Coordinating Markets via Perturbed Prices

Süleyman Kerimov, Pengyu Qian*

August 21, 2026

Abstract

We study the online allocation of finite, non-replenishable resources to self-interested buyers with private types who demand bundles of goods. The planner knows the type distribution and associated valuations but does not observe each arriving buyer's type. This informational constraint creates an implementability gap: type-observing allocation policies achieve constant regret, whereas conventional bid-price mechanisms, though implementable under private type information, can incur regret on the order of $\sqrt{T}$, where T is the number of buyers. The gap arises because posted prices can induce indifference, leaving the buyer's action not uniquely determined.

We close this gap with *Perturbed Posted Price* (`PPP`), an online and anonymous mechanism that posts a single vector of item prices together with a type-independent payment adjustment that depends on the chosen lottery over bundles. The adjustment induces the randomization needed when buyers are indifferent while preserving decentralized buyer self-selection. We prove that `PPP` achieves constant regret relative to the hindsight first-best for online packing (network revenue management), with a bound independent of both the horizon and initial inventory. We also extend our construction to online matching.

To the best of our knowledge, `PPP` is the first online, anonymous mechanism based on a single price vector to achieve constant regret in these private-information allocation settings. The mechanism is derived from a nonsmooth regularization of the fluid relaxation, revealing a broader role for regularization as a mechanism design instrument: it selects implementable behavior at indifference while stabilizing the allocation target against stochastic demand fluctuations.

*S. Kerimov is with the Jones Graduate School of Business, Rice University, Houston TX, USA, `kerimov@rice.edu`. P. Qian is with the Questrom School of Business, Boston University, Boston MA, USA, `pqian@bu.edu`.

1 Introduction

A central problem in online resource allocation is how to allocate scarce capacity among heterogeneous buyers who arrive sequentially. The planner faces two sources of uncertainty. First, allocation decisions are irrevocable and must be made before future demand is realized, requiring the planner to preserve capacity for valuable future uses. Second, in many markets, the planner does not observe the current buyer's type. She may know the distribution of buyer types and the valuation function associated with each type, but the buyer's realized type remains private.

Posted prices are the prevailing decentralized interface under such private information. The planner announces prices without observing the current buyer's type, and the buyer self-selects a utility-maximizing item or bundle. Cloud computing illustrates this informational structure: providers post prices for heterogeneous configurations of computing power, memory, storage, and network capacity, while customers privately know the value of workloads such as model training, fine-tuning, and inference. AWS Spot Instances (Amazon Web Services), for example, post prices that vary across instance types and availability zones. Similar informational constraints arise in applications such as airline seats, hotel rooms, and display advertising.

We study a finite-horizon online allocation problem with m types of indivisible resources and finite, non-replenishable initial inventories. A sequence of T self-interested buyers arrive online. There are finitely many buyer types, with each type associated with a publicly known valuation function over bundles of resources; buyers have quasilinear utility. Buyer types are drawn independently from a known distribution, but each realized type is privately known by the buyer. The planner seeks to maximize the expected social welfare using an online, anonymous mechanism that relies on a single vector of item prices and decentralized buyer self-selection. Our framework encompasses online packing (network revenue management) and online matching.

1.1 The implementability gap

We evaluate the performance of a mechanism relative to the *hindsight-first-best*, which observes the entire realized sequence of buyer types and selects the welfare-maximizing feasible allocation. Regret is the difference between the expected hindsight-first-best welfare and the mechanism's expected welfare. The strongest possible performance guarantee in this framework is *constant regret*, meaning that the regret bound is independent of both the horizon T and the initial inventory vector $\mathbf{b}$.

When the arriving buyer's type is observable, constant regret is attainable. In particular, Vera and Banerjee (2021) introduce compensated coupling and construct constant-regret allocation policies for online packing and matching. Related guarantees have been established for broader online control problems (Arlotto and Gurvich 2019, Bumpensanti and Wang 2020, Vera et al. 2021). These results show that uncertainty about future arrivals, by itself, does not preclude constant regret. These policies condition allocation decisions directly on the realized buyer type. They therefore do not immediately yield mechanisms in our private-information environment. To implement such allocations without observing private information, the planner must induce each buyer to select the desired allocation through the terms of the mechanism itself.

On the other hand, classical bid-price controls satisfy this informational requirement. They solve a deterministic fluid relaxation and post the resulting dual variables as resource prices (Talluri and Van Ryzin 1998, Jasin and Kumar 2012, Agrawal et al. 2014); each buyer's utility maximization then determines the allocation. This decentralizability can come at a performance cost. Standard bid-price controls can incur $\Omega(\sqrt{T})$ regret on network revenue management instances (Jasin and Kumar 2013). Thus, private information creates an *implementability gap.*

1.2 Why posted prices alone fail: indifference

A one-resource example isolates the obstruction. Suppose there are T periods, an initial inventory of $3T/4$, and two equally likely buyer types. A high type values one unit at 2, and a low type values it at 1. The fluid allocation serves all high types and exactly one half of the low types. A policy that observes types can regulate low-type admissions around this target while preserving capacity for later high types.

The unique optimal resource price is 1. At this price, a high type strictly prefers to buy, whereas a low type is indifferent between buying, not buying, and any lottery between the two. The posted price therefore implements the desired action for the high type but provides no instruction that induces the required one-half purchase probability for low types.

Under conventional re-solving bid-price controls, diffusion-scale fluctuations in early demand can lead to excessive admissions of low types and subsequently exhaust capacity needed by high types. The resulting displacement can produce regret on the order of $\sqrt{T}$ (Jasin and Kumar 2013). The problem is that a single linear price cannot communicate the desired interior action to an indifferent buyer. We simulate this example in Figure 1.

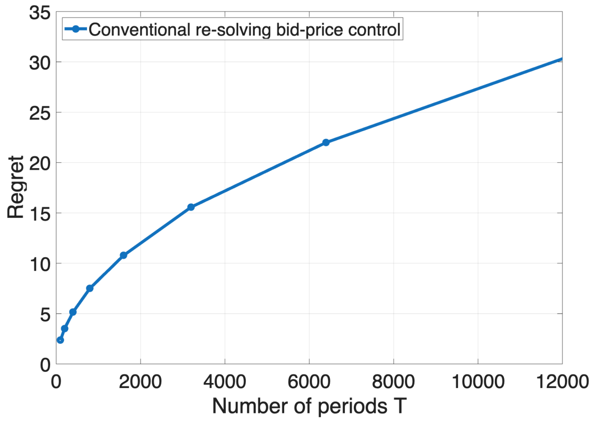

Figure 1: Illustration of the $\Omega(\sqrt{T})$ regret of the conventional re-solving bid-price control.

The example exposes a gap between the performance attainable by a centralized policy that observes buyer types and the performance attainable through posted prices. We ask:

> *Can an online, anonymous mechanism based on a single vector of item prices achieve constant regret without observing buyers' private types?*

1.3 Mechanism and contributions

We answer this question in the affirmative with *Perturbed Posted Price* (`PPP`), a minimal augmentation of the standard posted-price interface. In each period, `PPP` posts one vector of item prices and a publicly announced, type-independent payment adjustment that depends on the buyer's chosen lottery over bundles. The buyer pays the expected posted price payment together with the adjustment and chooses the lottery that maximizes her expected utility. The payment adjustment induces the randomization needed when posted prices alone create indifference.

The mechanism has a useful optimization interpretation. We show that `PPP` is action compatible with an ℓ_1-regularized fluid relaxation. The kink in the ℓ_1 term turns the knife-edge indifference condition in complementary slackness into a finite prescription: purchase, not purchase, or use a designated lottery, depending on whether the buyer's surplus is above, below, or within a small band around the perturbed dual price. This piecewise structure makes the regularized solution implementable through a single vector of item prices and a type-independent payment adjustment.

Our main performance result establishes that `PPP` achieves constant regret. The regret bound depends on the model primitives and the regularization budget but is independent of both the horizon T and the initial inventory $\mathbf{b}$. We establish an analogous result for online matching.

The analysis also reveals a broader methodological role for regularization in mechanism design. Our proof combines ideas from compensated coupling (Vera and Banerjee 2021) with Lipschitz stability of linear programs (Mangasarian and Shiau 1987). Regularization serves two functions simultaneously: it selects an implementable action when posted prices induce indifference, and it stabilizes the allocation target against stochastic perturbations in residual demand. To the best of our knowledge, `PPP` is the first online, anonymous mechanism based on a single item-price vector and a type-independent payment adjustment to achieve constant regret in these private-information online allocation problems.

1.4 Other related work

This paper contributes to the literature on online resource allocation, revenue management, and market design through prices.

Revenue management, bid prices, and re-solving. Classical revenue management replaces stochastic demand by its expectation and uses a deterministic program to derive capacity controls and shadow prices (Gallego and Van Ryzin 1994, 1997). Bid-price controls decentralize these dual variables as resource prices: a request is accepted when its reward exceeds the opportunity cost of the resources it consumes. Talluri and Van Ryzin (1998) establish asymptotic optimality in a large-market regime, and subsequent work develops dynamic re-solving methods with stronger performance guarantees under structural conditions (Jasin and Kumar 2012). Standard bid-price methods are attractive because they are decentralized, but known instances exhibit $\Omega(\sqrt{T})$ regret that frequent re-solving does not eliminate (Jasin and Kumar 2013). We introduce a minimal modification of the posted-price interface that achieves $O(1)$ regret.

Wang and Wang (2022) establish a closely related constant-regret result for price-based revenue management. Despite the common use of dynamically updated prices, the economic role of prices differs fundamentally across the two settings. In their model, prices are controls chosen by a seller to maximize expected revenue by influencing aggregate demand. We instead study prices in a Walrasian coordination role: an anonymous vector of item prices must decentralize a welfare-maximizing allocation among buyers with private preferences over bundles. Payments are transfers and do not enter our welfare objective. This distinction creates an implementation problem absent from the model of Wang and Wang (2022): when a buyer is indifferent, a linear item-price vector need not induce the randomized action required by the welfare target.

Constant-regret online allocation. A recent literature studies when online allocation policies can track stochastic demand with only a constant additive loss. Arlotto and Gurvich (2019) obtain constant regret in the multisecretary problem; Vera and Banerjee (2021) develop compensated coupling for finite-type packing and matching; and related guarantees follow from re-solving and geometric methods (Bumpensanti and Wang 2020, Vera et al. 2025). These results provide an important performance benchmark for our analysis. However, their policies condition actions on an observed request type. Our mechanism instead posts a type-independent menu without knowing the current buyer's private type.

Vera et al. (2021) provide a related pricing result in a model in which the seller observes a customer category and chooses a category-specific price, while the customer's valuation is private. Our information structure is different: the entire buyer type is private, and the mechanism posts a single anonymous vector of linear resource prices rather than a category-contingent pricing rule.

Prices, indifference, and market coordination. Our implementability problem is closely related to the literature on whether prices alone can coordinate decentralized markets. Hsu et al. (2016) show that a Walrasian price vector need not select a feasible welfare-maximizing allocation when buyers break ties independently, and they identify conditions under which uncoordinated tie-breaking creates limited over-demand. In a gross-substitutes environment, Paes Leme and Wong (2020) characterize robust Walrasian prices under which every buyer has a unique demanded bundle and aggregate demand clears the market. Rather than imposing genericity or uniqueness to eliminate indifference, our mechanism minimally perturbs the Walrasian price interface to recover implementability.

Regularization in revenue management and online allocation. Regularization has been used to stabilize bid-price controls and optimization solutions (Akan and Ata 2009, Ata and Akan 2015, Mangasarian and Shiau 1987), and nonlinear regularizers can encode objectives such as fairness (Balseiro et al. 2025). In our paper, regularization plays a different role as an implementation device. The kink in the ℓ_1 penalty converts a knife-edge indifference condition into a finite action prescription that can be implemented through a single vector of item prices and a type-independent payment adjustment. This also explains why the ℓ_1 form is important for our construction. A smooth penalty generally induces an interior choice probability that varies continuously with surplus, whereas the ℓ_1 penalty generates the piecewise response used by our mechanism and analysis.

Notation. We use $\mathbb{R}_+$ to denote the set of nonnegative real numbers and $\mathbb{Z}_+$ to denote the set of nonnegative integers. We denote the set of positive integers $\{1, 2, \cdots, n\}$ by $[n]$.

2 Model

We study an online resource allocation problem over a finite horizon of T periods. Time is indexed backward: $t \in \{T, T-1, \ldots, 1\}$ denotes the number of periods remaining, including the current period. The planner observes the public history and remaining inventory but not the current buyer's type.

2.1 Model primitives

Resources. There are m types of indivisible resources, indexed by $i \in [m]$. The initial inventory vector is $\mathbf{b} = (b_i)_{i\in[m]} \in \mathbb{Z}_+^m$. Let $\mathbf{b}_t = (b_{i,t})_{i\in[m]} \in \mathbb{Z}_+^m$ denote the inventory available at the beginning of period t, with $\mathbf{b}_T = \mathbf{b}$. Resources are non-replenishable.

Buyers and arrivals. There are n buyer types, indexed by $j \in [n]$. Exactly one buyer arrives in each period. The period-t buyer's type j_t is drawn independently from a known distribution $\mathbf{p} = (p_j)_{j\in[n]}$, where $p_j > 0$ for every j and $\sum_{j\in[n]} p_j = 1$. The distribution $\mathbf{p}$ is common knowledge. The realized type j_t is privately observed by the buyer and is never directly reported to the planner. The planner may observe only the buyer's menu choice and the realized allocation.

Bundle, feasibility, and valuations. Let $\mathcal{S}$ be a finite set of admissible bundles. Each bundle $S \in \mathcal{S}$ is represented by a consumption vector $\mathbf{a}_S = (a_{i,S})_{i\in[m]} \in \mathbb{Z}_+^m$, where $a_{i,S}$ is the number of units of resource i consumed by bundle S. The null bundle $\emptyset$ belongs to $\mathcal{S}$ and satisfies $\mathbf{a}_\emptyset = \mathbf{0}$. We assume that every non-null bundle consumes at least one unit of some resource. Given inventory $\mathbf{b}_t \in \mathbb{Z}_+^m$, the set of feasible bundles is

$$\mathcal{S}(\mathbf{b}_t) \triangleq \{S \in \mathcal{S} : \mathbf{a}_S \leq \mathbf{b}_t\},$$

where the inequality is componentwise. For any finite set $\mathcal{A}$, let

$$\Delta(\mathcal{A}) \triangleq \left\{ \boldsymbol{\sigma} \in \mathbb{R}_+^{\mathcal{A}} : \sum_{S\in\mathcal{A}} \sigma_S = 1 \right\}$$

denote the probability simplex over $\mathcal{A}$.

Each type j has an extended-real valuation function $v_j : \mathcal{S} \to \mathbb{R}_+ \cup \{-\infty\}$, with $v_j(\emptyset) = 0$. We use $v_j(S) = -\infty$ to exclude bundles that type j does not desire; consequently, any lottery assigning positive probability to such a bundle yields utility $-\infty$. Valuation functions and consumption vectors are common knowledge. If a type-j buyer purchases bundle S and makes total payment q, her quasilinear utility is

$$u_j(S, q) = v_j(S) - q.$$

Buyers are risk neutral and therefore evaluate lotteries by expected utility.

Social welfare. Let $S_t \in \mathcal{S}(\mathbf{b}_t)$ denote the random bundle allocated in period t. Pathwise social welfare is

$$\mathrm{SW} \triangleq \sum_{t=1}^{T} v_{j_t}(S_t)\,.$$

Item-price payments and payment adjustments are transfers between buyers and the planner and therefore do not enter social welfare. If the adjustment is instead interpreted as an external implementation cost, we show later that its total magnitude under our proposed mechanism is uniformly bounded independently of the horizon T. Thus, this alternative interpretation affects welfare by at most a horizon-independent constant.

2.2 Generalized posted-price mechanisms

A conventional posted-price mechanism announces item prices and lets the arriving buyer choose a utility-maximizing bundle (see, e.g., Milgrom 1989). We consider a minimal generalization in which the mechanism also announces a signed payment adjustment that depends on the buyer's chosen lottery over bundles. A positive adjustment is a rebate to the buyer, whereas a negative adjustment is an additional charge. This generalization preserves the decentralized posted-menu interface while allowing the mechanism to induce the desired randomization when posted prices alone leave the buyer indifferent.

Definition 1 (Generalized posted price mechanism)**.** In each period t, the mechanism proceeds as follows.

1. Without observing the arriving buyer's type, the planner publicly announces a price vector $\boldsymbol{\lambda}_t = (\lambda_{i,t})_{i\in[m]} \in \mathbb{R}^m_+$, and an adjustment function

   $$y(\cdot) : \Delta(\mathcal{S}(\mathbf{b}_t)) \to \mathbb{R}\,.$$

   A lottery $\boldsymbol{\sigma} \in \Delta(\mathcal{S}(\mathbf{b}_t))$ has expected item-price payment $\sum_{S\in\mathcal{S}(\mathbf{b}_t)} \sigma_S \boldsymbol{\lambda}_t^\top \mathbf{a}_S$ and receives adjustment $y(\boldsymbol{\sigma})$. Equivalently, its expected net payment is the expected item-price payment minus $y(\boldsymbol{\sigma})$.

2. The buyer privately observes her type j_t and chooses a lottery $\boldsymbol{\sigma}_t \in \Delta(\mathcal{S}(\mathbf{b}_t))$ maximizing her expected utility

$$\boldsymbol{\sigma}_t \in \underset{\Delta(\mathcal{S}(\mathbf{b}_t))}{\arg\max} \left\{ \sum_{S \in \mathcal{S}(\mathbf{b}_t)} \sigma_S \left(v_{j_t}(S) - \boldsymbol{\lambda}_t^\top \mathbf{a}_S \right) + y(\boldsymbol{\sigma}) \right\}. \tag{1}$$

3. The planner draws $S_t \sim \boldsymbol{\sigma}_t$, allocates S_t, charges the realized item-price payment $\boldsymbol{\lambda}_t^\top \mathbf{a}_{S_t}$ and applies the adjustment $y(\boldsymbol{\sigma}_t)$. The inventory updates as $\mathbf{b}_{t-1} = \mathbf{b}_t - \mathbf{a}_{S_t}$.

We focus on posted price mechanisms that are *online* and *anonymous*. Let $\mathcal{H}_t$ denote the public history at the beginning of period t: it contains initial inventory $\mathbf{b}$, all previously posted prices $\{\boldsymbol{\lambda}_\tau\}_{\tau=t+1}^T$, buyers' previous chosen lotteries $\{\boldsymbol{\sigma}_\tau\}_{\tau=t+1}^T$, and realized allocation outcomes $\{S_\tau\}_{\tau=t+1}^T$. It does not contain any buyer's private type.

Definition 2 (Online and anonymous posted-price mechanism)**.** A generalized posted-price mechanism is *online* if the price $\boldsymbol{\lambda}_t$ is measurable with respect to $\mathcal{H}_t$ in every period. It is *anonymous* if the same menu is offered to every possible current buyer and cannot condition on the current buyer's private type.

This is an indirect self-selection mechanism. The buyer is never asked to announce her type; after observing the posted menu, she simply chooses a utility-maximizing lottery according to (1). Individual rationality holds whenever the null lottery is feasible and the corresponding adjustment satisfies $y(\boldsymbol{\delta}_\emptyset) \geq 0$, because opting out then yields utility at least zero.

2.3 Optimality criterion

For a mechanism π, define its expected social welfare by

$$\mathrm{SW}^\pi \triangleq \mathbb{E}\left[\sum_{t=1}^T v_{j_t}(S_t^\pi)\right],$$

where the expectation is over buyer types and any randomization in the mechanism. We benchmark SW^π against the *hindsight first-best*. For a realized type sequence $\omega \triangleq (j_T, \cdots, j_1)$, the hindsight first-best social welfare is the value of:

$$\mathrm{SW}^{\mathrm{FB}}(\omega) \triangleq \left\{ \max_{S_1,\cdots,S_T \in \mathcal{S}} \sum_{\tau=1}^T v_{j_\tau}(S_\tau) \quad \text{subject to} \quad \sum_{\tau=1}^T \mathbf{a}_{S_\tau} \leq \mathbf{b} \right\}.$$

Let $\mathrm{SW}^{\mathrm{FB}} \triangleq \mathbb{E}[\mathrm{SW}^{\mathrm{FB}}(\omega)]$. The regret of mechanism π is

$$\textsc{Reg}(\pi) \triangleq \; \mathrm{SW}^{\mathrm{FB}} - \mathrm{SW}^{\pi} \,. \tag{2}$$

We say that π achieves constant regret if $\textsc{Reg}(\pi)$ is bounded by a constant that is independent of the horizon length T and the initial inventory $\mathbf{b}$, while possibly depending on fixed model primitives such as the type distribution, valuation functions, and resource-consumption vectors.

Remark 1 (Relationship with online resource allocation policies)**.** *The model nests standard online packing (including network revenue management) and online matching, once rewards are interpreted as buyer valuations in quasilinear utility. Most online resource allocation algorithms observe the arriving request type and then prescribe an allocation. Such algorithms remain valid performance benchmarks, but they are not directly implementable as mechanisms unless their prescribed actions can be implemented through a menu offered without observing the current type.*

2.4 Examples

The following examples illustrate how familiar bid-price controls relate to Definition 1. The first two can be implemented using a single item-price vector and an adjustment function, whereas the third generally requires a richer, nonlinear bundle-price menu.

Example 1 (Ordinary posted-prices and classical bid-price control)**.** If $y \equiv 0$, Definition 1 reduces to an ordinary posted-price mechanism. At time t, the planner posts $\boldsymbol{\lambda}_t$, and a type-j buyer chooses a bundle in

$$\arg\max_{S \in \mathcal{S}(\mathbf{b}_t)} \{v_j(S) - \boldsymbol{\lambda}_t^\top \mathbf{a}_S\}.$$

Because the null bundle is available, the buyer selects a non-null bundle only when its maximal surplus is nonnegative. Thus, bid prices implement allocation decision through buyer self-selection without requiring the planner to observe j_t. This decentralizability can come at a performance cost: ordinary bid-price controls need not achieve constant regret, and known network revenue management instances have $\Omega(\sqrt{T})$ regret (see, e.g., Jasin and Kumar 2013).

Example 2 (Quadratically perturbed bid-price control)**.** Ata and Akan (2015) propose a quadratically perturbed bid-price control for network revenue management. A decentralized mechanism implementation of this policy satisfies Definition 1, where the adjustment function is chosen as

$$y(\boldsymbol{\sigma}) = -\frac{\delta}{2} \sum_{S \in \mathcal{S} \setminus \{\emptyset\}} \sigma_S^2$$

for some $\delta > 0$. The opt-out lottery yields zero utility, so the implementation remains individually rational. They prove a regret bound of $O(\delta T)$.

Example 3 (The *BudgetRatio* policy of Vera et al. (2025))**.** Vera et al. (2025) propose the *BudgetRatio* policy and show that it admits a max-bid-price representation and achieves constant regret. In general, however, this representation falls outside Definition 1. Implementing it as a menu at time t would require posting a collection of price vectors $(\boldsymbol{\lambda}_{t,1}, \ldots, \boldsymbol{\lambda}_{t,\kappa_t})$, where κ_t is determined by the collection of relevant bases of the residual allocation problem. A buyer selecting bundle S would then be charged $\max_{k \in [\kappa_t]} \boldsymbol{\lambda}_{t,k}^\top \mathbf{a}_S$. Hence, bundle prices are given by the upper envelope of several linear price schedules rather than by a single linear item-price vector, as required by Definition 1. In contrast, our mechanism posts a single item-price vector in every period.

3 Perturbed Posted Price for Online Packing

In the online packing (network revenue management) setting, each type-j buyer is single-minded: she values a single non-null bundle $s_j \in \mathcal{S}$ at $r_j \triangleq v_j(s_j) \geq 0$, values the null bundle at zero, and assigns value $-\infty$ to every other bundle. Define

$$\mathbf{a}_j \triangleq \mathbf{a}_{s_j} = (a_{i,j})_{i \in [m]} \in \mathbb{Z}_+^m, \qquad A \triangleq [\mathbf{a}_1 \ \cdots \ \mathbf{a}_n] \in \mathbb{Z}_+^{m \times n}.$$

By the assumption in Section 2, every column of A is nonzero. Define

$$d \triangleq \max_{j \in [n]} \|\mathbf{a}_j\|_1, \qquad r_{\max} \triangleq \max_{j \in [n]} r_j, \qquad \alpha \triangleq \epsilon / T.$$

Here $\epsilon > 0$ is fixed independently of both T and the initial inventory. The coefficient α will scale the fluid regularizer.

3.1 Intuition from the regularized fluid problem

For an inventory vector $\mathbf{b} \in \mathbb{R}_+^m$ and a vector of demand arrivals $\mathbf{u} \in \mathbb{R}_+^n$, define the value of the *fluid problem* (`FP`):

$$F(\mathbf{b},\mathbf{u}) \triangleq \max_{\mathbf{z}\in\mathbb{R}^n_+} \quad \sum_{j\in[n]} r_j z_j \tag{3}$$
$$\text{subject to} \quad A\mathbf{z} \le \mathbf{b}, \qquad 0 \le z_j \le u_j, \quad j \in [n].$$

At state $(t, \mathbf{b}_t)$, the value of the fluid problem is $F(\mathbf{b}_t, t\mathbf{p})$. In the presence of a social planner that can observe the types of the arriving buyers, at any $t \in [T]$, the optimal solution z_j^* can be interpreted as the target number of type-j requests that the system should accept over the remaining t periods, while ignoring the stochasticity and integrality of decisions in the system. In our analysis, `FP` will serve as an upper bound on the hindsight first-best.

Posted-price mechanisms respect the informational constraints of the setting because the mechanism need not observe the buyer's type; instead, the buyer self-selects from the posted menu. To the best of our knowledge, no existing posted-price mechanism achieves constant regret without additional assumptions. We take a novel approach and introduce ℓ_1-regularization into the fluid problem. The resulting *regularized fluid problem* (`RFP`) at time t is:

$$\Phi_\alpha(\mathbf{b},\mathbf{u}) \triangleq \max_{\mathbf{z}\in\mathbb{R}^n_+} \quad \sum_{j\in[n]} \left(r_j z_j - \alpha \left| z_j - \frac{u_j}{2} \right| \right) \tag{4}$$
$$\text{subject to} \quad A\mathbf{z} \le \mathbf{b}, \qquad 0 \le z_j \le u_j, \quad j \in [n].$$

Let $\mathcal{Z}(\mathbf{b},\mathbf{u})$ denote the set of optimizers of (4). Because the feasible region is nonempty and compact, this set is nonempty. Our mechanism is constructed from the dual of this regularized fluid problem:

$$\texttt{Dual-RFP}(\mathbf{b},\mathbf{u}) \triangleq \left\{ \min_{\boldsymbol{\lambda}\in\mathbb{R}^m_+} \quad \mathbf{b}^\top \boldsymbol{\lambda} + \sum_{j\in[n]} \max_{0\le z_j\le u_j} \left\{ \left(r_j - \mathbf{a}_j^\top \boldsymbol{\lambda} \right) z_j - \alpha \left| z_j - \frac{u_j}{2} \right| \right\} \right\}. \tag{5}$$

Let $\Lambda(\mathbf{b},\mathbf{u})$ be its set of optimal solutions. Strong duality holds because (4) is a convex optimization problem. The next lemma records the exact coordinate-wise implication of optimality.

Lemma 1 (Coordinate characterization of the regularized fluid optimum)**.** *Fix* $(\mathbf{b},\mathbf{u})$*,* $\mathbf{z}^* \in$

$\mathcal{Z}(\mathbf{b}, \mathbf{u})$*, and* $\boldsymbol{\lambda}^* \in \Lambda(\mathbf{b}, \mathbf{u})$*. For every* $j \in [n]$*,*

$$z_j^* \in \begin{cases} \{u_j\}, & r_j - \mathbf{a}_j^\top \boldsymbol{\lambda}^* > \alpha, \\ [u_j/2, u_j], & r_j - \mathbf{a}_j^\top \boldsymbol{\lambda}^* = \alpha, \\ \{u_j/2\}, & -\alpha < r_j - \mathbf{a}_j^\top \boldsymbol{\lambda}^* < \alpha, \\ [0, u_j/2], & r_j - \mathbf{a}_j^\top \boldsymbol{\lambda}^* = -\alpha, \\ \{0\}, & r_j - \mathbf{a}_j^\top \boldsymbol{\lambda}^* < -\alpha. \end{cases} \tag{6}$$

Proof of Lemma 1. For any primal-dual optimal pair, equality of the primal and dual values implies that each z_j^* maximizes the scalar term in (5). On $[0, u_j/2]$ that scalar objective has slope $r_j - \mathbf{a}_j^\top \boldsymbol{\lambda}^* + \alpha$, and on $[u_j/2, u_j]$ it has slope $r_j - \mathbf{a}_j^\top \boldsymbol{\lambda}^* - \alpha$. The five cases in (6) follow immediately. □

Equation (6) reveals the key effect of the nonsmooth regularizer. When the net surplus lies strictly inside $(-\alpha, \alpha)$, i.e., near the indifference region, the regularized fluid target calls for serving exactly one half of type-j demand. This observation motivates a randomized choice between the desired bundle and the null bundle. We next implement such choices through the generalized posted-price interface.

3.2 The Perturbed Posted Price mechanism

We now implement action choices that are compatible with the rule in (6) through the generalized posted-price interface of Definition 1. Recall that each buyer of type j_t's best response is supported on $\{\emptyset, s_{j_t}\}$.[1] Then one can represent buyer of type j_t's choice of randomization as a single parameter $\delta \in [0, 1]$, where $\delta \triangleq \sigma_{s_{j_t}}$ and $1 - \delta = \sigma_\emptyset$. In every period, the mechanism uses the type-independent adjustment

$$y(\boldsymbol{\sigma}) \triangleq \frac{\alpha}{2} \left(\frac{1}{2} - \left| \sigma_\emptyset - \frac{1}{2} \right| \right). \tag{7}$$

Mechanism 1 summarizes the `PPP` mechanism.

[1]Since buyers are single-minded, any bundle other than $\emptyset$ and s_{j_t} is valued $-\infty$.

Mechanism 1 PPP for online packing

for $t = T, T-1, \ldots, 1$ **do**
  Select any $\boldsymbol{\lambda}_t^* \in \Lambda(\mathbf{b}_t, t\mathbf{p})$ by solving (5).
  Post $\boldsymbol{\lambda}_t^*$ and the adjustment function (7) on $\Delta(\mathcal{S}(\mathbf{b}_t))$.
  The buyer selects any best response in Proposition 1; draw S_t from that lottery.
  Allocate S_t and update $\mathbf{b}_{t-1} = \mathbf{b}_t - \mathbf{a}_{S_t}$.
end for

The adjustment function (7) is time-homogeneous and type-independent. We next characterize the buyer's best response and show that it is compatible with the regularized fluid target in (6).

Proposition 1 (Buyer best response)**.** *Fix a time period $t \in [T]$ and posted prices $\boldsymbol{\lambda}_t$. Denote the arriving buyer j_t's single desired bundle by $s_{j_t} \in \mathcal{S}$. If $\mathbf{a}_{j_t} \not\leq \mathbf{b}_t$, the null bundle is the unique finite-utility feasible choice of a type-j_t buyer. If $\mathbf{a}_{j_t} \leq \mathbf{b}_t$, every finite utility best response is supported on $\{\emptyset, s_{j_t}\}$. Writing $\delta \triangleq \sigma_{s_{j_t}}$, the best response correspondence under (7) is*

$$\mathrm{BR}_{j_t}(\boldsymbol{\lambda}_t) = \begin{cases} \{1\}\,, & \text{if } r_{j_t} - \mathbf{a}_{j_t}^\top \boldsymbol{\lambda}_t > \alpha/2, \\ [1/2,\ 1]\,, & \text{if } r_{j_t} - \mathbf{a}_{j_t}^\top \boldsymbol{\lambda}_t = \alpha/2, \\ \{1/2\}, & \text{if } -\alpha/2 < r_{j_t} - \mathbf{a}_{j_t}^\top \boldsymbol{\lambda}_t < \alpha/2, \\ [0, 1/2] & \text{if } r_{j_t} - \mathbf{a}_{j_t}^\top \boldsymbol{\lambda}_t = -\alpha/2, \\ \{0\} & \text{if } r_{j_t} - \mathbf{a}_{j_t}^\top \boldsymbol{\lambda}_t < -\alpha/2\,. \end{cases} \tag{8}$$

We illustrate the best response correspondence in Figure 2.

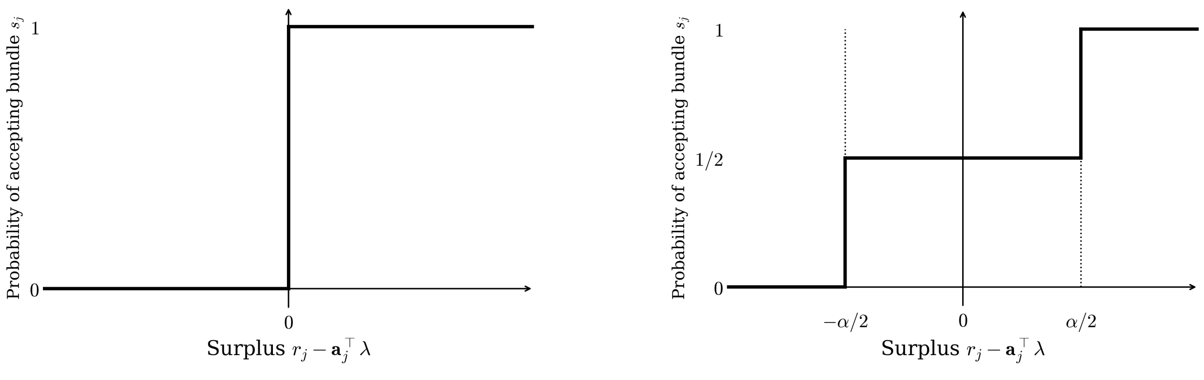

Figure 2: Buyer best response under conventional posted-price mechanism (left) and under PPP (right).

Proof of Proposition 1. Because all other non-null bundles have value $-\infty$, a finite-utility lottery is supported on $\{\emptyset, s_{j_t}\}$. The buyer therefore solves

$$\max_{0\le\delta\le 1}\ \delta(r_{j_t} - \mathbf{a}_{j_t}^\top \boldsymbol{\lambda}_t) - \frac{\alpha}{2}\left|\delta - \frac{1}{2}\right| + \frac{\alpha}{4}. \tag{9}$$

The objective has slope $r_{j_t} - \mathbf{a}_{j_t}^\top \boldsymbol{\lambda}_t + \alpha/2$ on $[0, 1/2]$ and slope $r_{j_t} - \mathbf{a}_{j_t}^\top \boldsymbol{\lambda}_t - \alpha/2$ on $[1/2, 1]$. This gives all five cases in (8). □

Two observations will be useful. First, the adjustment in any period is at most $\alpha/4 = \epsilon/(4T)$. Thus, if adjustments are interpreted as external implementation costs rather than transfers, the total cost incurred by PPP is at most $\epsilon/4$. Second, Proposition 1 allows multiple utility-maximizing actions at the boundary $r_{j_t} - \mathbf{a}_{j_t}^\top \boldsymbol{\lambda}_t^* = \pm\alpha/2$. Our performance guarantee is robust to arbitrary tie-breaking at these boundary points. The next lemma formalizes this robustness by showing that every decentralized best response remains aligned with the regularized fluid target.

Lemma 2 (Tie-breaking-robust action implications). *Fix* $(\mathbf{b}, \mathbf{u})$, $\mathbf{z}^* \in \mathcal{Z}(\mathbf{b}, \mathbf{u})$, *and* $\boldsymbol{\lambda}^* \in \Lambda(\mathbf{b}, \mathbf{u})$. *Suppose the type-$j$ desired bundle is feasible and the buyer selects any* $\delta \in \mathrm{BR}_j(\boldsymbol{\lambda}^*)$. *Then*

$$\delta > 0 \implies z_j^* \ge \frac{u_j}{2}, \qquad \delta < 1 \implies z_j^* \le \frac{u_j}{2}. \tag{10}$$

Consequently, if $X \sim \mathrm{Bernoulli}(\delta)$, *then almost surely*

$$X = 1 \implies z_j^* \ge \frac{u_j}{2}, \qquad X = 0 \implies z_j^* \le \frac{u_j}{2}.$$

Proof of Lemma 2. If $\delta > 0$, (8) implies $r_j - \mathbf{a}_j^\top \boldsymbol{\lambda}^* \ge -\alpha/2 > -\alpha$. Lemma 1 therefore gives $z_j^* \ge u_j/2$. If $\delta < 1$, (8) implies $r_j - \mathbf{a}_j^\top \boldsymbol{\lambda}^* \le \alpha/2 < \alpha$, so the same lemma gives $z_j^* \le u_j/2$. The realized-action implications follow because $X = 1$ has probability zero when $\delta = 0$, and $X = 0$ has probability zero when $\delta = 1$. □

We now state the performance guarantee for PPP.

Theorem 1 (Constant regret of PPP). *For online packing (network revenue management), under any selection from the buyer best-response correspondence, the Perturbed Posted Price mechanism PPP satisfies*

$$\textsc{Reg}(\texttt{PPP}) \le \epsilon + \left(dr_{\max} + \frac{\epsilon}{2}\right)\left(5n + 16(\kappa+1)^2 n^3 \sum_{j\in[n]} \frac{1}{p_j^2}\right),$$

where the constant $\kappa < \infty$ depends only on the consumption matrix A. The bound is independent of both the horizon T and the initial inventory $\mathbf{b}$.

3.3 Analysis

We compare PPP with two auxiliary benchmarks. Fix a realized type sequence $\omega = (j_T, \ldots, j_1)$. For $t \in \{0, 1, \ldots, T\}$, define

$$\Gamma_j(t) \triangleq \sum_{s=1}^{t} \mathbb{1}\{j_s = j\}, \qquad \mathbf{\Gamma}(t) \triangleq (\Gamma_j(t))_{j\in[n]}.$$

Thus $\|\mathbf{\Gamma}(t)\|_1 = t$ and $\mathbf{\Gamma}(t-1) = \mathbf{\Gamma}(t) - \mathbf{e}_{j_t}$. At the state reached by PPP, define the unregularized and regularized hindsight values by

$$F_t \triangleq F(\mathbf{b}_t, \mathbf{\Gamma}(t)), \qquad \Phi_t \triangleq \Phi_\alpha(\mathbf{b}_t, \mathbf{\Gamma}(t)),$$

with $F_0 = \Phi_0 = 0$. Since F_T allows fractional allocations, it upper-bounds the hindsight first best on every sample path:

$$\mathrm{SW}^{\mathrm{FB}}(\omega) \leq F_T. \tag{11}$$

Our first step bounds the pathwise cost of regularization, i.e., the gap between F_t and Φ_t.

Proposition 2 (Cost of regularization)**.** *For every $\mathbf{b} \in \mathbb{R}^m_+$ and $\mathbf{u} \in \mathbb{R}^n_+$,*

$$0 \leq F(\mathbf{b}, \mathbf{u}) - \Phi_\alpha(\mathbf{b}, \mathbf{u}) \leq \frac{\alpha}{2}\|\mathbf{u}\|_1. \tag{12}$$

Consequently, $0 \leq F_t - \Phi_t \leq \epsilon t/(2T)$.

Proof of Proposition 2. The first inequality follows because the regularization penalty is non-negative. Let $\widetilde{\mathbf{z}}$ be an optimizer of $F(\mathbf{b}, \mathbf{u})$. It is feasible for (4), and $0 \leq \widetilde{z}_j \leq u_j$ implies $|\widetilde{z}_j - u_j/2| \leq u_j/2$. Therefore,

$$\Phi_\alpha(\mathbf{b}, \mathbf{u}) \geq F(\mathbf{b}, \mathbf{u}) - \alpha \sum_{j\in[n]} \left|\widetilde{z}_j - \frac{u_j}{2}\right| \geq F(\mathbf{b}, \mathbf{u}) - \frac{\alpha}{2}\|\mathbf{u}\|_1. \qquad \square$$

We next compare the different actions taken under PPP and under the optimal solution of hindsight optimal regularized problem Φ_t. Let

$$X_t \triangleq \mathbb{1}\{S_t = s_{j_t}\}$$

be the realized purchase indicator under `PPP`, and recall that $\mathcal{Z}(\mathbf{b}_t, \mathbf{\Gamma}(t))$ is the hindsight regularized optimizer set (see (4)). Intuitively, we will compare these two decisions such that whenever `PPP` takes an online action, we will ask whether this action can be embedded in $\mathcal{Z}(\mathbf{b}_t, \mathbf{\Gamma}(t))$. If such an embedding is not possible, we say that there is a *disagreement*, and we will keep track of how much *compensation* is needed to force $\mathcal{Z}(\mathbf{b}_t, \mathbf{\Gamma}(t))$ to track the actions of `PPP`.

Given any $z_t^* \in \mathcal{Z}(\mathbf{b}_t, \mathbf{\Gamma}(t))$, note that $z_{j_t,t}^*$ is the total number of buyers of type j_t served among the remaining $\Gamma_{j_t}(t)$ buyers. Therefore, if there exists an optimal solution $z_t^* \in \mathcal{Z}(\mathbf{b}_t, \mathbf{\Gamma}(t))$ such that $z_{j_t}^* \geq 1$, then Φ_t can serve the current buyer without changing its value. Similarly, if there exists an optimal solution $z_t^* \in \mathcal{Z}(\mathbf{b}_t, \mathbf{\Gamma}(t))$ such that $z_{j_t}^* \leq \Gamma_{j_t}(t) - 1$, then Φ_t can reject the current buyer without changing its value. Define the disagreement event

$$E_t \triangleq \left\{ X_t = 1, \sup_{\mathbf{z} \in \mathcal{Z}(\mathbf{b}_t, \mathbf{\Gamma}(t))} z_{j_t} < 1 \right\} \cup \left\{ X_t = 0, \inf_{\mathbf{z} \in \mathcal{Z}(\mathbf{b}_t, \mathbf{\Gamma}(t))} z_{j_t} > \Gamma_{j_t}(t) - 1 \right\}. \tag{13}$$

Because $\mathcal{Z}(\mathbf{b}_t, \mathbf{\Gamma}(t))$ is compact, the supremum and infimum in (13) are attained. Outside E_t, we can select a hindsight regularized optimizer whose solution on the current buyer agrees with the realized `PPP` action.

To translate a disagreement into welfare loss, the next deterministic lemma bounds the loss caused by forcing one action in the unregularized hindsight problem F_t.

Lemma 3 (One-step trimming bound)**.** *Let* $\mathbf{u} \in \mathbb{Z}_+^n$ *with* $u_j \geq 1$*, let* $x \in \{0, 1\}$ *satisfy* $x\mathbf{a}_j \leq \mathbf{b}$*, and set* $\mathbf{b}' = \mathbf{b} - x\mathbf{a}_j$ *and* $\mathbf{u}' = \mathbf{u} - \mathbf{e}_j$*. Then*

$$F(\mathbf{b}, \mathbf{u}) - r_j x - F(\mathbf{b}', \mathbf{u}') \leq d r_{\max}. \tag{14}$$

Proof of Lemma 3. An incorrect rejection of type-j buyer decreases F by at most r_j, since any optimizer can be made feasible by removing at most one unit of type-j service. Similarly, reducing the capacity of any resource by one decreases F by at most $r_{\max}$: any resulting violation is at most one and, because A is integral, it can be eliminated by removing at most one unit of total service from classes that use that resource. Therefore, an incorrect acceptance of type j buyer decreases F by at most $\|\mathbf{a}_j\|_1 r_{\max}$. We have

$$\begin{aligned} F(\mathbf{b}, \mathbf{u}) - r_j x - F(\mathbf{b} - x\mathbf{a}_j, \mathbf{u} - \mathbf{e}_j) &\leq (1 - x) r_j + x \|\mathbf{a}_j\|_1 r_{\max} \\ &\leq d r_{\max}. \end{aligned}$$

□

Proposition 3 (Pathwise compensation bound)**.** *For any realized type sequence and every realization of the mechanism's lotteries,*

$$\Phi_T - \sum_{t=1}^{T} r_{j_t} X_t \le \frac{\epsilon}{2} + (dr_{\max} + \frac{\epsilon}{2}) \sum_{t=1}^{T} \mathbb{1}\{E_t\}. \tag{15}$$

Proof of Proposition 3. Define the one-period increment

$$\Delta_t \triangleq \Phi_t - r_{j_t} X_t - \Phi_{t-1}.$$

These increments telescope to the left-hand side of (15). Suppose first that E_t does not occur. If $X_t = 1$, select $\mathbf{z} \in \mathcal{Z}(\mathbf{b}_t, \mathbf{\Gamma}(t))$ with $z_{j_t} \ge 1$. Then $\mathbf{z} - \mathbf{e}_{j_t}$ is feasible for the hindsight regularized program at $t-1$. The reward terms cancel, and only the center of coordinate j_t in the regularization term changes. Hence it follows from triangle inequality that

$$\Delta_t \le \alpha \left(\left| z_{j_t} - 1 - \frac{\Gamma_{j_t}(t) - 1}{2} \right| - \left| z_{j_t} - \frac{\Gamma_{j_t}(t)}{2} \right| \right) \le \frac{\alpha}{2}.$$

If $X_t = 0$, select $\mathbf{z} \in \mathcal{Z}(\mathbf{b}_t, \mathbf{\Gamma}(t))$ with $z_{j_t} \le \Gamma_{j_t}(t) - 1$. The same vector is feasible at $t-1$, and the identical triangle inequality argument again gives $\Delta_t \le \alpha/2$.

If E_t occurs, Proposition 2 and Lemma 3 give

$$\begin{aligned} \Delta_t &\le F_t - r_{j_t} X_t - F_{t-1} + \frac{\alpha}{2}(t-1) \\ &\le dr_{\max} + \frac{\epsilon}{2}. \end{aligned}$$

Thus, for every t,

$$\Delta_t \le \frac{\alpha}{2} + \left(dr_{\max} + \frac{\epsilon}{2} \right) \mathbb{1}\{E_t\}.$$

Summing over t and using $T\alpha/2 = \epsilon/2$ proves the result. □

It remains to control the expected number of disagreements. The key ingredient is stability of the regularized LP with respect to the demand upper-bound vector.

Lemma 4 (Stability of the regularized optimizer set)**.** *There exists $\kappa < \infty$, depending only on the fixed constraint matrix determined by A, such that for every $\mathbf{b}$, every $\mathbf{u}, \mathbf{u}' \in \mathbb{R}^n_+$, and every*

$\mathbf{z} \in \mathcal{Z}(\mathbf{b}, \mathbf{u})$, *there is* $\mathbf{z}' \in \mathcal{Z}(\mathbf{b}, \mathbf{u}')$ *satisfying*

$$\|\mathbf{z} - \mathbf{z}'\|_\infty \leq \kappa \|\mathbf{u} - \mathbf{u}'\|_1. \tag{16}$$

Proof of Lemma 4. Introduce variables y_j and rewrite (4) as the linear program

$$\begin{aligned} \max_{\mathbf{z},\mathbf{y}} \quad & \sum_{j \in [n]} (r_j z_j - \alpha y_j) \\ \text{subject to} \quad & A\mathbf{z} \leq \mathbf{b}, \\ & y_j - z_j \geq -u_j/2, \qquad y_j + z_j \geq u_j/2, \\ & 0 \leq y_j \leq u_j/2, \qquad 0 \leq z_j \leq u_j, \quad j \in [n]. \end{aligned}$$

The added upper bound on y_j does not change the optimizer set because every feasible $z_j \in [0, u_j]$ satisfies $|z_j - u_j/2| \leq u_j/2$, and an optimum always takes $y_j = |z_j - u_j/2|$. The coefficient matrix is therefore fixed once A is fixed; only the right-hand side changes with $(\mathbf{b}, \mathbf{u})$. The Lipschitz stability theorem for linear-program solution sets in Mangasarian and Shiau (1987, Theorem 2.4), applied to this formulation and projected onto the $\mathbf{z}$ coordinates, gives (16). $\square$

We can now bound the expected number of disagreements.

Proposition 4 (Expected number of disagreements)**.** *The disagreement events in* (13) *satisfy*

$$\sum_{t=1}^{T} \mathbb{P}(E_t) \leq 5n + 16(\kappa + 1)^2 n^3 \sum_{j \in [n]} \frac{1}{p_j^2}. \tag{17}$$

Proof of Proposition 4. For each j, split the type-j disagreement events into

$$\begin{aligned} E_{t,j}^{+} &\triangleq E_t \cap \{j_t = j, X_t = 1\}, \\ E_{t,j}^{-} &\triangleq E_t \cap \{j_t = j, X_t = 0,\ \mathbf{a}_j \leq \mathbf{b}_t\}, \\ E_{t,j}^{0} &\triangleq E_t \cap \{j_t = j, X_t = 0,\ \mathbf{a}_j \not\leq \mathbf{b}_t\}. \end{aligned}$$

Then E_t is the disjoint union of these events over $j \in [n]$.

Let $D_t \triangleq \|\mathbf{\Gamma}(t) - t\mathbf{p}\|_1$. Fix j and suppose $p_j t \geq 4$. Take any optimizer $\mathbf{z}^{\mathrm{F}} \in \mathcal{Z}(\mathbf{b}_t, t\mathbf{p})$ and the corresponding $\mathbf{z}^{\mathrm{H}} \in \mathcal{Z}(\mathbf{b}_t, \mathbf{\Gamma}(t))$ supplied by Lemma 4. If $E_{t,j}^{+}$ occurs, the action implications in (10) give $z_j^{\mathrm{F}} \geq p_j t/2$, whereas the definition of disagreement gives $z_j^{\mathrm{H}} < 1$. Therefore,

$$\kappa D_t \geq |z_j^{\mathrm{F}} - z_j^{\mathrm{H}}| > \frac{p_j t}{2} - 1 \geq \frac{p_j t}{4}.$$

If $E_{t,j}^-$ occurs, then $z_j^{\mathrm{F}} \leq p_j t/2$ and $z_j^{\mathrm{H}} > \Gamma_j(t) - 1$. Since $\Gamma_j(t) \geq p_j t - D_t$,

$$\kappa D_t > \Gamma_j(t) - 1 - \frac{p_j t}{2} \geq \frac{p_j t}{2} - 1 - D_t,$$

so $(\kappa+1)D_t > p_j t/2 - 1 \geq p_j t/4$. In either case,

$$E_{t,j}^+ \cup E_{t,j}^- \subseteq \left\{ D_t > \frac{p_j t}{4(\kappa+1)} \right\}. \tag{18}$$

If $\|\mathbf{\Gamma}(t) - t\mathbf{p}\|_1 \geq a$, at least one coordinate differs from its mean by at least a/n. Hoeffding's inequality and a union bound therefore imply, for $p_j t \geq 4$,

$$\mathbb{P}(E_{t,j}^+ \cup E_{t,j}^-) \leq 2n \exp\left(-\frac{p_j^2 t}{8n^2(\kappa+1)^2} \right).$$

Summing the geometric series and using $e^x - 1 \geq x$ gives

$$\sum_{t: p_j t \geq 4} \mathbb{P}(E_{t,j}^+ \cup E_{t,j}^-) \leq \frac{16(\kappa+1)^2 n^3}{p_j^2}. \tag{19}$$

For the remaining periods,

$$\sum_{t: p_j t < 4} \mathbb{P}(E_{t,j}^+ \cup E_{t,j}^-) \leq \sum_{t: p_j t < 4} \mathbb{P}(j_t = j) < 4. \tag{20}$$

It remains to bound $E_{t,j}^0$. If this event occurs, every hindsight regularized optimizer has $z_j > \Gamma_j(t) - 1$. If $\Gamma_j(t) \geq 2$, then $z_j > 1$, and feasibility would imply $\mathbf{a}_j \leq \mathbf{b}_t$, contradicting the definition of $E_{t,j}^0$. Hence

$$E_{t,j}^0 \subseteq \{j_t = j,\ \Gamma_j(t) = 1\},$$

and independence gives

$$\sum_{t=1}^{T} \mathbb{P}(E_{t,j}^0) \leq \sum_{t=1}^{\infty} p_j (1-p_j)^{t-1} = 1. \tag{21}$$

Summing (19), (20), and (21) over j proves (17). $\square$

Proof of Theorem 1. By (11), Proposition 2, and Proposition 3,

$$\begin{aligned} \mathrm{SW}^{\mathrm{FB}}(\omega) - \sum_{t=1}^{T} r_{j_t} X_t &\le F_T - \Phi_T + \Phi_T - \sum_{t=1}^{T} r_{j_t} X_t \\ &\le \epsilon + \left(d r_{\max} + \frac{\epsilon}{2} \right) \sum_{t=1}^{T} \mathbb{1}\{E_t\}. \end{aligned}$$

Taking expectations and applying Proposition 4 yields the stated bound. □

Remark 2. *The proof of Theorem 1 follows the same broad intuition as the compensated-coupling approach of Vera and Banerjee (2021): force a hindsight benchmark to follow the online policy and charge a bounded compensation whenever their actions disagree. Their machinery, however, does not apply directly here because it requires the offline benchmark to be the value-to-go function of a sequential offline control problem. Our regularized hindsight benchmark generally does not satisfy such a Bellman recursion, because removing the current buyer also re-centers the regularization term. We therefore use the compensated-coupling intuition but establish the comparison through a direct pathwise compensation argument.*

4 Extensions and Final Remarks

The binary opt-in or opt-out structure of online packing is not essential to our methodology. We illustrate this by extending `PPP` to online matching, where a buyer may choose any one of the available resources.

Online matching. There are m resources and n buyer types. A type-j buyer obtains value $r_{ij} \ge 0$ from being matched to resource $i \in [m]$ and value zero from remaining unmatched. Every buyer therefore faces the same action set $\mathcal{S} = \{\emptyset, \{1\}, \ldots, \{m\}\}$. We have $v_j(\{i\}) = r_{ij}$, $\mathbf{a}_{\{i\}} = \mathbf{e}_i$, with $v_j(\emptyset) = 0$. For online matching, `PPP` uses the adjustment

$$y(\boldsymbol{\sigma}) = \frac{\alpha}{2} \left(1 - \|\boldsymbol{\sigma}\|_\infty\right), \qquad \boldsymbol{\sigma} \in \Delta^{m+1}. \tag{22}$$

Every buyer chooses a lottery over the same m resources and the null action. The corresponding *regularized fluid problem* at time t is:

$$\begin{aligned} \Phi_\alpha(\mathbf{b}, \mathbf{u}) \triangleq \max_{\{\mathbf{z}_j\}_{j\in[n]}} \quad & \sum_{j\in[n]} \left[\sum_{i\in[m]} r_{ij} z_{ji} - \alpha \left(\|\mathbf{z}_j\|_\infty - \frac{u_j}{m+1} \right) \right] \\ \text{subject to} \quad & \sum_{j\in[n]} z_{ji} \le b_i, \qquad i \in [m], \end{aligned} \tag{23}$$

$$\mathbf{z}_j \geq \mathbf{0}, \qquad \sum_{i \in [m] \cup \{\emptyset\}} z_{ji} = u_j, \qquad j \in [n].$$

Rather than repeating the entire online packing analysis, we isolate the step in which the richer action space matters. The following lemma is the matching analogue of Lemma 2: every action that the buyer may actually choose receives a nonnegligible amount of mass in the regularized fluid target.

Lemma 5 (Matching action implication). *Fix* $(\mathbf{b}, \mathbf{u})$ *with* $u_j > 0$*, let* $\mathbf{z}^*$ *be any regularized fluid optimizer, and let* $\boldsymbol{\lambda}^*$ *be any corresponding dual optimum. Every type-*j *buyer best response* $\boldsymbol{\sigma}$ *satisfies*

$$\sigma_i > 0 \quad \Longrightarrow \quad z^*_{j,i} \geq \frac{u_j}{m+1}, \qquad i \in [m] \cup \{\emptyset\}.$$

Proof of Lemma 5. Define the net surplus of resource i by $q_i \triangleq r_{ij} - \lambda^*_i$ and set $q_\emptyset = 0$. Ignoring the additive constant in (22), the buyer solves

$$\max_{\boldsymbol{\sigma} \in \Delta^{m+1}} \left\{ \mathbf{q}^\top \boldsymbol{\sigma} - \frac{\alpha}{2} \|\boldsymbol{\sigma}\|_\infty \right\}.$$

The normalized fluid allocation $\mathbf{z}^*_j / u_j$ solves the same problem with the larger regularization coefficient α. To compare these two problems, consider a generic coefficient $c > 0$. Order the surpluses as $q_{(1)} \geq \cdots \geq q_{(m+1)}$. Every extreme optimizer is uniform on a set of the k highest-surplus actions. Define

$$H_k = \sum_{\ell < k} \big(q_{(\ell)} - q_{(k)} \big), \qquad H_1 = 0.$$

Support size k is optimal exactly when $H_k \leq c \leq H_{k+1}$.

Now suppose $\sigma_i > 0$. Since every optimal lottery is a convex combination of optimal extreme points, some optimal extreme point for $c = \alpha/2$ contains action i. Let $b_i = \big|\{\ell : q_\ell \geq q_i\}\big|$ be the number of actions whose surplus is at least that of i. The preceding characterization, including ties, implies $H_{b_i} \leq \alpha/2$. Consider an extreme fluid optimizer, which corresponds to $c = \alpha$. If its support had size $k < b_i$, optimality would require $\alpha \leq H_{k+1} \leq H_{b_i} \leq \frac{\alpha}{2}$, a contradiction. Hence every extreme fluid optimizer contains action i. Because it is uniform on a support of at most $m+1$ actions, it assigns probability at least $1/(m+1)$ to i. Every fluid optimum is a convex combination of optimal extreme points, and therefore $z^*_{ji}/u_j \geq 1/(m+1)$. □

Lemma 5 supplies the key disagreement implication used in the online packing analysis.

The pathwise compensation argument then proceeds as in Proposition 3. Consequently, the same proof architecture yields constant expected regret for online matching, independent of the horizon and initial inventory.

Future work. Several directions remain open. The present guarantees rely on a finite and known type distribution, independent arrivals, and non-replenishable inventories. It would be valuable to determine whether the same implementation principle extends to unknown or non-stationary demand distributions. More broadly, our results suggest studying which nonsmooth regularizers both admit simple anonymous menu implementations and provide the stability required for horizon-uniform performance. Such a characterization could expand the range of online markets in which decentralized prices coordinate privately informed agents without sacrificing the strongest attainable regret guarantees.

References

Shipra Agrawal, Zizhuo Wang, and Yinyu Ye. A dynamic near-optimal algorithm for online linear programming. *Operations Research*, 62(4):876–890, 2014.

Mustafa Akan and Barış Ata. Bid-price controls for network revenue management: Martingale characterization of optimal bid prices. *Mathematics of Operations Research*, 34(4):912–936, 2009.

Amazon Web Services. Spot instances. *Amazon EC2 User Guide*. URL https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-spot-instances.html.

Alessandro Arlotto and Itai Gurvich. Uniformly bounded regret in the multisecretary problem. *Stochastic Systems*, 9(3):231–260, 2019.

Barış Ata and Mustafa Akan. On bid-price controls for network revenue management. *Stochastic Systems*, 5(2):268–323, 2015.

Santiago R. Balseiro, Haihao Lu, and Vahab Mirrokni. Regularized online allocation problems: Fairness and beyond. *Manufacturing & Service Operations Management*, 27(3):720–735, 2025. doi: 10.1287/msom.2022.0212.

Pornpawee Bumpensanti and He Wang. A re-solving heuristic with uniformly bounded loss for network revenue management. *Management Science*, 66(7):2993–3009, 2020.

Guillermo Gallego and Garrett Van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management science*, 40(8):999–1020, 1994.

Guillermo Gallego and Garrett Van Ryzin. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations research*, 45(1):24–41, 1997.

Justin Hsu, Jamie Morgenstern, Ryan Rogers, Aaron Roth, and Rakesh Vohra. Do prices coordinate markets? In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 440–453, 2016.

Stefanus Jasin and Sunil Kumar. A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research*, 37(2):313–345, 2012.

Stefanus Jasin and Sunil Kumar. Analysis of deterministic lp-based booking limit and bid price controls for revenue management. *Operations Research*, 61(6):1312–1320, 2013.

Olvi L Mangasarian and T-H Shiau. Lipschitz continuity of solutions of linear inequalities, programs and complementarity problems. *SIAM Journal on Control and Optimization*, 25(3):583–595, 1987.

Paul Milgrom. Auctions and bidding: A primer. *Journal of economic perspectives*, 3(3):3–22, 1989.

Renato Paes Leme and Sam Chiu-wai Wong. Computing walrasian equilibria: Fast algorithms and structural properties. *Mathematical Programming*, 179:343–384, 2020. doi: 10.1007/s10107-018-1334-9.

Kalyan Talluri and Garrett Van Ryzin. An analysis of bid-price controls for network revenue management. *Management science*, 44(11-part-1):1577–1593, 1998.

Alberto Vera and Siddhartha Banerjee. The bayesian prophet: A low-regret framework for online decision making. *Management Science*, 67(3):1368–1391, 2021.

Alberto Vera, Siddhartha Banerjee, and Itai Gurvich. Online allocation and pricing: Constant regret via bellman inequalities. *Operations Research*, 69(3):821–840, 2021.

Alberto Vera, Alessandro Arlotto, Itai Gurvich, and Eli Levin. Dynamic resource allocation: The geometry and robustness of constant regret. Technical Report 4, 2025.

Yining Wang and He Wang. Constant regret resolving heuristics for price-based revenue management. *Operations Research*, 70(6):3538–3557, 2022.