

Platform Design with Lemons: Ranking, Certification, and Endogenous Market Segmentation

Wenxiao Yang^{*} and Sihan Zhai[†]

^{*}University of California, Berkeley

[†]Harvard Business School

Abstract

Ranking algorithms are critical platform design mechanisms that determine seller visibility. However, existing research on platform design overlooks two key aspects: sellers' strategic signaling responses, especially the role of prices as quality signals, and how certification programs create seller heterogeneity. To address these gaps, we develop a theoretical framework examining how ranking design—specifically quality-based versus price-based—interacts with both strategic signaling and certification heterogeneity to shape market outcomes. Our analysis reveals two mechanisms. First, quality-based ranking creates a crowding-out effect: while revealing platform-held information, it inadvertently induces low-quality uncertified sellers to mimic high-quality uncertified sellers' pricing, creating pooling equilibria where prices lose informational value. Price-based ranking reverses this dynamic as low-quality uncertified sellers with cost advantages undercut competitors, generating separating equilibria where prices become informative. Second, compared to quality-based ranking, price-based ranking elevates these low-quality uncertified sellers to top positions, widening the perceived quality gap between certified and top-ranked uncertified sellers, and thus reducing competition between them. This creates endogenous market segmentation: certified sellers capture quality-sensitive consumers through premium pricing, while uncertified sellers expand market access to price-sensitive segments through competitive pricing. Empirical analysis of eBay's 2011 ranking algorithm redesign validates these predictions: after the redesign from quality-based to price-based ranking, prices of uncertified sellers declined due to intensified price competition, while prices of certified sellers slightly increased; consumers relied more on certifications and prices as quality signals; and high-quality uncertified sellers experienced significant welfare losses, while low-quality uncertified sellers and certified sellers potentially benefited.

Keywords: platform design, quality certificate, ranking algorithms, signaling

Email: ywenxiao@berkeley.edu; szhai@hbs.edu. Authors are ordered alphabetically. We thank J. Miguel Villas-Boas, Andrew Ching, Przemyslaw Jeziorski, and Song Lin for helpful comments and suggestions. We thank Michael Dinerstein for explaining the data set in detail to us. All errors are our own.

1 Introduction

Digital marketplaces facilitate a large share of retail transactions, with billions of products sold annually through platforms such as eBay and Amazon. These platforms employ two critical and interconnected mechanisms: ranking algorithms that determine the order in which products are displayed and certification programs that disclose the quality of sellers to consumers. Ranking algorithms profoundly influence consumer search behavior and seller strategic decisions, while certification programs serve dual functions: they reveal the quality of sellers to help consumers distinguish high-quality sellers from low-quality ones under asymmetric information, and they provide enhanced visibility that grabs consumer attention regardless of ranking positions. Major e-commerce platforms employ ranking algorithms (e.g., A9) that incorporate price and quality as critical determinants while simultaneously differentiating high-quality sellers through certification programs like Amazon's Choice and Seller Certification Program ([Amazon, 2024](#)). Similarly, eBay's Top-Rated Seller certification not only enhances search visibility and offers exclusive protections, but also increases transaction trust by reducing information asymmetry between buyers and sellers ([eBay, 2025](#)). Other examples include Airbnb's Superhost program and Uber's Pro designation. The proliferation of quality certification programs across digital platforms demonstrates the central role of these mechanisms in addressing the fundamental challenge of quality uncertainty in online marketplaces where physical inspection is impossible.

Despite their central roles and co-existence in digital marketplaces, ranking algorithms and certification programs have largely been studied in isolation. The literature on ranking algorithms has predominantly emphasized demand-side considerations, such as consumer search costs and attention allocation ([Ghose, Ipeirotis, and Li, 2012, 2014; Ursu, 2018; Ursu and Dzyabura, 2020; Compiani et al., 2024](#)), while research on certification programs has focused on reputation systems and buyer-seller matching ([Hui et al., 2016; Saeedi, 2019; Hui et al., 2023](#)). Even when both mechanisms are present, empirical studies typically pool certified and uncertified sellers together without examining their heterogeneous responses to changes in platform design. However, this heterogeneity is important. Our analysis of data from [Dinerstein et al. \(2018\)](#) uncovers that certified and uncertified sellers on eBay exhibit systematically different pricing patterns and responses to algorithmic changes, yet existing research does not account for these differential effects when

evaluating platform design.

The existing literature also overlooks a crucial supply-side phenomenon arising from information asymmetry in online marketplaces. When consumers cannot physically inspect products, they rely heavily on observable signals to infer quality, with pricing serving as a primary quality revelation mechanism for sellers who lack established reputations or other credible quality indicators. Research has extensively documented that consumers often interpret higher prices as signals of higher quality in the context of information asymmetry (Milgrom and Roberts, 1986; Rao and Monroe, 1989; Erdem, Keane, and Sun, 2008). However, existing research typically assumes that sellers' signaling behaviors are exogenous to platform design, neglecting how sellers strategically adjust their signaling responses when platform design changes. This oversight is particularly problematic because sellers' endogenous signaling behaviors can counteract platform design objectives, potentially reducing information transmission when platforms intend to enhance it.

To address these gaps, we develop a theoretical model of seller competition on a platform featuring vertically differentiated products where uncertified sellers possess private quality information while certified sellers benefit from credible quality signals and enhanced visibility provided by certifications. Our model captures how ranking algorithms trigger two mechanisms: changes in seller signaling incentives, and changes in market structure through the interaction between ranking and certification programs.

Our theoretical analysis examines how ranking design—specifically quality-based versus price-based—interacts with both strategic signaling and certification heterogeneity to shape market outcomes in vertically differentiated markets. Our analysis reveals two mechanisms that fundamentally shape market structure. The first operates through a counterintuitive crowding-out effect of ranking on seller signaling. Under quality-based ranking, the platform's visibility allocation based on quality assessments creates stronger incentives for low-quality sellers to mimic high-quality sellers' pricing strategies to gain top positions, leading to pooling equilibria where prices lose their informational value. Conversely, under price-based ranking, low-cost sellers gain systematic advantages by undercutting competitors, creating separating equilibria where prices become informative quality indicators. This demonstrates that quality-based ranking algorithms, designed to reveal the platform's quality information to consumers, can paradoxically reduce market information transmission by crowding out quality information from sellers' signaling.

The second mechanism operates through the interaction between ranking and certification. Compared to quality-based ranking, price-based ranking intensifies price competition among uncertified sellers, systematically elevating low-quality sellers with cost advantages to top positions, as they can profitably undercut high-quality competitors. This elevation widens the perceived quality gap between certified sellers and top-ranked uncertified sellers, reducing direct competition between these two groups and potentially enabling certified sellers to charge higher prices. This dynamic creates endogenous market segmentation: certified sellers capture quality-sensitive consumers through premium pricing strategies, while top-ranked uncertified sellers serve price-sensitive segments through competitive pricing. The resulting welfare redistribution systematically disadvantages high-quality uncertified sellers, who forfeit quality-based visibility advantages and are unable to compete on price, while potentially benefiting low-quality uncertified sellers with cost advantages who gain consistent top positioning and certified sellers who face reduced competitive pressure.

We validate these theoretical predictions using eBay's 2011 ranking algorithm redesign, where the platform transitioned from quality-based ranking to price-based ranking while holding certification programs unchanged. Our results on market responses to the redesign strongly support our theory: structural estimation reveals that consumers relied more on price and certifications as quality signals (validating the crowding-out effect through the signaling mechanism); data demonstrates that prices of uncertified sellers decreased, while prices of certified sellers slightly increased (validating the importance of certifications in competition landscapes); and we find empirical evidence for welfare redistribution and market expansion due to market segmentation (validating the certification mechanism).

The remainder of this paper is organized as follows. Section 2 reviews the related literature; Section 3 presents our theoretical model and develops predictions; Section 4 discusses the empirical context and model-free evidence for empirical predictions; Section 5 deploys a structural model to provide further empirical evidence for predictions; and Section 6 concludes with implications for business and society.

2 Literature Review

Our work intersects several streams of literature by examining how algorithmic design interacts with information asymmetry. We contribute to the literature on the economic implications of algorithms, which has examined their impact on minority welfare (Lambrecht and Tucker, 2019; Zhang et al., 2021), market collusion (Miklos-Thal and Tucker, 2019; Calvano et al., 2020), and political polarization (Levy, 2021). However, this literature has largely overlooked how algorithmic design changes participants' signaling strategies under information asymmetry, which endogenously affects the information available to algorithms. Within the literature on ranking algorithms, researchers have focused primarily on demand-side considerations (Ghose, Ipeirotis, and Li, 2012, 2014; Ursu, 2018; Ursu and Dzyabura, 2020; Compiani et al., 2024), typically treating product quality as exogenous or observable to algorithm designers. This approach neglects how algorithmic design choices alter sellers' signaling strategies when they possess private information. While Dinerstein et al. (2018) demonstrate that price-based ranking intensifies seller competition and Derakhshan et al. (2022) propose welfare-optimized ranking algorithms under complete information, neither study examines how ranking algorithms interact with quality signaling under information asymmetry. This gap is particularly significant because consumers' price-quality perceptions fundamentally shape sellers' signaling strategies, and evidence shows that ranking algorithms substantially influence these perceptions (Rao and Monroe, 1989; Erdem, Keane, and Sun, 2008).

We also contribute to the extensive literature on information disclosures and signaling mechanisms (Spence, 1978; Milgrom and Roberts, 1986; Bagwell and Riordan, 1991) by examining their behavioral and market effects. We summarize three branches in this stream of literature. First, researchers have examined the impact of information on behaviors of sellers and consumers. Research demonstrates that online reviews influence both consumer demand (Chevalier and Mayzlin, 2006; Sun, 2012) and seller quality investment (Farronato and Zervas, 2022; Jin and Leslie, 2003; Shi, Srinivasan, and Zhang, 2023). Jerath and Ning (2025) examine how signaling commitments to product inclusion affect consumer beliefs and demand patterns. Miklos-Thal and Zhang (2013) exploits the impact of information on consumer quality inferences to make profits. Second, studies of the effects on market competition are emerging, including Feng et al. (2024) and Armstrong and Zhou (2022). Among them, research on certifications reveals interesting market impacts. Hui

et al. (2023) find that certification expansion increases market efficiency with asymmetric effects across seller groups, and Hui, Jin, and Liu (2025) show that changes in certification design shape market structure by altering seller distributions and sales patterns. Wu, Huang, and Li (2023) investigate how consumer information disclosures affect market competition and derive profit-maximizing information design strategies. Third, welfare implications have also received attention, with Wu et al. (2015) measuring the economic value of reviews, Gandhi, Hollenbeck, and Li (2024) quantifying the costs of fake reviews, and Gardete (2013) discussing the welfare implications of advertising. Chen, Du, and Lei (2025) examine interactions between online reviews and pricing in signaling contexts, sharing conceptual similarities with our approach. However, our paper examines different interactions and draws distinct conclusions. Specifically, we demonstrate two unique mechanisms. Ranking interacts with quality signaling to cause the crowding-out effect, and interacts with certifications to cause market segmentation.

Our research identifies a novel form of market segmentation that emerges from the interaction between ranking algorithms and quality certifications, which is related to but differs fundamentally from traditional price discrimination mechanisms. The established literature on price discrimination (Mussa and Rosen, 1978) examines how firms directly segment consumers through pricing strategies, while recent platform research investigates algorithmic price discrimination through machine learning approaches (Johnson, Rhodes, and Wildenbeest, 2023). The literature covers various pricing strategies including input pricing (Miklos-Thal and Shaffer, 2021), dynamic pricing (Choe, King, and Matsushima, 2018), competitive pricing (Belleflamme, Lam, and Vergote, 2020; Do and Miklos-Thal, 2025), and two-sided pricing (Lin, 2020). Research on platform fee structures (Miklos-Thal and Shaffer, 2021) further illuminates how platforms optimize pricing across seller segments and strategically use pass-through rates to influence competitive dynamics. In contrast to these direct pricing approaches, our findings reveal that market segmentation can emerge through information design mechanisms rather than pricing manipulation.

3 Analytical Model

This section develops a theoretical model that combines signaling games and price competition to analyze how ranking algorithms affect quality information transmission and market structure. We examine how different ranking mechanisms interact with seller signaling and certification pro-

grams. Our analysis focuses on quality-based ranking and price-based ranking, characterizing their equilibrium properties and comparative implications for market segmentation and information provision.

We focus on these two ranking algorithms because quality and price represent the fundamental dimensions of competition in vertically differentiated markets. While we abstract from horizontal differentiation in our theoretical analysis, this simplification aligns with our empirical context: the platform in our dataset exclusively considered and implemented these two ranking approaches without personalization features, making quality-based and price-based algorithms the natural focus for both theoretical development and empirical validation.

3.1 Settings

Basic Market Setup We analyze a platform with both uncertified and certified sellers that offer vertically differentiated products to consumers. In such markets, quality differences are unambiguously agreed upon by all consumers, with seller quality measured by metrics such as fulfillment speed, packaging quality, customer service, and defect rates. The differences in quality of products reflect prevalent phenomena across major e-commerce platforms, where uncertified sellers exhibit substantial quality heterogeneity while certified sellers such as eBay's Top-Rated Sellers provide credible quality signals through rigorous verification processes.

Uncertified sellers offer products with either low quality q_L or high quality $q_H = q_L + \Delta_q$, where $\Delta_q > 0$. Each uncertified seller privately observes its own product quality, while the quality distribution across all uncertified sellers is common knowledge among market participants through platforms' aggregate statistics and consistent market expectations. The high-quality uncertified seller faces a constant marginal cost \bar{c} . The low-quality uncertified seller's marginal cost c_L is drawn from a continuous distribution with cumulative distribution function $G : [\underline{c}, \bar{c}) \rightarrow [0, 1]$ and full support on $[\underline{c}, \bar{c})$, where $\bar{c} > \underline{c} \geq 0$.¹

In contrast, certified sellers offer products with verified high quality $q_B := q_H + \Delta_B$ at marginal cost c_B , which is publicly observable, where $\Delta_B \geq 0$. For tractability, we focus on a setting with one certified seller (B) and two uncertified sellers comprising one low-quality uncertified seller (L) and

¹Cost heterogeneity of low-quality uncertified sellers naturally eliminates unreasonable equilibria where sellers set prices below their marginal cost under price-based ranking and enables capturing heterogeneous welfare effects across uncertified sellers with different cost structures.

one high-quality uncertified seller (H).²

Each seller sets its price based on its product quality and marginal cost. For sellers with constant quality and cost (H and B), their prices are denoted as $p_H \in \mathbb{R}_+$ and $p_B \in \mathbb{R}_+$, respectively. The pricing strategy of seller L is represented as a function of its cost: $p_L(\cdot) : [\underline{c}, \bar{c}] \rightarrow \mathbb{R}_+$.

Consumer Behavior and Platform Visibility Structure Consumers are vertically differentiated in quality sensitivity: $\alpha \in [\underline{\alpha}, \bar{\alpha}]$ where $0 \leq \underline{\alpha} < \bar{\alpha}$, with distribution $F \in \Delta[\underline{\alpha}, \bar{\alpha}]$. A consumer with parameter α who purchases product $j \in \{L, H, B\}$ obtains utility $u_j = u_0 + \alpha q_j - p_j$, where $u_0 \geq 0$ represents the base utility. The outside option o yields zero utility. Ties with the outside option are broken in favor of purchase.

Consumer consideration follows the random consideration set (RCS) framework of [Manzini and Mariotti \(2014\)](#), where each alternative j is included in the consideration set with attention probability γ_j . We adopt a specialized version with attention probabilities $(\gamma_B, \gamma_{(1)}, \gamma_{(2)}, \gamma_o) = (1, 1, 0, 1)$, where (k) denotes the k -th ranked uncertified product. This specification ensures that only the certified seller B , the top-ranked uncertified product, and the outside option enter the consideration set: $J = \{(1), B, o\}$. The consumer then makes a utility-maximizing choice within this consideration set.

Our model builds on the following stylized facts: certified sellers compete with enhanced visibility, whereas uncertified sellers compete for scarce visibility allocated through the platform's ranking algorithms. Certification programs (e.g., Amazon Prime / Amazon's Choice, eBay Top-Rated Sellers, Etsy Star Sellers) derive value from the resulting visibility advantage. We represent this via a parsimonious abstraction in which the certified product is always considered and uncertified sellers compete for a single visibility slot; this isolates how the ranking algorithm governs competition among uncertified sellers.

Platform's Ranking Algorithms The platform employs two distinct ranking algorithms to allocate visibility among uncertified sellers: quality-based ranking ($r = Q$) and price-based ranking ($r = P$).³

²An unintended consequence of this configuration with two uncertified sellers is that each uncertified seller can perfectly infer its competitor's quality by combining its private information with knowledge of the aggregate quality distribution. It is important to clarify that this perfect inference property does not drive our main results, and the framework readily extends to settings with $N \geq 2$ uncertified sellers while preserving the model's core insights.

³Ties between uncertified sellers are resolved by random assignment, with each tied seller receiving equal probability of being ranked first.

These algorithms operate exclusively within the uncertified segment, as certified sellers maintain guaranteed visibility through the certification mechanism.

The key distinction between these two algorithms lies in their information requirements and decision criteria. Under quality-based ranking, the platform leverages both private signals and observed pricing behavior to infer seller quality, while price-based ranking relies solely on publicly observable prices. This fundamental difference creates distinct competitive dynamics and strategic incentives for uncertified sellers.

Quality-based Ranking Algorithm ($r = Q$). The platform possesses private information about seller qualities, characterized by a signal s drawn from $\mathcal{S} := [q_L, q_H]$ representing its prior belief about seller H 's expected quality.⁴ This signal follows a known distribution with a cumulative distribution function $F_s : \mathcal{S} \rightarrow [0, 1]$. Only the distribution (not the realization) of s is publicly announced. The platform combines this private signal with observed prices (p_L, p_H) to form a belief about seller H 's expected quality, denoted as $b^0(p_L, p_H, s)$. It then ranks seller H first based on its perceived quality relative to the average quality of the uncertified seller: with probability 1 if $b^0(p_L, p_H, s) > q_L + \Delta_q/2$, with probability 1/2 if $b^0(p_L, p_H, s) = q_L + \Delta_q/2$, and with probability 0 otherwise. The expected probability of H being ranked first is $\Omega_Q(p_L, p_H; b^0(\cdot)) = \int_{\mathcal{S}} [1_{\{b^0(p_L, p_H, s) = q_L + \Delta_q/2\}}/2 + 1_{\{b^0(p_L, p_H, s) > q_L + \Delta_q/2\}}] dF_s(s)$, where $1_{\{\cdot\}}$ denotes the indicator function.

Price-based Ranking Algorithm ($r = P$). This algorithm ranks uncertified sellers purely based on their posted prices, with the lowest-priced seller receiving top visibility. The ranking rule is deterministic and independent of the platform's quality beliefs. Seller H is ranked first with probability 1 if $p_H < p_L$, with probability 1/2 if $p_H = p_L$, and with probability 0 if $p_H > p_L$. Formally, $\Omega_P(p_L, p_H; b^0(\cdot)) = 1_{\{p_H = p_L\}}/2 + 1_{\{p_H < p_L\}}$, where the platform's beliefs $b^0(\cdot)$ are retained for notational consistency but play no role in the ranking decision.

3.1.1 Equilibrium

Consumer Belief and Strategy The consumer observes the ranking algorithm r , the prices $(p_{(1)}, p_B)$ of the first-ranked uncertified and certified products, and the signal distribution F_s (but not its realization s). The consumer has perfect quality information about the certified product B (quality q_B), but must form beliefs about the quality of the first-ranked uncertified product based

⁴The corresponding prior for seller L 's expected quality is $q_L + q_H - s$ by Bayes' rule. A symmetric signal distribution around $q_L + \Delta_q/2$ represents the case where the platform does not have informative private signals

on its price and the ranking mechanism.

Specifically, consumers form a homogeneous belief $b^r(p_{(1)}; (1))$ about the quality of first-ranked uncertified product, which depends on the observed price and the ranking algorithm r and is independent of their quality sensitivity α . We focus on Perfect Bayesian Equilibrium (PBE) where consumers act sequentially rationally, choosing the product that maximizes their expected utility within the consideration set $J = \{(1), B, o\}$. The utility from purchasing product $j \neq o$ is $u_0 + \alpha b^r(p_j; j) - p_j$, where $b^r(p_j; j)$ represents the consumer's belief about product j 's quality, while the outside option yields zero utility. For simplicity of notations, we omit the explicit definition of consumers' strategy and characterize equilibrium in terms of sellers' pricing strategies $(p_L^*(\cdot), p_H^*, p_B^*)$, the platform's belief function $b^{0*}(\cdot)$, and consumers' belief function $b^{r*}(\cdot; (1))$, with consumers' strategies adhering to sequential rationality by default.

Timing The game proceeds in four stages:

1. The platform announces the ranking algorithm $r \in \{Q, P\}$ and the distribution of its private signal F_s (the realization s remains private to the platform).
2. Nature determines the low-quality uncertified seller L 's marginal cost c_L .
3. All sellers simultaneously set prices: L chooses $p_L(c_L)$, H chooses p_H , and B chooses p_B .
4. The platform implements its ranking algorithm. Consumers observe the ranking and prices, form beliefs, and make purchase decisions.

Figure 1 illustrates the game structure and information flow between the three key agents: sellers, platform, and consumers. The diagram depicts how sellers set prices, the platform implements either quality-based or price-based ranking algorithm, and consumers make purchase decisions based on the information provided through the ranking. The figure also highlights the belief structures of both the platform ($b^{0*}(\cdot)$) and consumers ($b^{r*}(\cdot; (1))$), showing how information flows through the system under different ranking algorithms.

Perfect Bayesian Equilibrium We analyze Perfect Bayesian Equilibrium (PBE), where strategic interactions are mediated by beliefs that endogenously emerge from observable actions. The equilibrium consists of pricing strategies $(p_L^*(\cdot), p_H^*, p_B^*)$ for all seller types and belief functions $(b^{0*}(\cdot), b^{r*}(\cdot; (1)))$ that govern platform ranking and consumer decisions.




Figure 1: Game Structure and Information Flow

A valid PBE requires three fundamental conditions: (i) *belief consistency*—beliefs are derived from seller strategies via Bayes’ rule whenever possible; (ii) *seller optimality*—each seller’s pricing strategy maximizes expected profits given the strategies of others and prevailing market beliefs; and (iii) *consumer rationality*—consumers make utility-maximizing choices given their beliefs about product quality.

To formalize seller optimization, we define the expected payoff functions that capture the strategic interaction between pricing, ranking, and consumer demand. An uncertified seller $j \in \{L, H\}$ with marginal cost c_j earns expected payoff $U_{(1)}(c_j, p_{(1)}, b^{r*}(p_{(1)}; (1)); p_B^*) = (p_{(1)} - c_j) \cdot D_{(1)}(p_B^*, p_{(1)}, b^{r*}(p_{(1)}; (1)))$ when ranked first at price $p_{(1)}$. The demand function $D_{(1)}(p_B^*, p_{(1)}, b^{r*}) = [F((p_B^* - p_{(1)})/(q_B - b^{r*})) - F((p_{(1)} - u_0)/b^{r*})]^+$ captures the measure of consumers who find the first-ranked uncertified product preferable to both the certified alternative and the outside option, where $[z]^+ := \max\{z, 0\}$ ensures non-negative demand. The certified seller faces the complementary demand structure, earning expected payoff $U_B(p_B, p_{(1)}, b^{r*}(p_{(1)}; (1))) = (p_B - c_B) \cdot D_B(p_B, p_{(1)}, b^{r*}(p_{(1)}; (1)))$, where $D_B(p_B, p_{(1)}, b^{r*}) = 1 - F(\max\{(p_B - u_0)/q_B, (p_B - p_{(1)})/(q_B - b^{r*})\})$ represents consumers who select the certified product as their preferred alternative. These demand functions embody the heterogeneity in consumer quality sensitivity $\alpha \sim F$ and the resulting utility-maximizing choices conditional on quality beliefs $b^{r*}(\cdot; (1))$. Complete PBE conditions and optimization details are provided in Appendix D.1.

Platform’s Belief Formation The platform’s ranking decisions depend on its ability to infer seller quality from two information sources: observable pricing patterns and private signals. When sellers choose different prices, the platform can directly infer seller types through price-based signaling. When sellers pool at identical prices, the platform relies on its private signal s to form quality beliefs $b^{0*}(\cdot)$.

To formalize this equilibrium on-path belief formation process, consider the equilibrium set $C_p := \{c \in [\underline{c}, \bar{c}) : p_L^*(c) = p_H^*\}$ representing cost realizations that lead seller L to match seller H ’s price. The platform’s belief formation follows a natural information aggregation rule consistent with Perfect Bayesian Equilibrium requirements: when prices differ ($c \notin C_p$), the platform perfectly identifies the high-quality seller and sets $b^{0*}(p_L^*(c), p_H^*, s) = q_H$; when prices are identical ($c \in C_p$), pricing provides no information, so the platform relies solely on its private signal: $b^{0*}(p_L^*(c), p_H^*, s) = s$.

Under quality-based ranking, this translates into ranking probabilities that depend on the informational environment. When prices separate, the platform deterministically ranks the high-quality seller first. When prices pool, ranking depends on the platform’s informational advantage, characterized by the platform’s ranking accuracy parameter $\omega := \Pr[s > q_L + \Delta_q/2] + \Pr[s = q_L + \Delta_q/2]/2 \in (1/2, 1]$ —the probability that the platform’s private signal generates a higher quality belief for seller H than for seller L when both sellers set identical prices. Values approaching $\omega = 1$ indicate the platform utilizes near-perfect quality information on quality-based ranking, while $\omega = 1/2$ corresponds to random assignment. The formal mathematical expressions for the belief function and ranking probabilities are provided in Appendix D.2.

3.2 Equilibrium Structure

Our analysis characterizes equilibrium outcomes under both ranking mechanisms. We focus on economically meaningful equilibria where uncertified sellers maintain positive profit opportunities. This focus requires that the certified seller does not dominate the market entirely and that there is sufficient consumer heterogeneity in quality sensitivity to support competitive interaction among all types of sellers. These conditions, formalized in Appendix A.1, ensure that ranking competition among uncertified sellers remains economically relevant and that strategic incentives for market participation are preserved across all types of sellers.

3.2.1 Quality-based Ranking Algorithm

Under the quality-based ranking algorithm ($r = Q$), platform ranking decisions incorporate private signal s alongside price observations. A crucial property emerges from this ranking approach: the equilibrium price of the first-ranked uncertified seller systematically equals that of the high-quality seller H .

Lemma 1. *In any equilibrium under the quality-based ranking algorithm $r = Q$, the price of the first-ranked uncertified seller always equals the price of seller H , i.e., $p_{(1)}^* = p_H^*$.*

This result emerges from the informational structure inherent in quality-based ranking. Remember the set $C_p := \{c \in [\underline{c}, \bar{c}) : p_L^*(c) = p_H^*\}$ that represents marginal cost realizations that induce seller L to match seller H 's equilibrium price. When L 's cost lies outside C_p , pricing differences enable the platform to identify seller types directly, leading to deterministically ranking H first. Conversely, when L 's cost falls within C_p , both sellers set identical prices ($p_H^* = p_L^*(c)$), so the first-ranked price equals p_H^* regardless of which seller the platform's algorithm selects for the top position.

Without additional restrictions, the model admits pathological equilibria where first-ranked sellers set prohibitively high prices, resulting in zero demand: $D_{(1)}(p_B^*(p_H^*), p_H^*, b^{Q*}(p_H^*; (1))) = 0$. Such outcomes arise when platforms form extreme off-equilibrium beliefs—treating any price deviation as definitive evidence of low quality, thereby discouraging any deviations from meaningful competition.

To focus on robust economic outcomes, we discipline platform belief formation through the following reasonable restriction:

Assumption 1 (Off-path Belief Restriction). *For any off-equilibrium price pair $(p_L, p_H) \neq (p_L^*(c), p_H^*)$ for any $c \in [\underline{c}, \bar{c})$ with $p_L, p_H \geq \bar{c}$, the platform's belief formation satisfies:*

1. *When $p_L = p_H$, beliefs equal the prior: $b^0(p_L, p_H, s) = s$;*
2. *When $p_L \neq p_H$, the platform cannot deterministically assign types: $\Omega_Q(p_L, p_H; b^0(\cdot)) \in (0, 1)$.*

This assumption reflects economic reality: platforms incorporating multiple quality indicators beyond price cannot achieve perfect type identification from off-path pricing alone, particularly when both prices exceed \bar{c} , where both seller types can operate profitably. The restriction eliminates

equilibria dependent on arbitrary belief specifications while preserving competitive dynamics driven by meaningful strategic considerations.

Proposition 1 (Equilibrium Structure under Quality-based Ranking). *Given Assumption 1, under the quality-based ranking algorithm ($r = Q$), there exist Perfect Bayesian Equilibria characterized by the following properties:*

1. **Pooling equilibrium structure:** $p_{(1)}^* = p_L^*(c) = p_H^* > \bar{c}$ for all $c \in [\underline{c}, \bar{c}]$, with both uncertified seller types choosing identical prices strictly above the high-quality marginal cost.
2. **Ranking and beliefs:** Seller H is ranked first with probability $\omega \in (1/2, 1]$. The resulting equilibrium belief is $b^{Q*}(p_H^*; (1)) = q_L + \omega\Delta_q$.
3. **Certified seller's response:** $p_B^*(p_{(1)}^*) = \operatorname{argmax}_{p_B \in \mathbb{R}_+} U_B(p_B, p_{(1)}^*, q_L + \omega\Delta_q)$.
4. **Equilibrium multiplicity and profitability:** Any price $p_{(1)}^* > \bar{c}$ satisfying $D_{(1)}(p_B^*(p_{(1)}^*), p_{(1)}^*, q_L + \omega\Delta_q) > 0$ constitutes a valid equilibrium, with all sellers earning strictly positive profits.

Proof. See Appendix A.2. □

The pooling equilibrium emerges because visibility generates sufficient profit advantages to make price mimicry optimal for low-quality sellers. When both uncertified seller types select identical prices, consumers cannot extract quality information from price signals and must rely on the platform's imperfect ranking algorithm. Consumers' beliefs consequently reflect a weighted average of seller types: $b^{Q*}(p_H^*; (1)) = q_L + \omega\Delta_q$, where the weighting depends on the platform's ranking accuracy ω . The equilibrium supports any pooling price $p_{(1)}^* > \bar{c}$ that maintains positive demand against certified competition, with the precise price level determined by off-equilibrium belief specifications.

3.2.2 Price-based Ranking Algorithm

The price-based ranking algorithm ($r = P$) operates through a fundamentally different competitive logic, prioritizing price competition over quality assessment. Under this mechanism, seller H achieves the first position in search ranks with probability 1 if $p_H < p_L$, probability 0 if $p_H > p_L$, and probability 1/2 when prices are equal. This ranking rule creates a direct trade-off between profit margins and visibility: gaining top position requires aggressive pricing, potentially at the expense of profitability.

This mechanism reverses the competitive dynamics relative to quality-based ranking. Low-quality sellers with cost advantages ($c < \bar{c}$) can secure top visibility by undercutting high-quality competitors, while high-quality sellers face the constraint that aggressive price competition erodes their already-thin profit margins. Consequently, cost-efficient low-quality sellers may systematically dominate visibility rankings under price-based ranking.

Multiple equilibria can emerge in signaling games depending on off-equilibrium belief specifications. To identify economically meaningful outcomes, we impose a natural restriction on consumer beliefs:

Assumption 2 (Continuous Consumer Belief). *The consumer's belief function $b^P(\cdot; (1))$ is continuous.*

This continuity requirement serves dual purposes. Economically, it embodies the realistic principle that consumer quality assessments should respond smoothly to price variations—small price differences should not trigger dramatic shifts in quality perceptions. Technically, this assumption eliminates equilibria that depend exclusively on discontinuous belief jumps without economic foundation, thereby focusing analysis on robust strategic outcomes.

Under this assumption, we characterize the equilibrium under the price-based ranking algorithm:

Proposition 2 (Equilibrium Structure under Price-based Ranking). *Given Assumption 2, under the price-based ranking algorithm ($r = P$), there exist Perfect Bayesian Equilibria characterized by the following properties:*

1. **Price separating among uncertified sellers:** $p_L^*(c) = \operatorname{argmax}_{p \in \mathbb{R}_+, p < \bar{c}} U_{(1)}(c, p, q_L; p_B^*)$ for all $c \in [\underline{c}, \bar{c}]$, while the high-quality seller sets $p_H^* = \bar{c}$.
2. **Ranking and beliefs:** The low-quality uncertified seller L is ranked first with probability 1, while H is never ranked first. The resulting equilibrium belief is $b^P(p_{(1)}^*; (1)) = q_L$.
3. **Certified seller's pricing strategy:** $p_B^* = \operatorname{argmax}_{p_B \in \mathbb{R}_+} \int_{\underline{c}}^{\bar{c}} U_B(p_B, p_L^*(c), q_L) dG(c)$.
4. **Profitability:** Low-quality uncertified sellers (L) and the certified seller (B) earn strictly positive expected profits, while the high-quality seller (H) earns zero profit.

Proof. See Appendix A.3. □

This separating equilibrium exhibits fundamentally different competitive dynamics from the

pooling outcome under quality-based ranking. High-quality uncertified sellers become constrained to marginal cost pricing ($p_H^* = \bar{c}$), earning zero economic profit, while cost-advantaged low-quality sellers ($c < \bar{c}$) can profitably secure top visibility through strategic underpricing.

The mechanism achieves perfect information revelation: since high-quality sellers cannot profitably compete at the aggressive price levels required for top ranking, the first-ranked position becomes systematically occupied by low-quality sellers. This generates the equilibrium belief $b^{P*}(p_{(1)}^*; (1)) = q_L$, providing perfect quality inference that contrasts sharply with the imperfect information environment under quality-based ranking where $b^{Q*}(p_H^*; (1)) = q_L + \omega\Delta_q$.

3.3 Analysis of Uncertified Sellers

This section compares the behaviors and welfare of sellers under quality-based and price-based ranking mechanisms. For brevity, we index equilibrium objects by superscripts Q and P for quality-based and price-based ranking algorithms, respectively (e.g., p_j^{*Q}, p_j^{*P}). Throughout the following analysis, $p_{(1)}^*$ denotes the pooling price of the uncertified seller who ranks first under $r = Q$. We define the profit difference between price-based and quality-based ranking as $\Delta\Pi_j := \Pi_j^{*P} - \Pi_j^{*Q}$ for seller $j \in \{L, H\}$.

3.3.1 Equilibrium Structure Comparison

The two ranking algorithms generate fundamentally different equilibrium structures. Under quality-based ranking, all uncertified sellers pool at identical prices strictly above the high-quality marginal cost ($p_L^{*Q}(c) = p_H^{*Q} > \bar{c}$ for all $c \in [\underline{c}, \bar{c}]$), as visibility-seeking incentivizes low-quality sellers to mimic high-quality pricing (Proposition 1). In contrast, price-based ranking creates separating equilibria where low-quality sellers leverage cost advantages to undercut high-quality competitors ($p_L^{*P}(c) < p_H^{*P} = \bar{c}$), thereby systematically securing top rankings (Proposition 2).

This structural transformation also alters the information environment. Under quality-based ranking, prices convey no quality information due to pooling, and consumers must rely on imperfect algorithmic signals. Under price-based ranking, separation implies perfect quality inference: the first-ranked uncertified seller is systematically low-quality because high-quality sellers cannot profitably undercut the equilibrium price level. This demonstrates a crowding-out effect, in which quality information provided through quality-based ranking algorithms crowds out quality information provided by sellers through signaling incentives.

Claim 1 (Crowding-Out Effect). *Under price-based ranking, prices serve as effective signals of product quality among uncertified sellers, whereas under quality-based ranking, they do not.*

3.3.2 Price Competition and Strategic Effects

Price-based ranking intensifies price competition among uncertified sellers through two channels. First, the visibility mechanism operates through pure price competition without incorporating quality signals, eliminating high-quality sellers' natural advantages. Second, low-cost sellers gain systematic ranking advantages, creating downward pricing pressure throughout the uncertified segment.

Claim 2 (Competition Intensification). *Under price-based ranking, uncertified sellers set lower prices than under quality-based ranking.*

3.3.3 Welfare Redistribution

Compared with quality-based ranking, price-based ranking fundamentally redistributes welfare among uncertified sellers. This section analyzes these welfare effects and demonstrates how they create systematic differential impacts across different types of sellers.

Our welfare analysis must account for equilibrium multiplicity under quality-based ranking (different pooling prices may arise), whereas price-based ranking yields a unique structure. Although absolute profit levels depend on the pooling price under quality-based ranking, our comparative welfare conclusions across seller types are robust to equilibrium selection.

Profit Changes Across Seller Types Under the quality-based ranking algorithm, low-quality sellers with cost $c \in [\underline{c}, \bar{c})$ earn expected profits $(1 - \omega) \cdot U_{(1)}(c, p_{(1)}^*, q_L + \omega\Delta_q; p_B^*(p_{(1)}^*))$ with equilibrium prices $p_{(1)}^* = p_L^{*Q}(c) = p_H^{*Q}$ and $p_B^*(p_{(1)}^*)$, where ω reflects the accuracy of the algorithm. High-quality sellers earn $\omega \cdot U_{(1)}(\bar{c}, p_{(1)}^*, q_L + \omega\Delta_q; p_B^*(p_{(1)}^*))$, benefiting from their higher probability of being ranked first.

Price-based ranking algorithm eliminates this quality-based advantage of high-quality sellers. Low-quality sellers can now guarantee top ranking by setting sufficiently low prices, earning $\max_{p < \bar{c}} U_{(1)}(c, p, q_L; p_B^{*P})$, where p_B^{*P} is the equilibrium price of the certified seller under price-based ranking. High-quality sellers, however, cannot profitably compete at these price levels and thus make a zero profit.

This stark contrast between quality-based and price-based ranking creates predictable welfare patterns. High-quality sellers unambiguously lose with price-based ranking, as they forfeit their advantage in quality-based ranking while becoming unable to compete on price due to their higher costs. The effects on low-quality sellers depend on their cost structure and the accuracy of the quality-based ranking algorithm. When the quality-based ranking algorithm is sufficiently accurate (ω is large), low-quality sellers rarely rank first under quality-based ranking, so the benefit of consistent top positioning under price-based ranking outweighs their losses due to lower prices.

Lemma 2 (Welfare Redistribution among Uncertified Sellers). *Compared with quality-based ranking, price-based ranking generates the following welfare effects:*

1. *High-cost uncertified sellers always experience welfare reductions: $\Delta\Pi_L(c) < 0$ and $\Delta\Pi_H < 0$ for all $c \in [\tilde{c}_H, \bar{c}]$, where $\tilde{c}_H \in [\underline{c}, \bar{c}]$ is a threshold cost level. Moreover, the welfare reduction is strictly less severe for low-quality sellers than for high-quality sellers: $\Delta\Pi_L(c) \in (\Delta\Pi_H, 0)$ for all $c > \tilde{c}_H$.*
2. *When the ranking algorithm's accuracy of the platform is sufficiently high ($\omega > \hat{\omega}$ for some threshold $\hat{\omega} \in (1/2, 1)$), low-cost uncertified sellers experience welfare gains while high-quality uncertified sellers face the most severe losses among all uncertified sellers. Formally, there exists $\tilde{c}_L \in [\underline{c}, \tilde{c}_H]$ such that $\Delta\Pi_L(c) > 0$ for all $c < \tilde{c}_L$, and $\Delta\Pi_H < \Delta\Pi_L(c)$ for all $c \in [\underline{c}, \bar{c}]$.*

Proof. See Appendix A.4. □

Note that this welfare redistribution is distortive, because high-quality sellers are hurt and low-quality sellers potentially benefit. If entry and exit are allowed, the market-level quality will decrease for uncertified sellers. We summarize these findings in the following claim:

Claim 3 (Distortive Welfare Redistribution among Uncertified Sellers). *Compared with quality-based ranking, price-based ranking generates the following welfare effects:*

- *Low-cost uncertified sellers can either experience an increase or a decrease in expected payoff (experience an increase if the quality-based ranking algorithm is relatively accurate).*
- *High-cost uncertified sellers experience a decrease in expected payoff.*
- *Among uncertified sellers, those offering high-quality products experience the largest decline in expected payoff if the quality-based ranking algorithm is relatively accurate.*

3.4 Analysis of Certified Sellers

We compare the pricing and welfare of certified sellers under quality-based and price-based ranking algorithms.⁵ To tractably illustrate how the pricing strategy of the certified seller differs under different signaling structures of uncertified sellers, we consider a special case in which the consumer's quality sensitivity, α , is uniformly distributed over $[\underline{\alpha}, \bar{\alpha}]$, with cumulative distribution function $F(\alpha) = (\alpha - \underline{\alpha}) / (\bar{\alpha} - \underline{\alpha})$. Additionally, we assume that $\bar{\alpha}$ is sufficiently high compared to $\underline{\alpha}$ to avoid corner solutions in equilibrium. Within this quality structure, the equilibrium price of the certified seller given the price of the first-ranked uncertified seller $p_{(1)}^*$ under the quality-based ranking algorithm is $p_B^*(p_{(1)}^*) = (c_B + \bar{\alpha}(\Delta_B + (1 - \omega)\Delta_q) + p_{(1)}^*)/2$. Similarly, under the price-based ranking algorithm, the certified seller's equilibrium price is $p_B^{*P} = (c_B + \bar{\alpha}(\Delta_B + \Delta_q) + \int_{\underline{c}}^{\bar{c}} p_L^{*P}(c)dG(c))/2$, where $p_L^{*P}(c) = \operatorname{argmax}_{p < \bar{c}} U_{(1)}(c, p, q_L; p_B^{*P})$, $\forall c \in [\underline{c}, \bar{c}]$.

Building on the equilibrium structures established in Propositions 1 and 2, we have $p_L^{*P}(c) < \bar{c} < p_{(1)}^*$, which implies $p_{(1)}^* > \int_{\underline{c}}^{\bar{c}} p_L^{*P}(c)dG(c)$. This price ordering enables us to decompose the certified seller's price change when switching from quality-based to price-based ranking into two countervailing effects:

$$p_B^{*P} - p_B^*(p_{(1)}^*) = \frac{1}{2} \left[\underbrace{\bar{\alpha}\omega\Delta_q}_{\text{Discrimination Effect } (>0)} - \underbrace{\left(p_{(1)}^* - \int_{\underline{c}}^{\bar{c}} p_L^{*P}(c)dG(c) \right)}_{\text{Competition Effect } (>0)} \right].$$

Both terms in this decomposition are strictly positive, but they affect the certified seller's price in opposite directions. The discrimination effect increases this difference, while the competition effect (entering with a negative sign) decreases it. We explain the economic intuitions of the two terms below:

1. **Discrimination Effect (Positive Impact on p_B):** Under the price-based ranking algorithm, the average quality of the first-ranked uncertified product is lower than under the quality-based ranking algorithm. Consequently, the perceived quality gap between the first-ranked uncertified product and the certified product increases. This reduces the competitive pressure

⁵We acknowledge a limitation of our stylized model: it considers only one certified seller, B . In reality, platforms typically host multiple certified sellers, and price competition among them could also intensify under the price-based ranking algorithm. This introduces potential price-reduction and market share increase effects that our model does not capture.

on the certified seller, allowing them to charge a higher price while attracting consumers with high-quality sensitivity.

2. **Competition Effect (Negative Impact on p_B):** The price of the first-ranked uncertified product decreases under the price-based ranking algorithm, intensifying price competition between the certified seller and the uncertified product. This increased competition drives the certified seller to lower its price.

Figure 2 illustrates discrimination effects. Under the quality-based ranking algorithm, consumers believe the first-ranked uncertified seller has quality $q_L + \omega\Delta_q$, while under the price-based ranking algorithm, they believe it has quality q_L . This shift in consumer beliefs about the first-ranked uncertified seller's quality separates q_B from their competitors as shown by Figure 2.




Figure 2: The Discrimination Effect: Quality Differentials Under Different Ranking Algorithms

The discrimination effect strengthens with the accuracy ω : a more accurate quality-based ranking algorithm enlarges the quality gap (the quality of certified sellers minus the quality of top-ranked uncertified sellers) under quality-based relative to price-based ranking. In contrast, the competition effect remains constant regardless of the accuracy ω when $p_{(1)}^*$ is fixed. Consequently, as the accuracy of the quality-based ranking algorithm increases, the certified seller's price becomes increasingly likely to rise if we change the ranking algorithms from quality-based to price-based ranking.

The welfare effects on certified sellers under quality-based ranking versus price-based ranking can vary in either direction. For instance, when the original quality-based algorithm is highly accurate, changing from quality-based ranking to price-based ranking significantly widens the quality gap between certified and top-ranked uncertified sellers, potentially benefiting certified sellers through reduced competitive pressure. Conversely, when the price of the top-ranked uncertified seller decreases substantially, intensified price competition under price-based ranking

reduces certified seller profits. A detailed explanation of certified seller welfare changes is provided in Appendix E.

Therefore, we propose the following claim regarding the price and payoff changes of certified sellers for empirical analysis:

Claim 4 (Price and Payoff Changes of Certified Sellers). *When the platform shifts from the quality-based ranking algorithm to the price-based ranking algorithm, certified sellers' prices can go either up or down (more likely to go up as the quality-based ranking algorithm is more accurate), and their expected payoffs can either increase or decrease.*

3.5 Market Segmentation

This section analyzes how ranking algorithms shape market structure through endogenous segmentation mechanisms. These algorithms create distinct competitive tiers that enable platform-level screening of consumers based on their quality sensitivity, transforming competition from direct rivalry across all sellers to a stratified structure where different seller types serve distinct consumer segments.

To illustrate the differences between market structures under quality-based and price-based ranking, we analyze a tractable example that extends our framework with several key simplifications: we assume $u_0 \geq \bar{c}$ to ensure all transactions are welfare-enhancing, and specify that consumer quality sensitivity α follows a uniform distribution over $[0, \bar{\alpha}]$ with cumulative distribution function $F(\alpha) = \alpha/\bar{\alpha}$ and probability density function $f(\alpha) = 1/\bar{\alpha}$. The upper bound $\bar{\alpha}$ is set sufficiently high to capture substantial consumer heterogeneity while ensuring interior solutions in equilibrium.

Platform-Level Market Structure Design The choice of ranking algorithm creates fundamentally different market architectures. Under quality-based ranking, the platform maintains a unified competitive tier where all uncertified sellers compete directly through pooling strategies. This creates largely overlapping market segments where certified and uncertified products compete for similar consumer types, resulting in head-to-head competition across quality levels with relatively modest perceived quality differentials between top uncertified and certified products.

Under price-based ranking, the platform systematically segments the market into distinct tiers: a low-price, low-quality tier dominated by cost-efficient uncertified sellers, and a premium tier

served by certified sellers with quality advantages. This algorithmic choice effectively implements a form of platform-level price discrimination by creating separate market segments with different price-quality combinations, allowing the platform to extract value from consumer heterogeneity in quality sensitivity.

Illustrative Example: Market Coverage Patterns The contrasting market structures under quality-based and price-based ranking can be characterized through market coverage patterns in our stylized setting. Under quality-based ranking, uncertified sellers pool at price $p_{(1)}^*$, generating market coverage of $1 - (p_{(1)}^* - u_0)^+ / [\bar{\alpha}(q_L + \omega\Delta_q)]$, where $(x)^+ := \max\{x, 0\}$. When pooling prices exceed base utility u_0 , consumers with low quality sensitivity may find participation unattractive, potentially leading to partial market coverage.

Under price-based ranking, intensified competition among uncertified sellers systematically drives pricing strategies where $p_L^{*P}(c) < \bar{c} \leq u_0$. This pricing pattern enables market expansion: the condition $p_L^{*P}(c) \leq u_0$ ensures that even consumers with the lowest quality sensitivity find participation attractive, achieving full market coverage. This market expansion effect demonstrates how price-based ranking can extend market access beyond what quality-based ranking achieves by creating specialized competitive tiers: a cost-efficient tier served by low-quality sellers with competitive pricing in top rankings that targets price-sensitive consumers, and a premium tier served by certified sellers that targets quality-sensitive segments.

Figure 3 visualizes these contrasting market segmentation patterns emerging from our illustrative example. The upper panel depicts the unified competitive structure under quality-based ranking, where uncertified sellers pool at identical prices $p_{(1)}^* > \bar{c}$, generating consumer beliefs $b^Q(p_{(1)}^*; (1)) = q_L + \omega\Delta_q$. This creates relatively intense competition between certified and uncertified sellers for overlapping consumer segments, as both target relatively quality-sensitive consumers.

The lower panel demonstrates the stratified market structure under price-based ranking. Here, the ranking algorithm creates two distinct competitive tiers: low-quality sellers dominate the low-price tier by setting $p_L^{*P}(c) < \bar{c}$, while certified sellers operate in the premium tier. This stratified structure eliminates direct competition between tiers and enables each to serve distinct consumer groups through specialized price-quality combinations.




Figure 3: Consumer Surplus and Market Segmentation Under Different Ranking Algorithms

The segmentation threshold α^* illustrates how price-based ranking enables natural consumer screening. Quality-sensitive consumers ($\alpha > \alpha^*$) self-select into the premium tier despite higher prices, while price-sensitive consumers ($\alpha < \alpha^*$) choose the low-price tier despite lower quality. This screening mechanism resembles second-degree price discrimination, where the ranking algorithm facilitates consumer surplus extraction through strategic market segmentation rather than direct pricing control.

The transition from quality-based to price-based ranking creates endogenous market segmentation that enables strategic value extraction from consumer heterogeneity. This market structure transformation allows the platform to capture value by creating distinct segments with different surplus extraction potential, and may facilitate overall market development by serving previously

excluded consumer segments.

3.6 Discussion




Figure 4: Core Mechanisms of Platform Ranking Redesign

Figure 4 summarizes the core theoretical mechanisms identified in our analysis. Compared with quality-based ranking, price-based ranking triggers two mechanisms that fundamentally shape market structure and information transmission.

The first mechanism reveals a counterintuitive crowding-out effect in seller signaling. Under quality-based ranking, the platform's visibility allocation based on quality assessments creates stronger incentives for low-quality uncertified sellers to mimic high-quality sellers' pricing strategies, leading to pooling equilibria where prices lose their informational value. Conversely, under price-based ranking, low-cost uncertified sellers gain systematic advantages by undercutting competitors, creating separating equilibria where prices become informative quality indicators. This demonstrates that quality-based ranking algorithms can paradoxically reduce price-based information transmission by altering sellers' strategic signaling incentives.

The second mechanism operates through the interaction between ranking and certification,

where the price-based ranking algorithms systematically elevate low-quality uncertified sellers to top positions, widening the perceived quality gap between certified and top-ranked uncertified products. This quality differentiation creates endogenous market segmentation: certified sellers capture quality-sensitive consumers through premium pricing strategies, while top-ranked uncertified sellers serve price-sensitive segments through competitive pricing.

The combined effect generates systematic welfare redistribution: if we change the ranking algorithms from quality-based to price-based ranking, high-quality uncertified sellers experience the largest losses as they forfeit quality-based visibility advantages while being unable to compete on price; low-cost uncertified sellers potentially benefit from guaranteed top positioning but may be hurt by decreased equilibrium prices; and certified sellers face mixed effects depending on whether discrimination or competition effects dominate.

This framework demonstrates that algorithmic choice constitutes active market structure design beyond mere information organization. Our subsequent empirical analysis validates these theoretical predictions using eBay's 2011 ranking algorithm redesign from quality-based ranking to price-based ranking, examining both the signaling mechanism through changes in price informativeness and consumer beliefs, and the certification mechanism through empirical welfare redistribution patterns, price changes, and market coverage.

Importantly, our two theoretical mechanisms do not depend on each other for their existence. The crowding-out effect would emerge purely from ranking algorithms' impact on signaling incentives, even without certification programs. Similarly, the market segmentation effect would arise from ranking-certification interactions, even without signaling considerations. However, when both mechanisms are present, signaling provides a buffering effect that moderates the market segmentation. This buffering operates through two channels. First, under quality-based ranking with imperfect platform information ($\omega < 1$), consumers perceive top uncertified sellers as having mixed quality ($q_L + \omega\Delta_q$) rather than pure high quality (q_H) as would occur with perfect platform information ($\omega = 1$), reducing the perceived quality change from the ranking redesign. Second, our subsequent empirical analysis incorporates consumers' belief adjustments where they use prices to infer quality—associating lower prices with lower quality—which discourages aggressive price competition among uncertified sellers that would otherwise intensify market segmentation.

4 Empirical Evidence

4.1 Empirical Context

To test the validity of our theoretical predictions, we exploit eBay's ranking algorithm redesign, which provides a natural experiment setting previously analyzed by [Dinerstein et al. \(2018\)](#). Regarding certification programs, eBay has a “top-rated seller” (TRS) flag shown to consumers, which serves as quality certifications. In order to become a “top-rated” seller, during the time period we studied, a seller needed at least 1,000 transactions and \$3,000 in sales in the previous 12 months and positive feedback scores more than 98 percent. In our theoretical framework, TRS-certified sellers correspond to the certified type B , while non-certified sellers represent the heterogeneous uncertified types L and H .

Regarding ranking algorithms, after a consumer enters a search query on eBay, eBay's algorithm filters listings and then ranks them. Before the redesign, the rank was generated by Best Match, which according to [Dinerstein et al. \(2018\)](#), corresponds to quality-based ranking ($r = Q$) in our analytical model. On May 19, 2011, eBay introduced a new two-stage product discovery process, which according to [Dinerstein et al. \(2018\)](#), corresponds to price-based ranking ($r = P$) in our analytical model. For more details on the ranking design, please refer to Appendix F. Following [Dinerstein et al. \(2018\)](#), we consider April 6, 2011 to May 18, 2011 as our sample before the redesign, and August 1, 2011 to September 20, 2011 as our sample after the redesign.

4.2 Model-Free Evidence

Due to the data availability constraint, we only study one product: the Microsoft Xbox 360 game *Halo Reach*, which is also the focus of the structural analysis in [Dinerstein et al. \(2018\)](#). This product has a large number of transactions and relatively stable demand and supply, making it an ideal focus for research. The data set that we borrow from [Dinerstein et al. \(2018\)](#) comprises both listing-level attributes and detailed individual search histories. The search records cover all visits to the *Halo Reach* product page and all visits to the search results page for queries containing the terms “xbox” (“x-box”), “halo”, and “reach”. The search records incorporate a total of 14,753 visits to the search results page and 6,733 visits to the product page.

Our theoretical predictions within this context can be summarized as follows:

Table 1: Model-Free Evidence

	Before	After	Change
<i>Panel A. Market Structure</i>			
Number of listings	270	218	-52
Number of listings from TRS	44	58	14
Number of listings from NTRS	226	160	-66
Percent of listings from TRS	16.3%	26.6%	10.3%
<i>Panel B. Pricing</i>			
Mean list price (+shipping)	\$39.73	\$37.88	-\$1.84
Mean list price (+shipping) for TRS	\$39.15	\$40.18	\$1.04
Mean list price (+shipping) for NTRS	\$39.84	\$37.05	-\$2.79
Median list price (+shipping)	\$37.00	\$35.00	-\$2.00
Median list price (+shipping) for TRS	\$38.00	\$39.47	\$1.47
Median list price (+shipping) for NTRS	\$37.00	\$34.95	-\$2.04
<i>Panel C. Consumer Search</i>			
Total number of views	137,544	237,688	100,144
Number of views from TRS	51,675	53,856	2,181
Number of views from NTRS	85,869	183,832	97,963
Average number of listings viewed in each session	14.59	19.71	5.12
<i>Panel D. Consumer Purchase</i>			
Total number of purchases	9,427	12,059	2,632
Number of purchases from TRS	3,107	3,309	202
Number of purchases from NTRS	6,320	8,750	2,430
Probability of purchase given view	6.85%	5.07%	-1.78%
Probability of purchase given view for TRS	6.01%	6.14%	0.13%
Probability of purchase given view for NTRS	7.36%	4.76%	-2.60%
Mean transacted price (+shipping)	\$34.56	\$33.30	-\$1.27
Mean transacted price (+shipping) for TRS	\$35.49	\$34.38	-\$1.11
Mean transacted price (+shipping) for NTRS	\$33.99	\$31.30	-\$2.69

- 1. Claim 1 (Crowding-Out Effect):** The signaling effect of prices becomes more pronounced among non-top-rated sellers (NTRS) following the redesign.
- 2. Claim 2 (Competition Intensification):** The redesign intensifies competition among NTRS, leading to price reductions in this group.
- 3. Claim 3 (Distortive Welfare Redistribution among Uncertified Sellers):** High-quality and high-cost NTRS experience the largest profit reductions due to the redesign, while low-quality and low-cost NTRS may experience profit increases.
- 4. Claim 4 (Price and Payoff Changes of Certified Sellers):** The prices and profits of top-rated sellers (TRS) can even increase following the redesign.

Table 1 presents descriptive statistics before and after eBay’s ranking algorithm redesign. We highlight three findings that can support our theoretical predictions in Section 3.

Price Change. Panel B reveals a 7.0% mean price decline among non-top-rated sellers (NTRS, $\Delta = -\$2.79$), consistent with Claim 2 that price-based ranking intensifies competition. Top-rated sellers (TRS), however, slightly increased list prices ($\Delta = +\$1.04$). This differential pricing pattern supports our theoretical prediction that the new ranking algorithm induced more intense competition and reduced prices among NTRS. It also aligns with Claim 4 that the prices of TRS may not decrease even with more intense competition.

The widening TRS price premium over NTRS in transacted prices following the redesign is consistent with increased market differentiation between certified and top-ranked uncertified sellers, so that consumers value certifications more. The magnitude of this premium expansion—from \$1.51 to \$3.08 in transacted prices—warrants a deeper investigation into the underlying changes in consumer behavior documented in Panels C and D.

Consumer Behavior. The redesign fundamentally altered consumers’ search and discovery patterns, as evidenced in Panel C. NTRS listings received 114% more views after the redesign, compared to only 4.2% more views for TRS listings. This dramatic shift in attention allocation reflects the mechanical effect of price-based ranking algorithms, which inherently favored lower-priced NTRS listings. Furthermore, we observe that the consideration set expanded by 35.1% (from 14.59 listings per browsing session to 19.71 listings per browsing session) due to the redesign. This is likely because consumers knew that products that ranked high in the search results were of low quality.

Comparing list prices and transacted prices reveals a notable pattern: transacted prices declined less than list prices did post-redesign, suggesting that consumers shifted toward higher-priced options within the available product range. This shift toward higher price percentiles could reflect consumers increasingly using prices as a quality signal, consistent with the crowding-out effect in Claim 1. Instead of guessing price sensitivities from observational data, we estimate a structural model in Section 5 to formalize this exercise.

An intriguing pattern emerges when the relationship between price changes and conversion rates (the probability of purchasing the product after viewing the product) is examined. Although economic intuition suggests that price reductions should increase conversion rates, we observe

the opposite for NTRS: a 7.0% price reduction coincides with a 35.3% decrease in conversion rate (from 7.36% to 4.76%). This apparent contradiction suggests that consumers' quality assessment mechanisms changed after the redesign. Two complementary changes in consumer behavior could explain this pattern. First, consumers may have increasingly relied on price as a quality signal for NTRS products, consistent with Claim 1, where lower prices signal lower quality in the absence of other quality indicators. Second, consumers may have placed greater weight on TRS certification as a quality assurance mechanism. These mechanisms work in tandem: as price-based quality inference becomes more pronounced for uncertified products, the value of certifications increases by providing an alternative quality signal that bypasses price-based judgments. In this process, prices, beliefs and the value of certifications all changed, so a formal demand model is needed to formalize the argument here. Therefore, we estimate a formal demand model in Section 5.

Market Structure. Panel A reveals a notable compositional shift that motivates a deeper investigation into the underlying competitive mechanisms. Despite the longer post-redesign observation period that would typically increase seller participation, the number of NTRS listings declined by 66 while the number of TRS listings increased by 14. This divergent pattern suggests that the algorithm redesign created heterogeneous welfare impacts across different types of sellers, primarily supporting Claim 3. It is likely that high-quality NTRS were hurt and thus exited the market, although the specific nature of the welfare redistribution requires further analysis.

We find that consumer behavioral adaptations are rational responses to the change in the market structure and seller behaviors, and can further weaken distortive welfare redistribution. First, consumers expanded their search intensity, viewing 35.1% more listings per session, which potentially reflects their awareness that top-ranked products under price-based algorithms may not represent the highest-quality options. Second, consumers shifted toward purchasing higher-priced products within the available range, suggesting an increased reliance on price as a quality signal, which is consistent with the crowding-out effect. These behavioral adjustments create buffering effects on market dynamics (see the last paragraph of Section 3): extended search reduces the competitive advantage of securing top ranking positions, while increased price-quality signaling creates costs for sellers who compete primarily through aggressive price competition. Such rational adaptations indicate that market participants respond intelligently to algorithmic changes, thus weakening incentives of low-cost and low-quality sellers to undercut and distort the market.

5 Empirical Model

The model-free evidence partially validates our theoretical predictions, but faces three limitations, especially in relation to Claim 1 and Claim 3. First, in terms of consumer behavior, we cannot comprehensively analyze all factors that shape consumer demand, nor can we formally identify quality beliefs or price sensitivity without further assumptions or models. Second, in terms of welfare redistribution, we cannot directly measure the costs of sellers or quantify the quality needed to demonstrate these effects and their underlying mechanisms. Third, when studying welfare redistribution, a new problem emerges. Specifically, we do not observe sellers who exited the market, so we cannot easily calculate profits of those who exited the market and thus are not in the observed data.

Our structural model in this section overcomes the three limitations in the following ways, respectively: (1) We formally specify a model to estimate consumer demand and beliefs (Table 5). (2) We impose the profit maximizing assumption to uncover the costs of sellers (Figure 5) and we make assumptions on how ranking algorithms work to uncover a proxy for the quality of products. (3) We run counterfactual profit simulations (Table 3) that output profit changes of all sellers.

5.1 Comparison with Dinerstein et al. (2018)'s Model

We would like to formally acknowledge that our empirical model is built on the basis of the structural model in Dinerstein et al. (2018), although we use our structural model to show different economic phenomena and arrive at different or somewhat opposite conclusions. We also want to highlight the following key improvements by standing on the shoulders of giants.

First, consistent with our theory, we consider asymmetric information. We assume that consumers do not know the true quality of products, so they have to infer the quality of products from information available to them. In contrast, Dinerstein et al. (2018) use a utility function of consumers that assumes consumers know the true quality of products. Alternatively, their utility function can also be regarded as the reduced form of our consumer utility function after we plug in the belief about quality formed by consumers in our model. This improvement allows us to estimate the demand on the data before and after the redesign, and equips us with tools to interpret the difference between parameters in the utility function before and after the redesign. Without

our model setting, the parameters in the utility function are policy-invariant and thus should not change so dramatically after the redesign. In our model, the change in these parameters has clear economic meanings and is consistent with both theoretical predictions and model-free evidence.

Second, our model differs from [Dinerstein et al. \(2018\)](#)'s in that we incorporate consumers' behavioral and belief adjustments in response to the algorithmic changes. In contrast, [Dinerstein et al. \(2018\)](#), if viewed as the reduced form of our model, assume that consumers had the same belief after the redesign although the market structure and seller actions changed substantially. Therefore, our approach should identify a more realistic counterfactual.

5.2 Model Specification

Consideration Set We adopt the consideration set framework developed by [Dinerstein et al. \(2018\)](#) to model consumer search. In this two-stage approach, each consumer i first forms a consideration set J_i and then makes a purchase decision by evaluating the utilities of products within this set. We operationalize consideration sets using consumers' observed search histories in the dataset. The consideration set J_i can be decomposed into J_i^J , representing targeted products (new *Halo Reach* items), and J_i^M , representing non-targeted products (other listings).

To model how targeted products are selected into consideration sets, we adapt the stochastic model of [Dinerstein et al. \(2018\)](#) with a key modification to accommodate information asymmetry. While they treated the pre-redesign ranking weight as the true quality, we instead interpret it as quality perceived by the platform (denoted as \hat{q}_j , where j indexes listings), which contains information about product quality but is distinct from the unobservable true quality (q_j). We assume that listings shown to consumers are drawn from a set of all active targeted products \mathcal{J}_i^J during consumer i 's session. \mathcal{J}_i^J represents the set of targeted products available to consumer i . Specifically, before the redesign, each active targeted product $j \in \mathcal{J}_i^J$ is associated with a sampling weight $\omega_j = \hat{q}_j$, which serves as a quality signal that partially reveals quality information to consumers. After the redesign, the sampling weight for each targeted product $j \in \mathcal{J}_i^J$ depends solely on its price and is defined as:

$$\omega_j = \exp \left[-\gamma \left(\frac{p_j - \min_{k \in \mathcal{J}_i^J} (p_k)}{\text{std}_{k \in \mathcal{J}_i^J} (p_k)} \right) \right]$$

See Appendix G for more details on the consideration set model.

Consumer Demand Before the redesign, the platform reveals a quality signal for each seller j 's product through the sampling weight $\hat{q}_j \in [0, 1]$ in the ranking algorithm. Given this quality signal, along with price p_j and TRS status TRS_j , consumers form beliefs about the quality of products \tilde{q}_{ij} according to:

$$\tilde{q}_{ij}(\hat{q}_j, p_j, TRS_j, \varepsilon_{ij}) := \beta_0 + \beta_1 p_j + \beta_2 TRS_j + \beta_3 p_j TRS_j + \beta_4 \hat{q}_j + \varepsilon_{ij}$$

The term ε_{ij} captures idiosyncratic information observable to consumer i about listing j but unobservable to researchers. We assume that it follows a Type-I extreme value distribution. Since the perfectly rational belief adjustment in economic theory does not exist in the real world, our way of modeling and understanding belief is the same as the empirical expectation (denoted by \mathcal{E} in [Pakes \(2010\)](#)) in industrial organization literature, which can be different from the rational expectation (denoted by E in [Pakes \(2010\)](#)). We use statistical methods to directly uncover belief (empirical expectation) from observational data and see how it changes across different environments.

The quality information available to consumers is tricky. It is not necessarily the ranking weights \hat{q}_j . Some consumers search multiple times, so the information they have is closer to the ranking weights \hat{q}_j . Other consumers may only come to eBay and buy the product immediately, so the information to them is more approximate to the rank in one specific search session. To solve this problem, we estimate multiple specifications with different proxies for quality information for robustness. In order to be consistent with [Dinerstein et al. \(2018\)](#), we use the specification with ranking weights \hat{q}_j as quality information in our main specification.

Given the consideration set J_i , consumer i evaluates the expected utility of each product $j \in J_i$. Since consumers cannot directly observe true quality q_j , they form beliefs based on observable information. We model consumer i 's expected utility from purchasing a targeted product j as shown in Equation (1). The structural parameters α_0 and α_1 are policy-invariant. Note that the price sensitivity parameter α_1 is not separably identifiable from the quality-price signaling relationship

parameter β_1 .

$$\begin{aligned} u_{ij} &:= \alpha_0 + \alpha_1 p_j + \tilde{q}_{ij} (\hat{q}_j, p_j, TRS_j, \varepsilon_{ij}) \\ &= (\alpha_0 + \beta_0) + (\alpha_1 + \beta_1) p_j + \beta_2 TRS_j + \beta_3 p_j TRS_j + \beta_4 \hat{q}_j + \varepsilon_{ij} := V_j + \varepsilon_{ij} \end{aligned} \quad (1)$$

After the redesign, consumers lose access to the platform's quality signal \hat{q}_j through the ranking algorithm. In this case, consumers must rely on observable characteristics to infer product quality. We model this adjustment by specifying that consumers' quality beliefs now depend only on price and certification status:

$$\tilde{q}_{ij} (p_j, TRS_j, \varepsilon_{ij}) := \beta_0 + \beta_1 p_j + \beta_2 TRS_j + \beta_3 p_j TRS_j + \varepsilon_{ij}$$

Consequently, the perceived utility becomes:

$$\begin{aligned} u_{ij} &:= \alpha_0 + \alpha_1 p_j + \tilde{q}_{ij} (p_j, TRS_j, \varepsilon_{ij}) \\ &= (\alpha_0 + \beta_0) + (\alpha_1 + \beta_1) p_j + \beta_2 TRS_j + \beta_3 p_j TRS_j + \varepsilon_{ij} := V_j + \varepsilon_{ij} \end{aligned}$$

Consistent with Dinerstein et al. (2018), we model the utility from purchasing a non-targeted product m as $u_{im} = \delta + \lambda \varepsilon_{im}$, where ε_{im} follows a Type-I extreme value distribution and λ parameterizes the degree of horizontal differentiation among non-targeted products as perceived by consumers. The value of the outside option is defined as $u_{io} = \varepsilon_{io}$, where ε_{io} also follows a Type-I extreme value distribution.

Pricing Behavior In equilibrium, sellers' pricing decisions represent optimal responses to consumer demand, with consumers' quality beliefs \tilde{q}_{ij} treated as fixed. This reflects the standard assumption that marginal deviations from equilibrium do not alter consumers' beliefs. Given the platform's transaction fee T and ad valorem fee t , seller j sets the price p_j to maximize the expected profit:

$$\max_{p_j} ((1 - t)p_j - c_j - T) D_j(p_j)$$

where c_j is the marginal cost and $D_j(p_j)$ represents the demand for seller j at price p_j , which follows logit choice probabilities:

$$D_j(p_j) = \sum_{J:j \in J \subseteq \mathcal{J}} \left[\frac{\exp(V_j)}{1 + \exp(\delta + \lambda \ln |J^M|) + \sum_{k \in J} \exp(V_k)} \right] \Pr(J \mid \mathcal{J})$$

5.3 Estimation

For demand side estimation, we utilize consumer search data to identify consideration sets before and after the redesign. From the composition of the consideration set, we identify the quality perceived by the platform (i.e., the sampling weight with quality-based ranking, \hat{q}_j) for each product. The estimation results of the consideration set model are presented in Table 5. We then estimate separate logit demand models for the pre- and post-redesign periods. The logit demand estimation results in Table 2 reveal several important changes. We will first discuss our conclusions based on our main specification that uses weights as quality information, and then discuss other specifications for robustness.

Table 2: Estimation Results of the Demand Model

	Before				After
	Main	Rank	Log Rank	No Signal	
Constant ($\alpha_0 + \beta_0$)	3.72 (1.07)	4.18 (1.13)	4.62 (1.18)	3.76 (1.03)	0.57 (0.75)
Price ($\alpha_1 + \beta_1$)	-0.24 (0.03)	-0.24 (0.03)	-0.24 (0.03)	-0.24 (0.03)	-0.19 (0.02)
Top-rated seller (TRS) (β_2)	4.13 (2.75)	3.00 (2.94)	2.35 (3.07)	4.72 (2.59)	10.85 (1.73)
Price \times TRS (β_3)	-0.10 (0.08)	-0.07 (0.08)	-0.06 (0.09)	-0.12 (0.07)	-0.24 (0.05)
Quality Information (β_4)	0.64 (0.36)	-0.03 (0.01)	-0.38 (0.08)		
Constant (Other listings) (δ)	-8.37 (0.41)	-8.36 (0.41)	-8.37 (0.41)	-8.37 (0.41)	-5.90 (0.26)
Size of epsilon (Other listings) (λ)	1.70 (0.14)	1.69 (0.14)	1.69 (0.14)	1.70 (0.14)	0.85 (0.09)

First, the price coefficient ($\alpha_1 + \beta_1$) increases from -0.24 to -0.19, indicating reduced price sensitivity. In economic terms, this means that consumers became less responsive to price changes after the redesign. This finding is consistent with the model-free evidence in Table 1, which shows that transacted prices declined less than listing prices post-redesign, suggesting consumers shifted

toward higher-priced options within the available range. Since the structural preference parameter α_1 is theoretically invariant across periods, this change must reflect an increase in β_1 , indicating that prices have become more informative signals of quality after the redesign. This supports the crowding-out effect in [Claim 1](#).

Second, the coefficient for TRS certification (β_2) increases dramatically from 4.13 to 10.85, with a decrease in estimation standard errors, indicating that certifications play a substantially more important role in consumer quality inference after the redesign. This provides empirical support for the theoretical prediction of a larger perceived quality gap between NTRS and TRS.

Third, the interaction term between price and Top-Rated Seller (TRS) certification (β_3) becomes more negative, changing from -0.10 to -0.24. This indicates that the relationship between price and quality beliefs differs significantly across seller types, with this differential effect becoming more pronounced following the redesign. Specifically, while higher prices signal higher quality for NTRS sellers (positive β_1), this price-quality signaling is attenuated for TRS sellers due to the negative interaction term. This coefficient pattern provides empirical support for [Claim 1](#): consumers increasingly rely on prices as quality signals for NTRS sellers after the redesign, while TRS sellers benefit from certifications as an alternative quality indicator that reduces their dependence on price signaling. This differential signaling mechanism explains the finding in [Table 1](#), where, despite price reductions, NTRS conversion rates decreased substantially, while TRS conversion rates remained stable despite higher list prices.

Finally, the observed changes in consumer belief parameters reveal effects that moderate the welfare impacts predicted by our theory. The strengthened price-quality signaling relationship (higher β_1) and the more negative interaction coefficient β_3 provide direct evidence of these adaptive responses. Specifically, consumers now interpret lower prices as signals of inferior quality, creating strategic costs for low-quality sellers attempting to gain ranking advantages through aggressive price competition. This effect operates most strongly for NTRS sellers, where price-based quality inference has become the primary signaling mechanism. These endogenous belief adjustments represent market responses that partially offset the welfare redistribution effects described in [Claim 3](#), consistent with the buffering mechanisms discussed in the final paragraph of [Section 3](#).

In [Table 2](#), in addition to sampling weights, we also use ranks and logs of ranks as quality signals. We also include a specification without any quality information in the column “No Signal”.

The coefficients do not vary significantly across specifications. The conclusion we draw from Table 2 remains exactly the same. Additionally, the fact that the result in the column “No Signal” is largely consistent with the results in other columns indicates that the change in the parameters after the redesign is not caused by removing the regressor, quality information.

For the supply side analysis, we employ the first-order conditions from sellers’ profit maximization problem to recover marginal costs for the pre-redesign period. Figure 5 presents the kernel density estimation of cost distributions separately for TRS and NTRS, and for high-quality NTRS and low-quality NTRS.⁶ The distributions reveal that TRS operate with systematically higher cost structures than NTRS, and high-quality NTRS operate with higher cost structures than low-quality NTRS. This is consistent with the assumption of our analytical model.




Figure 5: The Distribution of Estimated Marginal Costs

5.4 Welfare Analysis

We now test Claim 3, which proposes a welfare redistribution effect in which high-quality NTRS are hurt and low-quality/low-cost NTRS benefit. Since we have already estimated the costs of sellers, the simplest way is to calculate profits from data and compare them. However, this reduced-form method suffers from three key problems. First, many factors changed after the redesign and may confound the change in profits. A simple example is seasonal factors. Second, it has serious sample selection bias, because sellers whose welfare is hurt are less likely to remain in the market. So, they are less likely to be in our observed sample after the redesign. Third, we do not have information on the quality of the sellers after the redesign. Entry and exit are very frequent in this market, so we cannot merge sellers before the redesign and after the redesign to

⁶In this figure, high-quality NTRS are NTRS with \hat{q}_j above the average \hat{q}_j , and low-quality NTRS are NTRS with \hat{q}_j below the average \hat{q}_j .

retrieve information about quality from the observations before the redesign.

Therefore, we employ counterfactual simulations to quantify the welfare implications of eBay's ranking algorithm redesign. Specifically, following [Dinerstein et al. \(2018\)](#), we take the pre-redesign population of sellers, simulate their operation in the post-redesign environment with price-based ranking algorithms and adjusted beliefs and consumer browsing behaviors, and compute the resulting profit change (ΔU_j) for each seller. This approach allows us to measure how sellers of different types are differentially impacted by the redesign, while avoiding confounding problems and sample selection bias.

Our welfare analysis encounters two key methodological challenges that necessitate additional assumptions. First, consumer quality beliefs vary in different environments, but we do not have structural assumptions on how beliefs adjust. Our belief is purely empirical. Second, data limitations hinder direct observations of true product quality.

To deal with the first challenge, we employ a conservative belief. As we discover in Subsection [5.3](#), due to the buffering effects of signaling mechanisms discussed in Section [3](#), the adjusted belief after the redesign weakens the distortive welfare redistribution proposed in [Claim 3](#). Specifically, we notice that the adjusted belief especially hurts low-quality and low-cost NTRS, who are predicted by our theory to benefit the most after the redesign of ranking algorithms. First, consumers believe that sellers who charge low prices are of low quality. Low-cost and low-quality NTRS are the ones who charge low prices after the redesign, so they are especially harmed by this belief. Second, consumers believe that there is a larger quality gap between TRS and NTRS after the redesign. This also hurts low-quality and low-cost NTRS because they do not own TRS certifications. In a word, the adjusted belief reduces the probability that we can discover [Claim 3](#) in our empirical exercise. Therefore, we use this belief in the counterfactual. If we still find welfare redistribution in [Claim 3](#), it is likely that there is distortive welfare redistribution even if consumers know their existence and adjust their behaviors and beliefs in a way that weakens the distortive welfare impacts.

To deal with the second challenge, we assume that the platform's belief about the quality of products is informative of the true quality of products. In other words, we assume that the weights used in quality-based ranking algorithms capture meaningful information about quality of products. We utilize the quality perceived by the platform (i.e., the sampling weight under quality-based ranking, \hat{q}_j) as our quality measure.

This counterfactual is different from Dinerstein et al. (2018)'s in three important ways. First, we use this counterfactual to test distortive welfare redistribution proposed in Claim 3 despite the adjusted behaviors and beliefs of consumers that can weaken such welfare redistribution. Dinerstein et al. (2018) use this counterfactual to show the mechanism why prices decreased after algorithm redesign. Second, we use the adjusted belief that hurts low-quality and low-cost NTRS. In contrast, Dinerstein et al. (2018)'s model, if viewed as the reduced form of ours, uses the belief without adjustment from consumers. Third, we consider information asymmetry, which means that consumers know less about the quality of products than sellers. So, the ranking weights serve as a noisy signal of the true quality rather than the true quality itself.

To organize our welfare analysis results, we classify all 270 sellers into eight groups based on their welfare changes, with Group 1 experiencing the largest welfare losses and Group 8 experiencing the largest welfare gains (see Table 3 for group boundaries)⁷. The counterfactual analysis based on this classification reveals three empirical patterns that provide strong support for our theoretical predictions.

Table 3: Welfare Changes in Counterfactual Analysis

Groups	1	2	3	4	5	6	7	8
Upper Bound (\$)	-266.0	-177.1	-95.2	-66.6	-44.3	0.0	7.7	$+\infty$
Lower Bound (\$)	$-\infty$	-266.0	-177.1	-95.2	-66.6	-44.3	0.0	7.7
<i>Panel A. Numbers</i>								
Number of sellers	27	27	27	27	27	25	83	27
Number of TRS	4	2	7	4	1	6	6	14
Number of NTRS	23	25	20	23	26	19	77	13
<i>Panel B. Marginal Cost</i>								
Average marginal cost (\$)	35.53	40.39	41.81	35.67	31.40	37.31	36.16	26.53
Average marginal cost of TRS (\$)	36.63	36.03	40.05	39.80	36.57	35.66	42.27	30.69
Average marginal cost of NTRS (\$)	35.34	40.74	42.43	34.95	31.20	37.83	35.68	22.05
<i>Panel C. Quality</i>								
Average \hat{q}_j	0.155	0.045	0.029	0.019	0.013	0.008	0.016	0.052
Average \hat{q}_j of TRS	0.089	0.318	0.073	0.075	0.006	0.017	0.005	0.099
Average \hat{q}_j of NTRS	0.166	0.023	0.013	0.009	0.014	0.004	0.016	0.001

First, certification status strongly correlates with welfare changes due to the redesign. Panel A of Table 3 reveals that among sellers experiencing the most severe welfare losses (Groups 1-3,

⁷The lower bound is closed, and the upper bound is open.

with losses exceeding \$95.2), only 16.0% (13 out of 81) are TRS sellers. In contrast, TRS sellers constitute 51.9% (14 out of 27) of those experiencing welfare gains exceeding \$7.7 (Group 8). This asymmetric distribution cannot be attributed to random variation, as TRS sellers represent only 16.3% (44 out of 270) of the overall seller population. Furthermore, the conditional probability of experiencing nonnegative welfare effects is higher for TRS sellers at 45.5% (20 out of 44) compared to 39.8% for NTRS sellers (90 out of 226). Notably, 66 of the 90 NTRS sellers with nonnegative welfare effects generate zero sales under both ranking algorithms, indicating extremely low competitiveness. Excluding these inactive sellers, the vast majority of viable NTRS sellers experience welfare declines due to the redesign. This differential impact aligns with [Claim 4](#), which states that certification provides a buffer against increased price competition by offering an alternative quality signal that remains salient even when algorithmic quality information is reduced.

Second, cost structure emerges as a critical determinant of welfare outcomes within the NTRS category. Panel B of [Table 3](#) reveals a systematic pattern where high-cost NTRS sellers experience more severe welfare declines while low-cost NTRS sellers tend to achieve welfare gains: specifically, high-cost NTRS sellers (average marginal cost of \$42.43 in Group 3) experience substantial welfare declines, while low-cost NTRS sellers (average marginal cost of \$22.05 in Group 8) predominantly experience welfare gains. This inverse relationship between cost structure and welfare outcomes explains the differential impacts observed within the NTRS category, despite these sellers' common exposure to intensified price competition following the algorithmic redesign.

Third, product quality exhibits a consistent inverse relationship with welfare outcomes among NTRS sellers. Panel C of [Table 3](#) demonstrates this pattern: NTRS with the highest platform-perceived quality ($\hat{q}_j = 0.166$) are concentrated in Group 1, experiencing the largest welfare losses (exceeding \$266 per seller), while those with the lowest quality signals ($\hat{q}_j = 0.001$) predominate in Group 8, achieving welfare gains (exceeding \$7.7 per seller). This empirical pattern provides evidence supporting the welfare redistribution described in [Claim 3](#), demonstrating its persistence even under conditions where consumers recognize and adapt their beliefs and behaviors to mitigate it.

These findings have important implications for platform design and market efficiency. The counterfactual analysis reveals that algorithmic ranking changes generate redistributive effects, with welfare transferred from high-quality/high-cost sellers to low-quality/low-cost sellers. These

redistributive patterns may help explain why eBay eventually reverted to its original ranking algorithm—the negative welfare impacts on high-quality sellers may have outweighed efficiency gains from increased price competition.

5.5 Discussion: Comparison with Price Discrimination

As shown in our theory, the widened perceived quality difference after the redesign leads to endogenous market segmentation that resembles second-degree price discrimination. The endogenous market segmentation in our analysis has different mechanisms from market segmentation caused by price discrimination, but produces similar market outcomes.

Price discrimination is a profitable strategy only with heterogeneous consumers. However, because consumer IDs are not included in the data, it is difficult for us and for Dinerstein et al. (2018) to identify random coefficients. Despite this limitation, we still observe phenomena typically associated with second-degree price discrimination: price decrease and quality reduction of lower-tier products (NTRS) alongside price and quality premiums for higher-tier products (TRS). This validates our understanding that such market segmentation emerges endogenously rather than through deliberate design by any economic agent.

Price discrimination can expand market coverage and enhance profits with heterogeneous consumers. We find empirical evidence for both effects in our endogenously segmented market. From the model-free evidence presented in Panel D of Table 1, TRS sales increased by only 202 (6.5%), while NTRS sales increased by 2,430 (38.4%) after the redesign. This pattern suggests that the market expanded to serve consumers with lower willingness-to-pay following the redesign, providing evidence for market expansion.

Since we cannot directly observe profits in the data, we conduct counterfactual simulations with varying degrees of consumer heterogeneity by adding different variances to the price coefficient α_1 . The results are reported in Appendix C. We find that profit increases with consumer heterogeneity. When consumers are highly heterogeneous, the profit increase effect due to market segmentation can outweigh the profit decrease effect due to intensified price competition under price-based ranking.

6 Conclusion

This paper identifies two economic mechanisms through which the design of ranking algorithms interacts with seller signaling and certification programs to shape the competitive dynamics in digital marketplaces. First, we demonstrate a crowding-out effect in which quality information provided by the platform through ranking algorithms crowds out quality information provided by sellers through strategic signaling. Second, we identify a novel form of platform-mediated market segmentation that emerges from the interaction between ranking algorithms and certification programs: price-based rankings systematically allocate visibility to low-quality, low-cost sellers, creating a polarized market structure where certified sellers serve quality-sensitive consumers while top-ranked uncertified sellers primarily serve price-sensitive consumers.

Our empirical findings, summarized in Table 4, provide strong support for theoretical predictions. The transition to price-based ranking led to: (1) the crowding-out effect, as evidenced by the significant changes in price coefficients; (2) systematic price declines among uncertified sellers, with both list and transaction prices decreasing significantly; (3) distortive welfare redistribution, in which high-quality uncertified sellers experienced losses while low-quality/low-cost sellers benefited; and (4) mixed effects on certified sellers: prices slightly increased and a substantial proportion experienced welfare gains.

Table 4: Empirical Evidence Related to Theoretical Claims

Theoretical Claim	Empirical Observation & Quantification
Claim 1: Increased price-quality signaling of NTRS	NTRS price coefficient: $-0.24 \rightarrow -0.19$ TRS \times Price interaction: $-0.10 \rightarrow -0.24$
Claim 2: NTRS prices decline	Mean list price: $\$39.84 \rightarrow \37.05
Claim 3: High-quality NTRS hurt and Low-quality/low-cost NTRS benefit	Mean quality of most hurt NTRS: 0.166 Mean quality of most benefited NTRS: 0.001 Mean cost of most benefited NTRS: $\$22.05$
Claim 4: TRS prices can even increase	Mean list price: $\$39.15 \rightarrow \40.18 45.5% of TRS experienced welfare gains 51.9% of top gainers were TRS

All changes were measured before and after the platform redesign.

Our findings yield important implications for both platform design and regulation. For platform

designers, our results demonstrate the critical importance of considering (1) the informational consequences of algorithmic design choices, particularly how ranking algorithms interact with existing information disclosure mechanisms, and (2) the complex interdependencies among different platform design elements. The algorithmic allocation of visibility, in combination with quality certification programs, creates systematic winners and losers: low-quality uncertified sellers and certified sellers potentially benefit, while high-quality uncertified sellers experience welfare losses. Importantly, this redistribution of welfare occurs not through direct price discrimination but through the strategic manipulation of seller visibility and the resulting competitive dynamics. While these effects may enhance short-term platform profit by expanding demand-side coverage, they potentially undermine long-term market quality by discouraging high-quality sellers from market participation.

From a policy perspective, our work underscores the necessity of reconceptualizing ranking algorithms as active market-structuring mechanisms rather than neutral information-filtering tools. These algorithmic systems fundamentally reshape competitive dynamics, market structure, and welfare distribution through their embedded design choices. As platforms increasingly mediate economic transactions, the strategic implications of their algorithmic architectures warrant greater regulatory scrutiny, particularly regarding their effects on market concentration, quality provision, and competitive fairness. Furthermore, while society increasingly calls for algorithmic transparency in the regulation of digital platforms (Wang et al., 2023), our research demonstrates that even simple and transparent ranking algorithms generate complex, often counterintuitive market effects that challenge straightforward interpretation. This complexity suggests that algorithmic transparency, while valuable, represents only a partial solution to the broader challenges of platform governance and market regulation.

A Proofs

A.1 Assumption for positive demand

The first-ranked uncertified seller is able to capture a positive market share when its price is less than or equal to the marginal cost \bar{c} .

Assumption 3. $D_{(1)}(p_{Bm}, \bar{c}, q_L) > 0$, where p_{Bm} is the lowest possible best response price the certified seller can set, which is given by $p_{Bm} = \min_{\{p_{(1)} \geq \bar{c}, b \in [q_L, q_H]\} \cup \{p_{(1)} \in [\underline{c}, \bar{c}], b = q_L\}} \operatorname{argmax}_{p_B} (p_B - c_B)(1 - F(\max\{(p_B - u_0)/q_B, (p_B - p_{(1)})/(q_B - b)\}))$.

Given this assumption and the continuity of $D_{(1)}(p_{Bm}, \bar{c}, q_L)$, there exists a price $c > \bar{c}$ such that $D_{(1)}(p_{Bm}, c, q_L) > 0$. This ensures that sellers can obtain positive demand even when setting prices above the marginal cost \bar{c} .

This assumption is more likely to be satisfied when consumers display a wide range of quality sensitivities $[\underline{\alpha}, \bar{\alpha}]$ and when the base utility u_0 is sufficiently high. For example, consider the scenario in which consumers' quality sensitivity is distributed over a broad interval with $\underline{\alpha} = 0$ and a sufficiently large $\bar{\alpha}$ (or a large c_B), and $u_0 \geq \bar{c}$. In our analysis, we provide a tractable case where $\alpha \sim \text{Unif}[0, \bar{\alpha}]$ with sufficiently large $\bar{\alpha}$ and $u_0 \geq \bar{c}$, which ensures that the assumption is automatically satisfied.

A.2 Proof of Proposition 1

Under Assumption 1, H must obtain positive profit in equilibrium. This is because there exists an $\epsilon > 0$ such that H can profitably deviate to $p_H = \bar{c} + \epsilon$ to obtain positive profit, as guaranteed by Assumption 3. That is, a possible equilibrium must have $D_{(1)}(p_B^*(p_H^*), p_H^*, b^{Q*}(p_H^*; (1))) > 0$ and $p_H^* > \bar{c}$.

Then, we demonstrate that $C_p = [\underline{c}, \bar{c}]$. Suppose, by way of contradiction, that $\exists c' \in [\underline{c}, \bar{c}]$ such that $c' \notin C_p$. In this case, L with cost c' is always ranked second among uncertified sellers, resulting in zero profit. The price of the product ranked first is always the price of H , namely $p_{(1)}^* = p_H^*$. The existence of equilibrium requires that L 's expected profit from mimicking H 's price is non-positive for c' , that is, $U_{(1)}(c', p_H^*, b^{Q*}(p_H^*; (1)); p_B^*(p_H^*)) = (p_H^* - c')D_{(1)}(p_B^*(p_H^*), p_H^*, b^{Q*}(p_H^*; (1))) \leq 0$. Consequently, we have $p_H^* \leq c' < \bar{c}$, leading to a contradiction. Therefore, we must have a pooling equilibrium where $C_p = [\underline{c}, \bar{c}]$. Within the pooling equilibrium, the quality-based ranking algorithm ranks H first with a probability of ω . For the first-ranked product ((1)) on the equilibrium path, the consumer's belief about its quality is given by Bayes' rule, $b^{Q*}(p_{(1)}^*; (1)) = \omega q_H + (1 - \omega)q_L = q_L + \omega \Delta_q$. By the definition of PBE, the certified seller's optimal price is given by $p_B^* = \operatorname{argmax}_{p_B \in \mathbb{R}_+} U_B(p_B, p_{(1)}^*, q_L + \omega \Delta_q)$. In this equilibrium, both L and H obtain positive payoffs.

We construct an equilibrium with any pooling price $p^* > \bar{c}$ satisfying $D_{(1)}(p_B^*(p^*), p^*, q_L + \omega \Delta_q) > 0$. For any deviation $p' \neq p^*$: If $p' < \bar{c}$, the off-path belief gives quality belief q_L . If $p' \geq \bar{c}$, the off-path belief makes the deviating seller faces a low expected ranking probability that yields lower expected payoff than in pooling. With such off-path beliefs, no seller has profitable deviation, proving the existence of PBEs with any $p^* > \bar{c}$ such that $D_{(1)}(p_B^*(p^*), p^*, q_L + \omega \Delta_q) > 0$.

A.3 Proof of Proposition 2

Step 1: $p_H^* \geq \bar{c}$. Suppose by way of contradiction that there exists an equilibrium with $p_H^* < \bar{c}$. By Assumption 3, the corresponding demand is positive. For any L with marginal cost $c \in (p_H^*, \bar{c})$, setting $p_L = c > p_H^*$ would avoid negative profit. Consequently, H is ranked first when L 's cost falls in (p_H^*, \bar{c}) , resulting in negative profit for H . This contradicts equilibrium as H can profitably deviate to $p_H^* = \bar{c}$. Therefore, $p_H^* \geq \bar{c}$ must hold in equilibrium.

Step 2: $p_L^*(c) < p_H^*$ for any $c \in [\underline{c}, \bar{c}]$. Under the price-based ranking algorithm, L 's optimal price with cost c is given by $p_L^*(c) = \operatorname{argmax}_{p \in \mathbb{R}_+} [1 - \Omega_2(p, p_H^*)] U_{(1)}(c, p, b^{P*}(p; (1)); p_B^*)$. By Assumption 3, L can obtain positive demand while setting a profitable price. Therefore, in equilibrium, L must obtain positive profit, which implies $p_L^*(c) \leq p_H^*$ for all $c \in [\underline{c}, \bar{c}]$.

We now prove $p_L^*(c) < p_H^*$ for all $c \in [\underline{c}, \bar{c}]$: Suppose by contradiction that $p_L^*(c) = p_H^*$ for some c , i.e., $C_p \neq \emptyset$. In this case, p_H^* is on the equilibrium path, and by Bayes' rule, the consumer's belief is $b^{P*}(p_H^*; (1)) = q_L + \Delta_q / (1 + \int_{C_p} dG(c))$. The expected payoff for any L with cost $c' \in C_p$ is $(p_H^* - c') D_{(1)}(p_B^*, p_H^*, q_L + \Delta_q / (1 + \int_{C_p} dG(c))) / 2$. We consider two cases:

1. If $D_{(1)}(p_B^*, p_H^*, q_L + \Delta_q / (1 + \int_{C_p} dG(c))) = 0$, then L with cost $c' \in C_p$ can profitably deviate to the price $\operatorname{argmax}_{p \in (c', p_H^*)} U_{(1)}(c', p, b^{P*}(p; (1)); p_B^*(p_H^*))$. By Assumption 3, this deviation yields positive profit when setting a price in $(c', \bar{c}]$.
2. If $D_{(1)}(p_B^*, p_H^*, q_L + \Delta_q / (1 + \int_{C_p} dG(c))) > 0$, by the continuity of $b^{P*}(\cdot; (1))$ (Assumption 2), for $\epsilon \rightarrow 0^+$, we have $\lim_{\epsilon \rightarrow 0^+} b^{P*}(p_H^* - \epsilon; (1)) = b^{P*}(p_H^*; (1)) = q_L + \Delta_q / (1 + \int_{C_p} dG(c))$. This continuity property ensures that the belief remains consistent when L deviates to a price infinitesimally below p_H^* . Thus, by deviating to $p_H^* - \epsilon$ ($\epsilon \rightarrow 0^+$), L with cost $c' \in C_p$ obtains:

$$\begin{aligned} & \lim_{\epsilon \rightarrow 0^+} (p_H^* - \epsilon - c') D \left(p_B^*, p_H^* - \epsilon, b^{P*}(p_H^* - \epsilon; (1)) \right) \\ &= (p_H^* - c') D \left(p_B^*, p_H^*, q_L + \frac{\Delta_q}{1 + \int_{C_p} dG(c)} \right) > \frac{1}{2} (p_H^* - c') D \left(p_B^*, p_H^*, q_L + \frac{\Delta_q}{1 + \int_{C_p} dG(c)} \right). \end{aligned}$$

Therefore, L with cost $c' \in C_p$ can profitably deviate to $p_H^* - \epsilon$ ($\epsilon \rightarrow 0^+$).

Thus, we must have $p_L^*(c) < p_H^*$ for all $c \in [\underline{c}, \bar{c}]$.

Step 3: $p_L^*(c) = \operatorname{argmax}_{p < p_H^*} U_{(1)}(c, p, q_L; p_B^*)$ for any $c \in [\underline{c}, \bar{c}]$. Suppose by contradiction that there exists $c' \in [\underline{c}, \bar{c}]$ such that $p_L^*(c') \neq \operatorname{argmax}_{p < p_H^*} U_{(1)}(c, p, q_L; p_B^*)$ in equilibrium. By PBE consistency, the consumer's belief is $b^{P*}(p_L^*(c'); (1)) = q_L$. The profit from deviating to the price $p' = \operatorname{argmax}_{p < p_H^*} U_{(1)}(c, p, q_L; p_B^*)$ is:

$$U_{(1)}(c, p', b^{P*}(p'; (1)); p_B^*) \geq U_{(1)}(c, p', q_L; p_B^*) = \max_{p < p_H^*} U_{(1)}(c, p, q_L; p_B^*) > U_{(1)}(c, p_L^*(c'), q_L; p_B^*).$$

Therefore, L with cost c' can profitably deviate to p' , leading to a contradiction.

Step 4: $p_H^* = \bar{c}$. We have already shown $p_H^* \geq \bar{c}$. Suppose by contradiction that $p_H^* > \bar{c}$. Since $p_L^*(c) = \operatorname{argmax}_{p < p_H^*} U_{(1)}(c, p, q_L; p_B^*) < p_H^*$ for all $c \in [\underline{c}, \bar{c}]$, H is always ranked second and gets zero profit. Thus, we must have $p_L^*(c) \leq \bar{c}$ for all $c \in [\underline{c}, \bar{c}]$; otherwise, H could profitably deviate to a price in $(\bar{c}, p_L^*(c))$ to secure the first position and obtain positive profit.

By Assumption 3, there exists a price greater than \bar{c} that L can set while obtaining a positive demand. Therefore, we have $\operatorname{argmax}_{p < p_H^*} U_{(1)}(\bar{c}, p, q_L; p_B^*) = \operatorname{argmax}_{p < p_H^*} (p - \bar{c}) D_{(1)}(p_B^*, p, q_L) > \bar{c}$. Then, by the continuity of $\operatorname{argmax}_{p < p_H^*} U_{(1)}(c, p, q_L; p_B^*)$ about c , there exists $c' < \bar{c}$ such that $\operatorname{argmax}_{p < p_H^*} U_{(1)}(c', p, q_L; p_B^*) > \bar{c}$. In this case, H can profitably deviate to any price in $(\bar{c}, \operatorname{argmax}_{p < p_H^*} U_{(1)}(c', p, q_L; p_B^*))$, leading to a contradiction. Therefore, we must have $p_H^* = \bar{c}$.

Conclusion In conclusion, we have shown that:

1. $p_L^*(c) = \operatorname{argmax}_{p < \bar{c}} U_{(1)}(c, p, q_L; p_B^*)$ for all $c \in [\underline{c}, \bar{c}]$
2. $p_H^* = \bar{c}$
3. The certified seller's price is given by:

$$\begin{aligned} p_B^* &= \operatorname{argmax}_{p_B \in \mathbb{R}_+} \int_{\underline{c}}^{\bar{c}} U_B(p_B, p_L^*(c), q_L) dG(c) = \operatorname{argmax}_{p_B \in \mathbb{R}_+} (p_B - c_B) \int_{\underline{c}}^{\bar{c}} D_B(p_B, p_L^*(c), q_L) dG(c) \\ &= \operatorname{argmax}_{p_B \in \mathbb{R}_+} (p_B - c_B) \left(1 - \int_{\underline{c}}^{\bar{c}} F \left(\max \left\{ \frac{p_B - u_0}{q_B}, \frac{p_B - p_L^*(c)}{\Delta_B + \Delta_q} \right\} \right) dG(c) \right). \end{aligned}$$

In this equilibrium, both L (for any cost) and B obtain positive payoffs, whereas H always receives zero payoff since it is consistently ranked second. Furthermore, such a PBE can arise under a wide range of consumer beliefs. For instance, the belief $b^{P*}(p; (1)) = q_L$ for all $p \in \mathbb{R}_+$ constitutes a feasible belief supporting the existence of this type of PBE.

A.4 Proof of Lemma 2

For low-quality sellers with cost c , the payoff change is $\Delta\Pi_L(c) = \max_{p < \bar{c}} U_{(1)}(c, p, q_L; p_B^{*P}) - (1 - \omega) U_{(1)}(c, p_{(1)}^*, q_L + \omega\Delta_q; p_B^*(p_{(1)}^*))$. For high-quality sellers, the payoff change is $\Delta\Pi_H = -\omega U_{(1)}(\bar{c}, p_{(1)}^*, q_L + \omega\Delta_q; p_B^*(p_{(1)}^*))$. We establish the boundary condition at the upper bound. Given $\omega > 1/2$, we have $\lim_{c \rightarrow \bar{c}} \Delta\Pi_L(c) = -(1 - \omega) U_{(1)}(\bar{c}, p_{(1)}^*, q_L + \omega\Delta_q; p_B^*(p_{(1)}^*)) \in (\Delta\Pi_H, 0)$. Since $\Delta\Pi_L(c)$ is continuous and $\lim_{c \rightarrow \bar{c}} \Delta\Pi_L(c) > \Delta\Pi_H$, there must exist a threshold $\tilde{c}_H \in [\underline{c}, \bar{c}]$ such that $\Delta\Pi_L(c) \in (\Delta\Pi_H, 0)$ for all $c > \tilde{c}_H$.

For sufficiently large ω , we have $\lim_{c \rightarrow \underline{c}} \Delta\Pi_L(c) = \max_{p < \bar{c}} U_{(1)}(c, p, q_L; p_B^{*P}) - (1 - \omega) U_{(1)}(c, p_{(1)}^*, q_L + \omega\Delta_q; p_B^*(p_{(1)}^*)) > 0$ and $\Delta\Pi_H = -\omega U_{(1)}(\bar{c}, p_{(1)}^*, q_L + \omega\Delta_q; p_B^*(p_{(1)}^*)) < \Delta\Pi_L(c)$ for all $c \in [\underline{c}, \bar{c}]$. Since $\lim_{c \rightarrow \bar{c}} \Delta\Pi_L(c) \in (\Delta\Pi_H, 0)$, by the intermediate value theorem, there exists $\tilde{c}_L \in [\underline{c}, \tilde{c}_H]$ such that $\Delta\Pi_L(c) > 0$ for all $c < \tilde{c}_L$.

B Estimation Results of the Consideration Set Model

The results, presented in Table 5, reveal that the mean quality signal for TRS products (0.076) is substantially higher than for NTRS products (0.029), confirming that the certification status correlates with higher platform-perceived quality.


Table 5: Estimation Results of the Consideration Set Model

	Before	After
Average number of listings on the site	21	28
Average number of TRS listings on the site	3	10
Mean of quality information (\hat{q}_j normalized to 0 to 1)	0.037	
Mean of quality information (\hat{q}_j normalized to 0 to 1) for TRS	0.076	
Mean of quality information (\hat{q}_j normalized to 0 to 1) for NTRS	0.029	
Estimated gamma (γ)	0.80	
		(0.14)

Note: Estimated standard errors are in parentheses.

C Profit Increase

The average profit change of sellers is reported in Figure 6 (a). The vertical axis in the figure is the sum of profits of all sellers after the redesign minus the sum of profits of all sellers before the redesign. We can see that higher consumer heterogeneity is related to an increase in profit after redesign. This is consistent with the predictions of our theory. However, this group of simulations is too conservative because it does not allow sellers who suffer a large loss to exit the market. This motivates us to run a new group of counterfactuals that exogenously adjust the market structure. Specifically, we remove 66 NTRS that lose the most in our counterfactual in subsection 5.4 and duplicate 14 TRS that gain the most in our counterfactual. Note that this practice is also very conservative, because here entrants are perfect substitutes for TRS with positive profit gains from the redesign. However, in reality, we should never expect the entrants to be exactly the same as the incumbents. Then, we rerun the counterfactuals with heterogeneous consumers and with the new market structure. The results are shown in Figure 6 (b). We can see that when the heterogeneity of consumers is large enough, the profit increase effect due to market segmentation can outweigh the profit decrease effect due to increased price competition under price-based ranking. However, this only happens in extreme cases.

**Figure 6:** Average Profit Change Increases with Consumer Heterogeneity

D Model Details and Formal Definitions

D.1 Perfect Bayesian Equilibrium Definition

This section provides the complete formal definition of Perfect Bayesian Equilibrium referenced in the main text, building upon the payoff functions and demand structure already defined.

Definition 1 (Perfect Bayesian Equilibrium (PBE)). *A pure-strategy perfect Bayesian equilibrium (PBE) consists of:*

1. *A pricing strategy profile* $(p_L^*(\cdot) : [\underline{c}, \bar{c}] \rightarrow \mathbb{R}_+; p_H^* \in \mathbb{R}_+; p_B^* \in \mathbb{R}_+)$
2. *Beliefs* $(b^{0*}(\cdot) : \mathbb{R}_+^2 \times \mathcal{S} \rightarrow [q_L, q_H]; b^{r*}(\cdot; (1)) : \mathbb{R}_+ \rightarrow [q_L, q_H])$

such that:

1. *(Consistency)* The beliefs $b^{0*}(\cdot)$ and $b^{r*}(\cdot; (1))$ are derived from priors and seller strategies via Bayes' rule on the equilibrium path.
2. *(Seller L and H's optimality)* The pricing strategies maximize their expected payoffs given seller B's strategy and the beliefs of consumers and the platform:

$$p_L^*(c_L) = \underset{p \in \mathbb{R}_+}{\operatorname{argmax}} [1 - \Omega_r(p, p_H^*; b^{0*}(\cdot))] U_{(1)}(c_L, p, b^{r*}(p; (1)); p_B^*)$$

$$p_H^* = \underset{p \in \mathbb{R}_+}{\operatorname{argmax}} \int_{\underline{c}}^{\bar{c}} \Omega_r(p_L^*(c), p; b^{0*}(\cdot)) U_{(1)}(\bar{c}, p, b^{r*}(p; (1)); p_B^*) dG(c)$$

3. *(Seller B's optimality)* The pricing strategy maximizes its expected payoff given the strategies of sellers L and H and the beliefs of consumers and the platform:

$$p_B^* = \underset{p_B \in \mathbb{R}_+}{\operatorname{argmax}} \mathbb{E}_{p_{(1)}^*} \left[U_B(p_B, p_{(1)}^*, b^{r*}(p_{(1)}^*; (1))) \right]$$

$$= \underset{p_B \in \mathbb{R}_+}{\operatorname{argmax}} \left[U_B(p_B, p_H^*, b^{r*}(p_{(1)}^*; (1))) \int_{\underline{c}}^{\bar{c}} \Omega_r(p_L^*(c), p_H^*; b^{0*}(\cdot)) dG(c) \right.$$

$$\left. + \int_{\underline{c}}^{\bar{c}} U_B(p_B, p_H^*, b^{r*}(p_L^*(c); (1))) [1 - \Omega_r(p_L^*(c), p_H^*; b^{0*}(\cdot))] dG(c) \right]$$

where $\mathbb{E}_{p_{(1)}^*}$ denotes the expectation over the equilibrium price of the first-ranked uncertified product.

D.2 Platform's Belief Formation

This section provides additional technical details on the platform's belief formation process described in the main text.

Given a PBE $(p_L^*(\cdot), p_H^*, p_B^*; b^{0*}(\cdot), b^{r*}(\cdot; (1)))$ and the equilibrium set $C_p := \{c \in [\underline{c}, \bar{c}] : p_L^*(c) = p_H^*\}$, the formal belief function is:

$$b^{0*}(p_L^*(c), p_H^*, s) = \begin{cases} s, & \text{if } c \in C_p \\ q_H, & \text{if } c \notin C_p \end{cases}$$

The corresponding ranking probabilities under quality-based ranking are:

$$\Omega_Q(p_L^*(c), p_H^*; b^{0*}(\cdot)) = \begin{cases} \omega(F_s), & \text{if } c \in C_p \\ 1, & \text{if } c \notin C_p \end{cases}$$

where $\omega(F_s) := 1 - F_s(q_L + \Delta_q/2) + f_s(q_L + \Delta_q/2)/2$ represents the platform's ranking accuracy when prices provide no information. For analytical simplicity, we use ω to denote this accuracy parameter throughout the main analysis.

E Detailed Welfare Analysis for Certified Sellers

This section provides the formal analysis supporting the ambiguous welfare effects on certified sellers discussed in the main text. Following the uniform distribution case in Section 3.4, the change in certified seller's market share can be decomposed as:

$$\begin{aligned} & \int_{\underline{c}}^{\bar{c}} D_B(p_B^{*P}, p_L^{*P}(c), q_L) dG(c) - D_B(p_B^*(p_{(1)}^*), p_{(1)}^*, q_L + \omega\Delta_q) \\ &= \frac{1}{2(\Delta_B + (1 - \omega)\Delta_q)(\bar{\alpha} - \underline{\alpha})} \left[\underbrace{\frac{(c_B - \int_{\underline{c}}^{\bar{c}} p_L^{*P}(c) dG(c))\omega\Delta_q}{\Delta_B + \Delta_q}}_{\text{Discrimination Effect}} - \underbrace{\left(p_{(1)}^* - \int_{\underline{c}}^{\bar{c}} p_L^{*P}(c) dG(c) \right)}_{\text{Competition Effect } (>0)} \right] \end{aligned}$$

The discrimination effect dominates the competition effect only when c_B substantially exceeds $p_{(1)}^*$, resulting in an increased market share for the certified seller following the transition from quality-based to price-based ranking algorithm. This occurs because a certified seller with high costs has relatively modest quality advantages compared to the first-ranked uncertified seller that is perceived as having average quality under the quality-based ranking algorithm, thereby constraining the certified seller to serve only consumers with extremely high quality sensitivity. However, under the price-based ranking algorithm, the certified seller's quality advantage becomes more pronounced since the first-ranked seller is consistently a low-quality seller, enabling the certified seller to capture a broader segment of quality-sensitive consumers. Conversely, if the certified seller's cost is not excessively high, its market share may decrease as the competition effect becomes the dominant factor.

The welfare effects on certified sellers depend on the relative strength of these market share effects combined with price changes. In the case of a uniform distribution, when the certified seller's marginal cost equals that of a high-quality uncertified seller ($c_B = \bar{c}$), the certified seller always experiences a decline in demand, as the competition effect outweighs the discrimination effect. Consequently, if the certified seller's price decreases after the transition ($p_B^{*P} < p_B^*(p_{(1)}^*)$), their expected payoff under the price-based ranking algorithm will be lower than under the quality-based ranking algorithm.

Conversely, when \bar{c} , $p_{(1)}^*$, and $\int_{\underline{c}}^{\bar{c}} p_L^{*P}(c) dG(c)$ are all approximately equal, the change in equilibrium demand becomes negligible. In this case, the equilibrium price rises as the discrimination

effect dominates the competition effect in determining price changes. Under such circumstances, the certified seller's payoff increases when the platform switches from the quality-based ranking algorithm to the price-based ranking algorithm.

F Detailed Empirical Context

A screenshot of eBay webpage before and after the redesign is shown in Figure 7. Before the redesign, the rank was generated by Best Match, a relevance-based ordering mechanism introduced in 2008. According to Dinerstein et al. (2018), during the period of time we studied, this relevance score was mainly a quality measurement, it did not take into account the prices of the products, and it was not personalized to tastes of different consumers. Therefore, this setting corresponds to ranking by quality ($r = Q$) in our analytical model. We select April 6, 2011 to May 18, 2011 as our sample before the redesign.

On May 19, 2011, eBay introduced a new two-stage product discovery process. In the first stage, consumers are guided to select among various product models. In the second stage, they are shown a newly-designed product page. On the top of the product page, the top-rated seller with the lowest listing price is shown in a "Buy Box". Below the Buy Box, there are two columns, one for auctions and the other for posted prices. The column for posted prices corresponds to ranking by price ($r = P$) in our analytical model. We do not study auctions in this research.

In one week period from June 27, 2011 to July 2, 2011, this two-stage product discovery process became the default for five categories, including cell phones, digital cameras, textbooks, video games, and video game systems. Consistent with Dinerstein et al. (2018), we drop the data in July to allow sellers and consumers time to respond to the change and reach a new equilibrium. Therefore, we consider August 1, 2011 to September 20, 2011 as our sample after the redesign. Following Dinerstein et al. (2018), we pool sessions with default ranking approach and sessions in which consumers change the default ranking approach together.


G Details on Consideration Set Model

The search results may include some irrelevant listings. Therefore, Dinerstein et al. (2018) manually classify listings to "targeted listings" and "non-targeted listings", with "targeted listings" defined as new *Halo Reach* items listed either with a fixed price or as an auction with a Buy-It-Now price, and "non-targeted listings" defined as other listings, including used items and items that are not *Halo Reach* itself. We directly use their classifications.

To capture the post-redesign "Buy Box" feature, we reserve one position in the consideration set for a TRS product, following the approach in Dinerstein et al. (2018). The model first draws one product j_i^0 from the set of TRS sellers, $\mathcal{J}_i^{TRS,J}$, and then selects the remaining targeted products without replacement from the set $\mathcal{J}_i^J \setminus j_i^0$.

References

Amazon (2024). How seller certifications can help your business stand out. Retrieved from Amazon Seller Central Blog.



Categories: Video Games & Consoles (106,805), Video Game Consoles (3,278), Video Games (74,010), Video Game Accessories (20,497), Wholesale Lots (1,326), Replacement Parts & Tools (3,320), More. See all categories.

Format: Auction (568), Buy It Now (1). Platform: Sony PlayStation 1, Sony PlayStation 2, Sony PlayStation 3, Sony PlayStation Vita, Sony PSP. Region Code: NTSC-U/C (US/Canada), PAL, NTSC-J (Japan), NTSC-C (China), Not Specified. Bundled Items: Extra Cable(s), Extra Controller(s), Games(s), Hard Drive, Memory Card(s). Condition: New, Used, Not Specified. Price: \$0.01 - \$18,879.

143,517 active listings | sold listings | completed listings. Sort: Best Match. View: List.

Recently listed: 120GB SONY PS3 SLIM (WORKS GREAT) SYSTEM ONLY. \$139.99. 29d 23h left. 7/6, 9AM. Buy It Now. Free shipping.

SONY "BATMAN 2" PS3 SYSTEM LIMITED EDITION IN VERY GOOD CONDITION.... \$150.00. 1d left. Today 11:03AM. 23 bids. Free shipping.

SONY PS3 500GB God of War Ascension Legacy Bundle. \$274.99. 23d 1h left. 6/29, 12PM. Buy It Now. Free shipping.

SONY PS3 250GB MODEL CECH-4001B -- NO RESERVE SONY PLAYSTATION 3 SUPER SLIM. \$126.50. 7d left. Today 11:09AM. 24 bids.

Sony PlayStation 3 Slim (Latest Model)- 160 GB Charcoal Black Console (NTSC) \$199.99. 6d 23h left. +1 0. Like. See all listings (178).

Options: Charcoal Black, 160 GB (\$2...). Filters: New (\$267.68), Refurbished (\$209.99), New: other (\$225.00), Used (\$209.99), Bundles.

Best deal from a top-rated seller: *NEW* SONY PLAYSTATION 3 PS3 160GB Slim CONSOLE SYSTEM. Condition: New. dreambean_soft (2328) 99.7%   Location: Tampa, FL, USA. Returns: Accepted. Expedited Shipping available. eBay Buyer Protection. Learn more.

Your price: \$267.68. Free shipping. Buy It Now. Add to cart. Add Warranty from \$41.99.

Auction Time: ending soonest. Buy it Now. Price + Shipping: lowest first.

Sony PlayStation 3 Slim (Latest Model)- 160 G... jakandsig 1d 6h 46m \$218.50 0% (0) (30 bids) +\$15.00 Condition: New	Sony PlayStation 3 Slim (Latest Model)- 160 G... selle89 93.9% (28) \$235.00 93.9% (28) Condition: New Free shipping
BRAND NEW Sony PS3 SLIM 160 GB Console - w/ e... la2chicago 1d 8h 49m \$249.99 98% (54) (0 bids) +\$42.55 Condition: New	Sony PlayStation 3 Slim (Latest Model)- 160 G... redned07306911 100% (19) \$235.37 100% (19) Condition: New +\$15.00
ps3 with two games kendallgirl27 1d 16h 57m \$400.00 0% (-1) (0 bids) +\$15.00 Condition: New	** NEW** Sony PlayStation 3 Slim Latest Model... scuba diving minh 100% (1,721) \$255.99 100% (1,721) Condition: New Free shipping

Figure 7: eBay's Redesign. The upper panel shows the search results page before the redesign (Best Match), while the lower panel shows the search results page after the redesign (Buy Box and Rank by Price). These images are adapted from Figure 4 of Dinerstein et al. (2018).

Armstrong, M. and Zhou, J. (2022). Consumer information and the limits to competition. *American Economic Review*, 112(2):534–577.

Bagwell, K. and Riordan, M. H. (1991). High and declining prices signal product quality. *American Economic Review*, pages 224–239.

Belleflamme, P., Lam, W. M. W., and Vergote, W. (2020). Competitive imperfect price discrimination and market power. *Marketing Science*, 39(5):996–1015.

Calvano, E., Calzolari, G., Denicolo, V., and Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10):3267–3297.

Chen, Y., Du, J., and Lei, Y. (2025). The interactions of customer reviews and price and their dual roles in conveying quality information. *Marketing Science*, 44(1):155–175.

Chevalier, J. A. and Mayzlin, D. (2006). The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research*, 43(3):345–354.

Choe, C., King, S., and Matsushima, N. (2018). Pricing with cookies: Behavior-based price discrimination and spatial competition. *Management Science*, 64(12):5669–5687.

Compiani, G., Lewis, G., Peng, S., and Wang, P. (2024). Online search and optimal product rankings: An empirical framework. *Marketing Science*, 43(3):615–636.

Derakhshan, M., Golrezaei, N., Manshadi, V., and Mirrokni, V. (2022). Product ranking on online platforms. *Management Science*, 68(6):4024–4041.

Dinerstein, M., Einav, L., Levin, J., and Sundaresan, N. (2018). Consumer price search and platform design in internet commerce. *American Economic Review*, 108(7):1820–1859.

Do, J. and Miklos-Thal, J. (2025). Price discrimination against multi-clouders. *Available at SSRN*.

eBay (2025). Top rated seller program. Retrieved from eBay Seller Center.

Erdem, T., Keane, M. P., and Sun, B. (2008). A dynamic model of brand choice when price and advertising signal product quality. *Marketing Science*, 27(6):1111–1125.

Farronato, C. and Zervas, G. (2022). Consumer reviews and regulation: Evidence from nyc restaurants. *National Bureau of Economic Research Working Paper*.

Feng, X. F., Liu, X., Zhang, S., and Srinivasan, K. (2024). Sustainability and competition on amazon.

Gandhi, A., Hollenbeck, B., and Li, Z. (2024). Misinformation and mistrust: The equilibrium effects of fake reviews on amazon.com.

Gardete, P. M. (2013). Cheap-talk advertising and misrepresentation in vertically differentiated markets. *Marketing Science*, 32(4):609–621.

Ghose, A., Ipeirotis, P. G., and Li, B. (2012). Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content. *Marketing Science*, 31(3):493–520.

Ghose, A., Ipeirotis, P. G., and Li, B. (2014). Examining the impact of ranking on consumer behavior and search engine revenue. *Management Science*, 60(7):1632–1654.

Hui, X., Jin, G. Z., and Liu, M. (2025). Designing quality certificates: Insights from ebay. *Journal of Marketing Research*, 62(1):40–60.

Hui, X., Saeedi, M., Shen, Z., and Sundaresan, N. (2016). Reputation and regulations: Evidence from ebay. *Management Science*, 62(12):3604–3616.

Hui, X., Saeedi, M., Spagnolo, G., and Tadelis, S. (2023). Raising the bar: Certification thresholds and market outcomes. *American Economic Journal: Microeconomics*, 15(2):599–626.

Jerath, K. and Ning, Z. E. (2025). Inclusive product design and consumer inference.

Jin, G. Z. and Leslie, P. (2003). The effect of information on product quality: Evidence from restaurant hygiene grade cards. *The Quarterly Journal of Economics*, 118(2):409–451.

Johnson, J. P., Rhodes, A., and Wildenbeest, M. (2023). Platform design when sellers use pricing algorithms. *Econometrica*, 91(5):1841–1879.

Lambrecht, A. and Tucker, C. (2019). Algorithmic bias? an empirical study of apparent gender-based discrimination in the display of stem career ads. *Management Science*, 65(7):2966–2981.

Levy, R. (2021). Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review*, 111(3):831–870.

Lin, S. (2020). Two-sided price discrimination by media platforms. *Marketing Science*, 39(2):317–338.

Manzini, P. and Mariotti, M. (2014). Stochastic choice and consideration sets. *Econometrica*, 82(3):1153–1176.

Miklos-Thal, J. and Shaffer, G. (2021a). Input price discrimination by resale market. *The RAND Journal of Economics*, 52(4):727–757.

Miklos-Thal, J. and Shaffer, G. (2021b). Pass-through as an economic tool: on exogenous competition, social incidence, and price discrimination. *Journal of Political Economy*, 129(1):323–335.

Miklos-Thal, J. and Tucker, C. (2019). Collusion by algorithm: Does better demand prediction facilitate coordination between sellers? *Management Science*, 65(4):1552–1561.

Miklos-Thal, J. and Zhang, J. (2013). (de) marketing to manage consumer quality inferences. *Journal of Marketing Research*, 50(1):55–69.

Milgrom, P. and Roberts, J. (1986). Price and advertising signals of product quality. *Journal of Political Economy*, 94(4):796–821.

Mussa, M. and Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, 18(2):301–317.

Pakes, A. (2010). Alternative models for moment inequalities. *Econometrica*, 78(6):1783–1822.

Rao, A. R. and Monroe, K. B. (1989). The effect of price, brand name, and store name on buyers' perceptions of product quality: An integrative review. *Journal of Marketing Research*, 26(3):351–357.

Saeedi, M. (2019). Reputation and adverse selection: Theory and evidence from ebay. *The RAND Journal of Economics*, 50(4):822–853.

Shi, Z., Srinivasan, K., and Zhang, K. (2023). Design of platform reputation systems: Optimal information disclosure. *Marketing Science*, 42(3):500–520.

Spence, M. (1978). Job market signaling. In *Uncertainty in economics*, pages 281–306. Elsevier.

Sun, M. (2012). How does the variance of product ratings matter? *Management Science*, 58(4):696–707.

Ursu, R. M. (2018). The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science*, 37(4):530–552.

Ursu, R. M. and Dzyabura, D. (2020). Retailers' product location problem with consumer search. *Quantitative Marketing and Economics*, 18:125–154.

Wang, Q., Huang, Y., Jasin, S., and Singh, P. V. (2023). Algorithmic transparency with strategic users. *Management Science*, 69(4):2297–2317.

Wu, C., Che, H., Chan, T. Y., and Lu, X. (2015). The economic value of online reviews. *Marketing Science*, 34(5):739–754.

Wu, R., Huang, Y., and Li, N. (2023). Platform information design and competitive price targeting. *Available at SSRN 4619584*.

Zhang, S., Mehta, N., Singh, P. V., and Srinivasan, K. (2021). Frontiers: Can an artificial intelligence algorithm mitigate racial economic inequality? an analysis in the context of airbnb. *Marketing Science*, 40(5):813–820.