
Explaining the Law of Supply and Demand via Online Learning

Stratis Skoulakis
Aarhus University
stratis@cs.au.dk

Abstract

The *law of supply and demand* asserts that in a perfectly competitive market, the price of a good adjusts to a *market clearing price*. In a market clearing price p^* the number of sellers willing to sell the good at p^* equals the number of sellers willing to buy the good at price p^* . In this work, we provide a mathematical foundation on the law of supply and demand through the lens of online learning. Specifically, we demonstrate that if each seller employs a no-swap regret algorithm to set their individual selling price—aiming to maximize its individual revenue—the collective pricing dynamics converge to the market-clearing price p^* . Our findings offer a novel perspective on the law of supply and demand, framing it as the emergent outcome of an adaptive learning processes among sellers.

1 Introduction

The *law of supply and demand* is a fundamental economic principle explaining the price of a good in a perfectly competitive market [25]. A perfectly competitive market consists of a large number of sellers and buyers where each seller is interested in selling one unit of an indistinguishable good and each buyer is interested in buying one unit of the good. The market is thus described by a supply $\mathcal{S}$ and a demand curve $\mathcal{D}$ where $\mathcal{S}(p)/\mathcal{D}(p)$ is the number of sellers/buyers willing to sell/buy the good at price p . The law of supply and demand states that the price of the good will converge to the *market clearing price* p^* where $\mathcal{S}(p^*) = \mathcal{D}(p^*)$ —the number of sellers equals the number of buyers [22, 28].

The idea behind the law of supply and demand is very intuitive¹. If the price is higher than p^* then more sellers are willing to sell than buyers willing to buy, leading to a *market surplus*. As a result, some sellers are not able to sell good and will lower their prices to attract buyers. As prices fall, more buyers are willing to buy, and fewer sellers are willing to sell. Conversely, if the price is below p^* this will create a *market shortage* that will in turn cause an upward trend of the prices. These adjustments continue until all selling prices reach p^* where there is neither market surplus or shortage [25, 22, 28].

However, upon closer examination, the explanation above does not rule out the possibility of persistent price fluctuations—where a market surplus leads to a shortage, which then causes another surplus, and so on—preventing the selling prices from ever converging to the market-clearing price. The latter raises the following fundamental question:

Question 1. *Why do selling prices eventually stabilize to the market clearing price p^* instead of constantly oscillating around it?*

Surprisingly, despite the *law of supply and demand* being one of the most fundamental principles in modern economics, this question remains unanswered. In this work, we provide answers to Question 1 through the lens of *game theory* and *online learning*.

¹See also <https://www.investopedia.com/terms/l/law-of-supply-demand.asp>

Our Contribution and Techniques As already discussed, sellers choose their prices to maximize individual revenue—that is, to sell their goods at the highest possible price. We model this price competition through a suitable pricing game (see Definition 3). In this game, each seller selects a price for their good, after which buyers arrive sequentially and purchase from the lowest-priced available seller.

The main result of this work consists in establishing that any Correlated Equilibrium (CE) [3] of the pricing game effectively coincides with the market clearing price of the market.

Informal Theorem *In any correlated equilibrium, the selling price is the market clearing price p^* .*

To establish this result, we introduce a novel primal-dual technique. Specifically, we construct an appropriate linear program whose optimal value serves as a lower bound on the probability that a Correlated Equilibrium (CE) selects the market clearing price. We then show that this probability is at least 1 by demonstrating that the dual program admits a feasible solution with objective value 1. To the best of our knowledge, this is the first time the *dual-fitting technique* [29] is applied in the context of establishing convergence properties of learning dynamics. Thus our technique may be of independent interest.

It is well-known that in case agents use *no-swap regret algorithm* [7] to select their actions in a repeated game, then the overall joint time-average behavior converges to a Correlated Equilibrium [17, 19, 1]. We remark that assuming that sellers use no-swap regret learning is a natural assumption since these types of algorithms come with strong optimality guarantees no matter the actions of the other sellers. In view of the above the main take-away message of this work is the following:

Take-Away Message *If all sellers use no-swap regret algorithms to select their prices, the selling price of the good will converge to the market clearing price p^* .*

As a result, our results offer a novel perspective on the law of supply and demand, framing it as the emergent outcome of a learning process among sellers. The latter reinforces the classical understanding of market equilibrium but also bridges it with contemporary decision-making algorithms.

On the negative side, we show that the weaker solution concept of Coarse Correlated Equilibrium (CCE) does not necessarily align with the market clearing price. In Theorem 3, we construct an instance of the pricing game that admits a CCE which does not coincide with the corresponding market clearing price. This implies that convergence to the market clearing price cannot be guaranteed under general no-regret dynamics, which are known to converge to CCE [20]. That being said, in Section 4, we empirically evaluate both no-regret and no-swap regret algorithms. Our experiments show that both types of learning dynamics converge to the market clearing price, suggesting that the negative result of Theorem 3 may be circumvented by specific classes of no-regret algorithms.

1.1 Related Work

Our work relates with the line of research studying online learning dynamics in various market settings. However none of the prior work has not studied perfectly competitive markets and the law of supply and demand.

A significant body of work has investigated *Cournot competition* under the lens of online learning. Even-Dar et al.[15] were the first to establish that no-regret dynamics converge to the respective Nash Equilibrium in the case of linear Cournot competition. Nadav et al.[26] extended these convergence results in the case of product differentiation and they also studied *Bertrand Duopoly markets*. Fiat et al.[16] study best-response dynamics in a variant of Cournot competition where firms aim to either optimize either profit or revenue while Lin et al.[24] establish that best-response dynamics in Cournot competition converge either to a Nash equilibrium or to a periodic orbit of length two. Immorlica et al.[21] provide bounds on the Price of Anarchy in Cournot competition in case of coalitions among firms while Shit et al.[27] study the convergence properties of no-regret learning dynamics in case of limited information feedback.

Our work also relates with the line of research studying decentralized dynamics in *Fisher markets*. Bikhchandani et al. [6] showed that proportional response dynamics—an update rule closely related to online learning—converge to equilibrium in linear Fisher markets. Kolumbus et al. [23] extend these results in case of asynchronous updates. Cheung et al. [11] study the convergence properties of proportional response dynamics in Fisher markets with CES utilities while Branzei et al. [8] study

proportional dynamics in exchange economies. Finally [13, 14, 10, 9, 30] study the convergence properties of tâtonnement-type algorithms in various Fisher markets.

Babaioff et al. [4, 5] consider a pricing game very similar to ours and establish Price of Anarchy bounds on the social welfare as well as identifying structural conditions for pure Nash equilibrium in case of various types of valuations. Golrezaei [18] study the convergence properties of Online Mirror Descent (OMD) in case of firm-competition in case of consumer reference prices. Cheung et al. [12] show that in various types of markets, Online Gradient Ascent admits chaotic behavior. Zhu et al. [31] study online learning algorithms for betting markets.

2 Preliminaries and Results

Let n denote the number of seller and buyers. We assume k discrete prices $[k] = \{0, 1, \dots, k\}$. Each seller $i \in [n]$ can sell 1 unit of good at a price no smaller than its marginal cost $s_i \in [k]$. Each buyer $j \in [n]$ is interested in buying 1 unit of good at a price at most its valuation $b_j \in [k]$. Without loss of generality we assume that $s_1 \leq \dots \leq s_n$ ² and $b_1 \geq \dots \geq b_n$.

The supply curve $\mathcal{S} : [k] \mapsto [n]$ denotes the number of sellers willing to sell the good at a given price, $\mathcal{S}(p) = |\{i \in [n] : s_i \leq p\}|$. The demand curve $\mathcal{D} : [k] \mapsto [n]$ denotes the number of buyers willing to buy the good at a given price, $\mathcal{D}(p) = |\{j \in [n] : b_j \geq p\}|$.

Definition 1. A price $p \in [k]$ is a market clearing price if and only if $\mathcal{S}(p) = \mathcal{D}(p)$.

We remark market clearing prices are not necessarily unique.

Example 1. Consider $n = 3$ sellers and buyers. The marginal costs of the sellers are $(s_1, s_2, s_3) = (0, 2, 5)$ and the marginal prices of the buyers $(s_1, s_2, s_3) = (6, 4, 1)$. All prices $\{2, 3, 4\}$ are market clearing prices.

Lemma 1 establishes the fact that under minimal assumption, the set of market clearing prices is not empty and in fact forms a consecutive interval.

Lemma 1. In case the marginal costs and prices $s_i, b_j \in [k]$ lies in different places. Then the set of market clearing prices is not empty and is always an interval (a set of the form $\{p_1, p_1 + 1, \dots, p_2\}$).

Proof. Notice that $\mathcal{S}(k) = n$ and $\mathcal{D}(k) \leq 1$ due to the fact no two buyers' valuation can equal k . Similarly $\mathcal{S}(1) \leq 1$ and $\mathcal{D}(1) = n$. Thus, $\mathcal{S}(k) - \mathcal{D}(k) > 0$ and $\mathcal{S}(1) - \mathcal{D}(1) > 0$. Since marginal costs and valuations lie in different places, the difference $\mathcal{S}(p) - \mathcal{D}(p)$ can differ by at most 1 for consecutive prices. Thus, there exists $p^* \in [k]$ such that $\mathcal{S}(p^*) = \mathcal{D}(p^*)$. Since $\mathcal{S}(\cdot)$ is an increasing and $\mathcal{D}(\cdot)$ is a decreasing function, the set of market clearing prices is an interval $\{p_1, p_1 + 1, \dots, p_2\}$. $\square$

Despite the fact that strictly speaking market, clearing prices are not necessarily unique, if k and n are large enough and the marginal costs and prices are not very concentrated, the different market clearing prices will practical correspond to the exact same price.

Definition 2. We denote with p^* the highest market clearing price.

With some abuse of terminology in the rest of the paper we refer to p^* as the market clearing price. The law of supply and demand states that the selling price of the good will converge to the market clearing price.

2.1 Pricing Games

In this section we introduce the pricing game in order to provide theoretical foundations on the law of supply and demand. Our goal is to establish that the law of supply and demand is the outcome of the strategic pricing of the sellers in their attempt to maximize their revenue.

²With a slight abuse of notation when $s_i = \ell \in [k]$, we actually mean $s_i = \ell - \epsilon$ for an arbitrarily small $\epsilon > 0$. This convention is very convenient since it ensures that seller $i \in [n]$ strictly prefer selling the good at price s_i rather not selling at all.

In Definition 3 we introduce the pricing game that is an *one-shot game* capturing the competition between sellers. The strategy of each seller is the selling price that the agents selects. When selecting a price, an agent needs to balance between high revenue and the risk of not selling its good.

Definition 3. In a pricing game, each seller i selects a price $p_i \geq s_i$. The payoff $U_i(p_i, p_{-i})$ of seller $i \in [n]$ is defined as follows:

- Let S denote the set of sellers, $S = \{1, \dots, n\}$. All of them are initially available.
- for each buyer $j = 1$ to n
 - buyer j enters the market and finds the cheapest available seller, $i_{\text{low}} := \operatorname{argmin}_{i \in S} p_i^3$.
 - If $p_{i_{\text{low}}} \leq b_j$, then
 1. Buyer j buys the good from i_{low} at price $p_{i_{\text{low}}}$.
 2. Seller i_{low} gets utility $U_{i_{\text{low}}}(p_i, p_{-i}) := p_{i_{\text{low}}} - s_{i_{\text{low}}}$ and exits $S \leftarrow S \setminus \{i_{\text{low}}\}$.

A seller $i \in [n]$, that did not sell its good, gets utility $U_i(p_i, p_{-i}) = 0$.

We remark that in the pricing game of Definition 3 only sellers are strategic agents. When a buyer enters the market, it considers the lowest-priced seller who is still available at that moment. We emphasize that each seller can supply only 1 unit of the good. That is why once a seller sells its unit, then it exits the market.

We denote with $\mathcal{P}_i$ the strategy space of seller $i \in [n]$, $\mathcal{P}_i = \{s_i, \dots, k\}$. We also denote with $\mathcal{P} := \mathcal{P}_1 \times \mathcal{P}_2 \times \dots \times \mathcal{P}_n$ the set of pricing profiles. For a pricing profile $p := (p_1, \dots, p_n) \in \mathcal{P}$ we also use the notation $p = (p_i, p_{-i})$ to denote the price of seller $i \in [n]$ with the prices selected by the other sellers.

2.2 No-Swap Regret Minimization and Correlated Equilibrium

To model the market's behavior over time, we consider sellers repeatedly playing the pricing game of Definition 3 across multiple rounds. We remark that the beginning of each round all sellers admit 1 unit of good regardless of whether they performed a sale in the previous round.

Protocol 1: Pricing Game over time

At each round $t = 1, \dots, T$

- Each seller $i \in [n]$, (secretly) selects a price $p_i^t \in \mathcal{S}_i$.
 - Each seller $i \in [n]$, gets utility $U_i(p_i^t, p_{-i}^t)$ (see Definition 3).
 - Each seller $i \in [n]$, learns p_{-i}^t and uses this information to select its next price p_i^{t+1} .
-

A seller $i \in [n]$, needs to come up with a *pricing strategy* that at each round $t \in [T]$ selects a price p_i^t solely based on past prices $p_{-i}^1, \dots, p_{-i}^{t-1} \in [k]$. The online learning framework provides such decision-making algorithms that base their decision on prior observations [20].

In a pricing games an online learning algorithm $\mathcal{A}$, at round $t \in [T]$, produces a mixed strategy $x_i^t \in \Delta(\mathcal{P}_i)$ over a set of possible prices $\mathcal{P}_i$. The performance of an online learning algorithm $\mathcal{A}$ can be quantified through the notion of *regret* [20, 7].

Definition 4. The regret of an online learning algorithm $\mathcal{A}$ is defined as

$$\mathcal{R}_{\mathcal{A}}(T) := \max_{p_i^* \in \mathcal{P}_i} \sum_{t=1}^T \mathbb{E}_{p_i \sim x_i^t} [U_i(p_i^*, p_{-i}^t)] - \sum_{t=1}^T \mathbb{E}_{p_i \sim x_i^t} [U_i(p_i, p_{-i}^t)]$$

³In case of tie, buyer $j \in [n]$ selects the seller with the highest index.

The swap regret of an online learning algorithm $\mathcal{A}$ is defined as

$$\mathcal{R}_{\mathcal{A}}^{\text{swap}}(T) := \max_{\delta: \mathcal{P}_i \mapsto \mathcal{P}_i} \sum_{t=1}^T \mathbb{E}_{p_i \sim x_i^t} [U_i(\delta(p_i), p_{-i}^t)] - \sum_{t=1}^T \mathbb{E}_{p_i \sim x_i^t} [U_i(p_i, p_{-i}^t)]$$

Algorithm $\mathcal{A}$ is called *no-regret* if $\mathcal{R}_{\mathcal{A}}(T) = o(T)$ no matter the prices $p_{-i}^1, \dots, p_{-i}^T$ selected by the other sellers. Respectively if $\mathcal{R}_{\mathcal{A}}^{\text{swap}}(T) = o(T)$ the algorithm is called *no-swap regret*.

A no-regret algorithm $\mathcal{A}$ guarantees that its time-averaged utility converges to the time-averaged utility produced by the *best fixed price* $p_i^* \in \mathcal{P}_i$ regardless the prices selected by the other providers. A no-swap regret algorithm provides the stronger optimality guarantees that the time-average payoff is at least the time-average payoff of the *best fixed switching function* $\delta(\cdot)$.

It is well known that if all sellers use no-regret algorithms to select their prices, the overall behavior converges to a *Coarse Correlated Equilibrium* [2]. Similarly, if sellers use no-swap regret algorithms, the resulting dynamics converge to a *Correlated Equilibrium* [19, 17]. Both equilibrium concepts are formally defined in Definition 5.

Definition 5. A probability distribution $\mu \in \Delta(S)$ over pricing profiles $\mathcal{P} := \times_{i \in [n]} \mathcal{P}_i$, is a *Coarse Correlated Equilibrium* if for each seller $i \in [n]$,

$$\mathbb{E}_{p \sim \mu} [U_i(p_i, p_{-i})] \geq \max_{p'_i \in \mathcal{P}_i} \mathbb{E}_{p \sim \mu} [U_i(p'_i, p_{-i})].$$

A probability distribution $\mu \in \Delta(S)$ over pricing profiles $\mathcal{P} := \times_{i \in [n]} \mathcal{P}_i$, is a *Correlated Equilibrium* if for each seller $i \in [n]$,

$$\mathbb{E}_{p \sim \mu} [U_i(p_i, p_{-i})] \geq \max_{\delta: \mathcal{P}_i \mapsto \mathcal{P}_i} \mathbb{E}_{p \sim \mu} [U_i(\delta(p_i), p_{-i})].$$

2.3 Our Results and Paper Organization

To this end, we present the main contribution of our work: establishing that the law of supply and demand can be viewed as the limiting behavior of no-swap regret algorithms employed by sellers to maximize their individual payoffs.

Theorem 1. Let a pricing game with market clearing price p^* and $\mu \in \Delta(\mathcal{P})$ be a *Correlated Equilibrium*. If $(p_1, \dots, p_n) \sim \mu$ then with probability 1,

1. all sellers $i \in [n]$ with $s_i \leq p^*$, select $p_i = p^*$ and sell their good.
2. all buyers $j \in [n]$ with $b_j \geq p^*$, buy the good at price p^* .
3. all sellers $i \in [n]$ with $s_i > p^*$ do not sell anything and all buyers $j \in [n]$ with $b_j < p^*$ do not buy anything.

Theorem 1 establishes that the selling price of the good is the market clearing price $p^* \in [k]$. This is because executed sales occur at price p^* . Notice that Theorem 1 establishes that sellers $i \in [n]$ with marginal costs higher than $s_i > p^*$ do not sell their good. Similarly, buyers $j \in [n]$ with marginal prices $b_j < p^*$ never buy the good since p^* is the lowest price offered in the market.

Combining Theorem 1 with the known convergence results of no-swap regret dynamics [17, 19] we get Corollary 1, showing that the law of supply and demand emerges naturally when sellers use no-swap regret algorithms. This is the main takeaway of our work.

Corollary 1. If all sellers in a pricing game use a no-swap regret to select their prices, then the price of the good converges to the market-clearing price $p^* \in [k]$.

In Theorem 2 we show that Coarse Correlated Equilibria (CCE) is not always compatible with the market clearing price. In particular we show that there are instances of the pricing game admitting CCEs at which the probability of sellers selecting the market clearing price is arbitrarily close to zero.

Theorem 2. There exists a family of pricing games with $n = 2$ admitting a Coarse Correlated Equilibrium $\mu \in \Delta(\mathcal{P})$ such that the probability that any of the sellers plays a market clearing price is at most $\mathcal{O}(1/k)$ where $[k]$ is the set of prices.

The proof of Theorem 2 is deferred to Appendix A. In Section 3 we present the highlevel ideas behind the proof of Theorem 1. Finally in Section 4 we experimentally evaluate the convergence properties of both no-regret and no-swap regret algorithm in the context of pricing games.

3 Proof of Theorem 1

In this section, we present the proof of Theorem 1, which is based on a *dual-fitting* argument [29]. To the best of our knowledge, this is the first application of such techniques to establish convergence results, and it may be of independent interest.

We first introduce some necessary definitions. We remind that p^* denotes the maximum market clearing price, see Definition 2.

Definition 6. We call a pricing profile $p = (p_1, \dots, p_n) \in \mathcal{P}$ valid if and only if $p_i = p^*$ for all sellers $i \in [n]$ with $s_i \leq p^*$. We also denote with $\mathcal{V}$ the set of all valid pricing profiles.

Using Definition 6 we can restate Theorem 1 as of the fact that any Correlated Equilibrium $\mu \in \Delta(\mathcal{S})$, places all of its probability mass on *valid* pricing profiles. The latter is formally stated in Theorem 3 and consists the main technical contribution of our work.

Theorem 3. Let $\mu \in \Delta(\mathcal{S})$ be a Correlated Equilibrium of a pricing game. Then, $\mu(p) > 0$ if and only if $p \in \mathcal{V}$.

In the rest of the section we present the proof of Theorem 3 through our dual fitting technique.

Definition 7. Let the function $r : \mathcal{P} \mapsto \{0, 1\}$ on the pricing profiles: $r(p) = 1$ for all $p \in \mathcal{V}$ and $r(p) = 0$ for any $p \notin \mathcal{V}$.

Using the reward function defined above, we introduce the following linear program that we will play a key role in our proof.

Definition 8. Given a pricing game of Definition 3, consider the following linear program,

$$\begin{aligned} \min \quad & \sum_{p \in \mathcal{P}} r(p) \cdot \mu(p) \\ \text{s.t.} \quad & \sum_{p_{-i} \in \mathcal{P}_{-i}} \mu(p_i, p_{-i}) (U_i(p'_i, p_{-i}) - U_i(p_i, p_{-i})) \leq 0 \text{ for all } i \in [n] \text{ and } p_i, p'_i \in \mathcal{P}_i \\ & \sum_{p \in \mathcal{P}} \mu(p) = 1 \\ & \mu(p) \geq 0 \quad \forall p \in \mathcal{P} \end{aligned}$$

where $U_i(p_i, p_{-i})$ is the utility of seller $i \in [n]$. We also denote with Z_{LP}^* its optimal value.

The idea behind the linear program of Definition 8 is that any Correlated Equilibrium $\mu \in \Delta(\mathcal{P})$ will satisfies its constraints. One can show that $\mu \in \Delta(\mathcal{P})$ satisfies the first constraint by using the fact that $\mathbb{E}_{p \sim \mu} [U_i(p_i, p_{-i})] \geq \mathbb{E}_{p \sim \mu} [U_i(\delta(p_i), p_{-i})]$ for the switching function with $\delta(s_i) = s'_i$. The rest of the constraints are satisfied due to the fact that $\mu(\cdot)$ is a joint probability distribution. The latter is formally stated and established in Lemma 6.

Lemma 2. Let $\mu \in \Delta(\mathcal{P})$ be a Correlated Equilibrium. Then $\mu \in \Delta(\mathcal{P})$ satisfies the constraints of the linear program of Definition 8 and thus $\sum_{p \in \mathcal{P}} \mu(p) \cdot r(p) \geq Z_{LP}^*$.

Our cornerstone idea is that the optimal value of the linear program above acts as a lower bound on the probability of sampling a valid pricing profile. The latter is formally stated in Corollary 2.

Corollary 2. Let $\mu \in \Delta(\mathcal{P})$ be a Correlated Equilibrium. Then $\Pr_{p \sim \mu} [p \text{ is valid}] \geq Z_{LP}^*$.

Proof. By Definition 7, $r(p) = 1$ if p is valid and 0 otherwise. Thus, $\sum_{p \in \mathcal{P}} \mu(p) \cdot r(p) = \sum_{p \in \mathcal{V}} \mu(p) \cdot r(p) = \Pr[p \text{ is valid}]$. $\square$

To complete our proof we just need to establish that $Z_{LP}^* \geq 1$. We will show that the optimal value of the dual of the linear program of Definition 7 is at least 1. Then the claim follows by weak duality.

Lemma 3. The following LP is the dual of the program in Definition 8. We denote with D^* its optimal value.

$$\begin{aligned} \max \quad & \mu \\ \text{s.t.} \quad & \mu + \sum_{i \in [n]} \sum_{p'_i \in \mathcal{P}_i} \lambda_{p_i p'_i}^i (U_i(p_i, p_{-i}) - U_i(p'_i, p_{-i})) \leq r(p) \text{ for all } p \in \mathcal{P} \\ & \lambda_{p_i p_{-i}}^i \geq 0 \quad \forall i \in [n], p_i, p_{-i} \in \mathcal{P}_i \end{aligned}$$

Lemma 4. *The optimal value of the dual linear program is at least 1, $D^* \geq 1$.*

Lemma 4 is the main technical contribution of this section. In Section 3.1, we present a sketch of its proof. In Section 3.1. We complete this section with the proof of Theorem 3.

Proof of Theorem 3. By Corollary 2 we know that $\Pr_{p \sim \mu}[p \text{ is valid}] \geq Z_{LP}^* \geq D^* \geq 1$ where the fact that $Z_{LP}^* \geq D^*$ is due to weak duality. $\square$

3.1 The Dual-Fitting Argument

In this section we present the proof of Lemma 4 stating that the optimal value of the dual program is at least 1, $D^* \geq 1$.

Our approach is to select an assignment of the dual variables $\{\mu, \lambda_{s_i, s-i}^i\}$ that are dual feasible and at the same time $\mu = 1$. Our assignment is presented in Definition 9 that we denote with $\hat{\lambda}_{s_i s'_i}^i$ in order to differentiate it from the variables.

Definition 9. *For any seller $i \in [n]$ with $s_i \leq p^*$,*

$$\hat{\lambda}_{p_i p'_i}^i := \begin{cases} 0 & \text{if } p'_i \neq p^* \\ (2nk)^{2np_i - i} & \text{if } p'_i = p^* \end{cases}$$

For any seller $i \in [n]$ with $s_i > p^$, $\hat{\lambda}_{p_i p'_i}^i = 0$.*

Before proceeding, let us provide the high-level intuition behind the assignment of Definition 9. The dual variable $\lambda_{p_i p'_i}^i$ represents a deviation of seller $i \in [n]$ from price p_i to price p'_i . By setting $\hat{\lambda}_{p_i p'_i}^i > 0$ only when $p_i = p^*$ and 0 otherwise, we enforce that all agents $i \in [n]$ with $p_i \leq p^*$ are incentivized to deviate exclusively to price p^* , and not to any other price.

Next, we establish that by selecting the assignment $\hat{\lambda}_{s_i s'_i}^i$ as of Definition 9, we can select $\hat{\mu} = 1$ and the overall assignment $(1, \hat{\lambda})$ is dual feasible. To establish the latter, for any pricing profile $p \in \mathcal{P}$ we consider

$$\hat{b}_p := r(p) - \sum_{i \in [n]} \sum_{p'_i \in \mathcal{S}_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p_i, p_{-i}) - U_i(p'_i, p_{-i})).$$

We will establish that $\hat{b}_p \geq 1$ for all pricing profiles $p \in \mathcal{P}$. Then it directly follows that $(1, \hat{\lambda})$ is feasible for the dual, and thus $D^* \geq 1$, since the dual is a maximization linear program. The fact that $\hat{b}_p \geq 1$ for all $p \in \mathcal{P}$ is respectively established in Lemma 5 and Lemma 6.

Lemma 5. *Let a valid pricing profile $p \in \mathcal{V}$. Then $b_p = 1$.*

Proof. Since $p \in \mathcal{P}$ is a valid pricing profile, $r(p) = 1$. As a result, we only need to establish that $\sum_{i \in [n]} \sum_{p'_i \in \mathcal{S}_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p_i, p_{-i}) - U_i(p'_i, p_{-i})) = 0$. This follows by the fact that $\hat{\lambda}_{p_i p'_i}^i = 0$ for all sellers $i \in [n]$ with $s_i > p^*$ and for all sellers $i \in [n]$ with $s_i \leq p^*$, $p_i = p^*$ meaning that

$$\sum_{p'_i \in \mathcal{S}_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p_i, p_{-i}) - U_i(p'_i, p_{-i})) = \hat{\lambda}_{p_i p^*}^i \cdot (U_i(p_i, p_{-i}) - U_i(p^*, p_{-i})) = 0.$$

$\square$

We complete the section with Lemma 6-establishing the respective claim for non-valid pricing profiles $p \notin \mathcal{V}$. The full proof of Lemma 6 is presented to Appendix C.

Lemma 6. *Let a non valid pricing profile $p \notin \mathcal{V}$. Then $b_p \geq 1$.*

Sketch of Proof. Since the pricing profile $p \notin \mathcal{V}$ is not valid, by Definition 7 we get that $r(p) = 0$. As a result, we need to establish that

$$\sum_{i \in [n]} \sum_{p'_i \in \mathcal{S}_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p'_i, p_{-i}) - U_i(p_i, p_{-i})) \geq 1.$$

Since $p^* \in [n]$ is the highest market clearing price, $p^* + 1$ is not a market clearing price. Thus, $p^* = b_j$ for some buyer $j \in [n]$ or $p^* + 1 = s_i$ for some seller $i \in [n]$.

Let us here consider the case $p^* = b_j$. To simplify notation let $m := |\{i \in [n] : s_i \leq p^*\}|$ and i^* the highest price, $i^* := \operatorname{argmax}_{\{i \in [n] : s_i \leq p^*\}} p_i$.

Notice that in case $p_{i^*} \geq p^* + 1$ then seller i^* does not set its good and thus $U_{i^*}(p_{i^*}, p_{-i^*}) = 0$. This is because at $m - 1$ buyers are willing to buy the good at price $p_{i^*} \geq p^* + 1$ (recall that $p^* = b_j$) and they are at least $m - 1$ sellers with lower prices. At the same time, seller i^* always sells its good in case $p_{i^*} = p^*$. This is because there are m buyer willing to buy the good at price p^* . Thus, $U_{i^*}(p^*, p_{-i^*}) = p^* - s_{i^*} \geq 1$. As a result,

$$\hat{\lambda}_{p_{i^*} p^*}^i \cdot (U_{i^*}(p^*, p_{-i^*}) - U_{i^*}(p_{i^*}, p_{-i^*})) \geq (2nk)^{2np_{i^*} - i^*}. \quad (1)$$

Since $i^* := \operatorname{argmax}_{\{i \in [n] : s_i \leq p^*\}} p_i$ we know that $p_i \geq p_{i^*}$ for all sellers $i \in [n]$ with $s_i \leq p^*$. To simplify things here let us assume that $p_i > p_{i^*}$ (the case of ties introduces some additional complications and is presented in the full proof).

Let us now try to identify the worst-case for the quantity $U_i(p^*, p_{-i}) - U_i(p_i, p_{-i})$. First notice that $U_i(p_i, p_{-i}) \leq p_i - s_i$ since otherwise $U_i(p_i, p_{-i}) = 0$ and that $U_i(p^*, p_{-i}) = p^* - s_i$. Thus,

$$\begin{aligned} \sum_{i \neq i^*} \hat{\lambda}_{p_i p^*}^i \cdot \left(\underbrace{U_i(p^*, p_{-i})}_{p^* - s_i} - \underbrace{U_i(p_i, p_{-i})}_{\leq p_i - s_i} \right) &\geq \sum_{i \neq i^*} (2nk)^{2np_i - i} \cdot (p^* - p_i) \\ &\geq -n(2nk)^{2n(p_{i^*} - 1) - 1} k \geq -(2nk)^{2n(p_{i^*} - 1)} \quad (2) \end{aligned}$$

where the last inequality comes from the fact that $p_i \leq p_{i^*} - 1$ and that $p^* - p_i \leq k$. By Equations 1 and 2 we get that

$$\sum_{i \in [n]} \sum_{p'_i \in S_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p'_i, p_{-i}) - U_i(p_i, p_{-i})) \geq (2nk)^{2np_{i^*}} \left((2nk)^{-i^*} - (2nk)^{-2n} \right) \geq 1.$$

□

4 Experimental Evaluations

In this section we experimentally evaluate the well-known Hedge [20] no-regret algorithm and the no-swap regret algorithm proposed by Blum and Mansour [7]. We consider as the set of prices the $[0, 5]$ interval with 0.2 discretization—there are the following 30 possible prices $\{0, 0.2, \dots, 4.8, 5\}$.

We first consider the family of instances of the pricing game constructed to establish Theorem 2 (see also Appendix A). This instance is composed by $n = 2$ sellers and buyers where $(s_1, s_2) = (0, 0)$ and $(b_1, b_2) = (5, \lambda)$. As a result, the highest market clearing price is $p^* = \lambda$. Despite the fact that such instances admit CCEs that do not correspond to any market clearing price, our experimental evaluation reveal that the Hedge algorithm always converges to the market clearing price p^* , see Figure 1, 2 and 3.




Figure 1: $\lambda = 0.2$




Figure 2: $\lambda = 1$




Figure 3: $\lambda = 3$

Next we consider a more natural set-up with $n = 100$ sellers and buyers. We consider the linear demand curve $D(p) = -20p + 100$ for $p \in [0, 5]$ and three different supply curves, $\mathcal{S}_{\text{linear}}(p) :=$

$20p$, $\mathcal{S}_{\text{quad}}(p) := p^2/0.25$ and $\mathcal{S}_{\text{linear}}(p) := p/5$, $\mathcal{S}_{\text{sqrt}}(p) := 100\sqrt{p/5}$ (see Figure 4). Each supply curve intersects with the demand curve at a different price, resulting in different market clearing prices. As Figure 5 depicts, if the sellers use the Hedge algorithm, the average selling price converges fast to the respective market clearing price.




Figure 4: Supply and Demand Curves




Figure 5: Average Selling Price

The performance of the no-swap regret algorithm of Blum et al. [7] is very similar with the performance of Hedge. Since convergence to the market clearing price is ensured by Theorem 1, the respective figures are deferred to Appendix D. In Appendix D we also present additional experimental evaluations for other supply/demand curves.

5 Conclusion

In this paper, we establish a rigorous connection between the law of supply and demand and the dynamics of online learning within adequate pricing games. Our main contribution is establishing that any correlated equilibrium (CE) of the associated pricing game aligns with the market-clearing price. The latter implies that when all sellers employ no-swap regret algorithms, the price converges to the market clearing price. Thus, our results provide an interesting theoretical foundation on the law of supply and demand by framing it as the emergent behavior of a learning process among the sellers.

While we establish that coarse correlated equilibria (CCE) do not inherently guarantee convergence to the market-clearing price, our experimental evaluations indicate that once sellers use the well-known Hedge algorithm, the resulting price dynamics converge to the market clearing price. As a result, it is likely that a certain classes of no-regret algorithms, such as mean-based algorithms, may be able to always converge to the market clearing price. Providing formal theoretical convergence guarantees for specific classes of no-regret algorithms, is a very interesting research direction.

Limitations Establishing that general no-regret dynamics converge to the market-clearing price would offer a stronger theoretical foundation for the law of supply and demand. However, as Theorem 2 demonstrates, Coarse Correlated Equilibria do not always align with the market-clearing price, so such convergence properties cannot hold for all no-regret sequences of play. However our experiments suggest that the Hedge algorithm consistently reaches the market-clearing price. Extending our convergence results to specific classes of no-regret dynamics remains an open challenge and a limitation of this work.

Broader Impact We acknowledge that there are many potential societal consequences of our theoretical results, however none of which we feel must be specifically highlighted.

Acknowledgments This project was supported by the Villum Young Investigator Award (Grant no. 72091).

References

- [1] Ioannis Anagnostides, Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Noah Golowich, and Tuomas Sandholm. Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games. In Stefano Leonardi and Anupam Gupta, editors, *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 736–749. ACM, 2022.
- [2] Ioannis Anagnostides, Gabriele Farina, Christian Kroer, Chung-Wei Lee, Haipeng Luo, and Tuomas Sandholm. Uncoupled learning dynamics with $O(\log T)$ swap regret in multiplayer games. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022*, 2022.
- [3] Robert J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.
- [4] Moshe Babaioff, Noam Nisan, and Renato Paes Leme. Price competition in online combinatorial markets. In Chin-Wan Chung, Andrei Z. Broder, Kyuseok Shim, and Torsten Suel, editors, *23rd International World Wide Web Conference, WWW '14, Seoul, Republic of Korea, April 7-11, 2014*, pages 711–722. ACM, 2014.
- [5] Moshe Babaioff, Renato Paes Leme, and Balasubramanian Sivan. Price competition, fluctuations and welfare guarantees. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, page 759–776. Association for Computing Machinery, 2015.
- [6] Benjamin Birnbaum, Nikhil R. Devanur, and Lin Xiao. Distributed algorithms via gradient descent for fisher markets. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, EC '11, page 127–136, New York, NY, USA, 2011. Association for Computing Machinery.
- [7] Avrim Blum and Yishay Mansour. From external to internal regret. *J. Mach. Learn. Res.*, 8:1307–1324, December 2007.
- [8] Simina Brânzei, Nikhil R. Devanur, and Yuval Rabani. Proportional dynamics in exchange economies. In Péter Biró, Shuchi Chawla, and Federico Echenique, editors, *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, pages 180–201. ACM, 2021.
- [9] Yun Kuen Cheung, Richard Cole, and Nikhil Devanur. Tatonnement beyond gross substitutes? gradient descent to the rescue. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '13, page 191–200. Association for Computing Machinery, 2013.
- [10] Yun Kuen Cheung, Richard Cole, and Ashish Rastogi. Tatonnement in ongoing markets of complementary goods. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, page 337–354. Association for Computing Machinery, 2012.
- [11] Yun Kuen Cheung, Richard Cole, and Yixin Tao. Dynamics of distributed updating in fisher markets. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, page 351–368. Association for Computing Machinery, 2018.
- [12] Yun Kuen Cheung, Stefanos Leonardos, and Georgios Piliouras. Learning in markets: Greed leads to chaos but following the price is right. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 111–117. ijcai.org, 2021.
- [13] Bruno Codenotti, Benton McCune, and Kasturi Varadarajan. Market equilibrium via the excess demand function. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '05, page 74–83, New York, NY, USA, 2005. Association for Computing Machinery.
- [14] Richard Cole and Lisa Fleischer. Fast-converging tatonnement algorithms for one-time and ongoing market problems. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, page 315–324, New York, NY, USA, 2008. Association for Computing Machinery.

- [15] Eyal Even-Dar, Yishay Mansour, and Uri Nadav. On the convergence of regret minimization dynamics in concave games. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 523–532. ACM, 2009.
- [16] Amos Fiat, Elias Koutsoupias, Katrina Ligett, Yishay Mansour, and Svetlana Olonetsky. Beyond myopic best response (in cournot competition). *Games and Economic Behavior*, 113:38–57, 2019.
- [17] Dean P. Foster and Rakesh V. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1):40–55, 1997.
- [18] Negin Golrezaei, Patrick Jaillet, and Jason Cheuk Nam Liang. No-regret learning in price competitions under consumer reference effects. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [19] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- [20] Elad Hazan. Introduction to online convex optimization. *CoRR*, abs/1909.05207, 2019.
- [21] Nicole Immorlica, Evangelos Markakis, and Georgios Piliouras. Coalition formation and price of anarchy in cournot oligopolies. In *Proceedings of the 6th International Conference on Internet and Network Economics*, WINE’10, page 270–281, Berlin, Heidelberg, 2010. Springer-Verlag.
- [22] Sabiou Inoua and Vernon Smith. The classical theory of supply and demand, 03 2020.
- [23] Yoav Kolumbus, Menahem Levy, and Noam Nisan. Asynchronous proportional response dynamics: Convergence in markets with adversarial scheduling. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 25409–25434. Curran Associates, Inc., 2023.
- [24] Zhengyang Liu, Haolin Lu, Liang Shan, and Zihe Wang. On the oscillations in cournot games with best response strategies, 2024.
- [25] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
- [26] Uri Nadav and Georgios Piliouras. No regret learning in oligopolies: cournot vs. bertrand. In *Proceedings of the Third International Conference on Algorithmic Game Theory*, SAGT’10, page 300–311, 2010.
- [27] Yuanyuan Shi and Baosen Zhang. Multi-agent reinforcement learning in cournot games. In *59th IEEE Conference on Decision and Control, CDC 2020, Jeju Island, South Korea, December 14-18, 2020*, pages 3561–3566. IEEE, 2020.
- [28] Adam Smith. The wealth of nations, 1776].
- [29] David P. Williamson and David B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, USA, 1st edition, 2011.
- [30] Li Zhang. Proportional response dynamics in the fisher market. *Theoretical Computer Science*, 412(24):2691–2698, 2011.
- [31] Haiqing Zhu, Alexander Soen, Yun Kuen Cheung, and Lexing Xie. Online learning in betting markets: Profit versus prediction. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024.

A Proof of Theorem 2

Theorem 2. *There exists a family of pricing games with $n = 2$ admitting a Coarse Correlated Equilibrium $\mu \in \Delta(\mathcal{P})$ such that the probability that any of the sellers plays a market clearing price is at most $\mathcal{O}(1/k)$ where $[k]$ is the set of prices.*

Proof. Consider a pricing game of Definition 3 with $n = 2$ such that $s_1 = s_2 = 0$ and $b_1 = 1$ and $b_2 = k$. To this end notice that the market clearing prices of the game are either 0 or 1. We will show that there exists a Coarse Correlated Equilibrium $\mu \in \Delta(\mathcal{P})$ such that

$$\Pr_\mu[s_1 \in \{0, 1\}] + \Pr_\mu[s_2 \in \{0, 1\}] \leq \mathcal{O}\left(\frac{1}{k}\right).$$

To simplify notation let $k := 2\ell + 1$. Up next we define the joint probability distribution $\mu \in \Delta(\mathcal{P})$ as follows:

1. randomly select $s \sim \text{Unif}(1, \dots, k)$.
2. then $\Pr[(s, (s-1) \bmod k)] = \frac{\ell+1}{2\ell+1}$ and $\Pr[((s-1) \bmod k), s] = \frac{\ell}{2\ell+1}$.

Notice that due to the symmetry $\mu(\cdot)$, $\Pr_\mu[s_1 < s_2] = (\ell+1)/(2\ell+1)$ and $\Pr_\mu[s_2 > s_1] = \ell/(2\ell+1)$. The latter implies that for seller 1 that,

$$\mathbb{E}_\mu[U_1(s_1, s_2)] = \frac{\ell+1}{2\ell+1} \frac{\sum_{i=1}^{2\ell+1} i}{2\ell+1} = \frac{(\ell+1)^2}{2\ell+1}$$

Similarly for seller 2 we get that

$$\mathbb{E}_\mu[U_2(s_1, s_2)] = \frac{\ell}{2\ell+1} \frac{\sum_{i=1}^{2\ell+1} i}{2\ell+1} = \frac{\ell(\ell+1)}{2\ell+1}$$

Now let assume that seller 1 deviates to a fixed price i . In this case

$$\mathbb{E}_\mu[U_1(i, s_2)] = i \cdot \Pr_\mu[s_2 \geq i] = i \cdot \left(1 - \frac{i}{2\ell+1}\right) = i \cdot \frac{2\ell+1-i}{2\ell+1} \leq \frac{(\ell+1)^2}{2\ell+1} = \mathbb{E}_\mu[U_1(s_1, s_2)]$$

Now let assume that seller 2 deviates to fixed price i . In this case

$$\mathbb{E}_\mu[U_2(i, s_1)] = i \cdot \Pr_\mu[s_1 > i] = i \cdot \left(1 - \frac{i-1}{2\ell+1}\right) = i \cdot \frac{2\ell+2-i}{2\ell+1} \leq \frac{\ell(\ell+1)}{2\ell+1} = \mathbb{E}_\mu[U_2(s_2, s_1)]$$

As a result, the joint probability distribution $\mu \in \mathcal{P}$ is a Coarse Correlated Equilibrium (CCE) while probability that the price 0 or 1 (the set of market clearing prices) is played by either seller 1 or 2 is at most $\mathcal{O}(1/k)$. □

B Omitted Proofs of Section 3

Lemma 2. *Let $\mu \in \Delta(\mathcal{P})$ be a Correlated Equilibrium. Then $\mu \in \Delta(\mathcal{P})$ satisfies the constraints of the linear program of Definition 8 and thus $\sum_{p \in \mathcal{P}} \mu(p) \cdot r(p) \geq Z_{LP}^*$.*

Proof. Since $\mu \in \Delta(\mathcal{P})$ then $\sum_{p \in \mathcal{P}} \mu(p) = 1$ and $\mu(p) \geq 0$ for all $p \in \mathcal{P}$. Notice that by Definition 5 we know that for any seller $i \in [n]$ and any switching function $\delta : \mathcal{P}_i \mapsto \mathcal{P}_i$,

$$\sum_{p_i} \sum_{p_{-i}} \mu(p_i, p_{-i}) \cdot (U_i(p_i, p_{-i}) - U_i(\delta(p_i), p_{-i})) \geq 0$$

Let $p_i, p'_i \in \mathcal{P}_i$ and consider the switching function $\delta(p_i) = p'_i$ and $\delta(x) = x$ otherwise. Then we directly get that

$$\sum_{p_{-i}} \mu(p_i, p_{-i}) \cdot (U_i(p_i, p_{-i}) - U_i(p'_i, p_{-i})) \geq 0 \quad \text{for all } p_i, p'_i \in \mathcal{P}_i.$$

□

Lemma 4. *The optimal value of the dual linear program is at least 1, $D^* \geq 1$.*

Proof. By taking the Lagrangian

$$L := \sum_p \mu(p) \cdot r(p) - \sum_{i, p_i, p'_i} \lambda_{p_i, p'_i}^i \left(\epsilon + \sum_{p-i} \mu(p_i, p-i) \cdot (U_i(p_i, p-i) - U_i(p'_i, p-i)) \right) + \mu \left(1 - \sum_p \mu(p) \right) - \sum_p k_p \cdot \mu(p)$$

where $\lambda_{p_i, p-i}^i, k_p \geq 0$. By rearranging the terms we get that

$$L := \sum_p \mu(p) \left(r(p) + \sum_{p'_i} \lambda_{p_i, p'_i}^i (U_i(p'_i, p-i) - U_i(p_i, p-i)) - \mu - k_p \right) + \mu - \epsilon \sum_{i, p_i, p'_i} \lambda_{p_i, p'_i}^i$$

By setting $r(p) - \sum_{i \in [n]} \sum_{p'_i \in S_i} \lambda_{p_i, p'_i}^i (U_i(p_i, p-i) - U_i(p'_i, p-i)) - \mu - k_p = 0$ we get that

$$\mu + \sum_{i \in [n]} \sum_{p'_i \in \mathcal{P}_i} \lambda_{p_i, p'_i}^i (U_i(p_i, p-i) - U_i(p'_i, p-i)) \leq r(p)$$

since $k_s \geq 0$. □

C Omitted Proof of Section 3.1

Lemma 5. *Let a valid pricing profile $p \in \mathcal{V}$. Then $b_p = 1$.*

Proof. Since $p \in \mathcal{V}$ we know that $r(p) = 1$. Since p is a valid pricing profile we know that $p_i = p^*$ for each seller $i \in S(p^*)$ with $s_i \leq p^*$. As a result,

$$\begin{aligned} \sum_{p'_i \in S_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p_i, p-i) - U_i(p'_i, p-i)) &= \hat{\lambda}_{p_i p^*}^i \cdot (U_i(p_i, p-i) - U_i(p^*, p-i)) \quad (\hat{\lambda}_{p_i p'_i}^i = 0 \text{ for } p'_i \neq p^*) \\ &= 0 \quad \text{since } p_i = p^* \end{aligned}$$

For all the sellers i with $s_i > p^*$, by Definition 9 we have that $\sum_{p'_i \in S_i} \hat{\lambda}_{p_i p'_i}^i = 0$ since $\hat{\lambda}_{p_i p'_i}^i = 0$ for all prices p_i and p'_i . □

Lemma 6. *Let a non valid pricing profile $p \notin \mathcal{V}$. Then $b_p \geq 1$.*

Proof. Since the pricing profile $p \notin \mathcal{V}$ is not valid, by Definition 7 we get that $r(p) = 0$. As a result, we need to establish that $\sum_{i \in [n]} \sum_{p'_i \in S_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p'_i, p-i) - U_i(p_i, p-i)) \geq 1$. To simplify notation let $S(p^*)$ denote the set of sellers with marginal cost less than p^* , $S(p^*) = \{i \in [n] : s_i \leq p^*\}$. Notice that by Definition 9, $\hat{\lambda}_{s_i s'_i}^i = 0$ for all $i \notin S(p^*)$. To simplify notation we denote $m := |S(p^*)|$. Let $i^* \in [n]$ be the seller $i \in S(p^*)$ with the highest price p_i , $i^* := \operatorname{argmax}_{i \in S(p^*)} p_i$. In case there are multiple sellers with price p_{i^*} , we consider i^* to be the one with the lowest index.

Up next we show that for any non-valid pricing profile $p \notin \mathcal{V}$,

$$\sum_{i \in [n]} \sum_{p'_i \in S_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p'_i, p-i) - U_i(p_i, p-i)) \geq 1.$$

We establish the latter claim for the following mutually exclusive cases:

1. $p_{i^*} \leq p^*$
2. $p_{i^*} \geq p^* + 1$.

Let us start with the case $p_{i^*} \leq p^*$. Since $p \notin \mathcal{V}$ we know that there exists a seller $i \in [n]$ such that $p_i < p^*$ since otherwise $p \in \mathcal{V}$. Notice that any seller $i \in \mathcal{S}(p^*)$ that set p^* as its price, is ensured to sell its good. This is because all produce $i \notin \mathcal{S}(p^*)$ must essentially set a price $p_i > p^*$ while there are exactly m buyer willing to pay price p^* for the good. As a result, for any seller $i \in [n]$ we know that

$$\begin{aligned} \sum_{p'_i \in \mathcal{S}_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p'_i, p_{-i}) - U_i(p_i, p_{-i})) &= \hat{\lambda}_{p_i p^*}^i \cdot (U_i(p'_i, p_{-i}) - U_i(p^*, p_{-i})) \\ &\geq ((p^* - s_i) - (p_i - s_i)) = p^* - p_i \end{aligned}$$

Since there exists at least one seller $i \in [n]$ with $p_i < p^*$ and $\hat{\lambda}_{s_i s'_i}^i = 0$ for all $i \notin [n]$ we are ensured that

$$\sum_{i \in [n]} \sum_{p'_i \in \mathcal{S}_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p'_i, p_{-i}) - U_i(p_i, p_{-i})) \geq 1.$$

Let us now proceed with the case where $p_{i^*} \geq p^* + 1$. We first argue that $U_{i^*}(p_{i^*}, p_{-i^*}) = 0$ -the latter is formally established in Claim 1

Claim 1. *Let the seller $i^* = \operatorname{argmax}_{i \in [n]} p_i$ (in case of ties $i^* \in [n]$ is the one with lowest index). Then $U_{i^*}(p_{i^*}, p_{-i^*}) = 0$.*

Proof. Since p^* is the maximum market clearing price then $p^* + 1$ is not a market clearing price and thus $p^* = b_j$ for some buyer $j \in [n]$ or $p^* + 1 = s_i$ for some provider $i \in [n]$. Since all s_i and b_j lie in different positions the latter two cases are also mutually exclusive.

Let us start with the case $p^* = b_j$ for some buyer $j \in [n]$. Since $p_{i^*} \geq p^* + 1$ we are ensured that there are at most $m - 1$ buyers willing to buy the good at price $p_{i^*} \in [n]$. However since $i^* \in \operatorname{argmax}_i p_i$ and at the same time admits the lowest index, will not sell its good (notice that enter the market in decreasing order with respect to their b_j and break ties among sellers lexicographically). Thus $U_{i^*}(p_{i^*}, p_{-i^*}) = 0$. $\square$

Up next we consider the mutually exclusive cases $p^* + 1 = b_j$ and $p^* + 1 = s_i$ and separately establish that

$$\sum_{i \in [n]} \sum_{p'_i \in \mathcal{S}_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p'_i, p_{-i}) - U_i(p_i, p_{-i})) \geq 1.$$

We first start with the case $p^* + 1 = b_j$ for some buyer $j \in [n]$. Let $i^* \in [n]$ be the seller $i \in \mathcal{S}(p^*)$ with the highest price p_i , $i^* := \operatorname{argmax}_{i \in \mathcal{S}(p^*)} p_i$. In case there are multiple sellers with price $p_{i^*} \in [k]$ the i^* is the one with the lowest index among them.

Let us assume that $p_{i^*} > p^*$, by Claim 1 we are ensured that $U_{i^*}(p_{i^*}, p_{-i^*}) = 0$. In case seller $i^* \in [n]$ had selected price $p^* \in [n]$, it would have sold its good since at price $p^* \in [n]$ there are m buyers and m sellers willing to sell. Thus, $U_{i^*}(p^*, p_{-i^*}) = p^* - s_{i^*}$. As a result we get that

$$\hat{\lambda}_{p_i p^*}^i \cdot (U_{i^*}(p^*, p_{-i^*}) - U_{i^*}(p_{i^*}, p_{-i^*})) = (2nk)^{2np_{i^*} - i^*} \cdot (p^* - s_{i^*}) \geq (2nk)^{2np_{i^*} - i^*}$$

Now let a seller $i \in [n]$ such that $p_i = i^*$. By definition of i^* we get that $i \geq i^* + 1$. Thus,

$$\hat{\lambda}_{p_i p^*}^i \cdot (U_i(p^*, p_{-i}) - U_i(p_i, p_{-i})) \geq (2nk)^{2np_{i^*} - i^* - 1} \cdot (p^* - p_{i^*}) \geq -(2nk)^{2np_{i^*} - i^* - 1} k$$

Now let a seller $i \in [n]$ such that $p_i \leq p_{i^*} - 1$. In this case

$$\hat{\lambda}_{p_i p^*}^i \cdot (U_i(p^*, p_{-i}) - U_i(p_i, p_{-i})) \geq (2nk)^{2n(p_{i^*} - 1) - 1} \cdot (p^* - p_{i^*} + 1) \geq -(2nk)^{2np_{i^*} - 2n - 1} k$$

As a result, we overall get that

$$\begin{aligned} \hat{b}_p &:= r(p) - \sum_{i \in \mathcal{S}(p^*)} \sum_{p'_i \in \mathcal{S}_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p_i, p_{-i}) - U_i(p'_i, p_{-i})) \\ &= (2nk)^{2np_{i^*} - i^*} - n(2nk)^{2np_{i^*} - i^* - 1} k - n(2nk)^{2np_{i^*} - 2n - 1} k \\ &= (2nk)^{2np_{i^*} - i^*} \left(1 - \frac{1}{2} - nk(2nk)^{i^* - 2n - 1} \right) \geq 1 \end{aligned}$$

We now consider the case where $p^* + 1 = s_i$ for some seller $i \in [n]$. To simplify notation we denote this seller as Next. Following the exact same steps as in the previous case, we can prove that

$$\sum_{i \neq i^*} \sum_{p'_i \in S_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p_i, p_{-i}) - U_i(p'_i, p_{-i})) \geq -n(2nk)^{2np_{i^*} - i^* - 1}k - n(2nk)^{2np_{i^*} - 2n - 1}k$$

Let us now consider $\lambda_{p_i(p^*+1)}^{\text{next}} = (2nk)^{2np_i + \text{next}}/\epsilon$ where $\epsilon > 0$ is the small positive constant discussed in Section 2. Up next we establish Lemma 6 for the following mutually exclusive cases:

- $p_{i^*} = p^* + 1$ and $p_{\text{next}} = p^* + 1$

In this case $U_{i^*}(p) = 0$ since there are m sellers before seller $i^* \in [n]$ with higher priority (next has higher index than i^*). At the same time $U_{i^*}(p^*, p_{-i^*}) = p^* - s_{i^*}$. Finally in case next selects $p^* + 1$ as its price it gets the good and thus $U_{\text{next}}(p^* + 1, p_{-\text{next}}) = p^* + 1 - s_{\text{next}} = \epsilon$. Combining all the above we get that

$$\sum_{i \in \{i^*, \text{next}\}} \sum_{s'_i \in S_i} \hat{\lambda}_{p_i p'_i}^i (U_i(p_i, p_{-i}) - U_i(p'_i, p_{-i})) \geq (2nk)^{2np^* - i^*}$$

- $p_{i^*} = p^* + 1$ and $p_{\text{next}} \geq p^* + 2$

In this case $U_{i^*}(p) = p^* + 1 - s_{i^*}$ since $i^* \in [n]$ sells its good due to the fact that there are at least m buyers willing to buy the good at price $p^* + 1$. Similarly $U_{i^*}(p^*, p_{-i^*}) = p^* - s_{i^*}$. Also notice that next does not sell its good if it sets with price p_{next} but sell its if it sets price $p^* + 1$. As a result, $U_{\text{next}}(p_{\text{next}}, p_{-\text{next}}) = 0$ and $U_{\text{next}}(p^* + 1, p_{-\text{next}}) = \epsilon$. Combining all the above we get

$$\begin{aligned} \sum_{i \in \{i^*, \text{next}\}} \sum_{s'_i \in S_i} \hat{\lambda}_{p_i p'_i}^i (U_i(p'_i, p_{-i}) - U_i(p_i, p_{-i})) &\geq -(2nk)^{2np^* + 2n - i^*} + (2nk)^{2np_{\text{next}} - \text{next}} \\ &\geq -(2nk)^{2np^* + 2n - i^*} + (2nk)^{2np^* + 4n - \text{next}} \end{aligned}$$

- $p_{i^*} \geq p^* + 2$ and $p_{\text{next}} \leq p_{i^*}$

In this case notice that $U_{i^*}(p) = 0$ since there are m sellers with higher priority than seller $i \in [n]$. At the same time $U_{i^*}(p^*, p_{-i^*}) = p^* - s_{i^*}$ since seller next will never report price p^* ($s_{\text{next}} > p^*$) thus seller i^* always sells its good at price $p^* \in [n]$. Similarly as before we get that $U_{\text{next}}(p^* + 1, p_{-\text{next}}) = \epsilon$ and $U_{\text{next}}(p_{\text{next}}, p_{-\text{next}}) \leq p_{\text{next}} - s_{\text{next}} = p_{\text{next}} - (p^* + 1 - \epsilon)$. Combining all the above we get that

$$\begin{aligned} \sum_{i \in \{i^*, \text{next}\}} \sum_{s'_i \in S_i} \hat{\lambda}_{p_i p'_i}^i (U_i(p'_i, p_{-i}) - U_i(p_i, p_{-i})) &\geq (2nk)^{2np_{\text{next}} - i^*} - \frac{k}{\epsilon} (2nk)^{2np_{\text{next}} - \text{next}} \\ &= (2nk)^{2np_{\text{next}} - i^*} \left(1 - \frac{k}{\epsilon} (2nk)^{-1}\right) \\ &= (2nk)^{2np_{\text{next}} - i^*} \left(1 - \frac{1}{2\epsilon n}\right) \\ &= \frac{1}{2} (2nk)^{2np_{\text{next}} - i^*} \quad \text{for } \epsilon = 1/n \end{aligned}$$

- $p_{i^*} \geq p^* + 2$ and $p_{\text{next}} > p_{i^*}$

In this case $U_{i^*}(p) \leq p_{i^*} - s_{i^*}$ and $U_{i^*}(p^*, p_{-i^*}) = p^* - s_{i^*}$. At the same time notice that $U_{\text{next}}(p^* + 1, p_{-\text{next}}) = \epsilon$ since seller next is able to sell its good at price $p^* + 1$. While seller next can never sell its good at price $p_{\text{next}} \in [k]$ since there are m other sellers with higher priority. Thus, $U_{\text{next}}(p_{\text{next}}, p_{-\text{next}}) = 0$. Combining all the above we get

$$\begin{aligned}
\sum_{i \in \{i^*, \text{next}\}} \sum_{s'_i \in \mathcal{S}_i} \hat{\lambda}_{p_i p'_i}^i (U_i(p'_i, p_{-i}) - U_i(p_i, p_{-i})) &\geq -k(2nk)^{2n(p_{\text{next}}-1)-i^*} + \frac{1}{\epsilon} (2nk)^{2np_{\text{next}}-\text{next}} \cdot \epsilon \\
&= (2nk)^{2np_{\text{next}}} \left(-k(2nk)^{-2n-i^*} + (2nk)^{-\text{next}} \right) \\
&= \frac{1}{2} (2nk)^{2np_{\text{next}}-\text{next}}
\end{aligned}$$

As a result, we overall get that

$$\sum_{i \in \{i^*, \text{next}\}} \sum_{s'_i \in \mathcal{S}_i} \hat{\lambda}_{p_i p'_i}^i (U_i(p'_i, p_{-i}) - U_i(p_i, p_{-i})) \geq \frac{1}{2} (2nk)^{2np_{i^*}-i^*-1}$$

At the same time using the exact same arguments as in the case $p^* = b_j$ for some buyer $j \in [n]$, we establish that

$$\sum_{i \neq i^*, \text{next}} \sum_{p'_i \in \mathcal{S}_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p_i, p_{-i}) - U_i(p'_i, p_{-i})) \geq -n(2nk)^{2np_{i^*}-i^*-3}k - n(2nk)^{2np_{i^*}-2n-1}k$$

Putting everything together we get that

$$\sum_{i \in [n]} \sum_{p'_i \in \mathcal{S}_i} \hat{\lambda}_{p_i p'_i}^i \cdot (U_i(p_i, p_{-i}) - U_i(p'_i, p_{-i})) \geq \frac{1}{2} (2nk)^{2np_{i^*}-i^*-1} - n(2nk)^{2np_{i^*}-i^*-3}k - n(2nk)^{2np_{i^*}-2n-1}k \geq 1$$

□

D Additional Experimental Evaluations

All experiments were conducted in Apple M4 Pro and the Hedge algorithm was run with step-size $\gamma = 0.1$.

We also evaluate the Hedge algorithm in the following set-up. We consider the demand curve $\mathcal{D}(p) = -20p + 100$ for $p \in [1, 5]$ with a 0.2 discretization. We consider 4 different supply curves parametrized by $m \in \{0, 100\}$,

$$S_m(p) = \begin{cases} m & \text{if } p \leq 4.8 \\ \frac{100-m}{0.2} \cdot p + \frac{5m-480}{0.2} & \text{otherwise} \end{cases}$$

In Figure 6 we consider $m \in \{10, 30, 60, 80\}$. Each different curve S_m admits a different intersection point with the demand curve (see Figure 6). In Figure 7 we see that if all sellers use a no-swap regret algorithm the average selling price converges to the respective market clearing price. The average selling price at each round t , is the average price of the realized trades.




Figure 6: Supply and Demand Curves




Figure 7: Average Selling Price

In all of our experimental evaluation we used step-size $\gamma = 0.1$ for the Hedge algorithm.

D.1 No-Swap Regret Dynamics

We perform the exact same experimental evaluation for the no-swap regret algorithm of Blum et al. [7]. In our implementation we used the Hedge algorithm with step-size $\gamma = 0.1$ as our base no-regret algorithm. In all the above experimental evaluation the resulting no-swap regret dynamics converge to the market clearing price, something that is to be expected due to Theorem 1.

As in Section 4, we first consider the case $n = 2$ sellers and buyers where $(s_1, s_2) = (0, 0)$ and $(b_1, b_2) = (5, \lambda)$. In this instance the highest market clearing price is $p^* = \lambda$. Figures 8, 9 and 10 verify that the resulting no-swap regret dynamics converge to the respective market clearing price of each case.




Figure 8: $\lambda = 0.2$




Figure 9: $\lambda = 1$




Figure 10: $\lambda = 3$

Next we consider the case $n = 100$ sellers and sellers with demand curve $\mathcal{D}(p) = -20p + 100$ for $p \in [0, 5]$ and supply curves, $\mathcal{S}_{\text{linear}}(p) := p/5$, $\mathcal{S}_{\text{quad}}(p) := p^2/0.25$ and $\mathcal{S}_{\text{linear}}(p) := p/5$, $\mathcal{S}_{\text{sqrt}}(p) := 100\sqrt{p}/5$. Each supply curve intersects with the demand curve at a different price, resulting in different market clearing prices (see Figure 12). Figure 11 verifies that the resulting no-swap regret dynamics converge to the respective market clearing price of each case.




Figure 11: Supply and Demand Curves




Figure 12: Average Selling Price

Finally we consider the demand curve $\mathcal{D}(p) = -20p + 100$ for $p \in [1, 5]$ with a 0.2 discretization. We consider 4 different supply curves parametrized by $m \in \{10, 30, 60, 80\}$,

$$S_m(p) = \begin{cases} m & \text{if } p \leq 4.8 \\ \frac{100-m}{0.2} \cdot p + \frac{5m-480}{0.2} & \text{otherwise} \end{cases}$$

In Figure 14 we see that if all sellers use the no-swap regret algorithm of Blum et al. [7], the average price converges to the respective market clearing price.




Figure 13: Supply and Demand Curves




Figure 14: Average Selling Price

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: We described in detail the results of the paper both in abstract and the introduction.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We have included a limitation section describing the limitation our work.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: We explicitly state all of our assumptions in the beginning of the paper.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: Yes we describe in detail all the necessary information.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: We have included, the code used for our experimental evaluations along with detailed instructions on how to run it.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: Yes we have included all the necessary details.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[No\]](#)

Justification: All of our evaluated online learning dynamics converge to the market clearing price meaning that there is no variance on the final outcome.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: All of experimental evaluation have been conducted in Mac Pro 4. We have included this information in the experimental section.

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: All of our research aligns with the Code of Ethics.

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We have included a respective section.

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper deals with fundamental online learning algorithms.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: All the code, data and models were created by the authors.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new assets are introduced.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No experiments with human subjects were performed.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: No experiments with human subjects were performed.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [NA]

Justification: No LLM usage was used.