Electronic Companion for "Price Competition Under Platform-Mediated Search: A Consider-Then-Choose Framework"

Appendix A: Examples of CLC Choice Behavior

The following example illustrates the wide range of behaviors encompassed by CLC choice.

EXAMPLE EC.1. *Consider a set of third-party sellers listing athletic T-shirts for women, with attributes* $\{\texttt{brand}, \texttt{average rating}, \texttt{number of reviews}, \texttt{price}\}$. *The brand attribute takes values in the set* $\{\texttt{Nike}, \texttt{Adidas}, \texttt{Puma}\}$*; each T-shirt has an average rating in* $[0, 5]$*; and the number of reviews lies in* $\mathbb{N}$. *The following three customers belong to the CLC class:*

- **A price-sensitive customer.** *This customer is willing to pay at most \$25 and requires an average rating of at least 3. Of all such items, the customer chooses the cheapest one. If there are multiple lowest-priced items, she chooses the one with the highest rating, and breaks final ties by number of reviews.*
- **A quality-sensitive customer.** *This customer is willing to pay at most \$80 and requires an average rating of at least 4.5 with at least 100 reviews. Of all such items, she chooses the one with the highest average rating, breaking ties by the number of reviews, then by price.*
- **A loyal customer.** *This customer is willing to pay \$100 and only considers* `Nike` *products. Of all such products, she purchases the highest-rated one, breaking ties according to the number of reviews, followed by price.*

Appendix B: Section 4: Additional Results and Omitted Proofs

B.1. On the non-existence of equilibria

PROPOSITION EC.1. *Fix* $\epsilon \in (0, 1/2)$. *Consider a duopoly wherein* ϵ *fraction of the population only purchases from seller 1,* ϵ *fraction of the population only purchases from seller 2, and* $1 - 2\epsilon$ *fraction of the population purchases from the cheapest seller, breaking ties uniformly at random. Suppose moreover that customers' willingness-to-pay* $w_c \sim \text{Unif}[0, 1]$. *In this setting, no local Nash equilibrium exists.*

Proof. By Theorem 2, it suffices to analyze the VOPs associated with (i) seller 1 pricing above seller 2, and (ii) seller 2 pricing above seller 1.

Consider the case where seller 1 prices above seller 2. In this setting, seller 1 sets p_1 to maximize $R_1(p, 0) = \epsilon \cdot p \cdot (1 - p)$, whose maximum is attained at $p = 1/2$. Given $p_1 = 1/2$, the undercutting best-response of seller 2 is to set $p_2 = p_1 - \epsilon_0$, for arbitrarily small $\epsilon_0 > 0$, since there is an upward jump discontinuity at p_1. Since seller 2's undercutting optimum is not attained, there does not exist a LNE associated with this ordering. By symmetry, the same holds for seller 2 pricing above seller 1, which completes the proof of the claim. □

B.2. Proof of Theorem 1

Proof. We argue this at the customer choice level. Consider a seller i pricing at p and suppose that a seller j with $p_j < p$ changes its original price p_j to some $p'_j < p$. We show that this change does not impact whether or not seller i is chosen by a customer; hence, it does not impact the revenue of seller i.

Note that the only customers who could possibly impact seller i's revenue under this change are those for whom both i and j were in the consideration set when j priced at p_j. To see this, consider first a customer who did not include i in her consideration set under p_j. In this case, i will continue to be excluded under p'_j, and seller i will continue to earn zero revenue from such a customer. Consider next a customer who included i but not j in her consideration set under p_j. Since i was in the consideration set, the customer's willingness-to-pay exceeds p, and p_j as a result. Hence, j must have been excluded due to a non-price attribute, and j will continue to be excluded due to a non-price attribute under p'_j. The choice outcome from the perspective of i does not change in this case as well.

Consider finally customers who included both i and j in their consideration set under p_j. As argued above, since i was in the consideration set under p_j, $w_c \geq p > p'_j$, and thus both i and j remain in the consideration set under p'_j. Now, under j's old price, if the customer chose from the consideration set based on a non-price attribute ranked higher than price, she will continue to make the same choice under the new price. If she chose based on either the price attribute or some non-price attribute ranked lower than the price, then i is not chosen before or after the price change, since its price is not the lowest in the consideration set under p_j or p'_j. In summary, the outcome of whether or not i is chosen remains unchanged, which proves the result. □

B.3. Proof of Theorem 2

We first prove that the following property (which we refer to as *cross-monotonicity*) holds.

PROPOSITION EC.2. *Under CLC choice, for all $i \in \mathcal{N}$, $p_i \in \mathbb{R}$, $R_i(p_i, \mathbf{p}_{-i})$ is non-decreasing in $\mathbf{p}_{-i}$, i.e., if $\mathbf{p}'_{-i} \geq \mathbf{p}_{-i}$ (component-wise), then $R_i(p_i, \mathbf{p}'_{-i}) \geq R_i(p_i, \mathbf{p}_{-i})$.*

Proof. Fix seller $i \in \mathcal{N}$, and suppose there exists a seller $j \in \mathcal{N} \setminus \{i\}$ that strictly increased its price to $p'_j > p_j$. Note that this change only affects a customer c such that $\{i, j\} \in \mathcal{N}_c$. Consider any customer c for which this holds.

If c had not chosen i under p_j, then i's revenue from c is 0, and can only weakly increase after a price change of j. Thus, it suffices to argue that if c chooses i, then they continue to choose i after an increase in p_j, implying a weak increase in revenue for i.

If c had previously chosen i due to a non-price attribute ranked higher than price, then the price change of seller j does not change her previous decision. Suppose now that c chose i based on either the price attribute, or an attribute ranked lower than the price. Then, it must be that $p_i \leq p_j < p_{j'}$. In this case, even after the price increase in p_j, c would continue to choose i. This is because the customer cannot choose j

after this change since $p_i < p'_j$, and it cannot choose some other seller since if that were the case, then they would have chosen that seller under p_j as well, contradicting the fact that i was chosen under p_j. □

With this result in hand, we prove Theorem 2, deferring the proofs of all auxiliary lemmas to the end of the section.

Proof of Theorem 2. Suppose $\mathbf{p}$ is a LNE, and let π denote the corresponding ordering (breaking ties arbitrarily).

Suppose first that all prices in $\mathbf{p}$ are distinct. Then, these prices are necessarily valid with respect to π (i.e., $p_{\pi(i)}$ is a local optimum in the undercutting region of $\pi(i)$, for all i). This follows from the pseudo-competitive property: for any i, $R(p, \mathbf{p}^{S_i}_{-\pi(i)}) = R(p, \mathbf{p}_{-\pi(i)})$. Hence, the local optimality of a seller in the ordering is preserved even if the prices of lower-ranked sellers are reduced to 0, implying the validity of these prices with respect to π.

Hence, it remains to consider the case where at least two prices in $\mathbf{p}$ are equal. Consider the first set of sellers with equal prices (i.e., those ranked highest in the ordering), and let $l+1$ be the size of this set. We denote these sellers by $\{j_0, \ldots, j_l\}$, with $j_0, \ldots, j_l$ following the same order as π. By the same argument as the one used when all prices in $\mathbf{p}$ are distinct, the price of any seller ranked higher than j_0 is locally optimal with respect to $\mathbf{p}$. We argue that the local optimality of $p_{j_0} = \ldots = p_{j_l}$ is not perturbed.

Case 1: $p_{j_0} = p_{j_1} = \cdots = p_{j_l} = 0$. For any $j \in \{j_0, \ldots, j_l\}$, $R_j(p_j, \mathbf{p}_{-j}) = R_j(p_j, \mathbf{p}^{S_j}_{-j}) = 0$. Thus, if 0 is a local optimum for j over $[0, \bar{v}]$, it is necessarily a local optimum over j's relevant undercutting region.

Case 2: $p_{j_0} = p_{j_1} = \cdots = p_{j_l} > 0$. We will argue that p_{j_0} is locally optimal over its relevant undercutting region if any subset $\mathcal{S} \subseteq \{j_1, \cdots, j_l\}$ of sellers reduce their price to 0. By symmetry, then, the same will hold true for any $j \in \{j_0, \cdots, j_l\}$. Repeating this argument for the next set of equally priced sellers below $\{j_0, \ldots, j_l\}$, until no equally priced sellers remain (or all sellers are priced at zero), we will obtain the theorem.

Lemma EC.1 plays a key role in proving this claim. It establishes that sellers $j_1, \ldots, j_l$ reducing their prices does not impact the expected revenue of j_0.

LEMMA EC.1. *$R_{j_0}(p, \mathbf{p}_{-j_0}) = R_{j_0}(p, \mathbf{p}'_{-j_0})$ for all $\mathbf{p}'$ such that $p'_{j_i} \leq p_{j_i}$, $i \in \{1, \ldots, l\}$.*

We leverage Lemma EC.1 to show that p_{j_0} remains locally optimal under any price vector $\mathbf{p}'$ such that $p'_{j_i} = 0$ for some subset of sellers $\mathcal{S} \subseteq \{j_1, \ldots, j_l\}$. Lemma EC.2 establishes that it is never optimal for j_0 to ϵ-undercut its original price p_{j_0} under $\mathbf{p}'$.

LEMMA EC.2. *There exists $\epsilon_0 > 0$ such that, for all $0 < \epsilon \leq \epsilon_0$, $R_{j_0}(p_{j_0} - \epsilon, \mathbf{p}'_{-j_0}) \leq R_{j_0}(p_{j_0}, \mathbf{p}'_{-j_0})$.*

We conclude the proof by establishing that it is never optimal for j_0 to ϵ-overcut p_{j_0} under $\mathbf{p}'$.

LEMMA EC.3. *There exists $\epsilon_0 > 0$ such that, for all $0 < \epsilon \leq \epsilon_0$, $R_{j_0}(p_{j_0} + \epsilon, \mathbf{p}'_{-j_0}) \leq R_{j_0}(p_{j_0}, \mathbf{p}'_{-j_0})$.*

Putting Lemmas EC.2 and EC.3 together, the local optimality of p_{j_0} is preserved under any such $\mathbf{p}'$. □

Proof of Lemma EC.1. Consider the subset $\mathcal{S} \subseteq \{j_1, \ldots, j_l\}$ for which $p'_{j_i} < p_{j_i}$. As argued in the proof of Proposition EC.2, for j_0 to be affected by this price decrease, it must be that there exists a customer c such that $\{j_0, j\} \subseteq \mathcal{N}_c$ for some $j \in \mathcal{S}$. For any such customer c, it moreover must have been that j_0 and j were ranked equally on all attributes up until price, and that c had chosen j_0 (else, c chose j both before and after the price change, resulting in no change in revenue for j_0).

Given this, since j strictly decreased its price, all customers who chose j_0 based on an attribute k ranked lower than price under $\mathbf{p}$, now choose j under $\mathbf{p}'$. We claim that this occurs for a measure-zero set of customers. To see this, suppose there existed $j \in \{j_1, \ldots, j_l\}$ such that a strictly positive mass of customers chose j_0 over j based on an attribute ranked lower than price. Then, j would stand to gain a discrete mass of customers by infinitesimally undercutting p_{j_0}, thus contradicting the local optimality of $p_j = p_{j_0}$.

Since only a measure-zero set of customers move their choice from j_0 to $j \in \mathcal{S}$ under this price change, the expected demand (and hence expected revenue) of j_0 is invariant to weakly lower price changes of $j_1, \ldots, j_l$. □

Proof of Lemma EC.2. By local optimality of p_{j_0} under $\mathbf{p}$, there exists $\epsilon_0 > 0$ such that for $0 < \epsilon \leq \epsilon_0$:

$$\begin{aligned} 0 \leq R_{j_0}(p_{j_0}, \mathbf{p}_{-j_0}) - R_{j_0}(p_{j_0} - \epsilon, \mathbf{p}_{-j_0}) &\leq R_{j_0}(p_{j_0}, \mathbf{p}_{-j_0}) - R_{j_0}(p_{j_0} - \epsilon, \mathbf{p}'_{-j_0}) \\ &= R_{j_0}(p_{j_0}, \mathbf{p}'_{-j_0}) - R_{j_0}(p_{j_0} - \epsilon, \mathbf{p}'_{-j_0}). \end{aligned} \tag{EC.1}$$

Here, the second inequality follows from cross-monotonicity (Proposition EC.2), and the equality follows from Lemma EC.1. □

Proof of Lemma EC.3. By local optimality of p_{j_0} under $\mathbf{p}$, there exists $\epsilon_0 > 0$ such that, for all $0 < \epsilon \leq \epsilon_0$:

$$\begin{aligned} 0 \leq R_{j_0}(p_{j_0}, \mathbf{p}_{-j_0}) - R_{j_0}(p_{j_0} + \epsilon, \mathbf{p}_{-j_0}) &= R_{j_0}(p_{j_0}, \mathbf{p}_{-j_0}) - R_{j_0}(p_{j_0} + \epsilon, \mathbf{p}'_{-j_0}) \\ &= R_{j_0}(p_{j_0}, \mathbf{p}'_{-j_0}) - R_{j_0}(p_{j_0} + \epsilon, \mathbf{p}'_{-j_0}). \end{aligned} \tag{EC.2}$$

Here, the first equality follows from the pseudo-competitive property, and the second follows from Lemma EC.1. Thus, increasing its price is not locally beneficial to j_0 under price vector $\mathbf{p}'$. □

B.4. Unimodality: Special cases

Below, we provide a sufficient condition — a competitive extension of the notion of *regularity* in the monopolist pricing literature — for Assumption 1 to hold.

Let $\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})$ denote the probability that a customer chooses seller i at price p, given the prices of other sellers in set $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$. Moreover, let $c_i(p, \mathbf{p}^{\mathcal{S}}_{-i}) = \frac{\partial \mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})}{\partial p}$. We say that $\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})$ is *regular* if $\frac{\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})}{c_i(p, \mathbf{p}^{\mathcal{S}}_{-i})} + p$ exists and is increasing in p, for all $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$, and $\mathbf{p}^{\mathcal{S}}_{-i}$.

PROPOSITION EC.3. *Suppose $\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})$ is regular, for all $i \in \mathcal{N}$. Then, Assumption 1 holds.*

We will use Proposition EC.3 to show that the two special cases we present satisfy Assumption 1. We defer its proof to Appendix B.4.1.

In the remainder of this work, we let F denote the cumulative distribution function of customers' willingness-to-pay, and f the corresponding density. For ease of notation we define $\bar{F}(\cdot) = 1 - F(\cdot)$.

PROPOSITION EC.4. *Consider a setting in which (i) there is a mass $\lambda \in [0,1)$ of BICA-type customers, (ii) each seller $i \in \mathcal{N}$ has $\gamma_i > 0$ mass of loyal customers, and (iii) the remaining mass of customers are greedy, breaking ties according to quality. If customers' willingness-to-pay distribution is independent of other model primitives, and moreover has a monotone hazard rate[1] (MHR), then $\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})$ is regular for all $i \in \mathcal{N}$. Therefore, Assumption 1 holds.*

Examples of distributions with monotone hazard rate include the uniform distribution, the exponential distribution, and the Weibull distribution with shape $k \geq 1$.

Proof. For ease of notation, let $\gamma = \sum_{i=1}^{N} \gamma_i$. Then, the mass of greedy customers is given by $1 - \gamma - \lambda$. Now, fix seller $i \in \mathcal{N}$, and suppose $p_i = p$. We consider two cases.

Case 1: There exists $k \in \mathcal{N}$ such that $p_k \leq p$ and $q_k > q_i$ in $\mathbf{p}^{\mathcal{S}}_{-i}$. In this case, no BICA customer chooses i, since anyone who can afford p can also afford p_k, and will choose the higher-quality seller k as a result. Additionally, no greedy customer will choose i. Therefore, $\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i}) = \gamma_i \bar{F}(p)$, since only seller i's loyal customers purchase from i. We now verify that the regularity condition holds. We have:

$$\frac{\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})}{c_i(p, \mathbf{p}^{\mathcal{S}}_{-i})} + p = -\frac{\bar{F}(p)}{f(p)} + p,$$

which is increasing by the MHR assumption.

Case 2: $p_k > p$ for all $k \in \mathcal{N}$ such that $q_k > q_i$ in $\mathbf{p}^{\mathcal{S}}_{-i}$. In this case, the BICA-type customers who choose i are those for whom the prices of the higher-quality sellers exceed their willingness-to-pay. We further consider the following cases:

1. $p > 0$ and there exists $k \in \mathcal{N}$ such that $p_k < p$: In this case, no greedy customer chooses i. Therefore:

$$\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i}) = \gamma_i \bar{F}(p) + \lambda \mathbb{P}(p \leq w_c < \min_{k \in \mathcal{N}: q_k > q_i} p_k) = \bar{F}(p)(\lambda + \gamma_i) - \lambda \bar{F}(\min_{k \in \mathcal{N}: q_k > q_i} p_k) \tag{EC.3}$$

Therefore:

$$\frac{\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})}{c_i(p, \mathbf{p}^{\mathcal{S}}_{-i})} + p = -\frac{\bar{F}(p)(\lambda + \gamma_i) - \lambda \bar{F}(\min_{k \in \mathcal{N}: q_k > q_i} p_k)}{f(p)(\lambda + \gamma_i)} + p. \tag{EC.4}$$

We prove that the right-hand side of (EC.4) is increasing in p, which will establish the regularity of $\mathcal{C}_i$. For any interval over which $f(p)$ is non-decreasing, the right-hand side of (EC.4) is increasing, since $\bar{F}(p)$ is decreasing. For any interval over which $f(p)$ is decreasing, the right-hand side of (EC.4) is also increasing, since $\frac{\bar{F}(p)}{f(p)}$ is non-increasing by the MHR assumption, and $\frac{1}{f(p)}$ is non-decreasing.

[1] A distribution has a monotone hazard rate if its hazard rate $h(p) = \frac{f(p)}{1 - F(p)}$ is non-decreasing for all p, where $f(\cdot)$ and $F(\cdot)$ are respectively used to denote the probability density and cumulative distribution functions (Hartline 2013).

2. $\{p=0\}$, or $\{p>0$ and $p_k \geq p$ for all $k \in \mathcal{N}\}$: In this case, all greedy customers choose i since it is both the cheapest and the highest-quality seller. Adding the contribution of these greedy customers to (EC.3), we have:

$$\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i}) = \bar{F}(p)(\lambda + \gamma_i + (1-\gamma-\lambda)) - \lambda \bar{F}(\min_{k \in \mathcal{N}: q_k > q_i} p_k) = \bar{F}(p)(\gamma_i + 1 - \gamma) - \lambda \bar{F}(\min_{k \in \mathcal{N}: q_k > q_i} p_k). \tag{EC.5}$$

Regularity then follows from the same arguments as above.

□

We now provide sufficient conditions on the CLC model for Assumption 1 to hold in a duopoly. For any customer c, let $\overline{\mathcal{N}}_c$ denote the consideration set formed *without* accounting for the willingness-to-pay constraint, i.e., $\overline{\mathcal{N}}_c = \{i \in \mathcal{N} : (v_i^1, \ldots, v_i^K) \in \mathcal{Y}_c\}$. For $i, j \in \mathcal{N}$, conditional on the event $\overline{\mathcal{N}}_c = \{1,2\}$, we use $\overline{\lambda}_{i \succ j}$ to denote the probability that a customer strictly prefers i to j due to a non-price attribute ranked higher than price. Moreover, we let $\overline{\lambda}_{i=j}$ denote the probability that a customer equally prefers i and j on all attributes until price.

PROPOSITION EC.5. *Consider a duopoly, and suppose each seller $i \in \mathcal{N}$ has a positive fraction of customers who strictly prefer i to j based on an attribute ranked higher than price. Suppose moreover that customers' willingness-to-pay distribution is independent of all other CLC choice primitives, and has a monotone hazard rate. Then, $\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})$ is regular for all $i \in \mathcal{N}$. Therefore, Assumption 1 holds.*

Proof. Fix $i \in \mathcal{N}$, and suppose $p_i = p$. Moreover, for ease of notation, let $j = -i$. We derive a closed-form expression for $\mathcal{C}_i$ by considering the possible revenues obtained for each realization of $\overline{\mathcal{N}}_c$.

Case 1: $\overline{\mathcal{N}}_c = \{i\}$. In this case, seller i obtains revenue from all customers such that $w_c \geq p$.

Case 2: $\overline{\mathcal{N}}_c = \{j\}$. In this case, seller i obtains no revenue.

Case 3: $\overline{\mathcal{N}}_c = \{1,2\}$. In this case, customers purchase from i only if $w_c \geq p$. Of those customers, seller i obtains revenue from all customers who strictly prefer i based on an attribute ranked higher than price. Consider now the following sub-cases for the other sources of revenue:

1. $p > p_j$: In this case, seller i receives no additional revenue. This follows from the fact that customers who equally prefer i and j on all attributes up until price will break ties in favor of j. Moreover, customers who strictly prefer j based on an attribute ranked before price, and for whom $w_c < p_j$, will also have $w_c < p$.
2. $p < p_j$: In this case, seller i additionally receives revenue from (i) customers who equally prefer i and j on all attributes up until price, and (ii) customers who strictly prefer j based on an attribute ranked before price, but for whom $p \leq w_c < p_j$.
3. $p = p_j$: This would only occur if $j \notin \mathcal{S}$, since if $j \in \mathcal{S}$, we assume $p < p^{\min}(\mathcal{S}) = p_j$, by definition. Since $j \notin \mathcal{S}$, this implies that $p_j = 0$ in $\mathbf{p}^{\mathcal{S}}_{-i}$. This therefore implies that $p = 0$; we omit this case from

consideration as a result, since i can always make strictly positive revenue from the customers who strictly prefer i to j based on an attribute ranked higher than price. Therefore, the gradient cannot be nonpositive at zero.

Putting this case-by-case analysis together, we finally obtain the following two cases for the closed-form expression of $\mathcal{C}_i$.

Case a: $p > p_j$. Then,

$$\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i}) = \mathbb{P}(\overline{\mathcal{N}}_c = \{i\})\bar{F}(p) + \mathbb{P}(\overline{\mathcal{N}}_c = \{1,2\})\overline{\lambda}_{i\succ j}\bar{F}(p) \implies \frac{\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})}{c_i(p, \mathbf{p}^{\mathcal{S}}_{-i})} + p = -\frac{\bar{F}(p)}{f(p)} + p. \quad \text{(EC.6)}$$

Since F is MHR by assumption, $\frac{\bar{F}(p)}{f(p)}$ is non-increasing. Plugging this fact into the above, we obtain that $\mathcal{C}_i$ is regular in this case.

Case b: $p < p_j$. Then,

$$\begin{aligned}
\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i}) &= \mathbb{P}(\overline{\mathcal{N}}_c = \{i\})\bar{F}(p) + \mathbb{P}(\overline{\mathcal{N}}_c = \{1,2\})\big(\bar{F}(p)\big(\overline{\lambda}_{i\succ j} + \overline{\lambda}_{i=j}\big) + \overline{\lambda}_{j\succ i}\big(\bar{F}(p) - \bar{F}(p_j)\big)\big) \\
&= \mathbb{P}(\overline{\mathcal{N}}_c = \{i\})\bar{F}(p) + \mathbb{P}(\overline{\mathcal{N}}_c = \{1,2\})\big(\bar{F}(p) - \overline{\lambda}_{j\succ i}\bar{F}(p_j)\big) \\
&= \bar{F}(p)\big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{i\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big) - \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\bar{F}(p_j)\overline{\lambda}_{j\succ i}, \qquad \text{(EC.7)}
\end{aligned}$$

where the second equality follows from the fact that $\overline{\lambda}_{i\succ j} + \overline{\lambda}_{i=j} + \overline{\lambda}_{j\succ i} = 1$.

To see that $\mathcal{C}_i$ is regular, note that:

$$\frac{\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})}{c_i(p, \mathbf{p}^{\mathcal{S}}_{-i})} + p = p - \frac{\bar{F}(p)\big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{i\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big) - \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\bar{F}(p_j)\overline{\lambda}_{j\succ i}}{f(p)\big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{i\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big)}. \quad \text{(EC.8)}$$

For any interval over which $f(p)$ is non-decreasing, the right-hand side of the above is increasing, since $\bar{F}(p)\big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{i\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big) - \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\bar{F}(p_j)\overline{\lambda}_{j\succ i}$ is decreasing. On the other hand, for any interval over which $f(p)$ is increasing, $\bar{F}(p)/f(p)$ is non-increasing due to the MHR assumption; moreover, $\frac{\mathbb{P}\left(\overline{\mathcal{N}}_c=\{1,2\}\right)\bar{F}(p_j)\overline{\lambda}_{j\succ i}}{f(p)\left(\mathbb{P}\left(\overline{\mathcal{N}}_c=\{i\}\right)+\mathbb{P}\left(\overline{\mathcal{N}}_c=\{1,2\}\right)\right)}$ is decreasing. Plugging these facts in (EC.8), we obtain that $\mathcal{C}_i$ is regular. □

B.4.1. Proof of Proposition EC.3

Proof. By definition, $R_i(p, \mathbf{p}^{\mathcal{S}}_{-i}) = p \cdot \mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})$. Therefore, $g_i(p, \mathbf{p}^{\mathcal{S}}_{-i}) = \mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i}) + p \cdot c_i(p, \mathbf{p}^{\mathcal{S}}_{-i})$. Suppose now that $g_i(p, \mathbf{p}^{\mathcal{S}}_{-i}) \le 0$ for some $p \in [0, p^{\min}(\mathcal{S}))$. Then, we must show that $g_i(p', \mathbf{p}^{\mathcal{S}}_{-i}) < 0$ for all $p' \in (p, p^{\min}(\mathcal{S}))$. This latter inequality holds if, for all $p' \in (p, p^{\min}(\mathcal{S}))$:

$$\mathcal{C}_i(p', \mathbf{p}^{\mathcal{S}}_{-i}) + p' \cdot c_i(p', \mathbf{p}^{\mathcal{S}}_{-i}) < 0 \iff \frac{\mathcal{C}_i(p', \mathbf{p}^{\mathcal{S}}_{-i})}{c_i(p', \mathbf{p}^{\mathcal{S}}_{-i})} + p' > 0,$$

where the implication follows from the fact that a price increase can only decrease the likelihood that a customer purchases from seller i, therefore $c_i(p', \mathbf{p}^{\mathcal{S}}_{-i}) \le 0$. Since $\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})$ is regular and $p' > p$, we have:

$$\frac{\mathcal{C}_i(p', \mathbf{p}^{\mathcal{S}}_{-i})}{c_i(p', \mathbf{p}^{\mathcal{S}}_{-i})} + p' > \frac{\mathcal{C}_i(p, \mathbf{p}^{\mathcal{S}}_{-i})}{c_i(p, \mathbf{p}^{\mathcal{S}}_{-i})} + p \ge 0,$$

which proves the claim. □

B.5. Proof of Proposition 1

Proof. The proof is constructive. Let π denote an arbitrary ordering over sellers.

Initialization. Fix an arbitrary π, mark it as inspected, and proceed to Phase 1 below.

Phase 1: Valid order-price pair computation. Compute the VOP associated with π (see Definition 5). If a corresponding VOP exists, let $\mathbf{p}$ denote the resulting price vector, and proceed to Phase 2. Else, choose a new, uninspected ordering. Mark this new ordering as inspected, and go back to the beginning of Phase 1.

Phase 2: LNE / GNE verification. For $i \in [N]$, compute:

$$\max\left\{\sup_{p\in[p_{\pi(i)},p_{\pi(i-1)}]} R_{\pi(i)}(p,\mathbf{p}_{-\pi(i)}), \sup_{p\in[p_{\pi(i+1)},p_{\pi(i)}]} R_{\pi(i)}(p,\mathbf{p}_{-\pi(i)})\right\}.$$

If this quantity strictly exceeds $R_{\pi(i)}(\mathbf{p})$ for some $i \in [N]$, then $\mathbf{p}$ is not a LNE (and hence, not a GNE); else, it is a LNE.

If $\mathbf{p}$ is a LNE, verify that it is a GNE by computing for $i \in [N]$, $j \in [N] \setminus \{i\}$:

$$\begin{cases} \sup_{p\in[p_{\pi(j)},p_{\pi(j-1)}]} R_{\pi(i)}(p,\mathbf{p}_{-\pi(i)}) & \text{if } j<i \text{ or } j>i+1 \\ \sup_{p\in[p_{\pi(i+1)},p_{\pi(i-1)}]} R_{\pi(i)}(p,\mathbf{p}_{-\pi(i)}) & \text{if } j=i+1. \end{cases}$$

If for $i \in [N]$, any of these revenues is strictly higher than $R_{\pi(i)}(\mathbf{p})$, then $\mathbf{p}$ is not a GNE; else, it is. Choose a new, uninspected ordering, mark it as inspected, and proceed to the beginning of Phase 1. If all $N!$ orderings have been traversed, either output all identified LNE/GNE, or return that none exists.

Correctness and termination. By the unimodality assumption, for each ordering π there exists at most one VOP $(\pi, \mathbf{p})$. Hence, by Theorem 2, if a LNE exists, it must have been found in Phase 1. Since a GNE is a LNE, it must also have been found in Phase 1, if it exists. Phase 2 is a simple consistency check to see if the local optimum computed for each seller is locally / globally optimal with respect to $\mathbf{p}$.

We conclude the proof by giving an upper bound on the number of undercutting best-response computations required for the procedure to terminate. For each ordering (of which there are $N!$), at most N undercutting best-response computations are performed during Phase 1. Verifying whether each of these VOPs is a LNE requires an additional two best-response computations per seller, per ordering. This yields the upper bound on the number of best-response computations to find all LNE, or to find that none exists.

The second GNE verification phase requires at most N undercutting best-response computations per seller, for a given LNE. This then yields the $\mathrm{O}(N^2)$ best-response computations across all sellers.

The bound in Part 2 of the statement follows from putting these two together. Namely, each ordering requires at most $N^2 + N$ local best-response computations to check if the LNE is a GNE. Doing this for all $N!$ orderings gives rise to an upper bound of $\mathrm{O}((N^2+N)N!)$ undercutting best-response computations.

□

B.6. Gradient dominance: Special cases

In this section we present special cases of the CLC class for which gradient dominance holds, in particular highlighting the connection between gradient dominance and the strength of a seller along a non-price attribute.

We first illustrate that gradient dominance can arise due to quality differences between sellers. In particular, we show that gradient dominance weakly holds under Best-I-Can-Afford behavior, allowing us to leverage Proposition 2 to characterize all non-trivial LNE for this special case. Recall, under BICA behavior, the customer chooses the highest-quality seller whose price is at most her willingness-to-pay. Let q_i denote the quality of seller i, for $i \in \mathcal{N}$.

PROPOSITION EC.6. *Under BICA behavior, for any two sellers $i, j \in \mathcal{N}$ with qualities $q_i > q_j$, seller i weakly gradient-dominates seller j. Moreover, if all sellers' qualities are distinct and customers' willingness-to-pay distribution has a monotone hazard rate, there exists a unique non-trivial LNE in which sellers' prices are strictly ordered according to their respective qualities.*

Relating this back to the market dominance described in Section 4.1, under BICA behavior, for any sellers i and j such that $q_i > q_j$, seller i has no incentive to undercut j, since customers who can afford both sellers will always purchase from i. Seller j's customer base, on the other hand, is entirely composed of customers for whom i is too expensive, which pushes seller j's revenue-maximizing price downward. Mathematically, this translates to gradient dominance. We defer the proof of Proposition EC.6 to Appendix B.6.1.

We now provide sufficient conditions for gradient dominance to hold in a duopoly. Recall, for any customer c, $\overline{\mathcal{N}}_c$ denotes the consideration set formed *without* accounting for the willingness-to-pay constraint, i.e., $\overline{\mathcal{N}}_c = \{i \in \mathcal{N} : (v_i^1, \ldots, v_i^K) \in \mathcal{Y}_c\}$. For $i, j \in \mathcal{N}$, $\overline{\lambda}_{i \succ j}$ denotes the probability that a customer strictly prefers i to j due to a non-price attribute ranked higher than price, conditional on the event $\overline{\mathcal{N}}_c = \{1, 2\}$.

PROPOSITION EC.7. *Consider a duopoly in which customers' willingness-to-pay distribution is independent of all other CLC choice primitives, and has a monotone hazard rate. Then, seller i gradient-dominates seller j if and only if the following hold:*

1. *Non-price attribute dominance of seller i:*

$$\overline{\lambda}_{i \succ j} \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{i\}\big) > \big(1 - \overline{\lambda}_{i \succ j}\big) \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{j\}\big)$$

2. *Relative competitiveness of seller j:*

$$\big(1 - \overline{\lambda}_{i \succ j}\big) \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{j\}\big) \geq \mathbb{P}\big(\overline{\mathcal{N}}_c = \{i\}\big)$$

Under these two conditions, there exists a LNE in which $p_i^* > p_j^*$.

We provide some high-level intuition for Proposition EC.7, deferring its proof to Appendix B.6.2. Condition 1 requires the market dominance of seller i: if more customers are loyal to i, or prefer i to j based on an attribute ranked higher than price, then (i) seller j can capture low-valuation customers who could not afford their preferred seller i, and (ii) seller i has a higher incentive to increase its price when it is below its optimum. Condition 2 simply guarantees that j has an incentive to undercut i by ensuring that not too many customers are loyal to i; this essentially states that the market has to indeed be competitive.

The following corollary of Proposition EC.7 further illustrates the relationship between market and gradient dominance.

COROLLARY EC.1. *Consider a duopoly in which customers' willingness-to-pay distribution is independent of all other CLC choice primitives, and has a monotone hazard rate. Suppose moreover that customers always evaluate both sellers, and strictly more than half of customers prefer seller 1 to seller 2 according to an attribute ranked higher than price (i.e.,* $\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big) = 1$ *and* $\overline{\lambda}_{1\succ 2} > 1/2$*). Then, seller 1 gradient-dominates seller 2. In this case, there exists a LNE in which* $p_1^* > p_2^*$.

B.6.1. Proof of Proposition EC.6

Proof. We first argue that weak gradient dominance holds. Consider $\mathbf{p}$ such that $p_i > p_j$. We consider two cases.

Case 1: There exists $k \in \mathcal{N} \setminus \{i,j\}$ such that $q_k > q_i$ and $p_k \le p_i$. This is a trivial price configuration since no customer chooses i (anyone who can afford p_i can also afford p_k, and will choose this higher-quality seller). Hence, for any such price vector seller i has a revenue of zero.

Case 2: $p_k > p_i$ for all $k \in \mathcal{N}$ such that $q_k > q_i$. We have:

$$\begin{aligned} & R_i(p_i, p_j, \mathbf{p}_{-\{i,j\}}) = p_i \cdot \mathbb{P}\left(p_i \le w_c < \min_{k:q_k > q_i} p_k\right) \\ \implies & g_i(p_i, p_j, \mathbf{p}_{-\{i,j\}}) = \mathbb{P}\left(p_i \le w_c < \min_{k:q_k > q_i} p_k\right) - p_i f(p_i). \end{aligned}$$

Similarly:

$$\begin{aligned} & R_j(p_i, p_j, \mathbf{p}_{-\{i,j\}}) = p_j \cdot \mathbb{P}\left(p_j \le w_c < \min_{k:q_k > q_j} p_k\right) \\ \implies & g_j(p_i, p_j, \mathbf{p}_{-\{i,j\}}) = \mathbb{P}\left(p_j \le w_c < \min_{k:q_k > q_j} p_k\right) - p_j f(p_j) \qquad \text{(EC.9)} \\ \implies & \lim_{p_j \uparrow p_i} g_j(p_i, p_j, \mathbf{p}_{-\{i,j\}}) = -p_i f(p_i), \qquad \text{(EC.10)} \end{aligned}$$

where the final implication follows from the fact that $q_i > q_j$, therefore:

$$\lim_{p_j \uparrow p_i} \mathbb{P}\left(p_j \le w_c < \min_{k:q_k > q_j} p_k\right) \le \lim_{p_j \uparrow p_i} \mathbb{P}(p_j \le w_c < p_i) = 0.$$

Since $\lim_{p_j \uparrow p_i} \mathbb{P}\big(p_j \le w_c < \min_{k:q_k > q_j} p_k\big) \ge 0$, the above implies $\lim_{p_j \uparrow p_i} \mathbb{P}\big(p_j \le w_c < \min_{k:q_k > q_j} p_k\big) = 0$. Hence, we obtain:

$$g_i(p_i, p_j, \mathbf{p}_{-\{i,j\}}) - \lim_{p_j \uparrow p_i} g_j(p_i, p_j, \mathbf{p}_{-\{i,j\}}) = \mathbb{P}\left(p_i \le w_c < \min_{k:q_k > q_i} p_k\right) > 0.$$

Consider now the ordering π that ranks sellers in increasing order of quality, and note that this is the unique ordering that gives rise to a non-trivial price vector. We now argue that $\pi(i)$ attains its unique maximum in $(0, p_{\pi(i-1)})$, for any price vector $\mathbf{p}$ aligned with π. By (EC.9):

$$g_{\pi(i)}(p, \mathbf{p}_{-\pi(i)}) = \bar{F}(p) - \bar{F}(p_{\pi(i-1)}) - pf(p). \tag{EC.11}$$

Moreover, by the first-order condition:

$$p = \frac{\bar{F}(p) - \bar{F}(p_{\pi(i-1)})}{f(p)}. \tag{EC.12}$$

For any interval over which $f(p)$ is non-decreasing, the right-hand side of (EC.12) is decreasing, since $\bar{F}(p) - \bar{F}(p_{\pi(i-1)})$ is decreasing. For any interval over which $f(p)$ is decreasing, the right-hand side of (EC.12) is decreasing, since $\frac{\bar{F}(p)}{f(p)}$ is non-increasing by the MHR assumption, and $-\frac{\bar{F}(p_{\pi(i-1)})}{f(p)}$ is decreasing. Hence, the solution to (EC.12) is unique and attained in $(0, p_{\pi(i-1)})$, since any price in this interval yields strictly positive revenue.

Observe that the result does not follow immediately from Proposition 2, since unimodality and gradient dominance fail to hold for trivial price configurations in which a lower-quality seller prices above a high-quality seller. However, one can apply the proof of Proposition 2 to π, and we obtain the result. □

B.6.2. Proof of Proposition EC.7

Proof. For ease of notation, we let $i = 1, j = 2$. To show that seller 1 gradient-dominates seller 2, it suffices to consider $\mathbf{p}$ such that $p_1 > p_2$. We partition our analysis on the realization of $\overline{\mathcal{N}}_c$, and let $R_k(\mathbf{p} \mid \overline{\mathcal{N}}_c)$ denote the expected revenue of seller $k \in \mathcal{N}$, given $\overline{\mathcal{N}}_c$.

Case 1: $\overline{\mathcal{N}}_c = \{1\}$. Then, $R_1(\mathbf{p} \mid \overline{\mathcal{N}}_c) = p_1 \bar{F}(p_1)$, $R_2(\mathbf{p} \mid \overline{\mathcal{N}}_c) = 0$.

Case 2: $\overline{\mathcal{N}}_c = \{2\}$. Then, $R_1(\mathbf{p} \mid \overline{\mathcal{N}}_c) = 0$, $R_2(\mathbf{p} \mid \overline{\mathcal{N}}_c) = p_2 \bar{F}(p_2)$.

Case 3: $\overline{\mathcal{N}}_c = \{1, 2\}$. Let $\overline{\lambda}_{2=1}$ be the mass of customers who rank 2 and 1 equally on all attributes up until price, given $\overline{\mathcal{N}}_c = \{1, 2\}$. Since $p_1 > p_2$, the only customers who purchase the item from seller 1 are customers who strictly preferred it based on a non-price attribute ranked higher than price, and for whom $w_c \ge p_1$. Hence: $R_1(\mathbf{p} \mid \overline{\mathcal{N}}_c) = \overline{\lambda}_{1 \succ 2} p_1 \bar{F}(p_1)$.

On the other hand, the customers who purchase the item from seller 2 are those who satisfy one of three conditions: (1) $w_c \ge p_2$, and seller 2 was strictly preferred to seller 1 on a non-price attribute ranked higher than price, (2) $w_c \ge p_2$, and seller 2 was ranked identically to seller 1 on all attributes up until price, (3)

$p_2 \le w_c < p_1$, and seller 1 was strictly preferred to seller 2 based on a non-price attribute ranked higher than price. Then:

$$R_2(\mathbf{p} \mid \overline{\mathcal{N}}_c) = p_2\big(\big(\bar{F}(p_2) - \bar{F}(p_1)\big)\overline{\lambda}_{1\succ 2} + \bar{F}(p_2)\big(\overline{\lambda}_{2\succ 1} + \overline{\lambda}_{2=1}\big)\big) = p_2\big(\bar{F}(p_2) - \bar{F}(p_1)\overline{\lambda}_{1\succ 2}\big),$$

where the second equality follows from the fact that $\overline{\lambda}_{1\succ 2} + \overline{\lambda}_{2\succ 1} + \overline{\lambda}_{2=1} = 1$.

Putting these three cases together, we obtain the following:

$$\begin{aligned} R_1(\mathbf{p}) &= p_1\bar{F}(p_1)\big[\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1\}\big) + \overline{\lambda}_{1\succ 2}\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big] \\ \implies g_1(p_1,p_2) &= \big(\bar{F}(p_1) - p_1 f(p_1)\big)\big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1\}\big) + \overline{\lambda}_{1\succ 2}\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big), \end{aligned} \tag{EC.13}$$

and

$$\begin{aligned} R_2(\mathbf{p}) &= p_2\big[\bar{F}(p_2)\mathbb{P}\big(\overline{\mathcal{N}}_c = \{2\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big(\bar{F}(p_2) - \bar{F}(p_1)\overline{\lambda}_{1\succ 2}\big)\big] \\ &= p_2\big[\bar{F}(p_2)\big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{2\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big) - \bar{F}(p_1)\overline{\lambda}_{1\succ 2}\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big] \end{aligned} \tag{EC.14}$$

$$\begin{aligned} \implies g_2(p_1,p_2) &= \bar{F}(p_2)\big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{2\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big) - \bar{F}(p_1)\overline{\lambda}_{1\succ 2}\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big) \\ &\quad - p_2 f(p_2)\big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{2\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big) \\ &= \big(\bar{F}(p_2) - p_2 f(p_2)\big)\big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{2\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big) - \bar{F}(p_1)\overline{\lambda}_{1\succ 2}\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big) \end{aligned} \tag{EC.15}$$

$$\implies \lim_{p_2\uparrow p_1} g_2(p_1,p_2) = \big(\bar{F}(p_1) - p_1 f(p_1)\big)\big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{2\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big) - \bar{F}(p_1)\overline{\lambda}_{1\succ 2}\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big) \tag{EC.16}$$

Then:

$$\begin{aligned} &g_1(p_1,p_2) - \lim_{p_2\uparrow p_1} g_2(p_1,p_2) \\ &= \bar{F}(p_1)\overline{\lambda}_{1\succ 2}\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big) + \big(\bar{F}(p_1) - p_1 f(p_1)\big)\left(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1\}\big) - \mathbb{P}\big(\overline{\mathcal{N}}_c = \{2\}\big) + (\overline{\lambda}_{1\succ 2} - 1)\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\right). \end{aligned}$$

Dividing by $\bar{F}(p_1)$, we then have that, for all $p_1 \in (0, \bar{v})$:

$$\begin{aligned} &g_1(p_1,p_2) - \lim_{p_2\uparrow p_1} g_2(p_1,p_2) > 0 \\ \iff &\overline{\lambda}_{1\succ 2}\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big) > \left(1 - \frac{p_1 f(p_1)}{\bar{F}(p_1)}\right)\left(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{2\}\big) - \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1\}\big) + (1 - \overline{\lambda}_{1\succ 2})\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\right). \end{aligned}$$

Suppose first that $\mathbb{P}\big(\overline{\mathcal{N}}_c = \{2\}\big) - \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1\}\big) + (1 - \overline{\lambda}_{1\succ 2})\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big) \ge 0$. Since F is MHR, $1 - \frac{p_1 f(p_1)}{\bar{F}(p_1)}$ is non-increasing. Thus, the maximum of the right-hand side is achieved at $p_1 = 0$. For seller 1 to gradient-dominate seller 2, we must then have:

$$\overline{\lambda}_{1\succ 2}\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big) > \mathbb{P}\big(\overline{\mathcal{N}}_c = \{2\}\big) - \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1\}\big) + (1 - \overline{\lambda}_{1\succ 2})\mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)$$

Suppose now that $\mathbb{P}\big(\overline{\mathcal{N}}_c=\{2\}\big)-\mathbb{P}\big(\overline{\mathcal{N}}_c=\{1\}\big)+(1-\overline{\lambda}_{1\succ 2})\mathbb{P}\big(\overline{\mathcal{N}}_c=\{1,2\}\big)<0$. Then, again by the MHR assumption, the supremum of the right-hand side is attained at $p_1=\bar{v}$, with

$$\lim_{p_1\uparrow\bar{v}}\left(1-\frac{p_1 f(p_1)}{\bar{F}(p_1)}\right)\left(\mathbb{P}\big(\overline{\mathcal{N}}_c=\{2\}\big)-\mathbb{P}\big(\overline{\mathcal{N}}_c=\{1\}\big)+(1-\overline{\lambda}_{1\succ 2})\mathbb{P}\big(\overline{\mathcal{N}}_c=\{1,2\}\big)\right)=+\infty.$$

Thus, there exist values of $p_1\in(0,\bar{v})$ for which $g_1(p_1,p_2)\leq\lim_{p_2\uparrow p_1}g_2(p_1,p_2)$, and seller 1 does not gradient-dominate seller 2.

To establish that there exists a LNE $\mathbf{p}^*$ under which $p_1^*>p_2^*$, it suffices to show that each seller's undercutting revenue curve is unimodal for price vectors $\mathbf{p}$ such that $p_1>p_2$. Inspecting (EC.13) and (EC.15), unimodality follows from the MHR assumption. One can then use the same arguments as those used to establish Proposition 2 for these price configurations and obtain the result. □

B.7. Proof of Proposition 2

We first introduce some useful terminology. If seller i gradient-dominates seller j and $p_i>p_j$, we say that sellers i and j are *correctly ordered*. Otherwise, i and j are said to be *incorrectly ordered*. The following fact emerges as a consequence of gradient dominance.

LEMMA EC.4. *Under gradient dominance, there exists an ordering π such that, for all $\mathbf{p}\in\mathcal{V}^N$ aligned with π such that $p_{\pi(N)}>0$ and all $l\in[N-1]$:*

$$g_{\pi(l)}\left(p_{\pi(l)},\mathbf{p}_{-\pi(l)}^{\pi(1:l-1)}\right)>\lim_{p\uparrow p_{\pi(l)}}g_{\pi(l+1)}\left(p,\mathbf{p}_{-\pi(l+1)}^{\pi(1:l)}\right). \tag{EC.17}$$

Proof. The proof is constructive. Consider any ordering π such that (EC.17) does not hold, and let l^* denote the first rank in the ordering such that $\exists\mathbf{p}\in\mathcal{V}^N$ such that $g_{\pi(l)}\left(p_{\pi(l)},\mathbf{p}_{-\pi(l)}^{\pi(1:l-1)}\right)\leq\lim_{p\uparrow p_{\pi(l)}}g_{\pi(l+1)}\left(p,\mathbf{p}_{-\pi(l+1)}^{\pi(1:l)}\right)$, with $p_{\pi(N)}>0$. By gradient dominance, then, $g_{\pi(l+1)}\left(p_{\pi(l+1)},\mathbf{p}_{-\pi(l+1)}^{\pi(1:l)}\right)>\lim_{p\uparrow p_{\pi(l+1)}}g_{\pi(l)}\left(p,\mathbf{p}_{-\pi(l)}^{\pi(1:l-1)}\right)$ for all $\mathbf{p}\in\mathcal{V}^N$ such that $p_{\pi(N)}>0$. Define ordering π', which flips the two incorrectly ordered sellers, i.e.:

$$\begin{cases}\pi'(l) & =\pi(l)\quad\forall\, l\in[N]\setminus\{l^*,l^*+1\}\\ \pi'(l^*) & =\pi(l^*+1)\\ \pi'(l^*+1) & =\pi(l^*).\end{cases}$$

Repeat this process with π', as long as there exist incorrectly ordered sellers in the newly generated ordering π'. To see why this process must terminate with all sellers correctly ordered, note that, once two sellers have been swapped, by gradient dominance they will never be swapped again. Thus, the algorithm never revisits the same ordering. Since there are finitely many orderings and, for each ordering, finitely many pairs to potentially swap, this algorithm must terminate in an ordering π such that (EC.17) holds. □

We use this fact to establish existence of LNE.

Proof of Proposition 2. For $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$, let $p_i^*(\mathbf{p}_{-i}^{\mathcal{S}}) \in \arg\max_{p \in [0, p^{\min}(\mathcal{S}))} R_i(p, \mathbf{p}_{-i}^{\mathcal{S}})$ denote the unique optimal price of seller i in the open undercutting region relative to the set of sellers $\mathcal{S}$. (Uniqueness follows from Assumption 1.)

We first argue that $p_i^*(\mathbf{p}_{-i}^{\mathcal{S}}) \neq 0$, if it exists. Suppose for contradiction that $p_i^*(\mathbf{p}_{-i}^{\mathcal{S}}) = 0$. Since $R_i(0, \mathbf{p}_{-i}^{\mathcal{S}}) = 0$, the optimal revenue is 0, which implies that $R_i(p, \mathbf{p}_{-i}^{\mathcal{S}}) = 0$ for all $p \in [0, p^{\min}(\mathcal{S})]$, contradicting the unimodality assumption. Thus, $p_i^*(\mathbf{p}_{-i}^{\mathcal{S}}) \in (0, p^{\min}(\mathcal{S}))$, and moreover $g_i(p_i^*(\mathbf{p}_{-i}^{\mathcal{S}}), \mathbf{p}_{-i}^{\mathcal{S}}) = 0$.

Now, let π be such that (EC.17) holds, and define $p_{\pi(1)}^* = \arg\max_p R_{\pi(1)}(p, \mathbf{p}_{-\pi(1)}^{\phi})$. Using similar arguments as above, $p_{\pi(1)}^* \in (0, \bar{v})$, and $g_{\pi(1)}(p_{\pi(1)}^*, \mathbf{p}_{-\pi(1)}^{\phi}) = 0$. By (EC.17), we then have

$$\lim_{p \uparrow p_{\pi(1)}} g_{\pi(2)}\left(p, \mathbf{p}_{-\pi(2)}^{\pi(1)}\right) < g_{\pi(1)}(p_{\pi(1)}^*, \mathbf{p}_{-\pi(1)}^{\phi}) = 0,$$

which implies that $R_{\pi(2)}(\cdot, \mathbf{p}_{-\pi(2)}^{\pi(1)})$ attains its unique maximum $p_{\pi(2)}^*$ in $(0, p_{\pi(1)}^*)$, with $g_{\pi(2)}\left(p_{\pi(2)}^*, \mathbf{p}_{-\pi(2)}^{\pi(1)}\right) = 0$, by Assumption 1. Iterating this argument for $l = 3, \ldots, N$, we obtain that $\mathbf{p}^*$ is a VOP with distinct prices, and hence a LNE.

Note that finding an ordering π such that seller i is ranked above seller j if and only if i gradient-dominates j requires $O(N^2)$ comparisons. Given π, it suffices to find the aligned price vector $\mathbf{p}^*$, which can be done in $O(N)$ computations. □

B.8. Non-Equivalence of LNE and VOPs

Given Theorem 2, a natural question that arises is whether a VOP is necessarily a LNE. By the pseudo-competitive property, if all prices in a VOP are distinct, the corresponding price vector is necessarily a LNE. The issue then arises for VOPs in which two sellers price equally. We next establish that there exist settings for which such VOPs need not be LNE; as a result, the converse of Theorem 2 does not hold in general.

PROPOSITION EC.8. *There exist instances of CLC choice for which a VOP is not a LNE.*

Proof. Consider two sellers selling an identical item, with the only attributes being price and brand. All customers form their consideration set $\mathcal{N}_c$ based only on whether their willingness-to-pay w_c exceeds the item's price. Moreover, customers are partitioned equally into three types:

- Type 1 customers: rank brand over price and prefer seller 1 to seller 2;
- Type 2 customers: rank brand over price and prefer seller 2 to seller 1;
- Type 3 customers: rank price over brand and prefer seller 1 to seller 2.

Finally, suppose the willingness-to-pay w_c of customers is independent of type, and $w_c \sim \text{Unif}[0, 1]$.

In this case, both sellers setting a price of 0.5 is the only VOP in this game, corresponding to the ordering $\pi = (2, 1)$. To see this, consider first the ordering $\pi = (1, 2)$ in which seller 1 comes first, assuming seller 2 prices at 0. In this case, the best seller 1 can do is price at 0.5. This is obtained by optimizing only over Type 1 customers who have $w_c \sim \text{Unif}[0, 1]$, since Type 2 and Type 3 customers will both choose seller 2 at

a price of 0. For any $p < 0.5$, Type 2 and Type 3 customers choose seller 2; optimizing over these customers over $[0, 0.5)$ yields an optimum of $0.5 - \epsilon$ for seller 2, for arbitrarily small $\epsilon > 0$. If seller 2 prices at 0.5, on the other hand, it loses all Type 3 customers. Hence, the supremum of seller 2's undercutting revenue is not attained over $[0, 0.5]$, and the valid order-price pair corresponding to $\pi = (1, 2)$ does not exist.

Consider now the ordering where seller 2 comes first, assuming seller 1 prices at 0, followed by seller 1 pricing in the undercutting region. In this case, the best seller 2 can do is to optimize over Type 2 customers, since Type 1 and 3 customers always choose seller 1 when it prices at 0. Then, seller 2 sets a price of 0.5. For any $p \leq 0.5$, Type 1 and Type 3 customers choose seller 1. Optimizing over these customers, seller 1 has a local optimum of 0.5 over its undercutting region.

This VOP is not a LNE since seller 2 stands to gain all Type 3 customers from seller 1 by marginally undercutting seller 1 and setting a price of $0.5 - \epsilon$ for some small $\epsilon > 0$, for the same reason as described above. Since the only VOP is not a LNE, Theorem 2 implies that there is no LNE in this game. □

Appendix C: Section 5: Additional Results and Omitted Proofs

C.1. Specification of simulation framework

In Algorithm 1 we provide pseudo-code for the simulation framework described in Section 5. Our framework assumes sellers implement the classical Kiefer-Wolfowitz algorithm for stochastic function optimization with zero-order feedback (Kiefer and Wolfowitz 1952).

C.2. Formal description of Distributed Gradient Ascent dynamics

We present the Distributed Gradient Ascent algorithm in Algorithm 2.

Note that, in the case where $p_{i,t} = p_{j,t}$ for some $j \neq i$, by the pseudo-competitive property, the right-derivative $\lim_{p_i \downarrow p_{i,t}} g_i(p_i, p_j = p_{j,t} \, \forall \, j \in \mathcal{S}, \mathbf{p}_{-\{\mathcal{S},i\},t}) = g_i(p_{i,t}, p_j = 0 \, \forall \, j \in \mathcal{S}, \mathbf{p}_{-\{\mathcal{S},i\},t})$.

C.3. Assumption 3: Sufficient Conditions

We first provide a sufficient condition on sellers' revenue functions for Condition 3 to hold. We defer its proof to Appendix C.3.1.

LEMMA EC.5. *Suppose $R_i(p, \mathbf{p}^{\mathcal{S}}_{-i})$ is strongly concave with respect to p, for all $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$, $\mathbf{p}^{\mathcal{S}}_{-i}$. Then, Condition 3 holds.*

We now leverage Lemma EC.5 to provide sufficient conditions on the CLC choice model for Assumption 3 to hold in a duopoly.

PROPOSITION EC.9. *Consider a duopoly in which customers' willingness-to-pay distribution is independent of all other CLC choice primitives, and has Lipschitz continuous pdf and cdf. Then, Conditions 1 and 2 of Assumption 3 hold. Moreover:*

- *If* $\min_i \mathbb{P}(\overline{\mathcal{N}}_c = \{i\}) + \mathbb{P}(\overline{\mathcal{N}}_c = \{1, 2\})\overline{\lambda}_{i \succ -i} > 0$, *and there exists* $c > 0$ *such that* $2f(p) + pf'(p) \geq c$ *for all* $p \in [0, \bar{v}]$, *then Condition 3 holds.*

Algorithm 1 Equilibrium simulator

Initialize: Customer population size M, estimated CLC choice model, time horizon T, finite difference width schedule $(c_\tau, \tau \in \mathbb{N})$, stepsize $(a_\tau, \tau \in \mathbb{N})$

Output: Final prices $\mathbf{p}_T$

Generate an initial set of prices $\mathbf{p}_0 \in \mathcal{V}^N$ uniformly at random.

Initialize $\tau_i = 0$, for $i \in \mathcal{N}$.

for $t = 0, \ldots, T-1$ **do**

Generate batch of M customers.

Choose a seller $i \in \mathcal{N}$ uniformly at random to update their price.

For all sellers $j \neq i$, all M customers observe prices $p_{j,t}$.

Randomly sample $M/2$ customers to observe $p_{i,t}^+ = p_{i,t} + c_{\tau_i}/2$; the remaining $M/2$ customers observe $p_{i,t}^- = p_{i,t} - c_{\tau_i}/2$.

Simulate purchase decisions of all M customers, and estimate average revenue per customer at each of these prices, denoted by $\widehat{R}_i(p_{i,t}^+, \mathbf{p}_{-i,t})$, $\widehat{R}_i(p_{i,t}^-, \mathbf{p}_{-i,t})$.

Estimate partial gradient of seller i at $p_{i,t}$, i.e.,

$$\hat{g}_{i,t}(p_{i,t}, \mathbf{p}_{-i,t}) = \frac{\widehat{R}_i(p_{i,t}^+, \mathbf{p}_{-i,t}) - \widehat{R}_i(p_{i,t}^-, \mathbf{p}_{-i,t})}{c_{\tau_i}}.$$

Update price of seller i:

$$p_{i,t+1} = [p_{i,t} + a_{\tau_i} \hat{g}_{i,t}(p_{i,t}, \mathbf{p}_{-i,t})]_{\mathcal{V}}.$$

Increment update counter of seller i: $\tau_i \leftarrow \tau_i + 1$.

Algorithm 2 Distributed gradient ascent (GA)

Initialize: Convex set $\mathcal{V}$, time horizon T, learning rate schedule $(\eta_t)_{0 \leq t \leq T-1}$, initial prices $\mathbf{p}_0 \in \mathcal{V}^N$

Output: Final prices $\mathbf{p}_T$

for $t = 0, \ldots, T-1$ **do**

for $i \in \mathcal{N}$ **do**

if $p_{i,t} \neq p_{j,t}$ for all $j \neq i$ **then** let $f_i(p_{i,t}, \mathbf{p}_{-i,t}) = g_i(p_{i,t}, \mathbf{p}_{-i,t})$

else

Let $\mathcal{S} = \{j \in \mathcal{N} \setminus \{i\} : p_{i,t} = p_{j,t}\}$, and

$$f_i(p_{i,t}, \mathbf{p}_{-i,t}) = \begin{cases} \lim_{p_i \uparrow p_{i,t}} g_i(p_i, p_j = p_{j,t} \, \forall \, j \in \mathcal{S}, \mathbf{p}_{-\{\mathcal{S},i\},t}) & \text{with probability } 1/2 \\ \lim_{p_i \downarrow p_{i,t}} g_i(p_i, p_j = p_{j,t} \, \forall \, j \in \mathcal{S}, \mathbf{p}_{-\{\mathcal{S},i\},t}) & \text{with probability } 1/2. \end{cases} \tag{EC.18}$$

$p_{i,t+1} = [p_{i,t} + \eta_t f_i(p_{i,t}, \mathbf{p}_{-i,t})]_{\mathcal{V}}$.

- *If* $\mathbb{P}(\overline{\mathcal{N}}_c = \{1,2\})\overline{\lambda}_{-i\succ i} > 0$ *for all* $i \in \{1,2\}$ *and* $\bar{F}(p) > 0$ *for all* $p \in [0,\bar{v})$*, then Condition 4 holds.*

In words, the first sufficient condition for Condition 3 to hold is that, for each seller, either a strictly positive fraction of customers is loyal, or prefers them on an attribute ranked higher than price. The second condition is a technical condition on the willingness-to-pay distribution: it holds for uniform distributions, truncated exponential distributions, and distributions with linearly decreasing density.

The first sufficient condition for Condition 4 to hold states that, for each seller, a strictly positive mass of customers must prefer the competing seller based on an attribute ranked higher than price. This condition implies that, when the competitor is pricing at its optimum in a VOP, the focal seller benefits from undercutting in order to capture customers who cannot afford the competitor. The second sufficient condition is an extremely mild technical condition stating that the willingness-to-pay distribution is supported over the entire interval $[0,\bar{v})$.

Proof. Fix $i \in \mathcal{N}$, $\mathbf{p}_{-i} \in \mathcal{V}$. Moreover, for ease of notation, let $j = -i$. In order to derive $R_i(p,\mathbf{p}_{-i})$, we leverage the closed-form expression of $\mathcal{C}_i$ derived in the proof of Proposition EC.5. In particular, recall the following cases:

Case a: $p > p_j$. Then, by (EC.6), we have:

$$\begin{aligned} R_i(p,\mathbf{p}_{-i}) &= p\bar{F}(p)\big(\mathbb{P}(\overline{\mathcal{N}}_c = \{i\}) + \mathbb{P}(\overline{\mathcal{N}}_c = \{1,2\})\overline{\lambda}_{i\succ j}\big) \\ \implies g_i(p,\mathbf{p}_{-i}) &= \big(\mathbb{P}(\overline{\mathcal{N}}_c = \{i\}) + \mathbb{P}(\overline{\mathcal{N}}_c = \{1,2\})\overline{\lambda}_{i\succ j}\big)\big(\bar{F}(p) - pf(p)\big). \end{aligned} \tag{EC.19}$$

Case b: $p < p_j$. Then, by (EC.7), we have:

$$\begin{aligned} R_i(p,\mathbf{p}_{-i}) &= p\big(\bar{F}(p)\big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{i\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big) - \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\bar{F}(p_j)\overline{\lambda}_{j\succ i}\big) \\ \implies g_i(p,\mathbf{p}_{-i}) &= \bar{F}(p)\big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{i\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big) - \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\bar{F}(p_j)\overline{\lambda}_{j\succ i} \\ &\qquad - pf(p)\big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{i\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big) \\ &= \big(\mathbb{P}\big(\overline{\mathcal{N}}_c = \{i\}\big) + \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\big)\big(\bar{F}(p) - pf(p)\big) - \mathbb{P}\big(\overline{\mathcal{N}}_c = \{1,2\}\big)\bar{F}(p_j)\overline{\lambda}_{j\succ i}. \end{aligned} \tag{EC.20}$$

Case c: $p = p_j$. As described in the proof of Proposition EC.5, in this case it must be that $p = p_j = 0$. In this case, the seller will only ever evaluate the right partial-gradient (since she is at the boundary), which is the expression given in (EC.19).

In all of the above cases, Conditions 1 and 2 follow from Lipschitz continuity of the pdf and cdf, as well as boundedness of the space of feasible prices.

We now show that Condition 3 holds under the sufficient conditions described in the statement of the proposition. Specifically, we prove that $R_i(p,\mathbf{p}^{\mathcal{S}}_{-i})$ is strongly concave with respect to p, for all $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}, \mathbf{p}^{\mathcal{S}}_{-i}$. We again consider the three cases.

Case a: $p > p_j$ in $\mathbf{p}^{\mathcal{S}}_{-i}$.

$$\begin{aligned} g_i'(p, \mathbf{p}^{\mathcal{S}}_{-i}) &= \big(\mathbb{P}(\overline{\mathcal{N}}_c = \{i\}) + \mathbb{P}(\overline{\mathcal{N}}_c = \{1,2\})\overline{\lambda}_{i \succ j}\big)(-2f(p) - pf'(p)) \\ &\leq -c\big(\mathbb{P}(\overline{\mathcal{N}}_c = \{i\}) + \mathbb{P}(\overline{\mathcal{N}}_c = \{1,2\})\overline{\lambda}_{i \succ j}\big) \leq -K, \end{aligned}$$

for some $K > 0$.

Case b: $p < p_j$ in $\mathbf{p}^{\mathcal{S}}_{-i}$.

$$\begin{aligned} g_i'(p, \mathbf{p}^{\mathcal{S}}_{-i}) &= \big(\mathbb{P}(\overline{\mathcal{N}}_c = \{i\}) + \mathbb{P}(\overline{\mathcal{N}}_c = \{1,2\})\big)(-2f(p) - pf'(p)) \\ &\leq -c\big(\mathbb{P}(\overline{\mathcal{N}}_c = \{i\}) + \mathbb{P}(\overline{\mathcal{N}}_c = \{1,2\})\big) \leq -K, \end{aligned}$$

for some $K > 0$.

Case c: $p = p_j$ in $\mathbf{p}^{\mathcal{S}}_{-i}$. As in Case a.

We conclude by proving that Condition 4 holds under the condition described in the statement of the proposition. It suffices to consider the case where $p < p_j$, since Condition 4 describes settings where the seller is undercutting their competitor at a partial VOP.

Since seller j is pricing at its optimum, we have that $\lim_{p \uparrow p_j} \bar{F}(p) - pf(p) = 0$. Plugging this fact into (EC.20), we have: $\lim_{p \uparrow p_j} g_i(p, p_j) = -\mathbb{P}(\overline{\mathcal{N}}_c = \{1,2\})\bar{F}(p_j)\overline{\lambda}_{j \succ i} < 0$, by assumption. Hence, the supremum of the lower-priced seller is always attained in a partial VOP, and Condition 4 holds. □

Finally, to illustrate the idea that the settings in which Condition 4 fails to hold are relatively extreme, we provide the following instance.

PROPOSITION EC.10. *Consider a setting in which each seller has its own pool of loyal customers. Let γ_i denote the mass of loyal customers for seller $i \in \mathcal{N}$, with $\sum_{i \in \mathcal{N}} \gamma_i = 1$. Suppose moreover that all customers share the same willingness-to-pay distribution. In this case, Condition 4 fails to hold.*

Proof. Fix any ordering π. It suffices to consider the top two sellers in this ordering.

The first seller in the ordering sets the price $p_{\pi(1)}$ that maximizes $\gamma_{\pi(1)} p\bar{F}(p)$, since $\pi(1)$ can only sell to its loyal customers. Consider now the problem $\sup_{0 \leq p < p_{\pi(1)}} R_{\pi(2)}(p, \mathbf{p}^{\pi(1)}_{-\pi(2)}) = \sup_{0 \leq p < p_{\pi(1)}} \gamma_{\pi(2)} p\bar{F}(p)$, where the equality follows from the fact that $\pi(2)$ can only sell to its loyal customers. Notice that the supremum is not attained in this latter problem, since $\arg\max_{0 \leq p \leq p_{\pi(1)}} \gamma_{\pi(2)} p\bar{F}(p) = p_{\pi(1)}$. However, this fact also implies that $\lim_{p \uparrow p_{\pi(1)}} g_{\pi(2)}(p, \mathbf{p}^{\pi(1)}_{-\pi(2)}) = 0$. Therefore, Condition 4 fails to hold. □

Note that in the above instance, Condition 4 fails to hold not because customers are loyal, but because, conditional on the seller to whom customers are loyal, the willingness-to-pay distribution is *identical*. It is easy to see, given the proof of Proposition EC.10, that if the revenue functions associated with each loyal pool of customers have maximizers that differ by any $\epsilon > 0$, then Condition 4 does hold.

C.3.1. Proof of Lemma EC.5

Proof. By definition, $R_i(p, \mathbf{p}^{\mathcal{S}}_{-i})$ is strongly concave if and only if there exists $K > 0$ such that $g_i'(p, \mathbf{p}^{\mathcal{S}}_{-i}) \leq -K$ for all $p \in [0, p^{\min}(\mathcal{S}))$.

Case 1: $p_i^*(\mathbf{p}^{\mathcal{S}}_{-i}) \leq p^{\min}(\mathcal{S})$. In this case, we seek to show that $|g_i(p, \mathbf{p}^{\mathcal{S}}_{-i})| \geq K|p - p_i^*(\mathbf{p}^{\mathcal{S}}_{-i})|$. By the Mean Value Theorem (MVT), there exists $c \in \big(\min\{p, p_i^*(\mathbf{p}^{\mathcal{S}}_{-i})\}, \max\{p, p_i^*(\mathbf{p}^{\mathcal{S}}_{-i})\}\big)$ such that:

$$|g_i(p, \mathbf{p}^{\mathcal{S}}_{-i}) - g_i(p_i^*(\mathbf{p}^{\mathcal{S}}_{-i}), \mathbf{p}^{\mathcal{S}}_{-i})| = |g_i'(c, \mathbf{p}^{\mathcal{S}}_{-i})| \cdot |p - p_i^*(\mathbf{p}^{\mathcal{S}}_{-i})| \geq K|p - p_i^*(\mathbf{p}^{\mathcal{S}}_{-i})|. \tag{EC.21}$$

Noting that $g_i(p_i^*(\mathbf{p}^{\mathcal{S}}_{-i}), \mathbf{p}^{\mathcal{S}}_{-i}) = 0$ proves the claim.

Case 2: $p_i^*(\mathbf{p}^{\mathcal{S}}_{-i}) > p^{\min}(\mathcal{S})$. In this case, we seek to show that $|g_i(p, \mathbf{p}^{\mathcal{S}}_{-i})| \geq K|p - p^{\min}(\mathcal{S})|$. This however again follows from the MVT and strong concavity, since $p < p^{\min}(\mathcal{S}) < p_i^*(\mathbf{p}^{\mathcal{S}}_{-i}) \implies |p - p^{\min}(\mathcal{S})| < |p - p_i^*(\mathbf{p}^{\mathcal{S}}_{-i})|$. Applying this to (EC.21), we obtain the result. □

C.4. Proof of Theorem 3

We introduce some terminology that will be used throughout the proof. If $\mathbf{p}$ is aligned with an ordering π such that (EC.17) holds, we say that $\mathbf{p}$ is in a *right* configuration. If $\mathbf{p}$ is *strictly* ordered but (EC.17) does not hold, $\mathbf{p}$ is said to be in a *wrong* configuration. Finally, if there exist $i, j \in \mathcal{N}$ such that $p_i = p_j$, then $\mathbf{p}$ is in an *equal* configuration.

Theorem 3 follows from the following facts:

1. Assumption 2 and Assumption 3 imply that the LNE for which the prices are in a right configuration exists.
2. There exists t_1^* such that, if sellers' prices are in a right configuration at t_1^*, then they remain in a right configuration for all $t > t_1^*$ (Corollary EC.2).
3. If there exists t_2^* such that sellers' prices are in the same strictly aligned configuration for all $t > t_2^*$, then the LNE aligned with the corresponding ordering exists, and sellers' prices converge to this LNE (Lemma EC.8).

We put these facts together to prove Theorem 3.

Proof of Theorem 3. Existence of LNE follows from the fact that strong unimodality implies Assumption 1. Putting this together with Assumption 2 and applying Proposition 2, we obtain existence.

We now prove convergence to a LNE. Throughout the proof we assume that t is large enough that $p_{i,t+1} = p_{i,t} + \eta_t g_i(p_{i,t}, \mathbf{p}_{-i,t})$ for all $i \in \mathcal{N}$. This is without loss of generality, since by Lemma EC.6 below, the partial gradients of sellers' revenue functions point away from the boundaries of the decision space. We defer the straightforward proof of this fact to Appendix C.4.1.

LEMMA EC.6. *Under strong unimodality, the following holds:*

1. *if* $p^{\min}(\mathcal{S}) > 0$, *then* $g_i(0, \mathbf{p}^{\mathcal{S}}_{-i}) > 0$;
2. *if* $p^{\min}(\mathcal{S}) = \bar{v}$, *then* $\lim_{p \uparrow \bar{v}} g_i(p, \mathbf{p}^{\mathcal{S}}_{-i}) < 0$.

Consider $t > t_1^*$. We consider two cases.

Case 1: Sellers' prices are in a right configuration at time t. By Corollary EC.2, they remain in the right configuration for all $t' > t$. Then, by Lemma EC.8, the LNE for which sellers' prices are in this right configuration exists, and they converge to this LNE.

Case 2: Sellers' prices are in a wrong configuration at time t. In this case, either:

1. Sellers' prices remain in this wrong configuration for all $t' > t$. Then, by Lemma EC.8 the LNE corresponding to this wrong configuration exists, and the sellers' prices converge to this LNE.
2. There exist $t' > t$ and $i, j \in \mathcal{N}$ such that $p_{i,t'} < p_{j,t'}$ and $p_{i,t'+1} > p_{j,t'+1}$. In this sub-case, either:
   (a) $\mathbf{p}_{t'+1}$ is in a right configuration, in which case we have established that sellers' prices will remain in this right configuration and converge to the corresponding LNE.
   (b) $\mathbf{p}_{t'+1}$ is in another wrong configuration. Then, by Corollary EC.3, $p_{i,t''} > p_{j,t''}$ for all $t'' > t'$. As a result, sellers will never re-enter the wrong configuration they just exited.
3. There exist $t' > t$ and $i, j \in \mathcal{N}$ such that $p_{i,t'} = p_{j,t'} =: p$. Again, by Lemma EC.7 it must be that seller i gradient-dominates seller j. Due to the randomized gradient oracle at equality (see Algorithm 2), the following events occur with strictly positive probability:
   (i) sellers i and j observe different gradients,
   (ii) i observes $g_i(p_i = p, p_j = 0, \mathbf{p}^{\mathcal{S}}_{-\{i,j\},t})$, and j observes $\lim_{q \uparrow p} g_j(p_i = p, p_j = q, \mathbf{p}^{\mathcal{S}}_{-\{i,j\},t})$, for some $\mathcal{S} \subseteq \mathcal{N} \setminus \{i, j\}$.

   Thus, with probability 1 there exists $t'' > t'$ such that $p_{i,t''} \neq p_{j,t''}$, and one of three facts holds at t'':
   (a) sellers' prices are in a right configuration (Case 1),
   (b) sellers' prices are in a wrong configuration and remain in this wrong configuration (Case 2.1), or
   (c) sellers i and j become correctly ordered, and will remain correctly ordered from then on (Case 2.2(b)).

Case 3: Sellers' prices are in an equal configuration at time t. See Case 2.3.

Thus, in the worst case, sellers keep exiting and entering into successive wrong configurations. However, since there are finitely many configurations and sellers never re-enter a wrong configuration they have exited, it must be that sellers' prices eventually enter a right configuration and converge to the corresponding LNE.

We now show each of these key component results.

LEMMA EC.7. *There exists large enough $t_1^* \in \mathbb{N}$ such that, for all $i, j \in \mathcal{N}$, if sellers i and j are correctly ordered at time t_1^*, they remain correctly ordered for all $t \geq t_1^*$.*

Proof. Without loss of generality, suppose i gradient-dominates j. Let $c > 0$ be defined as follows:

$$c = \min\left\{ g_i(p_i = p, p_j = 0, \mathbf{p}^{\mathcal{S}}_{-\{i,j\}}) - \lim_{q \uparrow p} g_j(p_i = p, p_j = q, \mathbf{p}^{\mathcal{S}}_{-\{i,j\}}) \,\middle|\, \mathcal{S} \subseteq \mathcal{N} \setminus \{i, j\}, \mathbf{p} \in \mathcal{V}^N \right\}.$$

Fix $0 < \delta < c/(2L)$. We first show that there exists large enough t such that, if $|p_{i,t} - p_{j,t}| \geq \delta$, then $p_{i,t} > p_{j,t} \implies p_{i,t+1} > p_{j,t+1}$. We have:

$$\begin{aligned} p_{i,t+1} - p_{j,t+1} &= (p_{i,t} + \eta_t g_i(p_{i,t}, \mathbf{p}_{-i,t})) - (p_{j,t} + \eta_t g_j(p_{j,t}, \mathbf{p}_{-j,t})) \\ &= p_{i,t} - p_{j,t} + \eta_t (g_i(p_{i,t}, \mathbf{p}_{-i,t}) - g_j(p_{j,t}, \mathbf{p}_{-j,t})) \\ &\geq \delta - 2G\eta_t > 0 \end{aligned} \tag{EC.22}$$

for large enough t, since η_t is decreasing and $\lim_{t\to\infty} \eta_t = 0$ by square-summability of η_t.

We now argue that, if $|p_{i,t} - p_{j,t}| < \delta$, then $p_{i,t} > p_{j,t} \implies p_{i,t+1} > p_{j,t+1}$. By smoothness, for all $\epsilon > 0$ such that $(p_{i,t} - \epsilon, \mathbf{p}_{-j,t})$ is aligned with $\mathbf{p}_t$, $g_j(p_{j,t}, \mathbf{p}_{-j,t}) \leq g_j(p_{i,t} - \epsilon, \mathbf{p}_{-j,t}) + L|p_{j,t} - (p_{i,t} - \epsilon)|$. Using $|p_{i,t} - p_{j,t}| < \delta$ and taking the limit as ϵ approaches 0, we obtain $g_j(p_{j,t}, \mathbf{p}_{-j,t}) \leq \lim_{p \uparrow p_{i,t}} g_j(p, \mathbf{p}_{-j,t}) + L\delta$. Plugging this into (EC.22), we have:

$$\begin{aligned} p_{i,t+1} - p_{j,t+1} &\geq (p_{i,t} - p_{j,t}) + \eta_t \left(g_i(p_{i,t}, \mathbf{p}_{-i,t}) - \lim_{p \uparrow p_{i,t}} g_j(p_{i,t}, \mathbf{p}_{-j,t}) - L\delta \right) \\ &\geq (p_{i,t} - p_{j,t}) + \eta_t \cdot c/2 > p_{i,t} - p_{j,t} > 0. \end{aligned}$$

Taking $t_1^* = \min\{t \mid \delta - 2G\eta_t > 0 \,\forall\, i, j \in \mathcal{N}\}$, we obtain the claim. □

The following corollaries emerge as a direct consequence of Lemma EC.7.

COROLLARY EC.2. *If $\mathbf{p}_t$ is in a right configuration for some $t \geq t_1^*$, then, for all $t' > t$, $\mathbf{p}_{t'}$ remains in this configuration.*

COROLLARY EC.3. *Fix $t \geq t_1^*$. If sellers i and j were incorrectly ordered at time t but became correctly ordered at time $t+1$, then they remain correctly ordered for all $t' \geq t+1$.*

We now present our key lemma.

LEMMA EC.8. *Suppose there exists $t^* \in \mathbb{N}$ such that $\mathbf{p}_t$ is aligned with the same ordering π for all $t > t^*$. Then, there exists a valid order-price pair to which π belongs, and $\mathbf{p}_t$ converges to the corresponding LNE.*

Proof. Let $l^* = \sup\{l \mid \exists \text{ a valid order-price pair } (\widetilde{\pi}, \widetilde{\mathbf{p}}) \text{ such that } \widetilde{\pi}(l') = \pi(l') \,\forall\, l' \leq l\}$. If l^* fails to exist, we let $l^* = 1$. We prove that $\lim_{t\to\infty} p_{\pi(l'),t} = \widetilde{p}_{\pi(l')}$ for all $l' \leq l^*$ by induction.

Base case: $l' = 1$. This simply follows from the fact that for all $t > t^*$, Seller $\pi(1)$'s revenue is independent of $\mathbf{p}_{-\pi(1),t}$, and thus seller $\pi(1)$ is simply running gradient ascent over $R_{\pi(1)}(p, \mathbf{p}^{\phi}_{-\pi(1)})$, which under Assumption 3 converges to the unique maximizer $p^*_{\pi(1)}$ (Broadie et al. 2011).

Inductive step. Suppose $\lim_{t\to\infty} p_{\pi(l'),t} = \widetilde{p}_{\pi(l')}$ for all $l' \leq l$, with $l < l^*$. We first argue that $\sup_{p < p_{\pi(l)}} R_{\pi(l+1)}(p, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)})$ is attained. By continuity of the partial gradients, Condition 4 of Assumption 3 implies there exists $c > 0$ such that $\lim_{p \uparrow \widetilde{p}_{\pi(l)}} g_{\pi(l+1)}\left(p, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)}\right) \geq c$.

Let $\delta > 0$ be such that

$$\delta < \frac{1}{2}\min\left\{\frac{c}{2(1+clL/K+lL)}, \min_{l'<N}\{|\widetilde{p}_{\pi(l')} - \widetilde{p}_{\pi(l'+1)}|\}\right\}. \tag{EC.23}$$

Define $t_\delta > t^*$ such that $|\widetilde{p}_{\pi(l')} - p_{\pi(l'),t}| < \delta$ for all $t > t_\delta$, $l' \le l$. Consider $t \ge t_\delta$, and suppose $p_{\pi(l+1),t} < \widetilde{p}_{\pi(l)} - \epsilon$, for $\epsilon \ge \frac{c+lL\delta}{K}$. Since $\mathbf{p}_t$ and $\widetilde{\mathbf{p}}$ are both aligned with π for all $l' \le l$ and $t \ge t_\delta$, we have:

$$\begin{aligned} g_{\pi(l+1)}(p_{\pi(l+1),t}, \mathbf{p}^{\pi(1:l)}_{-\pi(l+1),t}) \ge g_{\pi(l+1)}(p_{\pi(l+1),t}, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)}) - l\cdot L\cdot\delta &\overset{(a)}{\ge} K|p_{\pi(l+1),t} - \widetilde{p}_{\pi(l)}| - l\cdot L\cdot\delta \\ &\ge K\epsilon - l\cdot L\cdot\delta \ge c, \end{aligned}$$

where (a) follows from strong unimodality.

Consider now the case where $p_{\pi(l+1),t} \ge \widetilde{p}_{\pi(l)} - \epsilon$, where $\epsilon < \frac{clL\delta}{K}$. Since $p_{\pi(l+1),t} < p_{\pi(l),t} < \widetilde{p}_{\pi(l)} + \delta$, we have $|p_{\pi(l+1),t} - \widetilde{p}_{\pi(l)}| \le \max\{\delta, \frac{clL\delta}{K}\}$. Then, by smoothness and continuity:

$$g_{\pi(l+1)}(p_{\pi(l+1),t}, \mathbf{p}^{\pi(1:l)}_{-\pi(l+1),t}) \ge \lim_{p\uparrow\widetilde{p}_{\pi(l)}} g_{\pi(l+1)}(p, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)}) - \delta\left(\max\{1, \frac{clL}{K}\} + l\cdot L\right) \ge c/2.$$

By non-summability of (η_t), then, there exists $t' > t_\delta$ such that $p_{\pi(l+1),t'} > \widetilde{p}_{\pi(l)} + \delta > p_{\pi(l),t'}$, a contradiction.

Having established that $\sup_{p<\widetilde{p}_{\pi(l)}} R_{\pi(l+1)}(p, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)})$ is attained, we now prove $\lim_{t\to\infty} p_{\pi(l+1),t} = \widetilde{p}_{\pi(l+1)}$. Let $d_t = (\widetilde{p}_{\pi(l+1)} - p_{\pi(l+1),t})^2$. We have:

$$d_{t+1} = d_t - 2(\widetilde{p}_{\pi(l+1)} - p_{\pi(l+1),t})(p_{\pi(l+1),t+1} - p_{\pi(l+1),t}) + (p_{\pi(l+1),t+1} - p_{\pi(l+1),t})^2$$

Plugging in the fact that $p_{\pi(l+1),t+1} = p_{\pi(l+1),t} + \eta_t g_{\pi(l+1)}(p_{\pi(l+1),t}, \mathbf{p}_{-\pi(l+1),t})$, we obtain:

$$\begin{aligned} d_{t+1} &= d_t - 2\eta_t g_{\pi(l+1)}(p_{\pi(l+1),t}, \mathbf{p}_{-\pi(l+1),t})(\widetilde{p}_{\pi(l+1)} - p_{\pi(l+1),t}) + \eta_t^2 g_{\pi(l+1)}(p_{\pi(l+1),t}, \mathbf{p}_{-\pi(l+1),t})^2 \\ &\le d_t - 2\eta_t g_{\pi(l+1)}(p_{\pi(l+1),t}, \mathbf{p}_{-\pi(l+1),t})(\widetilde{p}_{\pi(l+1)} - p_{\pi(l+1),t}) + \eta_t^2 G^2. \end{aligned} \tag{EC.24}$$

Fix $\epsilon > 0$, and suppose $d_t \ge \epsilon$. We show that, for large enough t, $d_{t+1} < d_t - \gamma_t$, with $\gamma_t > 0$ and $\lim_{T\to\infty}\sum_{t=1}^T \gamma_t = \infty$.

By the inductive hypothesis, $\lim_{t\to\infty} p_{\pi(l'),t} = \widetilde{p}_{\pi(l')}$ for all $l' \le l$. Let $\delta > 0$ be such that:

$$\delta < \frac{1}{2}\min\left\{\frac{K\epsilon}{lL\bar{v}}, \min_{l'<N}|\widetilde{p}_{\pi(l')} - \widetilde{p}_{\pi(l'+1)}|, -\frac{1}{(l+1)L}\lim_{p\uparrow\widetilde{p}_{\pi(l)}} g_{\pi(l+1)}(p, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)})\right\}.$$

(The fact that $\sup_{p<\widetilde{p}_{\pi(l)}} R_{\pi(l+1)}(p, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)})$ is attained implies $\lim_{p\uparrow\widetilde{p}_{\pi(l)}} g_{\pi(l+1)}(p, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)}) = -c$, for some $c > 0$, and thus δ necessarily exists.)

Let $t_\delta > t^*$ be such that, for all $t > t_\delta$: (i) $|p_{\pi(l'),t} - \widetilde{p}_{\pi(l')}| < \delta$ for all $l' \le l$, and (ii) $\eta_t < \delta/G$. (Such a t exists by the summability assumptions.)

Lemma EC.9 establishes that there exists $t' > t_\delta$ such that, for all $t > t'$, $p_{\pi(l+1),t} < \widetilde{p}_{\pi(l)}$. We formally state the lemma below, and defer its proof to the end of this section.

LEMMA EC.9. *Consider an ordering π and rank $l \in [N]$, and the price vector $\widetilde{\mathbf{p}}$ defined as follows:*

$$\begin{cases} \widetilde{p}_{\pi(1)} & := \arg\max_p R_{\pi(1)}(p, \mathbf{p}^{\phi}_{-\pi(1)}) \\ \widetilde{p}_{\pi(l')} & := \arg\max_{p<\widetilde{p}_{\pi(l'-1)}} R_{\pi(l')}(p, \widetilde{\mathbf{p}}^{\pi(1:l'-1)}_{-\pi(l')}) \quad \textit{for } l' \in \{2,\dots,l+1\}. \end{cases}$$

(That is, assume each of these local suprema is attained.) Let $\delta > 0$ be such that

$$\delta < \frac{1}{2}\min\left\{\min_{l'<l} |\widetilde{p}_{\pi(l')} - \widetilde{p}_{\pi(l'+1)}|, -\frac{1}{(l+1)L} \cdot \lim_{p\uparrow\widetilde{p}_{\pi(l)}} g_{\pi(l+1)}(p, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)})\right\}.$$

Suppose moreover that there exists $t_\delta \in \mathbb{N}$ such that, for all $t > t_\delta$, (i) $\mathbf{p}_t$ is aligned with π, (ii) $|p_{\pi(l'),t} - \widetilde{p}_{\pi(l')}| < \delta$ for $l' \le l$, and (iii) $\eta_t < \delta/G$. Then, there exists $t' > t_\delta$ such that $p_{\pi(l+1),t} < \widetilde{p}_{\pi(l)}$ for all $t > t'$.

Thus, by strong unimodality:

$$g_{\pi(l+1)}(p_{\pi(l+1),t}, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)})(\widetilde{p}_{\pi(l+1)} - p_{\pi(l+1),t}) \ge K(\widetilde{p}_{\pi(l+1)} - p_{\pi(l+1),t})^2 \ge K\epsilon.$$

Re-arranging, we obtain:

$$-g_{\pi(l+1)}(p_{\pi(l+1),t}, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)})(\widetilde{p}_{\pi(l+1)} - p_{\pi(l+1),t}) \le -K\epsilon. \tag{EC.25}$$

Since $\mathbf{p}_t$ and $\widetilde{\mathbf{p}}_t$ are aligned with the same ordering for $l' \le l$, and $p_{\pi(l+1),t} < \widetilde{p}_{\pi(l)}$, by smoothness we have that, for $t > t_\delta$:

$$g_{\pi(l+1)}(p_{\pi(l+1),t}, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)}) - lL\delta \le g_{\pi(l+1)}(p_{\pi(l+1),t}, \mathbf{p}^{\pi(1:l)}_{-\pi(l+1),t}) \le g_{\pi(l+1)}(p_{\pi(l+1),t}, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)}) + lL\delta.$$

Plugging this into (EC.25):

$$-g_{\pi(l+1)}(p_{\pi(l+1),t}, \mathbf{p}^{\pi(1:l)}_{-\pi(l+1),t})(\widetilde{p}_{\pi(l+1)} - p_{\pi(l+1),t}) \le -K\epsilon + lL\delta\bar{v} \le -K\epsilon/2.$$

Thus, we obtain $d_{t+1} \le d_t - \eta_t K\epsilon + \eta_t^2 G^2$.

Let $\gamma_t = \eta_t K\epsilon - \eta_t^2 G^2$. For large enough t, $\gamma_t > 0$, and moreover $\sum_t \gamma_t = \infty$ by the summability assumptions on (η_t). Let $t_\epsilon = \min\{t \mid \gamma_t > 0, t > t'\}$.

We complete the proof by showing that there exists t'_ϵ such that for all $t > t'_\epsilon$, if $d_t < \epsilon$, then $d_{t+1} < 2\epsilon$. We use the facts that $|g_{\pi(l+1)}(p_{\pi(l+1),t}, \mathbf{p}_{-\pi(l+1),t})| \le G$ and $|\widetilde{p}_{\pi(l+1)} - p_{\pi(l+1),t}| < \sqrt{\epsilon}$ and plug into (EC.24), to obtain $d_{t+1} < \epsilon + 2\eta_t G\sqrt{\epsilon} + \eta_t^2 G^2 < 2\epsilon$ for large enough t, since η_t is decreasing and $\lim_{t\to\infty} \eta_t = 0$ by the square-summability assumption.

Thus, for $t > \max\{t_\epsilon, t'_\epsilon, t_\delta\}$, as long as $d_t \ge \epsilon$, d_t strictly decreases by γ_t. Since $\sum_t \gamma_t = \infty$, there exists a point past which $d_t < \epsilon$. When this occurs, we have $d_{t+1} < 2\epsilon$. Taking ϵ to be arbitrarily small, we obtain convergence of $p_{\pi(l+1),t}$ to $\widetilde{p}_{\pi(l+1)}$. This completes the proof of the inductive step.

Rolling this argument out, we must have that $l^* = N$, and (i) the valid order-price pair to which π belongs exists, and (ii) $\lim_{t\to\infty} p_{\pi(l),t} = \widetilde{p}_{\pi(l)}$ for all $l \in [N]$. □

We conclude with a proof of Lemma EC.9.

Proof of Lemma EC.9. The proof is similar to that of Lemma EC.7. By assumption, $p_{\pi(l+1),t} < p_{\pi(l),t} < \widetilde{p}_{\pi(l)} + \delta$ for all $t > t_\delta$. We consider two cases.

Case 1: $p_{\pi(l+1),t} \le \widetilde{p}_{\pi(l)} - \delta$**.** As in the proof of Lemma EC.7, we have:

$$p_{\pi(l+1),t+1} = p_{\pi(l+1),t} + \eta_t g_{\pi(l+1)}(p_{\pi(l+1),t}, \mathbf{p}^{\pi(1:l)}_{-\pi(l+1),t}) \le p_{\pi(l+1),t} + \eta_t G \overset{(a)}{<} p_{\pi(l+1),t} + \delta < \widetilde{p}_{\pi(l)},$$

where (a) follows from the fact that $\eta_t < \delta/G$.

Case 2: $|p_{\pi(l+1),t} - \widetilde{p}_{\pi(l)}| < \delta$**.** We show that, for all $t > t_\delta$, $p_{\pi(l+1),t+1} \le p_{\pi(l+1),t} - \gamma_t$, where $\gamma_t > 0$ for all t and $\sum_t \gamma_t = \infty$. If this holds, there exists $t' > t_\delta$ such that $p_{\pi(l+1),t'} < \widetilde{p}_{\pi(l)} - \delta$, and we are back in Case 1. We have:

$$\begin{aligned} p_{\pi(l+1),t+1} &= p_{\pi(l+1),t} + \eta_t g_{\pi(l+1)}(p_{\pi(l+1),t}, \mathbf{p}^{\pi(1:l)}_{-\pi(l+1),t}) \\ &\le p_{\pi(l+1),t} + \eta_t \cdot \left(\lim_{p \uparrow \widetilde{p}_{\pi(l)}} g_{\pi(l+1)}(p, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)}) + (l+1)L\delta \right) \\ &\le p_{\pi(l+1),t} + \eta_t \cdot \frac{1}{2} \cdot \lim_{p \uparrow \widetilde{p}_{\pi(l)}} g_{\pi(l+1)}(p, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)}) \qquad \text{(EC.26)} \end{aligned}$$

where (EC.26) follows from the definition of δ. Taking $\gamma_t = -\frac{1}{2}\eta_t \lim_{p \uparrow \widetilde{p}_{\pi(l)}} g_{\pi(l+1)}(p, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)})$, we obtain $\gamma_t > 0$ since $\lim_{p \uparrow \widetilde{p}_{\pi(l)}} g_{\pi(l+1)}(p, \widetilde{\mathbf{p}}^{\pi(1:l)}_{-\pi(l+1)}) < 0$ by the assumption that the supremum is attained. Moreover, $\sum_t \gamma_t = \infty$, by non-summability of η_t. □

C.4.1. Proof of Lemma EC.6

Proof. Strong unimodality implies that $p_i^*(\mathbf{p}^{\mathcal{S}}_{-i})$ is unique. We first argue that $p_i^*(\mathbf{p}^{\mathcal{S}}_{-i}) \neq 0$. Suppose for contradiction that $p_i^*(\mathbf{p}^{\mathcal{S}}_{-i}) = 0$. This implies that the optimal revenue is 0. Since the revenue is non-negative, this means that i obtains zero revenue for all $p \in [0, p^{\min}(\mathcal{S}))$, contradicting strong unimodality. Thus $p_i^*(\mathbf{p}^{\mathcal{S}}_{-i}) > 0$, implying $\min\{p_i^*(\mathbf{p}^{\mathcal{S}}_{-i}), p^{\min}(\mathcal{S})\} > 0$. The strong unimodality condition then implies that $g_i(0, \mathbf{p}^{\mathcal{S}}_{-i}) > 0$, thus showing the first claim.

We next turn to the second claim. We have already argued that $p_i^*(\mathbf{p}^{S}_{-i}) \neq 0$. We now argue that $\bar{v} \notin \arg\max_{p \in [0,\bar{v}]} R_i(p, \mathbf{p}^{S}_{-i})$ if $p^{\min}(\mathcal{S}) = \bar{v}$. Suppose again for contradiction that $\bar{v} \in \arg\max_{p \in [0,\bar{v}]} R_i(p, \mathbf{p}^{S}_{-i})$. Similarly, since $R_i(\bar{v}, \mathbf{p}^{S}_{-i}) = 0$, this implies that seller i obtains zero revenue for all $p \in [0, p^{\min}(\mathcal{S}))$, contradicting the strong unimodality of $R_i(p, \mathbf{p}^{\phi}_{-i})$ over $[0, \bar{v})$. Thus, $p_i^*(\mathbf{p}^{\phi}_{-i}) \in (0, \bar{v})$. Since $g_i(p_i^*(\mathbf{p}^{\phi}_{-i}), \mathbf{p}^{\phi}_{-i}) = 0$, strong unimodality again implies the second result. □

C.5. Empirical analysis of gradient-driven price-update dynamics

Motivated by the fact that price and rating are two of the main attributes used for purchase decisions on e-commerce platforms, we consider price competition amongst $N \in \{2, \cdots, 5\}$ sellers who differ on these two attributes. Given N, we assume that the quality of seller $i \in [N]$ is $q_i = i/N$.

Choice model. We assume that customers are either loyal or non-loyal. Loyal customers purchase from the seller to whom they are loyal as long as its price is at most their willingness-to-pay. Non-loyal customers are one of two types: quality-sensitive (QS) or price-sensitive (PS). QS and PS customers both form their consideration set based on a maximum willingness-to-pay and a minimum rating floor. However, they differ in the lexicographic choice phase: QS customers choose the highest-rated item in their consideration set, breaking ties based on price; PS customers, on the other hand, choose the cheapest item in their consideration set, breaking ties based on rating. QS and PS customers respectively generalize the BICA and CICT customers presented in Section 3.

We consider three settings for the distribution over these model primitives:

- **Setting A:** Customers are loyal, PS, or QS with an equal probability of $1/3$. A loyal customer is loyal to seller i with probability $1/N$. The willingness-to-pay and rating floor for each customer (where applicable) are drawn independently from a Beta(α, β) distribution.
- **Setting B:** A customer is loyal with probability $1/3$ and non-loyal with probability $2/3$. A loyal customer is loyal to seller i with probability $2i/(N(N+1))$, capturing the fact that higher-rated sellers have more loyal customers. The willingness-to-pay of a loyal customer is drawn from a Beta(α, β) distribution. Conditional on being non-loyal, the willingness-to-pay, rating floor, and probability of the customer being QS are all equal, and drawn from a Beta(α, β) distribution. This captures the reality that willingness-to-pay and quality sensitivity for non-loyal customers may be correlated, i.e., a higher willingness-to-pay may indicate a higher quality floor and a higher propensity to be quality-sensitive.
- **Setting C:** Same as Setting B, except that there are no loyal customers: all customers are non-loyal with probability 1, with the same parameters.

Price dynamics. Sellers deploy distributed gradient-driven pricing schemes over a fixed time horizon $T = 10^4 \cdot N$. At the beginning of the horizon, we generate an initial set of prices independently and uniformly in $[0, 1]$. To model the reality that sellers make price updates asynchronously, we choose a seller uniformly at random in each period to price experiment and subsequently make a price update. We assume that sellers implement the classical Kiefer-Wolfowitz algorithm for function optimization with zero-order feedback (Kiefer and Wolfowitz 1952). Specifically, the chosen seller in each timestep observes a batch of 10^3 customers, and randomly sets one of two prices for each customer in the batch: one above and one below its current incumbent price. The seller then estimates the gradient given the average revenue at each price and updates its incumbent price by taking a gradient-scaled step. Letting c_τ and a_τ respectively denote the finite difference width for gradient estimation and stepsize after τ updates, we let $c_\tau = 1/\log \tau$ and $a_\tau = 1/\tau$.

Observe that if the price trajectories converge to distinct prices under such a gradient-based pricing scheme, then the limiting price vector must be a LNE. We say that a price trajectory has converged if (i) for each seller, the difference between the maximum and minimum prices set in the last 10^3 updates is within

(a) Setting A				(b) Setting B				(c) Setting C				(d) Setting D				
(α, β)	N	C	OC	(α, β)	N	C	OC	(α, β)	N	C	OC	(α, β)	N	C	OC	E
(1, 1)	2	100	100	(1, 1)	2	99	100	(1, 1)	2	99	100	(1, 1)	2	99	100	97
(1, 1)	3	97	100	(1, 1)	3	96	100	(1, 1)	3	98	100	(1, 1)	3	97	100	96
(1, 1)	4	97	100	(1, 1)	4	98	100	(1, 1)	4	98	100	(1, 1)	4	99	100	97
(1, 1)	5	94	99	(1, 1)	5	97	100	(1, 1)	5	96	99	(1, 1)	5	99	100	95
(0.5, 0.5)	2	100	100	(0.5, 0.5)	2	100	100	(0.5, 0.5)	2	98	100	(0.5, 0.5)	2	100	100	91
(0.5, 0.5)	3	98	100	(0.5, 0.5)	3	99	99	(0.5, 0.5)	3	100	100	(0.5, 0.5)	3	98	100	96
(0.5, 0.5)	4	93	100	(0.5, 0.5)	4	98	100	(0.5, 0.5)	4	96	100	(0.5, 0.5)	4	98	100	95
(0.5, 0.5)	5	93	100	(0.5, 0.5)	5	95	99	(0.5, 0.5)	5	96	99	(0.5, 0.5)	5	96	100	93
(2, 2)	2	100	100	(2, 2)	2	100	100	(2, 2)	2	100	100	(2, 2)	2	100	100	93
(2, 2)	3	98	100	(2, 2)	3	99	100	(2, 2)	3	97	100	(2, 2)	3	97	99	93
(2, 2)	4	91	100	(2, 2)	4	98	100	(2, 2)	4	93	100	(2, 2)	4	99	100	90
(2, 2)	5	91	100	(2, 2)	5	95	99	(2, 2)	5	97	100	(2, 2)	5	95	100	94
(2, 4)	2	99	100	(2, 4)	2	100	100	(2, 4)	2	100	100	(2, 4)	2	100	100	95
(2, 4)	3	95	99	(2, 4)	3	100	100	(2, 4)	3	97	100	(2, 4)	3	99	100	95
(2, 4)	4	91	100	(2, 4)	4	99	99	(2, 4)	4	96	99	(2, 4)	4	95	99	93
(2, 4)	5	80	100	(2, 4)	5	92	99	(2, 4)	5	93	99	(2, 4)	5	91	99	87
(4, 2)	2	98	100	(4, 2)	2	100	100	(4, 2)	2	99	100	(4, 2)	2	99	100	96
(4, 2)	3	90	100	(4, 2)	3	97	100	(4, 2)	3	96	100	(4, 2)	3	96	99	93
(4, 2)	4	84	100	(4, 2)	4	95	99	(4, 2)	4	91	100	(4, 2)	4	88	99	85
(4, 2)	5	72	100	(4, 2)	5	89	99	(4, 2)	5	83	100	(4, 2)	5	85	100	78

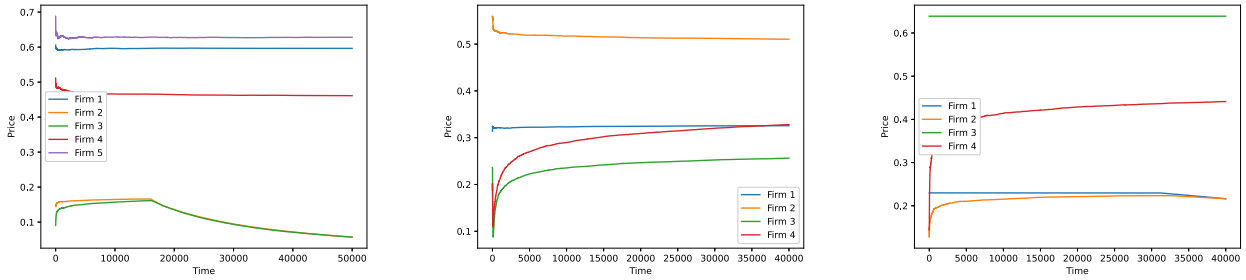
Table EC.1 Number of convergent (C) and order-convergent (OC) trajectories out of 100 random initializations in each setting

1% of the empirical average of the past 10^3 prices, and (ii) the ordering of all prices remains unchanged in the last 10^3 periods. If only (ii) is satisfied after T periods, we say that the trajectory is *order-convergent*. When the price trajectory is only order-convergent, we expect prices to eventually converge due to the pseudo-competitive property, by the same arguments as those provided in the proof sketch of Theorem 3.

Results. Table EC.1 shows the number of replications out of 100 random draws of initial prices that result in price and order convergence for Settings A-C. (Setting D shows a final set of robustness results, discussed at the end of the section.) We observe that the price dynamics converge in the majority of settings, with at least 72 convergent trajectories out of 100 in the worst case over all parameters in Setting A, 89 in Setting B, and 83 in Setting C. If we consider order-convergent trajectories, the numbers further improve, with at least 99 order-convergent trajectories in the worst case across all settings. Overall, our results suggest that LNE are frequently learnable via distributed gradient-driven pricing, under CLC models of practical interest.

We include sample non-order-convergent price trajectories in Figure EC.1. Two of the trajectories show sellers entering undercutting price wars as their prices approach each other (sellers 2-3 in Figure EC.1a, and sellers 1-2 in Figure EC.1c). These undercutting price wars are due to the existence of price-sensitive customers who choose the cheapest item in their consideration set. Indeed, if all customers are price-sensitive, the only equilibrium is one in which all sellers price at zero. Again, these results underscore the idea that

convergence arises in part due to the existence of quality-sensitive and loyal customers, who counteract the incentive to enter a price war.



(a) Setting A: $\alpha = 0.5$, $\beta = 0.5$ (b) Setting B: $\alpha = 1$, $\beta = 1$ (c) Setting C: $\alpha = 1$, $\beta = 1$

Figure EC.1 Sample non-convergent trajectories in Settings A-C

Robustness to non-CLC choice. We conclude our numerical experiments by investigating the robustness of these observations to non-CLC customers. Specifically, we modify Setting C such that 10% of customers behave according to the Multinomial Logit (MNL) choice model, with the remainder behaving as in Setting C. For an MNL customer, the utility of choosing a seller i with rating q and price p is assumed to be $q - p + \epsilon$, where ϵ is an independent noise term with a Gumbel distribution with location parameter 0 and scale parameter 1. The customer chooses the seller with the highest utility. We refer to this setting as Setting D. We choose the same price initializations for each of the 100 trajectories for both Settings C and D, and study the impact of MNL customers both on convergence and on the limit points themselves.

Convergence results for Setting D are shown in Table EC.1. We obtain consistent convergence to a LNE (in at least 85 random draws, in the worst case), as well as order-convergence (in at least 99 draws, in the worst case). The column labeled "E" for Setting D shows the number of convergent trajectories such that (i) the corresponding trajectories in Setting C converge as well (we call these *co-convergent* trajectories), and (ii) the convergence point is within a tolerance of 10% of the limit point under Setting C; this tolerance aligns with the fact that 10% of the population is of the MNL type. We observe that trajectories under Settings C and D are almost always co-convergent, and that these convergence points are frequently close, with both conditions holding in at least 78 random draws, in the worst case. Overall, these observations suggest that our equilibrium insights are robust to the limited inclusion of well-studied utility-maximizing behaviors.

Appendix D: Section 6 Omitted Proofs

D.1. Proof of Proposition 3

Proof. The coarse LNE follows from the fact that the coarse setting is equivalent to a Bertrand game of price competition.

We now consider the granular LNE. By Theorem 2, it suffices to study two possible VOP, corresponding to each seller being first in the ordering. Note that any price configuration for which $p_2 \geq p_1$ results in seller 2 making zero revenue. Therefore, this cannot be a non-trivial LNE. Therefore, it suffices to consider the case where $p_1 > p_2$. In this case, we have $R_1(p_1, p_2) = \lambda p_1(1-p_1)$, whose maximum is attained at $p_1^* = 1/2$. Moreover:

$$R_2(p_1^*, p_2) = p_2(\lambda(p_1^* - p_2) + (1-\lambda)(1-p_2)) \implies g_2(p_1^*, p_2) = \lambda(p_1^* - p_2) + (1-\lambda)(1-p_2) - p_2.$$

Solving the first-order condition, we obtain $p_2^* = \frac{1}{2} - \frac{\lambda}{4} < p_1^*$. □

D.2. Proof of Proposition 4

Proof. Revenue dominance of the granular LNE trivially follows from the fact that both sellers make zero revenue in the coarse LNE, as compared to strictly positive revenue in the granular LNE. Similarly, demand dominance of the coarse LNE trivially follows from the fact that all sellers' prices have increased. □

D.3. Proof of Proposition 5

Proof. The first part of the proposition follows immediately from Proposition 3.

For the second part of the proposition, fix $\lambda' \in (0,1)$. We have $R_1(p_1^*, p_2^*) = \lambda' \cdot p_1^*(1-p_1^*) = \frac{\lambda'}{4}$, and

$$\begin{aligned} R_2(p_1^*, p_2^*) = p_2^*(\lambda'(p_1^* - p_2^*) + (1-\lambda')(1-p_2^*)) &= \left(\frac{1}{2} - \frac{\lambda'}{4}\right)\left(\lambda' \cdot \frac{\lambda'}{4} + (1-\lambda')\left(\frac{1}{2} + \frac{\lambda'}{4}\right)\right) \\ &= \left(\frac{1}{2} - \frac{\lambda'}{4}\right)^2. \end{aligned}$$

Summing these two, we obtain $\Pi^g = \frac{\lambda'}{4} + \left(\frac{1}{2} - \frac{\lambda'}{4}\right)^2$, which is increasing in λ'. Therefore, the Ratings-First default ranking, which is such that $\lambda' = \lambda + \delta$, yields higher total revenue than the Price-First default ranking, which is such that $\lambda' = \lambda - \delta$.

We now analyze total demand on the platform. Since all customers for whom $w_c \geq p_2^*$ are served, we have $\mathcal{D}^g = 1 - p_2^* = \frac{1}{2} + \frac{\lambda'}{4}$, which is similarly increasing in λ'. Hence, the Ratings-First default ranking dominates the Price-First default ranking for this metric as well. □

D.4. Proof of Proposition 6

Proof. We first derive the Ratings-First LNE. Again, by Theorem 2, it suffices to consider the two VOPs corresponding to each seller pricing first in the ordering.

Suppose $p_1 > p_2$. In this case, seller 1 only sells to ratings-first customers with $w_c \geq p_1$. Then: $R_1(p_1, p_2) = \alpha_1 p_1(1-p_1)$, whose maximum is achieved at $p_1^* = 1/2$. Seller 2 sells to ratings-first customers for whom $w_c \in [p_2, p_1)$, and certification-first and price-first customers for whom $w_c \geq p_2$. Therefore:

$$R_2(p_1^*, p_2) = p_2\left(\alpha_1\left(\frac{1}{2} - p_2\right) + (1-\alpha_1)(1-p_2)\right) \implies g_2(p_1^*, p_2) = \alpha_1\left(\frac{1}{2} - p_2\right) + (1-\alpha_1)(1-p_2) - p_2. \tag{EC.27}$$

Solving the first-order condition, we obtain $p_2^* = \frac{1}{2} - \frac{\alpha_1}{4} < p_1^*$ for all $\alpha_1 \in (0,1)$.

The Certification-First LNE follows by symmetry, since seller 2 only sells to certification-first customers for whom $w_c \geq p_2$, and seller 1 sells to certification-first customers for whom $w_c \in [p_1, p_2)$ and ratings-first and price-first customers for whom $w_c \geq p_1$. □

D.5. Proof of Proposition 7

Proof. We analyze the two LNE separately.

Ratings-First LNE: $p_1^* = \frac{1}{2}$, $p_2^* = \frac{1}{2} - \frac{\alpha_1}{4}$. In this case, ratings-first customers choose seller 1 if $w_c \geq p_1^*$, and seller 2 if $w_c \in [p_2^*, p_1^*)$; certification-first and price-first customers choose seller 2 if $w_c \geq p_2^*$.

We first analyze the fraction of demand served in this LNE. By the above reasoning, we have $\mathcal{D}^{rf} = \alpha_1(1-p_2^*) + (\alpha_2+\gamma)(1-p_2^*) = 1 - p_2^* = \frac{1}{2} + \frac{\alpha_1}{4}$, which is increasing in α_1. Since $\alpha_1 \leq \lambda$, and $\mathcal{D}^g = \frac{1}{2} + \frac{\lambda}{4}$, we have that $\mathcal{D}^{rf} \leq \mathcal{D}^g$.

We now analyze seller revenues under this LNE. Noting that this LNE reduces to the setting with α_1 mass of ratings-first customers and $1-\alpha_1$ mass of price-first customers, we have: $\mathcal{R}_1^{rf} = \frac{\alpha_1}{4}$, $\mathcal{R}_2^{rf} = \left(\frac{1}{2} - \frac{\alpha_1}{4}\right)^2$. Then, the results follow from the fact that $\mathcal{R}_1^{rf}$ is increasing in α_1 and $\mathcal{R}_2^{rf}$ is decreasing in α_1, with $\alpha_1 \leq \lambda$.

We conclude by proving the bounds on ρ^{rf}. Using the closed-form expressions for $\mathcal{R}_1$ and $\mathcal{R}_2$ in each equilibrium, we have $\Pi^{rf} = \frac{\alpha_1^2}{16} + \frac{1}{4}$, and $\Pi^g = \frac{\lambda^2}{16} + \frac{1}{4}$. Hence, $\rho^{rf} \leq 1$ follows from the fact that $\alpha_1 \leq \lambda$. To prove the lower bounds, notice that, in the worst case, $\alpha_1 = 0$ and $\lambda = 1$, in which case $\rho^{rf} = \frac{1/4}{5/16} = \frac{4}{5}$. Additionally, $\rho^{rf} \geq \left(\frac{\alpha_1}{\lambda}\right)^2$ follows from algebra, and the fact that $\alpha_1 \leq \lambda$.

Certification-First LNE: $p_2^* = \frac{1}{2}$, $p_1^* = \frac{1}{2} - \frac{\alpha_2}{4}$. In this case, ratings-first and price-first customers choose seller 1 if $w_c \geq p_1^*$; certification-first customers choose seller 2 if $w_c \geq p_2^*$, and seller 1 if $w_c \in [p_1^*, p_2^*)$.

We first study the fraction of demand served in this LNE. The arguments are symmetric to those of the Ratings-First LNE. Namely, we have: $\mathcal{D}^{cf} = 1 - p_1^* = \frac{1}{2} + \frac{\alpha_2}{4}$, which is increasing in α_2. Therefore, $\mathcal{D}^{cf} \geq \mathcal{D}^g \iff \alpha_2 \geq \lambda$.

For seller revenues, by symmetry, $\mathcal{R}_2^{cf} = \frac{\alpha_2}{4}$, $\mathcal{R}_1^{cf} = \left(\frac{1}{2} - \frac{\alpha_2}{4}\right)^2$, $\Pi^{cf} = \frac{\alpha_2}{4} + \left(\frac{1}{2} - \frac{\alpha_2}{4}\right)^2$. Therefore:

$$\mathcal{R}_1^{cf} \geq \mathcal{R}_1^g \iff \left(\frac{1}{2} - \frac{\alpha_2}{4}\right)^2 \geq \frac{\lambda}{4} \iff \alpha_2 \leq 2 - 2\sqrt{\lambda},$$

and

$$\mathcal{R}_2^{cf} \geq \mathcal{R}_2^g \iff \frac{\alpha_2}{4} \geq \left(\frac{1}{2} - \frac{\lambda}{4}\right)^2 \iff \alpha_2 \geq 4\left(\frac{1}{2} - \frac{\lambda}{4}\right)^2.$$

Finally, $\Pi^{cf} = \frac{\alpha_2}{4} + \left(\frac{1}{2} - \frac{\alpha_2}{4}\right)^2$. Using the fact that Π^{cf} is increasing in α_2, and $\Pi^g = \frac{\lambda}{4} + \left(\frac{1}{2} - \frac{\lambda}{4}\right)^2$, we obtain $\Pi^{cf} \geq \Pi^g \iff \alpha_2 \geq \lambda$.

If $\alpha_2 \leq \lambda$, the same arguments as those used for ρ^{rf} apply to ρ^{cf}. If $\alpha_2 > \lambda$, on the other hand, the best case is when $\alpha_2 = 1$ and $\lambda = 0$, which gives us $\rho^{cf} \leq \frac{5}{4}$. Moreover, simple algebra yields $\rho^{cf} \leq \left(\frac{\alpha_2}{\lambda}\right)^2$. □