mechanism, with a principled method of constructing a universal "pricing function", that leads to the optimal competitive ratios for different setups including zero, linear and strictly-convex supply costs.

One of the drawbacks of the current design is that the optimal competitive ratio, as well as other key design parameters such as the optimal critical threshold ω_*, cannot be given in analytical forms. This brings difficulty in analyzing the properties of the competitive ratio w.r.t. important system parameters such as the lower and upper bounds of valuation densities (i.e., $\underline{p}$ and $\bar{p}$). It also remains unclear if there exists any analytical (and thus simple) pricing function so that the optimal competitive ratio can be approximately achieved. Meanwhile, extending the proposed techniques to other online optimization problems is also an interesting future direction.

8 ACKNOWLEDGEMENT

REFERENCES

[1] Ravi P. Agarwal and V. Lakshmikantham. 1993. *Uniqueness and Nonuniqueness Criteria for Ordinary Differential Equations.* World Scientific.

[2] Matthew. Andrews, Spyridon. Antonakopoulos, and Lisa. Zhang. 2016. Minimum-Cost Network Design with (Dis)economies of Scale. *SIAM J. Comput.* 45, 1 (2016), 49–66.

[3] Vladimir Arnold. 1973. *Ordinary differential equations.* MIT Press.

[4] Baruch Awerbuch, Yossi Azar, and Amir Epstein. 2005. The Price of Routing Unsplittable Flow. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing (STOC '05)*. ACM, New York, NY, USA, 57–66.

[5] Y. Azar, N. Buchbinder, T. H. Chan, S. Chen, I. R. Cohen, A. Gupta, Z. Huang, N. Kang, V. Nagarajan, J. Naor, and D. Panigrahi. 2016. Online Algorithms for Covering and Packing Problems with Convex Objectives. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, Vol. 00. 148–157.

[6] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. 2008. Item Pricing for Revenue Maximization. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC '08)*. ACM, New York, NY, USA, 50–59.

[7] Yair Bartal, Rica Gonen, and Noam Nisan. 2003. Incentive Compatible Multi Unit Combinatorial Auctions. In *Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge (TARK '03)*. ACM, New York, NY, USA, 72–87.

[8] Avrim Blum, Anupam Gupta, Yishay Mansour, and Ankit Sharma. 2011. Welfare and Profit Maximization with Production Costs. In *Proceedings of the 2011 IEEE 52Nd Annual Symposium on Foundations of Computer Science (FOCS '11)*. Washington, DC, USA, 77–86.

[9] Allan Borodin and Ran El-Yaniv. 1998. *Online Computation and Competitive Analysis.* Cambridge University Press, New York, NY, USA.

[10] Jonathan M. Borwein and Adrian S. Lewis. 2006. *Convex Analysis and and Nonlinear Optimization.* Springer.

[11] Niv Buchbinder and R. Gonen. 2015. Incentive Compatible Mulit-Unit Combinatorial Auctions: A Primal Dual Approach. *Algorithmica* 72 (2015), 167–190.

[12] Niv Buchbinder and Joseph Naor. 2006. Improved bounds for online routing and packing via a primal-dual approach. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*. 293–304.

[13] Niv Buchbinder and Joseph (Seffi) Naor. 2009. Online Primal-Dual Algorithms for Covering and Packing. *Mathematics of Operations Research* 34, 2 (May 2009), 270–286.

[14] Niv Buchbinder and Joseph (Seffi) Naor. 2009. The Design of Competitive Online Algorithms via a Primal-Dual Approach. *Found. Trends Theor. Comput. Sci.* 3, 2-3 (Feb. 2009), 93–263.

[15] Ying-Ju Chen and Jiawei Zhang. 2012. Design of price mechanisms for network resource allocation via price of anarchy. *Mathematical Programming* 131, 1 (01 Feb 2012), 333–364.

[16] George Christodoulou and Elias Koutsoupias. 2005. The Price of Anarchy of Finite Congestion Games. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing (STOC '05)*. ACM, New York, NY, USA, 67–73.

[17] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. 2017. Posted Price Mechanisms for a Random Stream of Customers. In *Proc. of the ACM Conference on Economics and Computation*.

[18] Sven de Vries and Rakesh Vohra. 2003. Combinatorial Auctions: A Survey. *INFORMS J. Comput.* 15, 3 (2003), 284–309.

[19] Nikhil R. Devanur and Thomas P. Hayes. 2009. The Adwords Problem: Online Keyword Matching with Budgeted Bidders Under Random Permutations. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC '09)*. ACM, New York, NY, USA, 71–78.

[20] Nikhil R. Devanur and Zhiyi Huang. 2017. Primal Dual Gives Almost Optimal Energy-Efficient Online Algorithms. *ACM Trans. Algorithms* 14, 1, Article 5 (Dec. 2017).

[21] Nikhil R. Devanur and Kamal Jain. 2012. Online Matching with Concave Returns. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing (STOC '12)*. ACM, New York, NY, USA, 137–144.

[22] Liran Einav, Chiara Farronato, Jonathan Levin, and Neel Sundaresan. 2018. Auctions versus Posted Prices in Online Markets. *Journal of Political Economy* 126, 1 (2018), 18.

[23] R. El-Yaniv, A. Fiat, R. M. Karp, and G. Turpin. 2001. Optimal search and one-way trading online algorithms. *Algorithmica (New York)* 30, 1 (2001), 101–139.

[24] Anupam Gupta, Ravishankar Krishnaswamy, and Kirk Pruhs. 2013. Online Primal-Dual for Non-linear Optimization with Applications to Speed Scaling. In *Approximation and Online Algorithms*. Springer Berlin Heidelberg, Berlin, Heidelberg, 173–186.

[25] Zhiyi Huang and Anthony Kim. 2015. Welfare Maximization with Production Costs: A Primal Dual Approach. In *Proceedings of the Twenty-sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '15)*. 59–72.

[26] Navendu Jain, Ishai Menache, Joseph (Seffi) Naor, and Jonathan Yaniv. 2014. A Truthful Mechanism for Value-Based Scheduling in Cloud Computing. *Theory of Computing Systems* 54, 3 (01 Apr 2014), 388–406.

[27] Bala Kalyanasundaram and Kirk R. Pruhs. 2000. An Optimal Deterministic Algorithm for Online B-matching. *Theor. Comput. Sci.* 233, 1-2 (Feb. 2000), 319–325.

[28] Qiulin Lin, Hanling Yi, John Pang, Minghua Chen, Adam Wierman, Michael Honig, and Yuanzhang Xiao. 2019. Competitive Online Optimization under Inventory Constraints. *Proc. ACM Meas. Anal. Comput. Syst.* 3, 1 (March 2019).

[29] Will Ma and David Simchi-Levi. 2019. Tight Weight-dependent Competitive Ratios for Online Edge-weighted Bipartite Matching and Beyond. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC '19)*. ACM, New York, NY, USA, 727–728.

[30] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory*. Oxford University Press.

[31] Aranyak Mehta. 2013. Online Matching and Ad Allocation. *Found. Trends Theor. Comput. Sci.* 8, 4 (Oct. 2013), 265–368.

[32] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. 2007. AdWords and Generalized Online Matching. *J. ACM* 54, 5, Article 22 (Oct. 2007).

[33] D. S. Mitrinović, J. E. Pečarić, and A. M. Fink. 1991. *Inequalities Involving Functions and Their Integrals and Derivatives*. Springer Netherlands.

[34] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press, Cambridge, UK.

[35] Lawrence Perko. 2001. *Differential Equations and Dynamical Systems*. Springer New York, New York, NY.

[36] Andrei D. Polyanin and Valentin F. Zaitsev. 2003. *Handbook of Exact Solutions for Ordinary Differential Equations*. Chapman & Hall/CRC.

[37] David Porter, Stephen Rassenti, Anil Roopnarine, and Vernon Smith. 2003. Combinatorial auction design. *Proceedings of the National Academy of Sciences* 100, 19 (2003), 11153–11157. https://doi.org/10.1073/pnas.1633736100

[38] Ryan Porter. 2004. Mechanism Design for Online Real-time Scheduling. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC '04)*. ACM, New York, NY, USA, 61–70.

[39] Xiaoqi Tan, Ablerto Leon-Garcia, and Danny .H.K. Tsang. 2019. Optimal Posted Prices for Online Resource Allocation with Supply Costs. In *Proceedings of ACM Workshop on Mathematical Performance Modeling and Analysis*. Phoenix, AZ, USA.

[40] Ruqu Wang. 1993. Auctions versus Posted-Price Selling. *The American Economic Review* 83, 4 (1993), 838–851.

[41] Z. Zhang, Z. Li, and C. Wu. 2017. Optimal Posted Prices for Online Cloud Resource Allocation. *Proc. ACM Meas. Anal. Comput. Syst.* 1, 1 (June 2017).

[42] Yunhong Zhou, Deeparnab Chakrabarty, and Rajan Lukose. 2008. Budget Constrained Bidding in Keyword Auctions and Online Knapsack Problems. In *Proceedings of the 17th International Conference on World Wide Web (WWW '08)*. ACM, New York, NY, USA, 1243–1244.

A PROOF OF THEOREM 4.1

The sufficient conditions in Theorem 4.1 are derived based on Proposition 3.1. Specifically, we need to prove that under PM_ϕ, i) the sequences of $\{P_n\}_{\forall n}$ and $\{D_n\}_{\forall n}$ are feasible, and ii) there exists some $k \in \mathcal{N}$ so that the initial inequality and the incremental inequality hold. Based on Proposition 3.2, the feasibility of $\{P_n\}_{\forall n}$ and $\{D_n\}_{\forall n}$ are trivial since the pricing function ϕ given by Eq. (19) is monotone. In the following we focus on the proof of the initial inequality and the incremental inequality.

(Initial Inequality) Let us assume $\omega = \sum_{n=1}^{k} r_n$, namely, k is the number of agents such that the total resource allocated equals ω. In the following we prove that the initial inequality $P_k \geq \frac{1}{\alpha} D_k$ holds. Based on Eq. (22), ω is the resource utilization threshold such that $F_{\underline{p}}(\omega) = \underline{p}\omega - \bar{f}(\omega) \geq \frac{1}{\alpha} h(\underline{p})$. Thus, we have

$$\underline{p} \cdot \left(\sum_{n=1}^{k} r_n \right) - \bar{f} \left(\sum_{n=1}^{k} r_n \right) \geq \frac{1}{\alpha} h(\underline{p}),$$

Since $\alpha \geq 1$ and $\hat{\gamma}_n \geq 0$ for all $n \in \mathcal{N}$, the above inequality leads to the following one

$$\left(1 - \frac{1}{\alpha} \right) \sum_{n=1}^{k} \hat{\gamma}_n + \sum_{n=1}^{k} \underline{p} r_n - \bar{f} \left(\sum_{n=1}^{k} r_n \right) \geq \frac{1}{\alpha} h(\underline{p}).$$

The pricing function given by Eq. (19) indicates that the requirements of all the requests will be satisfied as long as the resource utilization is below ω. For this reason, we have $\hat{y}_k = \sum_{n=1}^{k} r_n = \omega$, which represents the total resource utilization after processing request k. Therefore, we have

$$\sum_{n=1}^{k} \hat{\gamma}_n + \sum_{n=1}^{k} \underline{p} r_n - \bar{f}(\hat{y}_k) \geq \frac{1}{\alpha} \left(\sum_{n=1}^{k} \hat{\gamma}_n + h(\underline{p}) \right), \tag{54}$$

Note that the left-hand-side of Eq. (54) is equal to P_k since $\sum_{n=1}^{k} \hat{\gamma}_n + \sum_{n=1}^{k} \underline{p} r_n - \bar{f}(\hat{y}_k) = \sum_{n=1}^{k} v_n - \bar{f}(\hat{y}_k) = P_k$, where the first equality comes from $v_n = \hat{\gamma}_n + \hat{p}_{n-1} r_n = \hat{\gamma}_n + \underline{p} r_n$. Here, $\hat{p}_{n-1} = \underline{p}$ holds for all $n = \{1, 2, \cdots, k\}$ because the posted prices for all these agents are $\underline{p}$ (i.e., the flat-segment). Meanwhile, based on the objective of Problem (10), the terms inside the parenthesis of Eq. (54) equal D_k, namely, $D_k = \sum_{n=1}^{k} \hat{\gamma}_n + h(\underline{p})$. Therefore, Eq. (54) indicates that $P_k \geq \frac{1}{\alpha} D_k$ holds. We thus complete the proof of the initial inequality.

(Incremental Inequality) We now prove that the incremental inequality $P_n - P_{n-1} \geq \frac{1}{\alpha}(D_n - D_{n-1})$ holds for all $n \in \{k+1, \cdots, N\}$. Note that if agent n is rejected (i.e., the demand of agent n is not satisfied), then $P_n = P_{n-1}$ and $D_n = D_{n-1}$. In this case, the incremental inequality $P_n - P_{n-1} \geq \frac{1}{\alpha}(D_n - D_{n-1})$ holds for all $\alpha \geq 1$. Therefore, in the following we only focus on the case when the requirement of agent n is satisfied. Eq. (23) shows that the following differential inequality holds

$$\phi(\hat{y}_{n-1}) - \bar{f}'(\hat{y}_{n-1}) \geq \frac{1}{\alpha} \cdot h'\big(\phi(\hat{y}_{n-1})\big) \cdot \phi'(\hat{y}_{n-1}) \tag{55}$$

for all $\hat{y}_{n-1} \in [\omega, \bar{\rho})$. Since $\hat{y}_n - \hat{y}_{n-1} = r_n$, based on the Taylor series we have

$$\bar{f}(\hat{y}_n) = \bar{f}(\hat{y}_{n-1}) + \bar{f}'(\hat{y}_{n-1})(\hat{y}_n - \hat{y}_{n-1}) + O(r_n^2),$$

$$h\big(\phi(\hat{y}_n)\big) = h\big(\phi(\hat{y}_{n-1})\big) + h'\big(\phi(\hat{y}_{n-1})\big)\big(\phi(\hat{y}_n) - \phi(\hat{y}_{n-1})\big) + O\left(\big(\phi(\hat{y}_n) - \phi(\hat{y}_{n-1})\big)^2\right),$$

$$\phi(\hat{y}_n) = \phi(\hat{y}_{n-1}) + \phi'(\hat{y}_{n-1})(\hat{y}_n - \hat{y}_{n-1}) + O(r_n^2).$$

Based on Eq. (55) and the above Taylor series, we have

$$\phi\left(\hat{y}_{n-1}\right) - \frac{\bar{f}\left(\hat{y}_n\right) - \bar{f}\left(\hat{y}_{n-1}\right)}{\hat{y}_n - \hat{y}_{n-1}} - O(r_n)$$

$$\geq \frac{1}{\alpha} \cdot \left(\frac{h\left(\phi\left(\hat{y}_n\right)\right) - h\left(\phi\left(\hat{y}_{n-1}\right)\right)}{\phi\left(\hat{y}_n\right) - \phi\left(\hat{y}_{n-1}\right)} - O(r_n)\right) \cdot \left(\frac{\phi\left(\hat{y}_n\right) - \phi\left(\hat{y}_{n-1}\right)}{\hat{y}_n - \hat{y}_{n-1}} - O(r_n)\right),$$

Assumption 2 states that r_n is sufficiently small, i.e., $r_n \to 0$. Thus, the following inequality holds:

$$\phi\left(\hat{y}_{n-1}\right)\left(\hat{y}_n - \hat{y}_{n-1}\right) - \left(\bar{f}\left(\hat{y}_n\right) - \bar{f}\left(\hat{y}_{n-1}\right)\right) \geq \frac{1}{\alpha}\left(h\left(\phi\left(\hat{y}_n\right)\right) - h\left(\phi\left(\hat{y}_{n-1}\right)\right)\right).$$

Since $\hat{\gamma}_n \geq 0$ and $1 - \frac{1}{\alpha} \geq 0$, we have

$$\left(1 - \frac{1}{\alpha}\right)\hat{\gamma}_n + \left(\phi(\hat{y}_{n-1})(\hat{y}_n - \hat{y}_{n-1}) - \left(\bar{f}(\hat{y}_n) - \bar{f}(\hat{y}_{n-1})\right)\right) \geq \frac{1}{\alpha}\left(h\left(\phi\left(\hat{y}_n\right)\right) - h\left(\phi\left(\hat{y}_{n-1}\right)\right)\right), \quad (56)$$

Note that after processing agent n, the change in the primal objective is given by $P_n - P_{n-1} = v_n - \left(\bar{f}(\hat{y}_n) - \bar{f}(\hat{y}_{n-1})\right) = \hat{\gamma}_n + \phi(\hat{y}_{n-1})\left(\hat{y}_n - \hat{y}_{n-1}\right) - \left(\bar{f}(\hat{y}_n) - \bar{f}(\hat{y}_{n-1})\right)$, where the second equality is due to $v_n = \hat{\gamma}_n + \hat{p}_{n-1}r_n = \hat{\gamma}_n + \phi(\hat{y}_{n-1})(\hat{y}_n - \hat{y}_{n-1})$. Similarly, the change in the dual objective is given by $D_n - D_{n-1} = \hat{\gamma}_n + h(\hat{p}_n) - h(\hat{p}_{n-1}) = \hat{\gamma}_n + h\left(\phi\left(\hat{y}_n\right)\right) - h\left(\phi\left(\hat{y}_{n-1}\right)\right)$. Therefore, Eq. (56) indicates that the incremental inequality $P_n - P_{n-1} \geq \frac{1}{\alpha}\left(D_n - D_{n-1}\right)$ holds for all $n \in \{k+1, \cdots, N\}$ when $\hat{y}_{n-1} \in [\omega, \bar{\rho})$.

Note that we need to guarantee that the incremental inequality holds at each round for all $\hat{y}_{n-1} \in [\omega, \bar{\rho}]$, where $\bar{\rho}$ represents the maximum resource utilization level. Therefore, we need to consider the boundary condition of $\phi(y)$ when $y = \bar{\rho}$. First, $\phi(\omega) = \underline{p}$ is obvious as the pricing function ϕ must be continuous. Taking integration of both sides of Eq. (55) leads to

$$\int_{\omega}^{\rho} \phi(y)dy - \bar{f}(\rho) + \bar{f}(\omega) \geq \frac{1}{\alpha}\Big(h\big(\phi(\rho)\big) - h\big(\underline{p}\big)\Big).$$

Meanwhile, based on the flat-segment, we have $\underline{p}\omega - \bar{f}(\omega) \geq \frac{1}{\alpha}h(\underline{p})$. Thus, we have

$$\underline{p}\omega + \int_{\omega}^{\rho} \phi(y)dy - \bar{f}(\rho) \geq \frac{1}{\alpha}h\big(\phi(\rho)\big), \quad (57)$$

which is the same as Eq. (31). When $\rho = \bar{\rho}$, the above inequality also holds, and thus we have

$$\underline{p}\omega + \int_{\omega}^{\bar{\rho}} \phi(y)dy - f(\bar{\rho}) \geq \frac{1}{\alpha}h\big(\phi(\bar{\rho})\big). \quad (58)$$

On the other hand, PM_{ϕ} is α-competitive indicates that the following inequality must hold:

$$\underline{p}\omega + \int_{\omega}^{\bar{\rho}} \phi(y)dy - f(\bar{\rho}) \geq \frac{1}{\alpha}h(\bar{p}). \quad (59)$$

Therefore, if we have $\phi(\bar{\rho}) \geq \bar{p}$, then Eq. (58) guarantees that Eq. (59) always holds. This means, the incremental inequality $P_n - P_{n-1} \geq \frac{1}{\alpha}\left(D_n - D_{n-1}\right)$ holds at each round for all $\hat{y}_{n-1} \in [\omega, \bar{\rho}]$ as long as the pricing function satisfies

$$\begin{cases} \phi(y) - f'(y) \geq \frac{1}{\alpha} \cdot h'\left(\phi(y)\right) \cdot \phi'(y), y \in (\omega, \bar{\rho}), \\ \phi(\omega) = \underline{p}, \phi(\bar{\rho}) \geq \bar{p}. \end{cases}$$

We thus complete the proof of the incremental inequality in Eq. (23).

Summarizing our above proofs regarding the initial and incremental inequalities, we complete the proof of Theorem 4.1.

B PROOF OF THEOREM 4.2

Our proof of Theorem 4.2 is organized as follows. We first give a set of less restrictive conditions in Theorem B.1, based on which we give Theorem B.2 to show the existence of strictly-increasing allocation functions and pricing functions. Based on Theorem B.1 and Theorem B.2, Theorem 4.2 directly follows.

B.1 General Necessary Conditions

Theorem B.1 below gives a set of necessary conditions to guarantee the existence of an α-competitive online algorithm.

THEOREM B.1. *Given a convex setup $\mathcal{S}$, if there exists an α-competitive online algorithm, then there must exist a parameter ω which satisfies*

$$F_{\underline{p}}(\omega) \geq \frac{1}{\alpha} h(\underline{p}) \text{ and } 0 \leq \omega \leq \underline{\rho} \tag{60}$$

so that the following claims hold simultaneously:

- ***NC1)**: There exists a case when the total resource utilization is ω and all the accepted agents have the same valuation density $\underline{p}$.*
- ***NC2)**: There exists a non-decreasing function $\psi(p)$ that satisfies*

$$\begin{cases} \underline{p}\omega + \int_{\underline{p}}^{p} \eta\psi'(\eta)d\eta - f(\psi(p)) \geq \frac{1}{\alpha}h(p), p \in (\underline{p}, \bar{p}), \\ \psi(\underline{p}) = \omega, \psi(\bar{p}) \leq \bar{\rho}, \end{cases} \tag{61}$$

where $\psi(p)$ is a monotone mapping from $p \in [\underline{p}, \bar{p}]$ to $\psi(p) \in [\omega, \bar{\rho}]$.

PROOF. The proof of this theorem is based on constructing a resource utilization level ω and a monotone function ψ for any α-competitive online algorithm under two special arrival instances. We organize our proofs of **NC1** and **NC2** as follows.

(**Proof of NC1**) By definition, an α-competitive online algorithm must achieve an online social welfare that is no less than $\frac{1}{\alpha}$ of its offline counterpart for all possible arrival instances. To prove this, we first construct an arrival instance $\mathcal{A}_{\underline{p}}$ as follows: *there are K agents in $\mathcal{A}_{\underline{p}}$ with the same valuation density $\underline{p}$, and all the requests ask for resources with the same requirement Δ, where Δ is infinitely small and $K\Delta = \underline{\rho}$*. Based on Eq. (20), the optimal offline strategy under this arrival instance is to accept all the requests and the resulting optimal social welfare is given by $S_{\text{offline}}(\mathcal{A}_{\underline{p}}) = K\underline{p}\Delta - \bar{f}(K\Delta) = \underline{p} \cdot \underline{\rho} - \bar{f}(\underline{\rho}) = h(\underline{p})$. For any α-competitive online algorithm (including PM_ϕ), let us denote the final resource utilization level by $\omega \in [0, 1]$, then the online social welfare is given by $S_{\text{online}}(\mathcal{A}_{\underline{p}}) = \underline{p}\omega - \bar{f}(\omega) = F_{\underline{p}}(\omega)$. Based on $S_{\text{offline}}(\mathcal{A}_{\underline{p}})$ and $S_{\text{online}}(\mathcal{A}_{\underline{p}})$, the online algorithm is α-competitive indicates that $S_{\text{online}}(\mathcal{A}_{\underline{p}}) = F_{\underline{p}}(\omega) \geq \frac{1}{\alpha} S_{\text{offline}}(\mathcal{A}_{\underline{p}}) = \frac{1}{\alpha} h(\underline{p})$. Meanwhile, ω must be less than or equal to $\underline{\rho}$ since otherwise the marginal cost is larger than the valuation density $\underline{p}$, i.e., $\omega \in [0, \underline{\rho}]$. Combining these results, we can conclude that under the arrival instance $\mathcal{A}_{\underline{p}}$, if there is an α-competitive online algorithm, then this algorithm must accept all the agents so that the total resource utilization is ω. We thus complete the proof of **NC1**.

(**Proof of NC2**) Our proof of **NC2** is based on constructing a special arrival instance $\mathcal{A}_p$ such that any α-competitive online algorithm must satisfy the inequality in Eq. (61) in order to achieve at least $\frac{1}{\alpha}$ of the offline optimal social welfare in hindsight. Specifically, for any $p \in [\underline{p}, \bar{p}]$, $\mathcal{A}_p$ is constructed as follows. *We assume that $\mathcal{A}_p$ consists of two groups of arrivals. First, there is a group of identical agents whose valuation densities are $\underline{p}$ and the total resource requirements of these identical agents are ω. In the following, we refer to the identical agents in this group by **Group-$(\underline{p}, \omega)$**. Second, let us assume for each $\eta \in [\underline{p}, p]$, there is a group of agents parameterized by η, where all the agents in*

group η are identical and has a total requirement of $h'(\eta)$ (i.e., each customer's requirement is infinitely small and the total of them is $h'(\eta)$). Meanwhile, the valuation densities of all the agents in group η are the same value of η, namely, the total valuation of all the agents in this group is given by $v_\eta = \eta h'(\eta)$. We refer to the agents in the second continuum of groups by **Groups**-$(\underline{p}, p)$. Recall that based on Eq. (12), h' is given by

$$h'(\eta) = \begin{cases} f'^{-1}(\eta) & \text{if } \eta \in [\underline{c}, \bar{c}], \\ 1 & \text{if } \eta \in (\bar{c}, +\infty), \end{cases} \tag{62}$$

Therefore, $h'(\eta)$ is the maximum units of resource that can be supplied when the marginal cost is η per unit.

For a given arrival instance $\mathcal{A}_p$ with any $p \in [\underline{p}, \bar{p}]$, the social welfares achieved by the optimal offline algorithm and the α-competitive online algorithm are given as follows:

- **Offline**: the optimal offline result in hindsight is to allocate $h'(p)$ units of resources to the requests in the last group, i.e., group p in **Groups**-$(\underline{p}, p)$, and none to all the previous requests. The optimal social welfare is thus

$$S_{\text{offline}}(\mathcal{A}_p) = p h'(p) - f(h'(p)) = h(p),$$

  where the second equality comes from the Fenchel duality of h. In fact, if we substitute h' from Eq. (62) into the above $S_{\text{offline}}(\mathcal{A}_p)$, the equation follows.

- **Online**: for the α-competitive online algorithm, let $y = \psi(\eta)$ denote the total resource utilization after processing the customers in group $\eta \in [\underline{p}, p]$, and thus $\psi(\eta)$ represents the resources allocated to all the groups of customers in $[\underline{p}, \eta]$. According to **NC1**, all the requests in **Group**-$(\underline{p}, \omega)$ must be satisfied. Therefore, we have $\psi(\underline{p}) = \omega$, meaning that when $\eta = \underline{p}$, the total resource utilization is ω, which equals the total resource requirements of all the requests in **Group**-$(\underline{p}, \omega)$. Meanwhile, $\psi(\eta)$ is monotonically non-decreasing in $\eta \in [\underline{p}, p]$ since the resource allocations are irrevocable. The social welfare achieved by this online algorithm is thus the total valuation minus the total cost, namely,

$$S_{\text{online}}(\mathcal{A}_{\underline{p}}) = \underline{p}\omega + \int_{\psi(\underline{p})}^{\psi(p)} \eta d(\psi(\eta)) - f(\psi(p)) = \underline{p}\omega + \int_{\underline{p}}^{p} \eta \psi'(\eta) d\eta - f(\psi(p)).$$

Based on $S_{\text{offline}}(\mathcal{A}_p)$ and $S_{\text{online}}(\mathcal{A}_p)$, the online algorithm is α-competitive means that the following inequality

$$\underline{p}\omega + \int_{\underline{p}}^{p} \eta \psi'(\eta) d\eta - f(\psi(p)) \geq \frac{1}{\alpha} h(p) \tag{63}$$

holds for all $p \in [\underline{c}, \bar{p}]$. According to the definition of ψ and $\bar{\rho}$, we have $\psi(p) \leq \bar{\rho}, \forall p \in [\underline{c}, \bar{p}]$, and thus $\psi(\bar{p}) \leq \bar{\rho}$ holds as well. Therefore, if there exists an α-competitive online algorithm, then there must exist a non-decreasing function $\psi(p)$ that satisfies Eq. (61). □

B.2 Existence of Strictly-Increasing Allocation Functions and Pricing Functions

Based on **NC2** in Theorem B.1, Theorem B.2 below shows that it is necessary to have a strictly-increasing segment $[\omega, \bar{\rho}]$ for ϕ.

THEOREM B.2. *Given a convex setup $\mathcal{S}$, if there exists an α-competitive online algorithm, then there must exist a positive value $\omega \in [F_{\underline{p}}^{-1}\left(\frac{1}{\alpha} h(\underline{p})\right), \underline{\rho}]$ such that the following claims hold:*

- There exists a strictly-increasing function $\psi(p)$ that satisfies

$$
\begin{cases}
\psi'(p) = \frac{1}{\alpha} \cdot \frac{h'(p)}{p - f'(\psi(p))}, p \in (\underline{p}, \bar{p}), \\
\psi(\underline{p}) = \omega, \psi(\bar{p}) \le \bar{\rho}.
\end{cases}
\tag{64}
$$

- There exists a strictly-increasing function $\varphi(y)$ that satisfies

$$
\begin{cases}
\varphi'(y) = \alpha \cdot \frac{\varphi(y) - f'(y)}{h'(\varphi(y))}, y \in (\omega, \bar{\rho}), \\
\varphi(\omega) = \underline{p}, \varphi(\bar{\rho}) \ge \bar{p}.
\end{cases}
\tag{65}
$$

- $\psi(p)$ and $\varphi(y)$ are inverse to each other, i.e., $\psi = \varphi^{-1}$ or $\varphi = \psi^{-1}$.

PROOF. We organize our proofs of Eq. (64) and Eq. (65) as follows.

(**Proof of Eq.** (64)) We first prove the first bullet of Theorem B.2, namely, there exists a strictly-increasing function $\psi(p)$ that satisfies Eq. (64) as long as there exists an α-competitive online algorithm. Note that when $\psi(\underline{p}) = \omega$, we can derive based on the left-hand-side of Eq. (63) that

$$
\underline{p}\omega + \int_{\underline{p}}^{p} \eta \psi'(\eta) d\eta - f(\psi(p)) = p\psi(p) - \int_{\underline{p}}^{p} \psi(\eta) d\eta - f(\psi(p)).
$$

Based on Theorem B.1 and the above equivalent transformation, we know that there always exists a non-decreasing function $\psi(p)$ that satisfies the following inequality

$$
\psi(p) \ge \frac{h(p) + \alpha f(\psi(\underline{p}))}{\alpha p} + \frac{1}{p} \int_{\underline{p}}^{p} \left(\psi(\eta) + f'(\psi(\eta)) \right) d\eta
\tag{66}
$$

and the boundary conditions: $\{\psi(\underline{p}) = \omega, \psi(\bar{p}) \le \bar{\rho}\}$. For ease of exposition, let us write the right-hand-side of Eq. (66) in a more compact form as follows:

$$
\psi(p) \ge A(p) + B(p) \int_{\underline{p}}^{p} Q(\psi(\eta)) d\eta, p \in [\underline{p}, \bar{p}],
\tag{67}
$$

where $A(p)$, $B(p)$, and $Q(\psi(\eta))$ are given by

$$
A(p) = \frac{h(p) + \alpha f(\psi(\underline{p}))}{\alpha p},
$$
$$
B(p) = \frac{1}{p},
$$
$$
Q(\psi(\eta)) = \psi(\eta) + f'(\psi(\eta)).
$$

It can be observed that Eq. (67) is a nonlinear generalization of the standard Gronwell inequality [33]. Thus, there must exist a tightest lower bound function $\psi_L(p)$ such that

$$
\psi(p) \ge \psi_L(p), p \in [\underline{p}, \bar{p}],
$$

namely, ψ_L is defined to be the tightest lower bound of all the ψ's that satisfy Eq. (67) with $\{\psi(\underline{p}) = \omega, \psi(\bar{p}) \le \bar{\rho}\}$. We emphasize that the function ψ_L is well-defined as there exists at least one non-decreasing function $\psi(p)$ that satisfies Eq. (67) with $\{\psi(\underline{p}) = \omega, \psi(\bar{p}) \le \bar{\rho}\}$. Meanwhile, when Q is in some special forms (e.g., affine), ψ_L can be analytically given based on A, B, and Q [33]. Here, the nonlinearity of Q makes it difficult (if not impossible) to derive the exact form of ψ_L. To aid our following proof, let us define Ω_ψ by

$$
\Omega_\psi \triangleq \{ \psi | \psi \text{ is a strictly-increasing function that satisfies Eq. (64)} \}.
$$

The first bullet in Theorem B.2 argues that Ω_ψ is well-defined and non-empty. To prove this, in what follows we show that $\psi_L \in \Omega_\psi$, namely, ψ_L is a strictly-increasing function that satisfies Eq. (67) by equality and $\{\psi_L(\underline{p}) = \omega, \psi_L(\bar{p}) \le \bar{\rho}\}$.

Recall that ψ_L is defined to be the tightest lower bound of all the ψ's that satisfy Eq. (67) with $\{\psi(\underline{p}) = \omega, \psi(\bar{p}) \le \bar{\rho}\}$, and thus $\psi_L(\bar{p}) \le \bar{\rho}$ always holds. On the other hand, if $\psi_L(\bar{p}) \le \bar{\rho}$ does not hold, then it is obvious that there will be no solution to Eq. (61) since $\psi(\underline{p}) \ge \psi_L(\underline{p}) > \bar{\rho}$ holds for all $p \in [\underline{p}, \bar{p}]$, which thus contradicts the necessary condition **NC2** given by Theorem B.1. For this reason, we only need to prove that $\psi_L(p)$ is strictly-increasing in $p \in [\underline{p}, \bar{p}]$ and $\psi_L(\underline{p}) = \omega$. In fact, the definition of ψ_L indicates that the following equation of $\psi_L(p)$ holds for all $p \in [\underline{p}, \bar{p}]$:

$$\psi_L(p) = A(p) + B(p) \int_{\underline{p}}^{p} Q\big(\psi_L(\eta)\big) d\eta. \tag{68}$$

This is because, if there exists any $\hat{p} \in [\underline{p}, \bar{p}]$ such that Eq. (68) does not hold at $p = \hat{p}$, then we can always lower the value of $\psi_L(\hat{p})$ to reach the equality and make the smaller value as the new output of $\psi_L(\hat{p})$. Based on Eq. (68), we can prove that $\psi_L(p)$ is strictly-increasing in $p \in [\underline{p}, \bar{p}]$. Our proof is based on contradiction and the details are as follows. As illustrated by the dashed line in Fig. 4, if we assume that ψ_L is not strictly monotone in $p \in [\underline{p}, \bar{p}]$, then there must exist a point, say $p_0 \in [\underline{p}, \bar{p}]$, so that $\psi_L'(p_0) = 0$. Recall that $h'(p) > 0$ holds for all $p \in [\underline{p}, \bar{p}]$ as the conjugate function $h(p)$ is strictly-convex in $p \in [\underline{p}, \bar{p}]$. Thus, we can derive from Eq. (68) that

$$\psi_L'(p) = \frac{1}{\alpha} \cdot \frac{h'(p)}{p - f'(\psi_L(p))} \ne 0, p \in [\underline{p}, \bar{p}],$$

which contradicts $\psi_L'(p_0) = 0$. Therefore, ψ_L is either strictly-increasing or strictly-decreasing. Note that our design of $\omega \in [F_p^{-1}(\frac{1}{\alpha}h(\underline{p})), \underline{\rho}]$ guarantees that $\underline{p} - f'(\omega) > 0$. Therefore, based on the Picard–Lindelöf theorem [3, 35], there always exists a unique strictly-increasing function ψ_L that satisfies the following first-order initial value problem (IVP)[8]:

$$\begin{cases} \psi_L'(p) = \frac{1}{\alpha} \cdot \frac{h'(p)}{p - f'(\psi_L(p))} \ne 0, p \in (\underline{p}, \bar{p}), \\ \psi_L(\underline{p}) = \omega. \end{cases} \tag{69}$$

As a result, ψ_L must be strictly-increasing. Recall that $\psi_L(\bar{p}) \le \bar{\rho}$ always holds, together with Eq. (69), we can conclude that $\psi_L \in \Omega_\psi$, which means that Ω_ψ is well-defined and non-empty. Thus, we complete the proof of Eq. (64).

REMARK 6. *Our above proof shows that the tightest lower bound function $\psi_L(p)$ is strictly-increasing in $p \in [\underline{p}, \bar{p}]$ and satisfies Eq. (64). If we define similarly that ψ_U is the tightest upper bound of all the ψ's that satisfy Eq. (67) with $\{\psi(\underline{p}) = \omega, \psi(\bar{p}) \le \bar{\rho}\}$, then we can prove a stronger result: for any given $\alpha \ge 1$, if there exists an α-competitive online algorithm, then there exist infinitely-many strictly-increasing functions that satisfy Eq. (64), unless $\psi_L = \psi_U$ since in this case Ω_ψ is a singleton, i.e., $\Omega_\psi = \{\psi_L\}$. Geometrically, when α is larger, the shaded area in Fig. 4 will be expanded as $\psi_U(\bar{p}) - \psi_L(\bar{p})$ becomes larger. In contrast, a smaller α indicates that ψ_U and ψ_L are closer to each other, i.e., $\psi_U(\bar{p}) - \psi_L(\bar{p})$ is smaller, and in the extreme case when $\psi_U = \psi_L$, the corresponding value of α is the minimum competitive ratio that any online algorithm could possibly achieve. For the detailed discussion of the optimal/minimum competitive ratio, please refer to the proof of Theorem 4.3 in Section 6.*

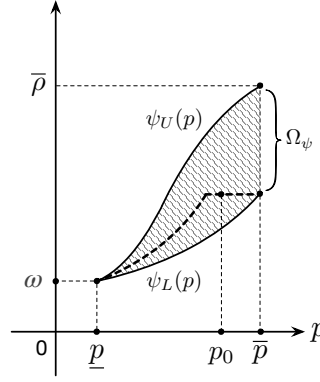[8]We give a brief overview of first-order IVPs in Appendix E.1 and Appendix E.2.

Fig. 4. Illustration of Ω_ψ for given $\underline{p}, \bar{p}, \omega$, and $\underline{\rho}$. In the figure, ψ_L and ψ_U denote the tightest lower and upper boundaries of Ω_ψ, respectively.

(**Proof of Eq.** (65)) Now we are ready to prove Eq. (65). Since $\psi(p)$ is strictly-increasing, the inverse of $\psi(p)$ is well-defined. Let us denote the inverse of $y = \psi(p)$ by $p = \psi^{-1}(y) \triangleq \varphi(y)$. The derivative of $\psi(p) = \varphi^{-1}(p)$ can be written as

$$\psi'(p) = \frac{1}{\varphi'\left(\varphi^{-1}(p)\right)} = \frac{1}{\varphi'(y)}. \tag{70}$$

Substituting Eq. (70) into Eq. (64) leads to the following ODE

$$\frac{1}{\varphi'(y)} = \frac{1}{\alpha} \cdot \frac{h'(\varphi(y))}{\varphi(y) - f'(y)}.$$

Therefore, we have the following ODE in terms of φ:

$$\varphi'(y) = \alpha \cdot \frac{\varphi(y) - f'(y)}{h'\left(\varphi(y)\right)}, y \in (\omega, \bar{\rho}),$$

which is exactly the ODE in Eq. (65). Based on the two boundary conditions of $\psi(p)$ given in Eq. (61), the BVP in Eq. (65) directly follows. Based on the above proofs of Eq. (64) and Eq. (65), we thus complete the proof of Theorem B.2. □

Based on **NC1** in Theorem B.1 and the two BVPs in Theorem B.2, Theorem 4.2 follows.

C PROOF OF COROLLARY 4.5

The proof is simple and intuitive by following our previous analysis of the ODE with boundary conditions.

(**Zero Cost**: $q = 0$) When $f(y) = 0$, i.e., there is no supply cost, the extended cost function is given as follows:

$$\bar{f}(y) = \begin{cases} 0 & \text{if } y \in [0, 1], \\ +\infty & \text{if } y \in (1, +\infty), \end{cases}$$

Recall that the conjugate function $h(p)$ is given by $h(p) = \max_{y \geq 0} py - \bar{f}(y) = \max_{y \geq 0} F_p(y)$, which indicates that $h(p) = p, \forall p \in [0, +\infty)$. Therefore, the BVP given by Eq. (25) can be written as

follows:

$$\begin{cases} \varphi'(y) = \alpha \cdot \varphi(y), y \in (\omega, \bar{\rho}), \\ \varphi(\omega) = \underline{p}, \varphi(1) \geq \bar{p}. \end{cases}$$

Based on Eq. (60), we know that $\omega \geq \frac{1}{\alpha}$. Solving the above ODE leads to the following solution

$$\varphi(y) = \delta \cdot \exp(\alpha y),$$

where δ can be any real scalar. Based on the two boundary conditions, to minimize α, we set

$$\delta \exp(\alpha \omega) = \underline{p}, \delta \exp(\alpha) = \bar{p}, \alpha = \frac{1}{\omega}.$$

Solving the above equations leads to $\delta = \underline{p}/e$, where e denotes the base of the natural logarithm. Therefore, the final optimal pricing function is

$$\phi_*(y) = \begin{cases} \underline{p} & \text{if } y \in [0, \omega_*), \\ \underline{p} \cdot \exp(y/\omega_* - 1) & \text{if } y \in [\omega_*, 1], \\ +\infty & \text{if } y \in (1, +\infty), \end{cases}$$

where ω_* and α_* are given by

$$\alpha_* = 1 + \ln\left(\frac{\bar{p}}{\underline{p}}\right), \omega_* = \frac{1}{1 + \ln(\bar{p}/\underline{p})}.$$

(**Linear Cost:** $q > 0$) When $f'(y) = q$ and $q > 0$, namely $f(y) = qy$ when $y \in [0, 1]$. The conjugate function $h(p)$ is given by

$$h(p) = \begin{cases} 0 & \text{if } p \in [0, q], \\ p - q & \text{if } p \in (q, +\infty). \end{cases}$$

Based on Eq. (60), we have $\underline{p}\omega - f(\omega) \geq \frac{1}{\alpha}(\underline{p} - f(1))$, which thus indicates that $\omega \geq \frac{1}{\alpha}$. The differential equation of interest is thus given by

$$\begin{cases} \varphi(y) - f'(y) \geq \frac{1}{\alpha} \cdot \varphi'(y) \cdot h'(\varphi(y)), \forall y \in (\omega, 1), \\ \varphi(\omega) = \underline{p}, \varphi(1) = \bar{p}, \end{cases}$$

which leads to the following general solution: $\varphi(y) = \delta \cdot \exp(\alpha y) + q$. Substituting the two boundary conditions into the general solution leads to $\varphi(\omega) = \delta \cdot \exp(\alpha \omega) + q = \underline{p}, \varphi(1) = \delta e^\alpha + q = \bar{p}$. To minimize α, it suffices to set $\omega = 1/\alpha$ and thus we have

$$\phi_*(y) = \begin{cases} \underline{p} & \text{if } y \in [0, \omega_*), \\ (\underline{p} - q) \cdot \exp(y/\omega_* - 1) + q & \text{if } y \in [\omega_*, 1], \\ +\infty & \text{if } y \in (1, +\infty), \end{cases}$$

where ω_* and α_* are given by

$$\alpha_* = 1 + \ln\left(\frac{\bar{p} - q}{\underline{p} - q}\right), \omega_* = \frac{1}{1 + \ln\left(\frac{\bar{p} - q}{\underline{p} - q}\right)}.$$

Combining the results regarding $q = 0$ and $q > 0$, we complete the proof of Corollary 4.5.

D PROOF OF COROLLARY 5.1

Based on Theorem 4.1 and Eq. (32), the worst-case competitive ratio of PM_ϕ satisfies:

- When the final resource utilization $\rho \in [0, \omega]$, the worst-case competitive ratio is

$$\alpha_1 = \frac{h(\underline{p})}{F_{\underline{p}}(\omega)}.$$

  This corresponds to the arrival instance that all agents have the minimum valuation density $\underline{p}$.
- When the final resource utilization $\rho = \bar{\rho}_\varphi$, the worst-case competitive ratio is

$$\alpha_2 = \frac{h(\bar{p})}{\underline{p}\omega + \int_\omega^{\bar{\rho}_\varphi} \varphi(y)dy - f\left(\bar{\rho}_\varphi\right)}.$$

  This corresponds to the worst-case scenario when the first group of agents are all satisfied by PM_ϕ and all of them obtain zero utilities (i.e., the valuation equals the payment), while the offline optimal social welfare is to satisfy all the agents in the second group whose valuation densities are all $\bar{p}$.
- When the final resource utilization $\rho \in [\omega, \bar{\rho}_\varphi)$, the worst-case competitive ratio is

$$\alpha_3 = \max_{\rho \in [\omega, \bar{\rho}_\varphi)} \frac{h(\varphi(\rho))}{\underline{p}\omega + \int_\omega^\rho \varphi(y)dy - f(\rho)},$$

  which comes from Eq. (32). The rationality of α_3 is based on analyzing the competitive ratio of PM_ϕ under a series of worst-case scenarios constructed in Section 5.2.

Combining the above three cases, we have

$$\alpha(\omega, \varphi) = \max\{\alpha_1, \alpha_2, \alpha_3\}$$

$$= \max\left\{\frac{h(\underline{p})}{F_{\underline{p}}(\omega)}, \frac{h(\bar{p})}{\underline{p}\omega + \int_\omega^{\bar{\rho}_\varphi} \varphi(y)dy - f\left(\bar{\rho}_\varphi\right)}, \max_{\rho \in [\omega, \bar{\rho}_\varphi]} \frac{h(\varphi(\rho))}{\underline{p}\omega + \int_\omega^\rho \varphi(y)dy - f(\rho)}\right\}.$$

We thus complete the proof of Corollary 5.1. Note that when $\varphi(1) \geq \bar{p}$, the second term (i.e., α_2) is not needed as it is contained by α_3.

E PROOF OF PROPOSITION 6.2

E.1 Background of First-Order Two-Point BVPs

For a better reference, we revisit the two BVPs in Corollary 6.1 as follows:

$$\text{BVP}_1(\omega, u, \alpha) \begin{cases} \varphi_1'(y) = \alpha \cdot \frac{\varphi_1(y) - f'(y)}{f'^{-1}(\varphi_1(y))}, y \in (\omega, u); \\ \varphi_1(\omega) = \underline{p}, \varphi_1(u) = \bar{c}, \end{cases} \tag{71}$$

$$\text{BVP}_2(u, \alpha) \begin{cases} \varphi_2'(y) = \alpha \cdot (\varphi_2(y) - f'(y)), y \in (u, 1); \\ \varphi_2(u) = \bar{c}, \varphi_2(1) \geq \bar{p}. \end{cases} \tag{72}$$

A solution to the above first-order two-point BVPs is a function $\varphi_1(y)$ or $\varphi_2(y)$ that satisfies the corresponding ODE and also the two boundary conditions.

In the field of differential equations, when there is only one additional condition other than the ODE, the resulting problem is an initial value problem (IVP). For example, the following IVP is

derived from $\text{BVP}_1(\omega, u, \alpha)$ by excluding the first boundary condition:

$$\text{IVP}_1(\omega, u, \alpha) \begin{cases} \varphi_1'(y) = \alpha \cdot \frac{\varphi_1(y) - f'(y)}{f'^{-1}(\varphi_1(y))}, y \in (\omega, u); \\ \varphi_1(u) = \bar{c}, \end{cases} \tag{73}$$

where $\varphi_1(u) = \bar{c}$ is usually termed as the initial condition. We denote the solution to the above IVP by $\varphi_{\text{ivp}}(y; \omega, u, \alpha)$ (if it exists). Meanwhile, we denote the first derivative of $\varphi_{\text{ivp}}(y; \omega, u, \alpha)$ w.r.t $y \in (\omega, u)$ by $\varphi_{\text{ivp}}'(y; \omega, u, \alpha)$.

The key to the analysis of IVPs and BVPs is the existence and uniqueness of solutions. For first-order IVPs, the existence and uniqueness theorem is well understood. In particular, the Picard–Lindelöf theorem [3, 35] guarantees the unique exsitence of solutions as long as the ODE satisfies a certain Lipschitz continuity conditions. Meanwhile, there are plenty of iterative methods off-the-shelf that can solve IVPs numerically. However, for BVPs, there is no general uniqueness and existence theorem. As argued by [35], it is even non-trivial to obtain numerical solutions for some BVPs in the most basic two-point case as Eq. (71) and Eq. (72).

Our optimal pricing function design in Theorem 4.3 relies on the uniqueness and existence property of first-order two-point BVPs. On the one hand, based on the existence and uniqueness property, we can prove that there always exists a competitive pricing function as long as the competitive ratio α satisfies a certain condition. On the other hand, we can prove that there exists no other online algorithms that can outperform PM_ϕ with a smaller competitive ratio since the corresponding BVP has no solution.

E.2 Basic Lemmas

Our pricing function design is related to the existence, uniqueness and monotonicity property of solutions to $\text{IVP}_1(\omega, u, \alpha)$. Note that to simplify the analysis below, we temporarily assume that $\omega \in [0, \rho]$ is a given parameter. Later in our proofs of the existence of $\Gamma_1(\omega, u)$ and $\Gamma_2(u)$, we will include the discussions of ω as a design variable.

LEMMA E.1. *For each* $(\alpha, u) \in [1, +\infty) \times (\omega, 1]$, $\text{IVP}_1(\omega, u, \alpha)$ *has a unique solution* $\varphi_{\text{ivp}}(y; \omega, u, \alpha)$ *that is defined over* $y \in [\omega, u]$.

Lemma E.1 follows the Picard–Lindelöf theorem for the existence and uniqueness of solutions to IVPs. We refer the details to [1, 3, 35]. Basically the Picard–Lindelöf theorem guarantees that there always exists a unique solution to $\text{IVP}_1(\omega, u, \alpha)$, defined on a small neighbourhood of the initial point $\varphi(u) = \bar{c}$, as long as the ODE in $\text{IVP}_1(\omega, u, \alpha)$ is Lipschitz continuous within that neighbourhood. Moreover, this unique solution extends to the whole region of $y \in [\omega, u]$. Based on Lemma E.1, we can prove the following Lemma E.2 and Lemma E.3.

LEMMA E.2. *Given* $\alpha \geq 1$, $\varphi_{\text{ivp}}(y; \omega, u, \alpha)$ *is increasing in* $y \in [\omega, u]$ *and lower bounded by* $f'(y)$ *for all* $y \in [\omega, u]$.

PROOF. We first prove that $\varphi_{\text{ivp}}(y; \omega, u, \alpha) \geq f'(y)$ holds for all $y \in [\omega, u]$. Note that when $y = u \in (0, 1)$, we have $\varphi_{\text{ivp}}(u; \omega, u, \alpha) > f'(u)$ and $\varphi_{\text{ivp}}'(u; \omega, u, \alpha) > 0$. Therefore, if $\varphi_{\text{ivp}}(y; \omega, u, \alpha) \geq f'(y)$ does not hold for all $y \in [\omega, u]$, then there must exists at least one point within (ω, u), say $y_0 \in (\omega, u)$, such that $\varphi_{\text{ivp}}(y; \omega, u, \alpha) > f'(y)$ for all $y \in (y_0, u]$, $\varphi_{\text{ivp}}(y_0; \omega, u, \alpha) = f'(y_0)$, and $\varphi_{\text{ivp}}(y; \omega, u, \alpha) < f'(y)$ for all $\omega \in (y_0 - \epsilon, y_0)$, where ϵ is a small positive value. However, when $\varphi_{\text{ivp}}(y; \omega, u, \alpha) < f'(y)$, $\varphi_{\text{ivp}}'(y; \omega, u, \alpha)$ is negative according to the ODE, and thus $\varphi_{\text{ivp}}(y; \omega, u, \alpha)$ is decreasing in $(y_0 - \epsilon, y_0)$. This means that $\varphi_{\text{ivp}}(y; \omega, u, \alpha) > \varphi_{\text{ivp}}(y_0; \omega, u, \alpha) = f'(y_0) > f'(y)$ for all $y \in (y_0 - \epsilon, y_0)$, leading to a contradiction. Therefore, $\varphi_{\text{ivp}}(y; \omega, u, \alpha) \geq f'(y)$ always holds for all $y \in [\omega, u]$, and the monotonicity directly follows. □

Lemma E.2 guarantees that for any $\alpha \geq 1$ and $u \in (\omega, 1)$, the unique solution to $\text{IVP}_1(\omega, u, \alpha)$ is a monotone pricing function. Below we give Lemma E.3, which states that $\varphi_{\text{ivp}}(y; \omega, u, \alpha)$ is also monotonic in $(u, \alpha) \in (\omega, 1) \times [1, +\infty)$.

LEMMA E.3. $\varphi_{\text{ivp}}(y; \omega, u, \alpha)$ is continuous and non-increasing in $(u, \alpha) \in (\omega, 1) \times [1, +\infty)$.

PROOF. The continuity directly follows since $\varphi_{\text{ivp}}(y; \omega, u, \alpha)$ is well defined for all $(u, \alpha) \in (\omega, 1) \times [1, +\infty)$.

We first prove the monotonicity in $u \in (\omega, 1)$ by contradiction. Suppose we have $u_1 \in (\omega, 1)$ and $u_2 \in (\omega, 1)$, and assume w.l.o.g. that $u_1 > u_2$, we can prove that $\varphi_{\text{ivp}}(y; \omega, u_1, \alpha) < \varphi_{\text{ivp}}(y; \omega, u_2, \alpha)$ holds for all $y \in (\omega, u_2)$. The idea is that these two functions cannot have any intersection point, since otherwise the IVP with the same ODE as $\text{IVP}_1(\omega, u, \alpha)$ but with the initial condition defined at the intersection point will have at least two solutions, namely $\varphi_{\text{ivp}}(y; \omega, u_1, \alpha)$ and $\varphi_{\text{ivp}}(y; \omega, u_2, \alpha)$, which is impossible due to the uniqueness property given by Lemma E.1. Note that it is also impossible to have $\varphi_{\text{ivp}}(y; \omega, u_1, \alpha) > \varphi_{\text{ivp}}(y; \omega, u_2, \alpha)$ since if this is the case, then $\varphi_{\text{ivp}}(y; \omega, u_1, \alpha)$ is not monotonic in $y \in (\omega, u_1)$. Therefore, when $u_1 > u_2$, we always have $\varphi_{\text{ivp}}(y; \omega, u_1, \alpha) < \varphi_{\text{ivp}}(y; \omega, u_2, \alpha)$.

We now prove the monotonicity in $\alpha \in [1, +\infty)$. Suppose we have α_1 and α_2, and assume w.l.o.g. that $\alpha_1 > \alpha_2$. We need to prove that $\varphi_{\text{ivp}}(y; \omega, u, \alpha_1) < \varphi_{\text{ivp}}(y; \omega, u, \alpha_2)$ for all $y \in (\omega, u]$. Based on the ODE of (73), when $\alpha_1 > \alpha_2$, the derivative of φ at $y = u$ satisfies

$$\varphi'_{\text{ivp}}(u; \omega, u, \alpha_1) > \varphi'_{\text{ivp}}(u; \omega, u, \alpha_2).$$

Therefore, there must exist a small interval on the left-hand-side of u, say $[u - \epsilon, u]$, where ϵ is a small positive value, such that $\varphi_{\text{ivp}}(y; \omega, u, \alpha_1) < \varphi_{\text{ivp}}(y; \omega, u, \alpha_1)$ for all $y \in [u - \epsilon, u]$. This can be easily proved based on the definition of the derivative, which is omitted for brevity.

Now suppose $\varphi_{\text{ivp}}(y; \omega, u, \alpha_1) < \varphi_{\text{ivp}}(y; \omega, u, \alpha_2)$ does not hold for all $y \in [\omega, u]$, then there must exist an intersection point u_0, such that $\varphi_{\text{ivp}}(y; \omega, u, \alpha_1) < \varphi_{\text{ivp}}(y; \omega, u, \alpha_2)$ when $y \in (u_0, u]$, and $\varphi_{\text{ivp}}(y; \omega, u, \alpha_1) \geq \varphi_{\text{ivp}}(y; \omega, u, \alpha_2)$ when $y \in (u_0 - \epsilon, u_0]$, where ϵ is a very small positive value. Now let us consider two new IVPs with the same initial condition defined at point $y = u_0$, and denote the unique solutions to these two new IVPs by $\varphi_{\text{new}}(y; \omega, u_0, \alpha_1)$ and $\varphi_{\text{new}}(y; \omega, u_0, \alpha_2)$, according to the uniqueness property, we must have

$$\varphi_{\text{new}}(y; \omega, u_0, \alpha_1) = \varphi_{\text{ivp}}(y; \omega, u, \alpha_1), \forall y \in (\omega, u_0),$$
$$\varphi_{\text{new}}(y; \omega, u_0, \alpha_2) = \varphi_{\text{ivp}}(y; \omega, u, \alpha_2), \forall y \in (\omega, u_0).$$

Since $\varphi_{\text{ivp}}(y; \omega, u, \alpha_1) \geq \varphi_{\text{ivp}}(y; \omega, u, \alpha_2)$ when $y \in (u_0 - \epsilon, u_0]$, which means that we cannot find a small interval on the left-hand-side of u_0, say $[u_0 - \hat{\epsilon}, u_0]$, such that $\varphi_{\text{new}}(y; \omega, u_0, \alpha_1) < \varphi_{\text{new}}(y; \omega, u_0, \alpha_2)$. However, this contradicts with the fact that

$$\varphi'_{\text{new}}(u_0; \omega, u_0, \alpha_1) > \varphi'_{\text{new}}(u_0; \omega, u_0, \alpha_2).$$

Therefore, we have $\varphi_{\text{ivp}}(y; \omega, u, \alpha_1) < \varphi_{\text{ivp}}(y; \omega, u, \alpha_2)$ for all $y \in (\omega, u]$. We thus complete the proof. □

E.3 Proof of the Existence of $\Gamma_1(\omega, u)$

We now prove the unique existence of a well-defined function $\Gamma_1(\omega, u)$ in the first bullet of Proposition 6.2.

Our proof is based on the idea of "shooting method" in ODE [3, 35]. Here we briefly explain the intuition. For each given $\omega \in [0, \rho]$ and $u \in (\omega, 1)$, based on the ODE in Eq. (73), when α approaches infinity, $\varphi_1(y)$ must be very close to $f(y)$ when y approaches ω from the right, namely, $\varphi_1(\omega)$ approaches $f'(\omega)$. In this case, $\varphi_1(\omega) \leq p$. Similarly, when α is very small, say approaches zero (our definition of α requires that $\alpha \geq 1$, here we purely focus on the numerical analysis of
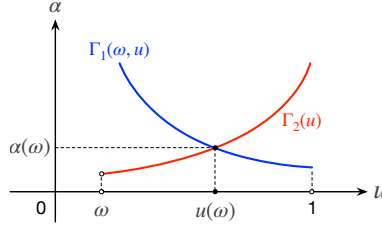
Fig. 5. Illustration of the unique existence of the intersection point $u(\omega)$ between $\Gamma_1(\omega, u)$ and $\Gamma_2(u)$ for a given $\omega \in (0, \underline{\rho})$.

the IVP given by Eq. (73)), then $\varphi_1(\omega)$ must be very close to $\bar{c}$ since $\varphi_1'(y)$ will be close to 0 and thus $\varphi_1(y)$ will be almost horizontal. In this case, $\varphi_1(\omega) \geq \underline{p}$. The monotonicity property in Lemma E.3 indicates that there must exist a unique α in between so that $\varphi(\omega)$ is exactly equal to $\underline{p}$. In the "shooting method", this indicates the value setting of α hits the target of $\varphi(\omega) = \underline{p}$. In summary, for a given $\omega \in [0, \underline{\rho}]$ and $u \in (\omega, 1)$, there exists a unique α so that the unique solution to the IVP in Eq. (73) is the unique solution to the BVP in Eq. (71). We thus complete the proof of the existence of $\alpha = \Gamma_1(\omega, u)$.

E.4 Proof of the Existence of $\Gamma_2(u)$

Proving the unique existence of $\alpha = \Gamma_2(u)$ is more straightforward as the ODE in Eq. (72) can be solved in analytical forms. To be more specific, the general solution can be written as follows:

$$\varphi_2(y) = \exp(\alpha y) \cdot \left(\int^y \frac{\alpha f'(\eta)}{\exp(\alpha \eta)} f'(\eta) d\eta + C \right),$$

where C is any real constant. Substituting the two boundary conditions into the above general solution leads to the following

$$\varphi_2(u) = \exp(\alpha \cdot u) \cdot \left(\int^u \frac{\alpha f'(\eta)}{\exp(\alpha \eta)} d\eta + C \right) = \bar{c},$$

$$\varphi_2(1) = \exp(\alpha \cdot 1) \cdot \left(\int^1 \frac{\alpha f'(\eta)}{\exp(\alpha \eta)} d\eta + C \right) \geq \bar{p},$$

which thus indicates that

$$\int_u^1 \frac{\alpha f'(\eta)}{\exp(\alpha \eta)} d\eta \geq \frac{\bar{p}}{\exp(\alpha)} - \frac{\bar{c}}{\exp(u\alpha)} \tag{74}$$

We can prove that for each given $u \in (0, 1)$, there exists a unique α that satisfies Eq. (74) by equality as long as $\bar{p} > \bar{c}$. The proof is similar to that of $\alpha = \Gamma_1(\omega, u)$ as any given setup of $\bar{p} > \bar{c}$ will guarantee the existence of a monotone function $\varphi_2(y)$ in the interval $[u, 1]$ with an appropriate setting of α. Moreover, this α must be unique as otherwise the uniqueness property in Lemma E.1 would be contradicted. Therefore, if we denote this unique α by $\alpha = \Gamma_2(u)$, then we complete the proof of the existence of $\Gamma_2(u)$. However, note that in the second bullet of Proposition 6.2, we argue that for any $\alpha \geq \Gamma_2(u)$, the BVP in Eq. (72) will have a unique monotone solution. This can be proved by the monotonicity of $\varphi_2(y)$ w.r.t. α, similar to our proof of Lemma E.3, and thus we skip the details for brevity. In summary, for each given $u \in (0, 1)$, there exists a unique solution to $\text{BVP}_2(u, \alpha)$ as long as $\alpha \geq \Gamma_2(u)$. We thus complete the proof of Proposition 6.2.

F PROOF OF THEOREM 6.3

The proof of Theorem 6.3 heavily relies on the monotonicity properties of ϕ w.r.t ω, u, and α. To achieve the minimum α, we make the following claims:

- **Claim 1**: ω must be as small as possible. Since a larger ω will shrink the interval of $[\omega, 1]$, and the value of α must be increased as the overall slope the pricing function will be steeper. For any α, to minimize the choice of ω, we set

$$\alpha = \frac{h(\underline{p})}{F_{\underline{p}}(\omega)} \text{ or } \omega = F_{\underline{p}}^{-1}\left(\frac{1}{\alpha}h(\underline{p})\right).$$

- **Claim 2**: for each given $\omega \in [F_{\underline{p}}^{-1}(\frac{1}{\alpha}h(\underline{p})), \rho]$, $\Gamma_1(\omega, u)$ is strictly decreasing in $u \in (\omega, 1)$. Meanwhile, $\Gamma_2(u)$ is strictly increasing in $u \in (\omega, 1)$. The proof of the monotonicity of $\Gamma_1(\omega, u)$ and $\Gamma_2(u)$ is elementary based on our analysis in Appendix E above. To visualize the intuition, Fig. 5 illustrates $\Gamma_1(\omega, u)$ and $\Gamma_2(u)$ w.r.t. $u \in (\omega, 1)$ for a given ω. We can prove that for any $\omega \in (0, \bar{\rho}]$, there exists a unique $u(\omega)$ such that $\Gamma_1(\omega, u(\omega)) = \Gamma_2(u(\omega))$, as illustrated by the intersection point in Fig. 5. The proof is straightforward by evaluating the values of $\Gamma_1(\omega, u)$ and $\Gamma_2(u)$ at their two boundaries, e.g., $\Gamma_1(\omega, u)$ and $\Gamma_2(u)$ approaches infinity if u approaches ω and 1, respectively, Therefore, the monotonicity of $\Gamma_1(\omega, u)$ and $\Gamma_2(u)$ indicates that there must exist a unique intersection point $u(\omega)$ in the interval $[\omega, 1]$. To minimize α, we simply choose $u = u(\omega)$ so that $\alpha = \alpha(\omega)$ is minimized for a given $\omega \in [\omega, \rho]$.

Based on the above two claims, to minimize $\alpha(\omega)$, we must set $\omega = F_{\underline{p}}^{-1}\left(\frac{1}{\alpha}h(\underline{p})\right)$. In this case the unique intersection point between $\Gamma_1(\omega, u)$ and $\Gamma_2(u)$ is the optimal design of u, and the corresponding α is the minimum competitive ratio, i.e., $\alpha_*(\mathcal{S})$. In summary, the optimal design of ω and u is the unique solution to the following system of equations

$$\frac{h(\underline{p})}{F_{\underline{p}}(\omega_*)} = \Gamma_1(\omega_*, u_*) = \Gamma_2(u_*).$$

We thus complete the proof of Theorem 6.3.

G REMAINING PROOFS OF THEOREM 4.3 IN CASE-2 AND CASE-3

In this section, we provide the remaining proofs of Theorem 4.3 in **Case-2** and **Case-3**.

G.1 Case-2: $\underline{c} < \bar{c} \le \underline{p} \le \bar{p}$

In this case, $\bar{c} \le \underline{p} \le \bar{p}$ indicates that $\rho = \bar{\rho} = 1$. Meanwhile, there exist no such a resource utilization threshold u so that $\phi(u) = \bar{c}$, as illustrated in Fig. 1(c). Based on Theorem 4.1 and Theorem 4.2, we give the following Corollary G.1 which summarizes the sufficient and necessary conditions in **Case-2**.

COROLLARY G.1. *Given a convex setup $\mathcal{S}$ in **Case-2**, PM_ϕ is IC and α-competitive if there exists a critical threshold $\omega \in [F_{\underline{p}}^{-1}(\frac{1}{\alpha}h(\underline{p})), 1]$ such that ϕ is given by*

$$\phi(y) = \begin{cases} \underline{p}, & \text{if } y \in [0, \omega], \\ \hat{\varphi}_2(y), & \text{if } y \in [\omega, 1], \\ +\infty, & \text{if } y \in (1, +\infty), \end{cases} \tag{75}$$

where $\hat{\varphi}_2(y)$ is a solution to the following BVP:

$$\widehat{\text{BVP}}_2(\omega, \alpha) \begin{cases} \hat{\varphi}'_2(y) = \alpha \cdot (\hat{\varphi}_2(y) - f'(y)), y \in (\omega, 1); \\ \hat{\varphi}_2(\omega) = \underline{p}, \hat{\varphi}_2(1) \geq \bar{p}. \end{cases} \tag{76}$$

On the other hand, if there exists an α-competitive online algorithm, then there must exist a strictly-increasing solution to $\widehat{\text{BVP}}_2(\omega, \alpha)$ for some $\omega \in [F_{\underline{p}}^{-1}(\frac{1}{\alpha}h(\underline{p})), 1]$.

PROOF. The proof is similar to our previous analysis regarding Corollary 6.1 in **Case-1**.   □

Similar to Proposition 6.2, we give the following Proposition G.2 to show the existence of function $\hat{\Gamma}_2(\omega)$.

PROPOSITION G.2. *For each given $\omega \in [0, 1)$, $\widehat{\text{BVP}}_2(\omega, \alpha)$ has a unique strictly-increasing solution if and only if $\alpha \geq \widehat{\Gamma}_2(\omega)$, where $\widehat{\Gamma}_2(\omega)$ satisfies the following equation in variable $\widehat{\Gamma}_2$:*

$$\int_\omega^1 \frac{\widehat{\Gamma}_2 f'(y)}{\exp(y\widehat{\Gamma}_2)} dy = \frac{\bar{p}}{\exp(\widehat{\Gamma}_2)} - \frac{\underline{p}}{\exp(\omega\widehat{\Gamma}_2)}. \tag{77}$$

PROOF. Eq. (77) is similar to Eq. (46) and thus we omit the proof for brevity.   □

Based on Proposition G.2, we give the following Theorem G.3 to show the unique existence of ω_* as well as the calculation of $\alpha_*(S)$ in **Case-2**.

THEOREM G.3. *Given a convex setup S in **Case-2**, there exists a unique $\omega_* \in (0, 1)$ such that the optimal competitive ratio $\alpha_*(S)$ is given by*

$$\alpha_*(S) = \frac{h(\underline{p})}{F_{\underline{p}}(\omega_*)} = \widehat{\Gamma}_2(\omega_*), \tag{78}$$

where $\omega_* \in [0, 1)$ is the unique root that satisfies Eq. (78). Meanwhile, PM_ϕ is $\alpha_*(S)$-competitive if and only if ϕ is given by Eq. (75) with $(\omega, \alpha) = (\omega_*, \alpha_*(S))$.

PROOF. The result follows similarly as Theorem 6.3.   □

Note that in **Case-2**, the calculation of the unique ω_* is equivalent to finding the unique root to the following equation

$$\int_{\omega_*}^1 \frac{h(\underline{p})}{F_{\underline{p}}(\omega_*)} \cdot \exp\left(-\frac{h(\underline{p})}{F_{\underline{p}}(\omega_*)} \cdot y\right) \cdot f'(y) \, dy = \bar{p} \cdot \exp\left(-\frac{h(\underline{p})}{F_{\underline{p}}(\omega_*)}\right) - \underline{p} \cdot \exp\left(-\frac{h(\underline{p})}{F_{\underline{p}}(\omega_*)} \cdot \omega_*\right).$$

G.2 Case-3: $\underline{c} < \underline{p} \leq \bar{p} \leq \bar{c}$

In this case, $\underline{c} < \underline{p} \leq \bar{p} \leq \bar{c}$ indicates that $\underline{\rho} \leq \bar{\rho} \leq 1$. Meanwhile, there exist no such a threshold $u \in (0, 1)$ so that $\bar{\phi}(u) = \bar{c}$, as illustrated in Fig. 2(d). Based on Theorem 4.1 and Theorem 4.2, we give the following Corollary G.4 which summarizes the sufficient and necessary conditions in **Case-3**.

COROLLARY G.4. *Given a convex setup S in **Case-3**, PM_ϕ is IC and α-competitive if there exists a pair of resource utilization thresholds $(\omega, \rho) \in [F_{\underline{p}}^{-1}(\frac{1}{\alpha}h(\underline{p})), \underline{\rho}] \times [\omega, \bar{\rho}]$ such that ϕ is given by*

$$\phi(y) = \begin{cases} \underline{p}, & \text{if } z \in [0, \omega], \\ \hat{\varphi}_1(y), & \text{if } z \in [\omega, \rho], \\ +\infty, & \text{if } z \in (\rho, 1], \end{cases} \tag{79}$$

where $\hat{\varphi}_1(y)$ is a solution to the following BVP:

$$\widehat{\text{BVP}}_1(\omega, \rho, \alpha) \begin{cases} \hat{\varphi}'_1(y) = \alpha \cdot \frac{\hat{\varphi}_1(y) - f'(y)}{f'^{-1}(\hat{\varphi}_1(y))}, y \in (\omega, \rho); \\ \hat{\varphi}_1(\omega) = \underline{p}, \hat{\varphi}_1(\rho) \geq \bar{p}. \end{cases} \tag{80}$$

On the other hand, if there exists an α-competitive online algorithm, then there must exist a strictly-increasing solution to $\widehat{\text{BVP}}_1(\omega, \rho, \alpha)$ for some pair of $(\omega, \rho) \in [F_{\underline{p}}^{-1}\left(\frac{1}{\alpha}h(\underline{p})\right), \underline{\rho}] \times [\underline{\omega}, \bar{\rho}]$.

PROOF. The proof is similar to our analysis of Theorem 6.3. □

Similar to Proposition 6.2, we give the following Proposition G.5 to show the existence of function $\hat{\Gamma}_1(\omega, \rho)$.

PROPOSITION G.5. *Given a convex setup $\mathcal{S}$ in **Case-3**, for each given pair of $(\omega, \rho) \in [0, \underline{\rho}) \times (\underline{\omega}, \bar{\rho}]$, there exists a well-defined function $\widehat{\Gamma}_1(\omega, \rho)$ so that $\widehat{\text{BVP}}_1(\omega, \rho, \alpha)$ has a unique strictly-increasing solution if and only if $\alpha \geq \widehat{\Gamma}_1(\omega, \rho)$.*

PROOF. The result follows the first bullet in Proposition 6.2. □

Based on Proposition G.5, we give the following Theorem G.6 to show the calculation of $\alpha_*(\mathcal{S})$ in **Case-3**.

THEOREM G.6. *Given a convex setup $\mathcal{S}$ in **Case-3**, there exists a unique $\omega_* \in [0, \underline{\rho}]$ such that the optimal competitive ratio $\alpha_*(\mathcal{S})$ is given by*

$$\alpha_*(\mathcal{S}) = \frac{h(\underline{p})}{F_{\underline{p}}(\omega_*)} = \widehat{\Gamma}_1(\omega_*, \bar{\rho}), \tag{81}$$

where ω_* is the unique root that satisfies Eq. (81). Meanwhile, PM_ϕ is $\alpha_*(\mathcal{S})$-competitive if and only if ϕ is given by Eq. (79) with $(\omega, \rho, \alpha) = (\omega_*, \bar{\rho}, \alpha_*(\mathcal{S}))$.

PROOF. The proof is similar to our analysis of Theorem 6.3. □

Based on Theorem 6.3, Theorem G.3, and Theorem G.6, Theorem 4.3 follows. We thus complete the proof of Theorem 4.3 in all the three cases.

REMARK 7. *Based on Theorem 6.3, Theorem G.3, and Theorem G.6, we conclude that the upper bound of the optimal pricing function is always equal to the maximum valuation density, namely, $\phi_*(\bar{\rho}) = \bar{p}$. Note that compared to $\phi_*(\bar{\rho}) > \bar{p}$, by setting $\phi_*(\bar{\rho}) = \bar{p}$ we can have a higher resource utilization level while still guarantee the competitive ratio of PM_ϕ. The intuition here is that it is always beneficial to have a higher resource utilization level as long as the resources are sold at the right price.*

H PROOF OF COROLLARY 4.4

The properties of $\alpha_*(\mathcal{S})$ directly follow the monotonicity of $\Gamma_1(\omega, u)$ and $\Gamma_2(u)$. We briefly explain them as follows:

- For a given $\bar{p} \in (\underline{c}, +\infty)$, we claim that $\Gamma_1(\omega, u)$ is strictly decreasing in $\underline{p} \in (\underline{c}, \bar{p}]$. We use **Case-1** to explain the intuition. A larger $\underline{p}$ indicates a smaller value of $\bar{c} - \underline{p}$, and thus for a given ω and u, the pricing function φ_1 can increase from $\underline{p}$ to $\bar{c}$ with a smaller rate. A smaller increasing rate of φ_1 indicates that the value of α can be smaller as well. Therefore, $\Gamma_1(\omega, u)$ is decreasing in $\underline{p} \in (\underline{c}, \bar{p}]$. As we can see from Fig. 5, a smaller $\Gamma_1(\omega, u)$ indicates a smaller value of $u(\omega)$ (note that $\Gamma_2(u)$ will not change w.r.t. the increasing of $\underline{p}$), as well as a smaller value of $\alpha(\omega)$, which thus leads to a smaller value of $\alpha_*(\mathcal{S})$ in the end. Therefore, $\alpha_*(\mathcal{S})$ is strictly decreasing in $\underline{p} \in (\underline{c}, \bar{p}]$ for a given $\bar{p} \in (\underline{c}, +\infty)$.

- We can prove similarly that $\alpha_*(\mathcal{S})$ is strictly increasing in $\bar{p} \in [\underline{p}, +\infty)$ for a given $\underline{p} \in (\underline{c}, +\infty)$. We skip the details for brevity.
- Note that $\underline{p} = \bar{p} \in (\underline{c}, +\infty)$ happens only in **Case-2** and **Case-3**. For instance, in **Case-2**, $\underline{p} = \bar{p} \in (\bar{c}, +\infty)$ indicates that $\hat{\varphi}_2(\omega_*) = \hat{\varphi}_2(1) = \underline{p} = \bar{p}$, where $\hat{\varphi}_2$ is the unique solution to the BVP in Eq. (76). We argue that such a solution occurs only when $\omega_* = 1$ and in this case the optimal competitive ratio $\alpha_*(\mathcal{S}) = 1$ since both the online and offline strategies are to satisfy all the request until reaching the total capacity limit. We can prove similarly that in **Case-3**, $\omega_* = \rho$ and $\alpha_*(\mathcal{S}) = 1$ as long as $\underline{p} = \bar{p} \in (\underline{c}, \bar{c}]$. Combining the results in both cases, we have $\alpha_*(\mathcal{S}) = 1$ and $\omega_* = \rho$ when $\underline{p} = \bar{p} \in (\underline{c}, +\infty)$.

We thus complete the proof of Corollary 4.4.

I PROOF OF THEOREM 5.2

Before presenting the proof of Theorem 5.2, we first give the extension of PM_ϕ in Algorithm 2.

Algorithm 2: Posted Price Mechanism (PM_ϕ)

1: **Inputs:** A given setup $\mathcal{S} = \{f_t, \underline{p}_t, \bar{p}_t\}_{\forall t \in \mathcal{T}}$ and $\boldsymbol{\phi} = \{\phi_t\}_{\forall t \in \mathcal{T}}$.
2: **Initialize**: $\hat{y}_t^{(0)} = 0$ and $\hat{p}_t^{(0)} = \phi(\hat{y}_t^{(0)})$ for all $t \in \mathcal{T}$.
3: **while** a new agent n arrives **do**
4:     Supplier publishes the price $\hat{p}_t^{(n-1)}$ for $t \in \mathcal{T}$.
5:     **if** $v_n - \sum_{t \in \mathcal{T}_n} r_n^t \cdot \phi_t\big(y_t^{(n-1)}\big) < 0$ **then**
6:         Agent n leaves (i.e., set $\hat{x}_n = 0$)
7:     **else if** $\hat{y}_t^{(n-1)} + r_n^t > 1$ for any $t \in \mathcal{T}$ **then**
8:         Request n is rejected (i.e., set $\hat{x}_n = 0$)
9:     **else**
10:         Request n is satisfied (i.e., set $\hat{x}_n = 1$)
11:         Set the payment $\hat{\pi}_n$ by $\hat{\pi}_n = \sum_{t \in \mathcal{T}_n} r_n^t \cdot \phi_t\left(y_t^{(n-1)}\right)$.
12:         Update the total resource utilization by $\hat{y}_t^{(n)} = \hat{y}_t^{(n-1)} + r_n^t, \forall t \in \mathcal{T}$.
13:         Update the price by $\hat{p}_t^{(n)} = \phi_t\left(\hat{y}_t^{(n)}\right), \forall t \in \mathcal{T}$.
14:     **end if**
15: **end while**

Based on Algorithm 2, the sufficient conditions of Theorem 5.2 can be derived based on Proposition 3.1 as follows. We first prove that there indeed exists an index $k \in \mathcal{N}$ such that the initial inequality hold. The major argument is that in the worst-case, each agent may only require resource for a single time slot, leading to the existence of a critical threshold ω_t for each $t \in \mathcal{T}$. Specifically, suppose $k \in \mathcal{N}$ and after processing all the agents within $\{1, 2, \cdots, k\}$, we have

$$P_k = \sum_{n=1}^{k} v_n - \sum_{t \in \mathcal{T}} \bar{f}_t(\hat{y}_t^{(k)}) = \sum_{n=1}^{k} \left(\hat{\gamma}_n + \sum_{t \in \mathcal{T}} r_n^t \cdot \phi_t(y_t^{(n-1)}) \right) - \sum_{t \in \mathcal{T}} \bar{f}_t(\hat{y}_t^{(k)}).$$

The dual objective after processing request k is given by

$$D_k = \sum_{n=1}^{k} \hat{\gamma}_n + \sum_{t \in \mathcal{T}} h_t(\hat{p}_t^{(k)}).$$

Based on the pricing function given by Eq. (38), the initial inequality $P_k \geq \frac{1}{\alpha_t} D_k$ leads to the following

$$\left(1 - \frac{1}{\alpha_t}\right) \sum_{n=1}^{k} \hat{\gamma}_n + \sum_{t \in \mathcal{T}} \underline{p}_t \cdot \left(\sum_{n=1}^{k} r_n^t\right) - \sum_{t \in \mathcal{T}} \bar{f}_t \left(\sum_{n=1}^{k} r_n^t\right) \geq \frac{1}{\alpha_t} \sum_{t \in \mathcal{T}} h_t\left(\underline{p}_t\right).$$

Since $\alpha_t \geq 1$ and $\hat{\gamma}_n \geq 0$, it suffices to have the following inequality for each time slot:

$$\underline{p}_t \cdot \left(\sum_{n=1}^{k} r_n^t\right) - \bar{f}_t \left(\sum_{n=1}^{k} r_n^t\right) \geq \frac{1}{\alpha_t} h_t\left(\underline{p}_t\right). \tag{82}$$

The above inequality Eq. (82) holds if the following one holds:

$$\underline{p}_t \omega_t - \bar{f}_t \left(\omega_t\right) = F_{\underline{p}_t}(\omega_t) \geq \frac{1}{\alpha_t} h_t\left(\underline{p}_t\right),$$

where $\omega_t = \sum_{n=1}^{k} r_n^t$. Meanwhile, for the same reason as the previous basic resource allocation model, it is obvious that ω_t must be less than or equal to ρ_t because the selling price must be larger than or equal to the marginal cost, i.e,. $\phi_t(\omega_t) = \underline{p}_t \geq f'_t(\omega_t)$. Therefore, Eq. (39) indicates that the initial inequality $P_k \geq \frac{1}{\alpha_t} D_k$ holds for some $k \in N$

We next prove the ordinary differential inequality in Eq. (40). The change in the primal objective is given as follows:

$$P_n - P_{n-1} = v_n - \sum_{t \in \mathcal{T}} \left(\bar{f}_t(\hat{y}_t^{(n)}) - \bar{f}_t(\hat{y}_t^{(n-1)})\right)$$

$$= \hat{\gamma}_n + \sum_{t \in \mathcal{T}} \phi_t(\hat{y}_t^{(n-1)}) \left(\hat{y}_t^{(n)} - \hat{y}_t^{(n-1)}\right) - \sum_{t \in \mathcal{T}} \left(\bar{f}_t(\hat{y}_t^{(n)}) - \bar{f}_t(\hat{y}_t^{(n-1)})\right).$$

The change in the dual objective is given as follows:

$$D_n - D_{n-1} = \hat{\gamma}_n + \sum_{t \in \mathcal{T}} \left(h_t(\hat{p}_t^{(n)}) - h_t(\hat{p}_t^{(n-1)})\right).$$

To guarantee that $P_n - P_{n-1} \geq \frac{1}{\alpha} (D_n - D_{n-1})$ holds, it is equivalent to having the following

$$\left(1 - \frac{1}{\alpha_t}\right) \hat{\gamma}_n + \sum_{t \in \mathcal{T}} \left(\phi_t(\hat{y}_t^{(n-1)})(\hat{y}_t^{(n)} - \hat{y}_t^{(n-1)}) - (\bar{f}_t(\hat{y}_t^{(n)}) - \bar{f}_t(\hat{y}_t^{(n-1)}))\right)$$

$$\geq \frac{1}{\alpha_t} \sum_{t \in \mathcal{T}} \left(h_t(\hat{p}_t^{(n)}) - h_t(\hat{p}_t^{(n-1)})\right),$$

where the first term $\hat{\gamma}_n \geq 0$ always holds. Therefore, if the following inequality holds at each round, then the incremental inequality holds as well:

$$\sum_{t \in \mathcal{T}} \left(\phi_t(\hat{y}_t^{(n-1)})(\hat{y}_t^{(n)} - \hat{y}_t^{(n-1)}) - (\bar{f}_t(\hat{y}_t^{(n)}) - \bar{f}_t(\hat{y}_t^{(n-1)}))\right) \geq \frac{1}{\alpha_t} \sum_{t \in \mathcal{T}} \left(h_t(\hat{p}_t^{(n)}) - h_t(\hat{p}_t^{(n-1)})\right). \tag{83}$$

To guarantee that Eq. (83) holds, it suffices to have the following per-slot inequality

$$\phi_t(\hat{y}_t^{(n-1)}) \left(\hat{y}_t^{(n)} - \hat{y}_t^{(n-1)}\right) - \left(\bar{f}_t(\hat{y}_t^{(n)}) - \bar{f}_t(\hat{y}_t^{(n-1)})\right) \geq \frac{1}{\alpha_t} \left(h_t(\hat{p}_t^{(n)}) - h_t(\hat{p}_t^{(n-1)})\right),$$

which is equivalent to the following:

$$\phi_t(\hat{y}_t^{(n-1)}) - \frac{\bar{f}_t(\hat{y}_t^{(n)}) - \bar{f}_t(\hat{y}_t^{(n-1)})}{\hat{y}_t^{(n)} - \hat{y}_t^{(n-1)}} \geq \frac{1}{\alpha_t} \cdot \frac{h_t\left(\phi_t(\hat{y}_t^{(n)})\right) - h_t\left(\phi_t(\hat{y}_t^{(n-1)})\right)}{\phi_t(\hat{y}_t^{(n)}) - \phi_t(\hat{y}_t^{(n-1)})} \cdot \frac{\phi_t(\hat{y}_t^{(n)}) - \phi_t(\hat{y}_t^{(n-1)})}{\hat{y}_t^{(n)} - \hat{y}_t^{(n-1)}}.$$

Therefore, we have the following inequality in the differential form:

$$\phi_t(y) - \bar{f}'_t(y) \geq \frac{1}{\alpha_t} \cdot h'_t\big(\phi_t(y)\big) \cdot \phi'_t(y), \forall y \in [\omega_t, 1). \tag{84}$$

This means, if Eq. (84) holds for all $y \in [\omega_t, 1)$, then the incremental inequality holds at each round for all $y \in [\omega_t, 1)$.

Similar to our previous analysis, we need to guarantee that the incremental inequality holds at each round for all $y \in [\omega_t, 1]$. Therefore, we need to consider the boundary condition of $\phi_t(y)$ when $y = 1$. First, it is obvious that $\phi_t(\omega_t) = \underline{p}_t$ as the pricing function ϕ_t must be continuous. Taking integration of both sides of Eq. (84) leads to

$$\int_{\omega_t}^{\rho_t} \phi_t(y)dy - \bar{f}_t(\rho_t) + \bar{f}_t(\omega_t) \geq \frac{1}{\alpha_t}\Big(h_t\big(\phi_t(\rho_t)\big) - h_t\big(\underline{p}_t\big)\Big),$$

where ρ_t denotes the final utilization level at time slot $t \in \mathcal{T}$. Meanwhile, based on the flat-segment, we have $\underline{p}_t \omega_t - \bar{f}_t(\omega_t) \geq \frac{1}{\alpha_t} h_t(\underline{p}_t)$. Thus, we have

$$\underline{p}_t \omega_t + \int_{\omega_t}^{\rho_t} \phi_t(y)dy - \bar{f}_t(\rho_t) \geq \frac{1}{\alpha_t} h_t\big(\phi_t(\rho_t)\big),$$

which is the same as Eq. (57). When $\rho_t = 1$, the above equality also holds, and thus we have

$$\underline{p}_t \omega_t + \int_{\omega_t}^{1} \phi_t(y)dy - \bar{f}_t(1) \geq \frac{1}{\alpha_t} h_t\big(\phi_t(1)\big).$$

On the other hand, $\mathsf{PM}_{\boldsymbol{\phi}}$ is α_t-competitive indicates that the following inequality must hold:

$$\underline{p}_t \omega_t + \int_{\omega_t}^{1} \phi_t(y)dy - \bar{f}_t(1) \geq \frac{1}{\alpha_t} \sum_{t \in \mathcal{T}} h_t(\bar{p}_t), \tag{85}$$

where the right-hand-side represents the optimal social welfare in the offline setting when the resource at time t has been used up. *Note that here the worst-case scenario is different from the single time slot case. The rationality of Eq.* (85) *is as follows. When the resource at time t is used up by agents whose valuation equals the payment, let us assume that there is a follow-up arrival instance which consists of two groups of agents as follows: i) the first group of agents ask for resource at time t with the maximum valuation density $\bar{p}_t$, and ii) the second group of agents ask for resources at time t as well as at time $\hat{t}$ with the maximum valuation density $\bar{p}_{\hat{t}}$, $\forall \hat{t} \in \mathcal{T} \backslash \{t\}$. Meanwhile, we assume that the number of agents in these two groups are large enough so that the total capacity of each time slot can be used up. In the online setting, the social welfare is given by the left-hand-side of Eq.* (85)*, meaning that the valuation of the accepted agents are all equal to their payments, while the follow-up agents are all rejected due to the capacity limit at time t. In comparison, in the offline setting, the optimal social welfare is to allocate all the capacity at time t to those agents in the follow-up instance whose valuation densities are $\bar{p}_t$, $\forall t \in \mathcal{T}$, as a result, the optimal optimal social welfare in hindsight is $\sum_{t \in \mathcal{T}} h_t(\bar{p}_t)$, as given by the right-hand-side of Eq.* (85). *Therefore,* to guarantee that the above Eq. (85) holds, it suffices to have the following boundary condition: $\phi_t(1) - f_t(1) \geq \sum_{t \in \mathcal{T}} h_t(\bar{p}_t) \Longrightarrow \phi_t(1) \geq \bar{p}_t + \sum_{\hat{t} \in \mathcal{T} \backslash \{t\}} h_{\hat{t}}(\bar{p}_{\hat{t}})$.

Summarizing our above proofs regarding the critical threshold, the differential inequality and the boundary conditions, we complete the proof of Theorem 5.2.

(*Impacts of Multiple Time Slots on the Competitive Ratio*) Based on the second boundary condition in Eq. (40), $\phi_t(1)$ follows roughly linearly w.r.t. the total number of time slots. For example, when we assume $\bar{p}_t \geq \bar{c}_t$ for all $t \in \mathcal{T}$, then the second boundary condition in Eq. (40) indicates that $\phi_t(1) \geq \bar{p}_t + \sum_{\hat{t} \in \mathcal{T} \backslash \{t\}} h_{\hat{t}}(\bar{p}_{\hat{t}}) = \sum_{t \in \mathcal{T}} \bar{p}_t - \sum_{\hat{t} \in \mathcal{T} \backslash \{t\}} f_{\hat{t}}(1)$. Our previous analysis regarding the basic model shows that the calculation of the optimal competitive ratio $\alpha_*(\mathcal{S})$ depends on the boundary

condition of ϕ_* (i.e., the monotonicity of $\alpha_*(\mathcal{S})$ in Corollary 4.4). Therefore, in the case with multiple time slots, the final competitive ratio $\alpha = \max_t\{\alpha_t\}$ depends on $|\mathcal{T}|$ as well. As we mention in the conclusion of the paper, it is difficult to analyze the property of $\alpha_*(\mathcal{S})$ w.r.t. $\underline{p}$ and $\bar{p}$. Similarly, here it is also difficult to analyze the dependency of α on the total number of time slots. We leave this for future research.

(*Necessary Conditions*) The necessary conditions in Theorem 4.2 can also be generalized to consider the case with multiple time slots. We skip the details for brevity as most of the proofs are similar to our previous analysis in the paper, except that the worst-case analysis regarding the boundary conditions should be performed in a different manner, as demonstrated by Eq. (85).

J COMPUTATION OF ω_*, u_*, AND $\alpha_*(\mathcal{S})$

In this section, we present the algorithms for computing ω_*, u_*, and $\alpha_*(\mathcal{S})$ in **Case-1**. We skip the computational details regarding **Case-2** and **Case-3** for brevity.

We give Algorithm 3 below to show the calculation of $\Gamma_1(\omega, u)$ for a given ω and u. We also give Algorithm 4 to show the calculation of $\Gamma_2(u)$ for a given u. Based on $\mathsf{GAMMA}_1(\omega, u)$ in Algorithm 3 and $\mathsf{GAMMA}_2(u)$ in Algorithm 4, Algorithm 5 shows the calculation of ω_*, u_*, and $\alpha_*(\mathcal{S})$ for a given convex setup $\mathcal{S}$.

Algorithm 3: $\mathsf{GAMMA}_1(\omega, u)$

1: **Inputs:** A given setup $\mathcal{S} = \{f, \underline{p}, \bar{p}\}$, ω, u, and ϵ. Set $\alpha_1 = 1$, $\alpha_2 = R_1$, and $\varphi_1(u) = R_2$.
2: **while** $|\varphi_1(u) - \bar{c}| > \epsilon$ **do**
3:      Compute $\alpha_0 = \frac{\alpha_1 + \alpha_2}{2}$.
4:      Solve the following IVP and obtain the solution $\varphi_1(y)$:

$$\mathsf{IVP}(\omega, u, \alpha_0) \begin{cases} \varphi_1'(y) = \alpha_0 \cdot \frac{\varphi_1(y) - f'(y)}{f'^{-1}(\varphi_1(y))}, y \in (\omega, u); \\ \varphi_1(\omega) = \underline{p}. \end{cases}$$

5:      **if** $\varphi_1(u) > \bar{c}$ **then**
6:          $\alpha_2 = \alpha_0$.
7:      **else**
8:          $\alpha_1 = \alpha_0$.
9:      **end if**
10: **end while**
11: $\Gamma_1(\omega, u) = \alpha_0$.
12: **Outputs:** $\Gamma_1(\omega, u)$.

We can see that Algorithm 3 performs a bisection search for the unique output of $\Gamma_1(\omega, u)$. Note that ϵ is a constant that is sufficiently small. In both Algorithm 3 and Algorithm 4, R_1 and R_2 are two large positive constants that can be initialized appropriately based on the setup $\mathcal{S}$. In line 4 of Algorithm 3, we need to solve the IVP parameterized by ω, u, and α_0. We emphasize that solving $\mathsf{IVP}(\omega, u, \alpha_0)$ in this step is light-weight and various off-the-shelf methods/solvers can be used. Algorithm 5 searches between $[0, 1]$ for the unique u_*. Based on the unique u_*, the optimal critical threshold ω_* can be given by line 5 of Algorithm 5. Recall that in Algorithm 5, $F_{\underline{p}}^{-1}$ denotes the inverse of $F_{\underline{p}}$, which is defined by Eq. (8).

Before ending this section, we give the following remark regarding the implementation of PM_{ϕ_*}.

REMARK 8. *It is worth pointing out that the calculation of ϕ_* in a numerical manner should not be a concern for the online implementation of PM_{ϕ_*}. On the one hand, our above algorithms shows that*

Algorithm 4: $\mathsf{GAMMA}_2(u)$

1: **Inputs:** A given setup $\mathcal{S} = \{f, \underline{p}, \bar{p}\}$, u, and ϵ. Set $\alpha_1 = 1$, $\alpha_2 = R_1$, and $GAP = R_2$.
2: **while** $|GAP| > \epsilon$ **do**
3:     Compute $\alpha_0 = \frac{\alpha_1 + \alpha_2}{2}$.
4:     Set $GAP = \int_u^1 \frac{\alpha_0 f'(y)}{\exp(y\alpha_0)} dy - \left(\frac{\bar{p}}{\exp(\alpha_0)} + \frac{\bar{c}}{\exp(u\alpha_0)} \right)$.
5:     **if** $GAP > 0$ **then**
6:         $\alpha_2 = \alpha_0$.
7:     **else**
8:         $\alpha_1 = \alpha_0$.
9:     **end if**
10: **end while**
11: $\Gamma_2(u) = \alpha_0$.
12: **Outputs:** $\Gamma_2(u)$.

Algorithm 5: Calculation of ω_*, u_*, and $\alpha_*(\mathcal{S})$

1: **Inputs**: A given setup $\mathcal{S} = \{f, \underline{p}, \bar{p}\}$ and ϵ. Set $u_1 = 0$, $u_2 = 1$, and $GAP = 1$.
2: **while** $|GAP| > \epsilon$ **do**
3:     Compute $u_* = \frac{u_1 + u_2}{2}$.
4:     $\Gamma_2(u_*) = \mathsf{GAMMA}_2(u_*)$.
5:     Set $\omega_* = F_{\underline{p}}^{-1}\left(\frac{h(p)}{\Gamma_2(u_*)} \right)$.
6:     $\Gamma_1(\omega_*, u_*) = \mathsf{GAMMA}_1(\omega_*, u_*)$.
7:     Set $GAP = \Gamma_2(u_*) - \Gamma_1(\omega_*, u_*)$.
8:     **if** $GAP > 0$ **then**
9:         $u_2 = u_*$.
10:     **else**
11:         $u_1 = u_*$.
12:     **end if**
13: **end while**
14: Set $\alpha_*(\mathcal{S}) = \mathsf{GAMMA}_1(\omega_*, u_*) = \mathsf{GAMMA}_2(u_*)$.
15: **Outputs:** $\omega_*, u_*, \alpha_*(\mathcal{S})$.

---

the computation of ω_*, u_* and $\alpha_*(\mathcal{S})$ can be performed efficiently via various numerical root-finding methods such as bisection searching. On the other hand (and perhaps more importantly), for a given convex setup $\mathcal{S}$, the parameters $\alpha_*(\mathcal{S}), \omega_*$, and u_* can be computed offline before the start of PM_{ϕ_*}. This means, we only need to solve either $BVP_1(\omega_*, u_*, \alpha_*)$ or $BVP_2(u_*, \alpha_*)$ "on-the-fly", depending on which segment the total resource utilization level is located.