Scarcity by Design: How Recommendation Algorithms Impact Prices, Participation, and Customer Welfare

Ömer Sarıtaç
Lee Kong Chian School of Business, Singapore Management University

Nur Kaynar
Johnson Graduate School of Management, Cornell University

Özge Şahin
Johns Hopkins Carey Business School

Recommendation systems are central to digital platforms, but their design often focuses on maximizing immediate match quality. This myopic view overlooks the long-term equilibrium effects driven by strategic responses from market participants. We develop an infinite-horizon equilibrium model of a two-sided marketplace that incorporates strategic sequential search by customers, strategic pricing by suppliers, and endogenous entry from both, to characterize the platform's long-term revenue-maximizing recommendation strategy. We find that maximizing immediate recommendation quality by always recommending the best-fit suppliers can be suboptimal; instead, it may be optimal to concentrate a portion of the recommendations on the very worst supplier types. This strategy, which we term "Scarcity by Design," exploits strategic feedback loops. We identify conditions (high customer patience, sufficient market thickness) under which this strategy is profitable, noting it typically benefits the platform and suppliers at the expense of customer surplus. Leveraging the structure of the optimal policy, we illustrate the revenue benefit of the optimal policy against benchmarks widely used in practice. We then extend the analysis to personalized recommendations across heterogeneous customer segments by leveraging a structural property that renders the high-dimensional optimization problem tractable. We decompose the Value of Information (VoI) from personalization into an efficiency gain (Matching Gain) and a strategic feedback gain (Equilibrium Price Manipulation), which reveals the underlying mechanism and the role customer patience plays in mediating these two gains. While personalization always increases platform and supplier revenues, it may increase or decrease the aggregate customer welfare. Interestingly, we demonstrate that personalization can create asymmetric distributional welfare effects, often harming vulnerable customer segments. This suggests a role for regulatory oversight to ensure fairness.

Key words: Recommendation, Pricing, Market Design, Stationary Equilibrium.

1. Introduction

Labor and freelance platforms (e.g., Toptal, TaskRabbit, Thumbtack, Angi) rely heavily on recommendation systems to connect customers with providers. These systems create substantial economic value; for example, algorithmic recommendations have been shown to increase job vacancy fill rates by 20% (Horton 2017). By directing customer attention, recommendation systems shape engagement and hiring decisions and play a central role in overall platform performance. As the variety of available services and providers continues to expand, steering users toward the most relevant options becomes even more critical. Platforms, therefore, invest heavily in advanced recommendation systems to generate precise matches (MacKenzie et al. 2013). For instance, Toptal combines AI-driven algorithms with a rigorous human screening process to recommend freelancers who match a client's project needs (Fuller et al. 2020). TaskRabbit not only recommends "Taskers" tailored to specific tasks but also reinforces these suggestions through targeted notifications, which actively guide customer choices at the moment of decision (Kim et al. 2016). Empirical research consistently shows that such high-quality recommendations improve immediate customer outcomes, such as click-through and conversion rates (Ursu 2018a, Pei et al. 2019).

While maximizing immediate outcomes appears beneficial, it can trigger strategic responses that lead to unintended consequences over time. For instance, customers may adjust their search behavior—how long they search or how selective they are. Indeed, recommendations can alter how customers perceive the options they encounter, sometimes making them more willing to accept offers they might otherwise reject (Zhang and Bockstedt 2016, Dean et al. 2025). Suppliers, in turn, adjust their pricing strategies or participation based on their exposure and the level of competition created by the recommendations. For the platform, these reactions by the market participant ultimately determine customer and supplier engagement, commission revenue, and the sustainability of the marketplace. The long-term equilibrium effects of recommendation design, and how they interact with strategic behavior, are not well understood. Therefore, moving beyond short-term metrics and understanding these long-term interactions is critical for optimal design.

This paper formally analyzes how recommendations influence these long-term strategic interactions within a two-sided marketplace. Specifically, we ask: How should platforms optimally steer customers with heterogeneous preferences toward suppliers of varying attractiveness? What are the mechanisms driving the optimal strategy, and what are its implications for platform revenue, social welfare, and the value of personalized recommendations? To address these questions, we develop a dynamic, infinite-horizon model of a two-sided market that explicitly incorporates endogenous entry, strategic pricing by suppliers, and strategic sequential search by customers. This framework allows us to characterize the platform's revenue-maximizing recommendation strategy in a stationary equilibrium, capturing the long-term feedback loops often overlooked in static or myopic models. Our results provide novel insights into when maximizing immediate match quality aligns with long-term revenue goals and when strategically worsening recommendations may yield superior outcomes.

1.1. Model and Results

Model. We study an infinite-horizon marketplace mediated by a central platform, consisting of two strategic sides: suppliers and customers. Agents on each side aim to maximize their expected long-term payoffs based on private costs (for suppliers) and valuations (for customers). Each period, potential customers arrive with heterogeneous preferences and private valuations drawn from a known distribution, while potential suppliers enter with privately observed product/service-customer fit quality. Entry decisions for both sides are endogenous and depend on expected profitability. Upon entry, the platform recommends active suppliers to customers based on a probability distribution over supplier types.

Customers are exposed to suppliers sequentially. After observing a supplier's price and fit quality, customers decide to either purchase immediately or wait, balancing urgency against potentially better future offers. Customers who do not purchase exit with some probability each period, reflecting impatience and the opportunity cost of time. Unlike standard choice models, which often assume customers know their valuations for all options ex ante, sequential search reflects settings common in labor and service markets where customers must inspect suppliers (e.g., interview a consultant, evaluate a B2B provider), balancing the returns from continued search against their outside option. Suppliers remain active until their product is sold. They may discount future revenues, setting prices to trade off lower prices for quicker sales against higher prices for greater profits per transaction.

The platform's recommendation system directly influences supplier exposure, entry, and pricing strategies, and indirectly affects customer purchasing decisions. The platform collects a commission fee on each transaction and selects the recommender to maximize steady-state revenue. Importantly, the platform fully internalizes the long-term implications of its recommendation policy, including equilibrium entry decisions by both sides. This contrasts with models focused on short-term objectives that typically prioritize immediate recommendation outcome. We focus on stationary equilibria, where agent strategies are best responses to prevailing market conditions—including the number of active agents, their strategic pricing and purchase-delay decisions, and their entry behaviors.

Results. Our main contribution is characterizing the platform's revenue-maximizing recommendation strategy in this dynamic equilibrium setting and analyzing its implications for revenue, welfare, and personalization.

To this end, we first derive several key structural results that transform this analytically challenging model into a tractable analytical framework. We establish the existence and uniqueness of the stationary equilibrium (Theorem 1), which ensures that comparisons across recommendation policies are well-posed. We then characterize the equilibrium strategies. We show that customer purchase thresholds and willingness-to-pay are strictly increasing in valuations (Proposition 1). On the supplier side, we show that posted prices increase with fit quality, while, interestingly, the *effective price* faced by customers—posted price plus misfit disutility—rises with the misfit level (Proposition 2). Furthermore, we prove that the equilibrium distribution of customer willingness-to-pay preserves the shape of the arrival distribution of valuations (Theorem 2), a powerful result that simplifies the analysis.

Building on these foundational results, we investigate the optimal recommendation and its broader implications, starting with a single customer segment to isolate the fundamental tradeoffs. The central insight, formalized in Theorem 3, is that a revenue-maximizing platform does not always benefit from steering customers exclusively toward their best-matched suppliers. Instead, the optimal recommendation policy concentrates exposure on the best and the worst supplier types while excluding the middle of the spectrum. This result overturns the conventional wisdom that maximizing immediate recommendation outcome invariably increases platform performance.

We then unpack the mechanism behind this counterintuitive result. Propositions 3 and 4 demonstrate that the platform's optimal strategy exploits the feedback loop in supplier and customer strategies. When customers are occasionally shown less-preferred suppliers, they anticipate a riskier search environment, which lowers their acceptance thresholds and makes them less selective. Simultaneously, limiting exposure to the most-preferred suppliers diminishes direct competition. Together, these forces allow suppliers to post higher prices. Thus, strategically introducing scarcity in customers' search environment elevates prices and enhances platform profits, provided that customers are sufficiently patient and the supply is sufficiently thick (Proposition 5).

We also study the welfare implications of implementing the revenue-maximizing recommender. We show that steering traffic toward niche suppliers generally reduces customer surplus while increasing supplier profits whenever platform revenue rises (Propositions 6 and 7), highlighting a fundamental tradeoff. Incentive alignment between customers and the platform occurs only when customers are relatively impatient or when the market has limited supplier capacity (Corollary 1).

Finally, Section 6 extends the analysis to a setting with an arbitrary number of heterogeneous customer segments. This section studies the value of personalization of the recommender by segment when customers vary in preferences, arrival rates, patience, and willingness-to-pay. Theorem 4 generalizes the structure of the optimal recommendation, showing that it continues to concentrate exposure at the extremes. Key to Theorem 4 is a property we term the Segmental-Exclusivity (SE). This property dictates that for any type of supplier, the platform's optimal strategy is to assign that type's entire available exposure capacity to just one customer segment. In essence, the optimal allocation becomes locally exclusive. This finding is powerful because it dramatically simplifies a complex, high-dimensional allocation problem. Instead of a single, highly coupled decision, the problem elegantly collapses into a collection of simpler, one-dimensional choices, one for each supplier type. Leveraging this key property and the structure of optimal recommendations, we decompose the Value of Information (VoI) from personalization into an efficiency gain (due to better recommendations) and a strategic feedback gain (due to changes in equilibrium prices) (Proposition 8). While personalization always increases platform revenue, its impact on customer surplus is ambiguous due to the strategic feedback effect. Our numerical analysis illustrates that personalization allows platforms to tailor diversion based on customer segment, which can lead to substantial distributional welfare effects. For instance, it can disadvantage customers with lower valuations[1] or higher patience (bargain hunting).

2. Literature

There is a substantial body of empirical research demonstrating that the way products are recommended and ranked shapes customer behavior in various ways. Recommendations guide how customers explore and navigate online marketplaces: the way products are ranked and recommended shapes which items users view next and how long they continue browsing, thereby influencing overall engagement (De los Santos and Koulayev 2017). They also affect click and purchase behavior: higher-ranked or frequently recommended items draw more user attention and clicks, though the extent to which these interactions lead to purchases depends on factors such as product quality, customer trust, and search costs (Ghose et al. 2014, Ursu 2018b, Lee and Musolff 2025). Finally, recommendations impact customers' perceptions of value and willingness-to-pay: empirical findings reveal that recommendations can recalibrate how much customers are willing to pay for products, altering their perceived value independently of direct product experience (Adomavicius et al. 2018).

Building on these empirical findings, theoretical studies have also examined how recommendations influence customer behavior. One stream models recommendations as assortment or ranking optimization problems under exogenous attention, explicitly incorporating how product position and multi-page framing affect choices. In these frameworks, customers inspect a random subset of products and choose according to random-utility or multinomial-logit models (Gallego et al. 2020, Lei et al. 2022). The platform then optimizes revenue or engagement by selecting an ordering or framing: page-based framing with provable guarantees under MNL (Abeliuk et al. 2016, Gallego et al. 2020); sequential ranking formulated as submodular maximization (Asadpour et al. 2023); learning-to-rank policies that directly target engagement from clickstream data (Ferreira et al. 2022); and, in operations settings, Lei et al. (2022) explicitly coordinate display, pricing, and order fulfillment under the MNL model. This modeling approach captures how ranking and page placement shape demand and engagement but abstracts from strategic interactions on the supply side and from the equilibrium feedback between recommendations, prices, and customer search.

Recent studies have considered *strategic supplier decisions* from multiple angles. Some papers treat entry/participation as endogenous and use rankings as an information–design tool: platforms

[1] Although we do not directly model socioeconomic status (SES); in many markets, lower observed valuations can reflect tighter budgets and may correlate with lower SES. Hence, our findings suggest that personalization may lead to exploitation of customers with lower SES.

commit to personalized orders so that customers optimally follow the shown path, which sustains niche sellers and balances the long tail; prices remain exogenous and seller strategy operates through the participation margin (Papanastasiou and Sekar 2024). A distinct stream of work on curated two-sided matching studies how display policies allocate attention and determine which users compete for exposure. User 'like' behaviors are probabilistic rather than price-based, and the platform balances present versus future exposure across sequential and same-period matching designs, achieving constant-factor guarantees for practical heuristics (Rios and Torrico 2023). Others study *price competition for algorithmic exposure* in one-shot environments: sellers choose prices to win recommendation slots, the platform applies a recommendation rule, and customers take the induced recommendation as given Zhou and Zou (2023). They show that greater personalization can have a non-monotonic effect on equilibrium prices—initially intensifying competition and lowering prices, but eventually leading to higher prices as sellers exploit better targeting. (Shi 2024) also studies recommendation-driven pricing under capacity constraints but assumes myopic customer behavior, showing that the welfare-maximizing policy ranks options by expected value net of price and can be implemented through impression probabilities. Prices adjust endogenously as the recommendation policy changes which providers are shown together and to whom—greater overlap intensifies competition and lowers prices. Relative to these findings, our key insight is that deliberately inducing scarcity in recommendations reduces the incentive to continue searching, softens competition among similar sellers, and raises prices and platform revenue. [2] This follows from the fact that our model internalizes the feedback between the recommendation policy, sellers' pricing and participation, and customers' stopping decisions in a stationary equilibrium, with the platform optimizing the reccomendation policy knowing it shifts both competitive intensity and the continuation value of search.

3. Model

We study a marketplace mediated by a central platform over an infinite discrete-time horizon. We refer to the two sides of the market as 'supply" and 'demand", and the agents associated with each side as 'suppliers" and 'customers", respectively. Suppliers and customers make strategic decisions to maximize their expected long-term payoffs based on their private costs and valuations, given the platform recommendation policy.

At the beginning of each period, a continuum of potential customers arrives, each endowed with a private valuation for the product drawn from a known distribution. Simultaneously, a mass of

[2] The role of randomization in our exposure policy echoes Balseiro et al. (2024): when agents are not perfectly optimizing, small behavioral slack can justify randomized designs; although their environment is single-seller and ours is a market equilibrium with pricing and stopping, both results point to randomized policies as first-order tools.

potential suppliers enters, each characterized by a privately observed quality level reflecting how closely their product matches customers' preferences. Customers and suppliers decide independently and endogenously whether to enter the platform based on expected profitability.

Among customers who enter, each is directed by the platform's recommendation algorithm to inspect a supplier drawn according to a probability distribution over active suppliers. This reflects a sequential search process, where customers evaluate options one at a time. While parallel browsing (often modeled via discrete choice models) is common in settings with standardized products and negligible evaluation costs (e.g., online retail), sequential search is more appropriate when evaluation is costly. In labor and service markets (e.g., hiring a consultant, evaluating a B2B provider), assessing fit requires significant time and cognitive investment, making simultaneous evaluation impractical or prohibitively costly. After observing the recommended supplier's posted price and quality (equivalently, misfit disutility), each customer decides either to accept the offer (generating an immediate transaction) or to remain unmatched for another period. Customers who do not purchase leave the platform with some probability at the end of each period, reflecting impatience and the opportunity cost.[3] Thus, customers strategically balance between accepting an available offer quickly and holding out for a better match and lower-price suppliers in the future. Suppliers, on the other hand, remain active in the market across periods until they successfully sell a product, after which they exit immediately. Each supplier with an unsold product stays in the market for future periods, but discounts expected future revenues using a per-period discount rate. This reflects the fact that, under discounting, immediate sales are more valuable than delayed ones. Thus, suppliers face a trade-off between setting lower prices to achieve a faster sale versus higher prices to maximize profit per transaction. Importantly, the platform's choice of recommendation algorithm affects supplier prices through two paths: (i) directly, by determining how often each supplier is exposed to potential customers, and (ii) indirectly, by shaping customers' purchase decision threshold upon inspecting products.

The platform charges suppliers an ad-valorem (commission) fee proportional to each successful transaction's value. The platform's objective is to choose its recommendation distribution optimally to maximize its total fee revenues. We analyze the system in steady state, focusing on stationary equilibria in which the inflows and outflows of customers and suppliers are balanced, and each agent's strategy represents a best response given aggregate market conditions (Hopenhayn 1992). These aggregate conditions include the total number of active customers (C) and suppliers (S), as

[3] We model the departure probability (formally μ_j in Section 3.1) as exogenous. This allows us to present a clean characterization of the unique steady state and WTP distributions in our dynamic equilibrium problem that underpin the analysis of the platform's optimization problem. In Appendix K, we extend our model and allow the departure decision to be endogenous. We show that the structural properties of the equilibrium is preserved, and our main qualitative insights are robust to this generalization.

well as the distributions governing the strategies of suppliers and customers that also determine their entry, as made formal in Section 3.4.

3.1. Customers' Problem

We consider a general setting with J customer segments indexed by $j \in \{1,\dots,J\}$. In each period, a mass λ_j of segment-j customers arrives. Each such customer has a valuation v drawn from the (segment-specific) uniform distribution on $[\underline{v}_j, \bar{v}_j]$ where $0 \le \underline{v}_j < \bar{v}_j \le 1$. The platform's recommendation profile is $\mathbf{D} := (D_j)_{j=1}^{J}$, where each D_j is a probability distribution over supplier types $\delta \in [0,1]$ (with $\int_0^1 D_j(\delta)\, d\delta = 1$). A supplier type δ is drawn from D_j to determine which supplier is shown to a segment-j customer. Supplier type δ indicates how well the supplier's product fits the preferences of segment j customers. The customer observes the supplier's posted price $p(\delta;\mathbf{D})$ (which is determined in equilibrium and may depend on the full recommendation profile $\mathbf{D}$) and a segment-specific misfit disutility $d_j(\delta) \in [0,1]$, where we assume a linear form $d_j(\delta) = a_j\delta + b_j$ with $a_j \neq 0$. The realized net utility from accepting an offer from a supplier of type δ is therefore $u_{c,j}(v,\delta) = v - p(\delta;\mathbf{D}) - d_j(\delta)$.

After observing this utility, the customer decides whether to purchase the inspected product or continue searching[4], comparing the immediate payoff to the continuation value (reward-to-go) from waiting for future opportunities. Customers who do not purchase exit the market at the end of the period with probability μ_j, reflecting limited patience and the implicit cost of waiting. With probability $1-\mu_j$, they remain active and are recommended another supplier in the next period according to D_j.

Given this setting, the customer's best response can be formulated as a dynamic program. Since each individual customer is infinitesimal, their actions do not affect the aggregate steady-state. It follows that the continuation value depends on the customer's own valuation v and the full recommendation profile $\mathbf{D}$ (through prices and market conditions). We represent this by the threshold $\theta_j(v,\mathbf{D})$. This threshold structure implies a simple cutoff rule: customers accept an offer if their realized utility exceeds $\theta_j(v,\mathbf{D})$. Hence, the strategy space can be reduced without loss of generality to valuation-dependent thresholds. Formally, given a threshold $\theta_j(v,\mathbf{D})$ and recommendation distribution D_j, the expected lifetime surplus $\pi_{c,j}(v,\theta_j(v,\mathbf{D}))$ consists of two components. First, the *immediate surplus*, realized if utility exceeds the threshold. Second, if the offer is rejected, the *continuation surplus*, discounted by the probability $(1-\mu_j)$ that the customer remains active in the market. Accordingly, the Bellman equation is

$$\begin{aligned}\pi_{c,j}(v,\theta_j(v,\mathbf{D})) = \mathbb{E}_{\delta\sim D_j}\Big[&(v-d_j(\delta)-p(\delta;\mathbf{D}))\mathbf{1}\big(v-d_j(\delta)-p(\delta;\mathbf{D}) \ge \theta_j(v,\mathbf{D})\big)\Big]\\ &+(1-\mu_j)\Pr_{\delta\sim D_j}\big(v-d_j(\delta)-p(\delta;\mathbf{D}) < \theta_j(v,\mathbf{D})\big)\,\pi_{c,j}(v,\theta_j(v,\mathbf{D})),\end{aligned}$$

[4] One can trivially show that allowing recall of previously examined options does not alter any of our results.

which can be written explicitly as

$$\pi_{c,j}(v,\theta_j(v,\mathbf{D})) = \frac{\mathbb{E}_{\delta\sim D_j}\left[(v-d_j(\delta)-p(\delta;\mathbf{D}))\,\mathbf{1}\big(v-d_j(\delta)-p(\delta;\mathbf{D})\geq\theta_j(v,\mathbf{D})\big)\right]}{\mu_j+(1-\mu_j)\,\Pr_{\delta\sim D_j}\big(v-d_j(\delta)-p(\delta;\mathbf{D})\geq\theta_j(v,\mathbf{D})\big)}.$$

Thus, each customer of valuation v chooses a threshold $\theta_j(v,\mathbf{D})$ to maximize this expected lifetime surplus and enters the market only if this maximal surplus is non-negative. Aggregating across all customer valuations, we define $\mathcal{V}_j=\{v\in[0,1]:\pi_{c,j}(v,\theta_j(v,\mathbf{D}))\geq 0\}$ as the set of customers of segment j who enter the market.

3.2. Suppliers' Problem

At the beginning of each period, a mass m of suppliers arrive. Upon arrival, suppliers decide whether to enter the market and, if they do, which price to post. Each supplier has a known fit type $\delta\in[0,1]$, distributed according to F_Δ. Suppliers also observe the customer acceptance thresholds $\{\theta_j(\cdot,\mathbf{D})\}_{j\in\{1,\dots,J\}}$, which are induced by the platform's recommendation profile $\mathbf{D}:=(D_j)_{j\in\{1,\dots,J\}}$. Let $s(\delta,\mathbf{D})$ denote the steady–state *supplier density* by type, and let $C_j(\mathbf{D})$ denote the steady–state mass of active customers in segment j given $\mathbf{D}$. Then, the per–period exposure probability that a type–δ listing is shown (to any segment) is

$$\psi(\delta,\mathbf{D})=\min\left\{1,\ \frac{\sum_{j\in\{1,\dots,J\}}C_j(\mathbf{D})\,D_j(\delta)}{s(\delta,\mathbf{D})}\right\}.$$

Note that the term $\sum_{j\in\{1,\dots,J\}}C_j(\mathbf{D})\,D_j(\delta)$ is the total number of per–period display slots that the policy $\mathbf{D}$ allocates to type–δ suppliers across segments $j\in\{1,\dots,J\}$. A higher supplier density $s(\delta,\mathbf{D})$ lowers the probability that any individual type–δ listing is shown, since a fixed measure of display slots must be divided across more suppliers.

Similarly, let $\psi_j(\delta,\mathbf{D})$ denote the per–period probability that a type–δ supplier is shown to a customer in segment $j\in\{1,\dots,J\}$. Then we have

$$\psi_j(\delta,\mathbf{D})=\frac{C_j(\mathbf{D})\,D_j(\delta)}{\sum_{j'\in\{1,\dots,J\}}C_{j'}(\mathbf{D})\,D_{j'}(\delta)}\,\psi(\delta;\mathbf{D})\ \text{for}\ j\in\{1,\dots,J\}.$$

This expression states that the probability a type–δ supplier is shown to segment j equals the overall probability of being shown, $\psi(\delta,\mathbf{D})$, multiplied by segment j's proportional weight among all segments.

Next, we characterize the probability of sale for a supplier of type δ. Let F_{V_j} denote the steady–state distribution of valuations among active customers in segment j. If a type–δ supplier charging price p is recommended to a segment–j customer, a sale occurs if and only if $v-p-d_j(\delta)\geq\theta_j(v,\mathbf{D})$. Thus, conditional on recommendation to segment j, the probability of sale is $\Pr_{v\sim F_{V_j}}\big(v-p-d_j(\delta)\geq$

$\theta_j(v,\mathbf{D})\big)$. Since a type–$\delta$ supplier is recommended to segment j with probability $\psi_j(\delta,\mathbf{D})$, the overall per–period sale probability is the weighted average of these conditional probabilities across segments

$$\varphi_S(p,\delta,\mathbf{D}) = \sum_{j\in\{1,\dots,J\}} \psi_j(\delta,\mathbf{D}) \Pr_{v\sim F_{V_j}}\big(v-p-d_j(\delta) \geq \theta_j(v,\mathbf{D})\big).$$

Each supplier supplies a single unit. This framework readily accommodates multi-unit suppliers (i.e., firms selling finitely many units) by treating each unit as a replica supplier of the same type. Under this interpretation, the arrival mass m represents the total arriving *capacity* (mass of units). This replica approach does not alter the subsequent analysis. The platform charges an ad-valorem commission $\tau\in[0,1)$, leaving the supplier with revenue $(1-\tau)p$ upon a sale. Suppliers who do not sell their unit remain active, discounting future profits by a factor $(1-\rho)$ per period, where $\rho\in(0,1]$ (higher ρ corresponds to a more myopic supplier). Given these elements, a type δ supplier, posting price p, earns a lifetime expected discounted profit satisfying the Bellman equation

$$\pi_s(\delta,p,\mathbf{D}) = (1-\tau)\,p\,\phi_S(p,\delta,\mathbf{D}) + (1-\phi_S(p,\delta,\mathbf{D}))(1-\rho)\pi_s(\delta,p,\mathbf{D}),$$

which can be written explicitly as

$$\pi_s(\delta,p,\mathbf{D}) = \frac{(1-\tau)\,p\,\phi_S(p,\delta,\mathbf{D})}{\rho+(1-\rho)\,\phi_S(p,\delta,\mathbf{D})}.$$

Each supplier of type δ chooses a price to maximize $\pi_s(\delta,p,\mathbf{D})$ and enters the market if and only if the lifetime profit is positive, that is, if $\pi_s(\delta,p,\mathbf{D})>0$. Hence, the equilibrium entry set is determined endogenously as $\mathcal{S}=\{\delta\in[0,1]:\pi_s(\delta,p,\mathbf{D})>0\}$.

Since supplier profit $\pi_s(\delta,p,\boldsymbol{D})$ depends on the recommendation policy $\boldsymbol{D}$ through the sale probability $\phi_S(p,\delta,\boldsymbol{D})$, the algorithm itself becomes a key determinant of suppliers' strategies. First, by allocating recommendation weights $D_j(\delta)$, the platform directly determines the showing probabilities $\psi_j(\delta,\boldsymbol{D})$. Second, because the distribution of suppliers that customers expect to encounter depends on $\boldsymbol{D}$, the recommendation policy indirectly shapes customer acceptance thresholds $\theta_j(v,\mathbf{D})$. Together, these direct and indirect paths explain how the platform's recommendation policy influences equilibrium supplier entry and pricing decisions.

3.3. Equilibrium Concept

We now formalize the notion of a stationary equilibrium by combining optimal agent strategies and market-consistency conditions. Let $c_j(v)$ denote the steady-state density of active customers in segment j with valuation $v\in\mathcal{V}_j$. The overall density of active customers at valuation v is then given by $c(v)=\sum_{j=1}^{J}c_j(v)$. Define the total customer mass as $C=\int_{\mathcal{V}}c(v)\,dv$. Similarly, note that $s(\delta,\mathbf{D})$ denotes the steady-state density of active suppliers with fit type $\delta\in\mathcal{S}$, and define $S=\int_{\mathcal{S}}s(\delta,\mathbf{D})\,d\delta$ as the total mass of active suppliers. Having defined the steady-state masses and distributions, we next describe transaction behavior. A segment-j customer with valuation v accepts an offer from a supplier of type δ if $v - p(\delta, \boldsymbol{D}) - d_j(\delta) \geq \theta_j(v, \boldsymbol{D})$, and otherwise rejects it. This condition is captured by the indicator function $\mathbf{1}\{v - p(\delta, \boldsymbol{D}) - d_j(\delta) \geq \theta_j(v, \boldsymbol{D})\}$, which equals 1 when the inequality holds (purchase occurs) and 0 otherwise (no purchase). Then the per-valuation purchase probability $q_{c,j}(v)$ is defined as $q_{c,j}(v) = \int_{\mathcal{S}} \mathbf{1}\{v - p(\delta, \mathbf{D}) - d_j(\delta) \geq \theta_j(v, \mathbf{D})\} D_j(\delta)\, d\delta$. This is the probability that a segment-j customer with valuation v makes a purchase under D_j.

Stationarity requires inflows to equal outflows for every customer valuation v and supplier type δ:

$$\text{(Customers)} \quad \frac{\lambda_j}{\bar{v}_j - \underline{v}_j} \mathbf{1}_{\{v \in \mathcal{V}_j\}} = \left[\mu_j(1 - q_{c,j}(v)) + q_{c,j}(v)\right] c_j(v), \quad \forall j, \forall v \in [\underline{v}_j, \bar{v}_j], \tag{1a}$$

$$\text{(Suppliers)} \quad m\, f_\Delta(\delta)\, \mathbf{1}_{\{\delta \in \mathcal{S}\}} = s(\delta, \mathbf{D})\, \phi_S(p, \delta, \mathbf{D}), \quad \forall \delta \in [0,1]. \tag{1b}$$

The customer balance equation (1a) equates the inflow of customers of valuation v to their outflow, due to either purchasing or exogenous exit. The supplier balance equation (1b) states that supplier inflow of type δ equals the outflow resulting from successful sales. With this, we next formally define the stationary equilibrium.

Definition 1 (Stationary equilibrium given $\mathbf{D}$). Given a recommendation policy $\mathbf{D}$, a stationary equilibrium is characterized by a tuple of strategies (θ^*, p^*), entry sets $(\mathcal{V}^*, \mathcal{S}^*)$, and steady-state distributions $(\{c_j^*\}_{j=1}^J, s^*)$ such that:

(i) **Customer Optimality.** For all j and $v \in \mathcal{V}_j^*$, the strategy $\theta_j^*(v)$ maximizes the customer's expected lifetime surplus $\pi_{c,j}$ (defined in 3.1), taking the equilibrium supplier price p^* and policy $\mathbf{D}$ as given.

(ii) **Supplier Optimality.** For all $\delta \in \mathcal{S}^*$, the strategy $p^*(\delta)$ maximizes the supplier's expected discounted profit π_s (defined in 3.2), taking the equilibrium customer strategies θ^* and the steady-state distributions (which determine exposure ψ_j and F_{V_j}) as given.

(iii) **Market Consistency (Steady State).** The distributions $(\{c_j^*\}_{j=1}^J, s^*)$ satisfy the balance equations (1a) and (1b) when evaluated at the strategies (θ^*, p^*).

(iv) **Endogenous Entry.** Entry occurs if and only if the expected payoff under the equilibrium strategies is non-negative: $\mathcal{V}_j^* = \{v \mid \pi_{c,j}(v, \theta_j^*(v)) \geq 0\}$ and $\mathcal{S}^* = \{\delta \mid \pi_s(\delta, p^*(\delta)) \geq 0\}$.

Given the definition above, next, we show the existence and uniqueness of a stationary equilibrium. Existence and uniqueness ensure that, for any fixed recommendation profile $\mathbf{D} = (D_j)_{j=1}^J$, the induced stationary outcome (acceptance thresholds $\{\theta_j^*\}_{j=1}^J$, optimal prices p^*, active densities (c_j^*, s^*), and exposure probabilities) is unique. This allows us to treat all equilibrium objects as functions of $\mathbf{D}$, making the platform's optimization problem defined in Section 3.4 well–posed and enabling clean comparative statics and welfare analysis in subsequent sections. A high–level proof idea (formalized in Appendix F) is to construct a contraction in the vector of mean effective prices, $\mathbf{M} = (M_1, \ldots, M_J)$ where $M_j = \mathbb{E}_{\delta \sim D_j}[p^*(\delta) + d_j(\delta)]$, and apply Banach's Fixed–Point Theorem.

Theorem 1 **(Existence & uniqueness of stationary equilibrium)**. *Given any recommendation policy* $\mathbf{D}$*, there exists a* unique *stationary equilibrium.*

The contraction mapping in Appendix F yields a unique fixed point $\mathbf{M}^*(\mathbf{D})$, which pins down $\{\theta_j^*\}$, p^*, and (c_j^*, s^*) for each policy $\mathbf{D}$. This single–valued mapping $\mathbf{D} \mapsto E(\mathbf{D})$ (the stationary equilibrium) is what allows us to optimize recommendations without equilibrium selection issues and to state the structural results (e.g., monotonicity of strategies and the distributional properties) globally rather than equilibrium–by–equilibrium.

3.4. Platform's Problem

The platform's objective is to maximize the commission revenue. Specifically, the platform selects recommendation policies $\mathbf{D} := (D_j)_{j \in \{1,\ldots,J\}}$, anticipating the resulting unique stationary equilibrium (as shown in Theorem 1). Given the commission rate $\tau \in (0,1)$, the platform's revenue in a stationary equilibrium equals the commission collected from completed transactions. The total mass of transactions occurring per period for suppliers of type δ is given by the steady-state mass of these suppliers, $s^*(\delta, \boldsymbol{D})$, multiplied by their per-period sale probability, $\phi_S(p^*(\delta, \boldsymbol{D}), \delta, \boldsymbol{D})$. Thus, the per-period platform revenue is given by

$$R(\boldsymbol{D}) = \tau \int_{\mathcal{S}} \quad s^*(\delta, \boldsymbol{D}) \quad \phi_S(p^*(\delta, \boldsymbol{D}), \delta, \boldsymbol{D}) \quad p^*(\delta; \boldsymbol{D})\, d\delta.$$

The integral aggregates the transaction value across all active supplier types $\delta \in \mathcal{S}$. Multiplying by the commission rate τ gives the platform's per-period commission revenue. This expression is mathematically equivalent to the customer-side aggregation of expected payments. Hence, the platform chooses $\boldsymbol{D}$ that maximizes $R(\boldsymbol{D})$.

While the platform has full control over the recommendation distribution $\mathbf{D}$, it must satisfy a feasibility condition: the number of customers directed to any supplier type δ cannot exceed the available mass of type-δ suppliers in the market, $s(\delta, \mathbf{D})$.

While the platform has full control over the recommendation distribution $\mathbf{D}$, it must ensure feasibility: the number of customers directed to any supplier type δ cannot exceed the available mass of type-δ suppliers in the market, $s^*(\delta, \mathbf{D})$. Specifically, $s^*(\delta, \mathbf{D}) = \frac{m\, f_\Delta(\delta)}{\phi_S(p^*, \delta, \mathbf{D})}$ where $m\, f_\Delta(\delta)$ is the inflow rate of type-δ suppliers and the denominator is the rate at which suppliers of type δ make a sale. This expression follows from the expected lifetime of a supplier under a geometric distribution with success probability $\phi_S(p^*, \delta, \mathbf{D})$. Formally, the recommendation policy $\mathbf{D}$ must satisfy $\sum_{j=1}^J C_j(\mathbf{D}) \cdot D_j(\delta) \leq s^*(\delta, \mathbf{D}) = \frac{m\, f_\Delta(\delta)}{\phi_S(p^*, \delta, \mathbf{D})}$.

Under this constraint, the platform solves the optimization problem OPT-G.

$$\max_{\boldsymbol{D}:\,[0,1]^J\to\mathbb{R}_{\geq 0}} \quad R(\boldsymbol{D}) \tag{2a}$$

$$\text{(OPT-G)} \quad \text{s.t.} \quad \int_0^1 D_j(\delta)\,d\delta \,\leq\, 1, \quad \forall\, j, \tag{2b}$$

$$\sum_{j=1}^{J} C_j(\boldsymbol{D})\, D_j(\delta) \,\leq\, s^*(\delta,\boldsymbol{D}), \quad \forall\, \delta\in[0,1]. \tag{2c}$$

The first constraint requires that, for each customer segment, D_j forms a valid probability distribution over supplier types. The second constraint enforces feasibility on the supply side by ensuring that the customer flow directed to any supplier type does not exceed the available mass of that type. The platform chooses a recommendation policy $\mathbf{D}$, denoted by $\mathbf{D}^*$, that solves this maximization problem. Importantly, although the platform maximizes current-period revenue; it is not myopic. The equilibrium approach ensures that the platform fully internalizes the long-term implications of its recommendation policy. Thus, unlike models in which myopic decisions typically discourage diversification and fairness, here an optimal steady-state recommendation incorporates the consequences of deliberately steering customers toward a diverse set of suppliers.

4. Equilibrium Analysis

In this section, we analyze the market equilibria for given recommendation distributions $\mathbf{D}$. To describe our result about customer-side strategies, we introduce the functional change of variable $\omega_j(v) = v - \theta_j(v)$, which we call the *willingness-to-pay* of customers of segment j with valuation v. As seen subsequently, equilibrium properties are expressed more conveniently in terms of $\omega_j(\cdot)$ rather than $\theta_j(\cdot)$. Given a stationary equilibrium for recommendation distribution $\mathbf{D}$, remember that $\mathcal{V}_j^* = \{v \in [0,1] : \pi_{c,j}(v, \theta_j^*(v;\mathbf{D})) \geq 0\}$ denote the set of valuations of segment-j customers who enter the market in equilibrium. We define the upper and lower bounds on the distribution of willingness-to-pay as follows:

$$\underline{\omega}_j = \inf\{v - \theta_j^*(v) : v \in \mathcal{V}_j^*\} \quad \text{and} \quad \bar{\omega}_j = \sup\{v - \theta_j^*(v) : v \in \mathcal{V}_j^*\},$$

where, for simplicity, our notation omits the dependence of $\underline{\omega}_j$ and $\bar{\omega}_j$ on the stationary equilibrium. The upper and lower bounds of willingness-to-pay, $\underline{\omega}_j$ and $\bar{\omega}_j$, are well-defined when a positive mass of customers enters the market. Hence, in our analysis, we disregard trivial equilibria in which no agent participates. The next two results demonstrate structural facts about the customer willingness-to-pay values and the supplier prices, which are instrumental in subsequent analysis.

Next, we present our findings on the structure of optimal strategies for suppliers and customers. These results provide insights into optimal strategies in the stationary equilibrium and are instrumental in later stages of the paper.

Proposition 1 **(Monotonicity of optimal customer strategies)**. *In the unique stationary equilibrium, the acceptance threshold $\theta_j^*(v;\mathbf{D})$ is strictly increasing in v on $\mathcal{V}_j^*$ for any j. The effective willingness to pay $\omega_j(v;\mathbf{D}) = v - \theta_j^*(v;\mathbf{D})$ is also strictly increasing in v for any j.*

Recall that an optimal acceptance threshold balances delaying purchase against the net surplus obtained from an immediate purchase. Proposition 1 shows that equilibrium customer strategies exhibit vertical differentiation. Customers with higher valuations have higher acceptance thresholds, meaning they demand a higher net surplus from successful transactions. Interestingly they also have higher willingness-to-pay for the product, which grants them a competitive advantage in the market.

We now turn to characterizing supplier prices. For ease of exposition of results and presenting fundamental trade-offs clearly, we first analyze the *single–segment* benchmark. The multi–segment case ($J > 1$) is presented in detail in Section 6. Recall that $\mathcal{S}^* \subseteq \mathcal{S}$ denote the set of supplier types that are active in equilibrium (those that receive positive exposure and post finite prices).

Proposition 2 **(Structure of optimal supplier strategies: single segment)**. *Consider the stationary equilibrium given a recommendation policy $\mathbf{D}$ when $J = 1$ and $d(\delta) = \delta$. The optimal price mapping $p^* : \mathcal{S}^* \to \mathbb{R}_+$ satisfies:*

(i) ***Price and profit monotonicity:*** *$p^*(\delta)$ and the resulting supplier profit $\pi_s^*(\delta;\mathbf{D})$ are strictly decreasing in δ on $\mathcal{S}^*$.*

(ii) ***Effective–price monotonicity:*** *the effective price paid by a customer, $p_e^*(\delta) := p^*(\delta) + \delta$, is strictly increasing in δ on $\mathcal{S}^*$.*

Intuitively, the proposition illustrates a compelling economic insight. When the level of misfit between a supplier's product and a customer's preferences is small, the product closely aligns with the customer's needs, and the supplier can post a relatively high price without distressing demand. As misfit disutility δ increases, the customer's perceived value decreases, and the supplier has to reduce her posted price to remain competitive in the market. Yet these reductions never fully offset the customer's growing misfit disutility. Consequently, the effective price, which is the posted price plus the misfit disutility, rises with increasing misfit disutility. Effectively, this result implies that competitive pressure marginalizes niche suppliers; customers naturally gravitate toward suppliers offering the lowest effective prices, thus favoring those with lower misfit disutilitys.

Finally, Theorem 2 shows that the equilibrium density of customers' willingness-to-pay preserves the shape of the arrival valuation distribution of customers.

Theorem 2 **(Distribution of willingness-to-pay)**. *Let $f_{W_j}(\omega)$ denote the steady-state density of willingness-to-pay among* active *customers in segment j. In the stationary equilibrium, this density is constant over its support $[\underline{\omega}_j, \bar{\omega}_j]$. The density is given by*

$$f_{W_j}(\omega) = \frac{\lambda_j}{\mu_j(\bar{v}_j - \underline{v}_j)}, \qquad \omega \in [\underline{\omega}_j, \bar{\omega}_j] \ .$$

Theorem 2 establishes that, at the stationary equilibrium, the shape of the willingness-to-pay distribution matches exactly the distribution of arriving customer valuations (both are uniform). To clarify this result, reconsider the model dynamics: each customer arrives with an initial valuation and either makes an immediate purchase or chooses to remain in the market to search further in subsequent periods. Consequently, at any given time, the population of customers of a specific type consists of those who entered in previous periods (or in the current period) and persisted in the marketplace until now. Since the duration that each customer type spends in the system is influenced by type-specific strategies—and these strategies vary across customer types as established in Proposition 1—the equilibrium number and valuations of customers belonging to any particular type are generally expected to differ. Theorem 2 demonstrates that, in the long run, the distribution of customers' willingness-to-pay (distinct from their initial valuations) replicates the initial valuation arrival distribution precisely, scaled by a factor related to μ_j and the valuation support length $\bar{v}_j - \underline{v}_j$. In fact, the result can be generalized for any arrival distribution. That is, for any arrival distribution, the equilibrium distribution of willingness-to-pay has the same shape as the arrival distribution, appropriately scaled. This is a powerful result simplifying the subsequent analysis, as knowledge of the valuation arrival distribution and the customer departure rate directly yields the willingness-to-pay distribution at equilibrium.

Having established these foundational monotonicity and distributional properties, we next examine how the platform can leverage them when designing recommendation policies.

5. Scarcity by Design is Optimal: The Case of a Single Customer Segment

In this section, we study a single customer segment where customers prefer low-δ suppliers with $d(\delta) = \delta$. Studying a single market segment first allows us to highlight the core trade-offs clearly before generalizing our results in Section 6. In this setting, we establish that the platform's optimal strategy involves designing a recommender system that concentrates customer traffic on two distinct groups of suppliers: the very best and worst matches, while avoiding intermediate types.

The key insight is that steering a fraction of customers toward "less-preferred" suppliers creates 'scarcity' within the search environment, prompting customers to become less selective. Moreover, exposing misfit suppliers poses a smaller competitive threat to best match suppliers than introducing other good match ones, thus softening price competition. These two factors enable suppliers to optimally increase their prices. We identify the specific market conditions, related to customer patience, supplier myopia, and market thickness, under which this strategy increases the platform's total commission revenues.

Furthermore, we demonstrate that improvements in supplier profitability and platform revenue often arise at the expense of customer surplus and reduce the total volume of transactions. In the following subsections, we formally characterize the optimal policy, the mechanism driving it, and the resulting welfare implications.

5.1. Revenue-maximizer recommendation policy

We begin by presenting our main result: the revenue-maximizing recommendation policy concentrates customer traffic on two distinct groups of suppliers positioned at opposite extremes of the supplier-type spectrum. We subsequently dissect the logic behind this result, providing detailed insights into the underlying mechanism, supported by formal results.

THEOREM 3 **(Structure of the revenue–maximizing recommendation)**. *Let* D^* *denote the optimal recommendation distribution for the single–segment model with* $d(\delta) = \delta$*. Then:*

(i) ***Structure.*** *There exist unique thresholds* $0 \le \delta_0 \le \delta_1 \le \bar{\delta}$ *such that*

$$D^*(\delta) = \begin{cases} \dfrac{s^*(\delta; D^*)}{C(D^*)}, & \delta \in [0, \delta_0] \cup [\delta_1, \bar{\delta}], \\ 0, & \delta \in (\delta_0, \delta_1). \end{cases}$$

Hence, the recommendation policy saturates the capacity of suppliers on two extreme blocks while allocating zero recommendations to the suppliers in the middle.

(ii) ***Thick-Market Limit.*** *As the supplier arrival rate* m *increases, both intervals* $[0, \delta_0]$ *and* $[\delta_1, \bar{\delta}]$ *monotonically shrink in length. In the limit as* $m \to \infty$*,* D^* *converges weakly to a two-point distribution (a mixture of two Dirac measures) concentrated at* 0 *and* $\bar{\delta}$*.*

Theoretical challenges and proof ideas for Theorem 3. Although the optimal recommendation algorithm has a clean structure, characterizing this recommendation strategy presents significant analytical challenges. The platform's optimization problem involves maximizing revenue over an infinite-dimensional space (the space of distributions D). Crucially, both the objective function (revenue) and the constraints (supplier capacity) are endogenous equilibrium outcomes of the policy itself. Changing the recommendation strategy alters the strategic behavior of customers and suppliers, which feeds back into the market equilibrium, reshaping both revenue potential and capacity utilization. Standard optimization techniques are insufficient to handle these feedback loops and endogenous (state-dependent) constraints. We detail how we overcome these challenges in Appendix C where we use variational analysis combined with a methodology we term "Freeze-then-Lift". We first derive the Gateaux derivative (marginal revenue density), which fully incorporates the equilibrium feedback. We then prove this derivative is strictly convex (U-shaped) in the supplier type δ. This approach allows us to establish the structural characterization in Theorem 3.

Theorem 3 establishes that the platform optimally assigns positive recommendation probability only to two groups of suppliers: the very best fits ($\delta \in [0, \delta_0]$) and the worst fits among those that still profitably participate ($\delta \in (\delta_1, \bar{\delta}]$). Suppliers in both groups fully utilize their capacities. Furthermore, as the market thickens ($m \to \infty$), the platform becomes more selective, concentrating recommendations entirely on the extremes ($\delta = 0$ and $\delta = \bar{\delta}$). On the other hand, for sufficiently small m, it may be optimal to have $\delta_0 = \delta_1$, meaning the gap disappears and all active suppliers are utilized.

It may seem paradoxical that *intentionally* diverting customers away from their best-match suppliers could increase revenue. The explanation hinges upon the strategic responses of customers and suppliers. To illustrate this mechanism, we formally introduce a baseline recommender and a shifted recommender.

Baseline and shifted recommenders. We define a *baseline recommender* D^B that efficiently uses capacity for low-misfit suppliers. Specifically, D^B completely *saturates* available supplier capacity within a best-fit segment $[0, \delta_0]$, i.e., $C(D^B)\, D^B(\delta) = s^*(\delta, D^B)$ for all $\delta \in [0, \delta_0]$. In contrast, capacity is *slack* for the remainder, i.e., $C(D^B)\, D^B(\delta) < s^*(\delta, D^B)$ for $\delta \in (\delta_0, \bar{\delta}]$. In other words, capacity is fully utilized among better-matched (low-δ) suppliers, while availability exists among undesirable (higher-δ) suppliers. Let $\omega_B(v)$ and $p_B^*(\delta)$ denote the induced equilibrium willingness-to-pay and supplier prices under D^B.

To examine the impact of steering customers toward less desirable suppliers, we select a *source set* $I_\ell \subseteq [0, \delta_0]$ and a *destination set* $I_h \subset (\delta_0, \bar{\delta}]$. For a sufficiently small $\varepsilon > 0$, we define the *shifted recommender* D^ε as:

$$D^\varepsilon(\delta) := D^B(\delta) - \varepsilon\, \frac{\mathbf{1}_{I_\ell}(\delta)}{|I_\ell|} + \varepsilon\, \frac{\mathbf{1}_{I_h}(\delta)}{|I_h|}.$$

This construction reallocates an ε-fraction of recommendation mass from best match suppliers (in I_ℓ) to less desirable suppliers (in I_h). Let $\omega_\varepsilon(v)$ and $p_\varepsilon^*(\delta)$ denote the equilibrium outcomes under D^ε. The following results characterize the strategic responses of customers (Proposition 3) and suppliers (Proposition 4) to this shift.

Proposition 3 **(Injecting scarcity softens search)**. *Starting from a baseline recommender D^B defined above, steering any positive mass of customers toward relatively less desirable suppliers increases each customer's willingness-to-pay:*

$$\omega_\varepsilon(v) > \omega_B(v) \qquad \text{for every } v \in \mathcal{V}^* \text{ and every sufficiently small } \varepsilon > 0.$$

Proposition 3 establishes that introducing a probability of encountering highly misfit suppliers makes the prospect of future search less attractive. Customers recognize that rejecting a current offer carries the risk of encountering a lower–quality fit—which, by Proposition 2, carries a higher effective price—in the next period. This injection of scarcity makes search riskier, reducing the customer's continuation value and thus increasing their willingness to accept the current offer. This, in turn, decreases the effective price elasticity of demand perceived by suppliers. The following proposition formalizes the suppliers' strategic response.

Proposition 4 **(Scarcity raises prices)**. *Starting from a baseline recommender D^B defined above, steering any positive mass of customers toward relatively less desirable suppliers increases each supplier's price:*

$$p^*_\varepsilon(\delta) > p^*_B(\delta) \qquad \text{for every } \delta \in \mathcal{S}^* \text{ and every sufficiently small } \varepsilon > 0.$$

Facing customers with a degraded outside option, suppliers optimally respond by raising their prices. This highlights the core tradeoff behind the strategic injection of scarcity. On one hand, diverting customers to higher misfit suppliers softens price competition, enabling higher markups (Propositions 3 and 4). On the other hand, this diversion increases the expected search duration, risking lost sales as impatient customers exit the market before purchasing.[5] The following proposition resolves this fundamental tradeoff, identifying how to perform the diversion and the conditions under which the revenue gained from higher margins outweighs the loss from increased churn.

Proposition 5 **(When and how to allocate high misfit recommendations)**. *Suppose diverting a positive fraction of customer traffic from a baseline recommendation policy D^B toward less desirable suppliers increases platform revenue. Then, revenue is maximized by allocating the entire diverted mass to the supplier with the highest available misfit disutility, i.e., to the supplier type $\delta = \sup\{\delta' : D^B(\delta') < s(\delta')/C(D^B)\}$.*

Moreover, such a diversion increases platform revenue if and only if the following conditions are satisfied simultaneously:

1. Customers are sufficiently patient*: $\mu_1 < \bar{\mu}(m,\rho)$, where $\bar{\mu}(m,\rho)$ is a critical threshold. In the limit of abundant supply ($m \to \infty$), the threshold converges to $\bar{\mu}(\infty,\rho) = \frac{\rho}{1+2\rho}$.*

[5] One might expect that if customer departure (μ_j) were endogenous (i.e., customer becomes more likely to exit upon receiving a bad recommendation, potentially due to an outside option), this mechanism would be significantly dampened by increased abandonment (*Attrition Effect*). However, our analysis in Appendix K reveals a counteracting effect. Indeed, when the platform worsens recommendations, the departure rate increases. This shortens the expected duration of search, making customers effectively more impatient. In a sequential search environment, however, higher impatience reduces the *option value of waiting* for a better future match (lower continuation value). Consequently, customers become less selective and increase their willingness-to-pay to secure a transaction quickly and avoid the heightened risk of the search terminating (the *Impatience Effect*) at equilibrium. This allows suppliers to raise prices even further when scarcity is induced. Our analysis proves that the revenue gain from this Impatience Effect dominates the revenue loss from the Attrition Effect, thereby reinforcing the platform's incentive to induce scarcity.

2. Supplier capacity is sufficiently large*: $m > m^\dagger$ for some critical market size $m^\dagger > 0$.* [6] *Equivalently, $m^\dagger$ is the smallest m such that the interval $(\delta_1, \bar{\delta}]$ in Theorem 3 has positive measure.*

Proposition 5 establishes that if scarcity injection is profitable, it should be done using the least desirable suppliers available (highest $\delta \in \mathcal{S}^*$). The mechanism driving this result, and the broader structure of Theorem 3, is the *curvature* of the effective prices $p^*(\delta) + \delta$ (Proposition 2 and Appendix B)[7]. Because the effective price increases at an increasing rate as the misfit δ grows, the impact on the equilibrium (the "return on scarcity") grows disproportionately when customers are diverted to the extremes. Consequently, the platform's marginal revenue density is also convex (U-shaped). Suppliers with moderate quality fit produce smaller price improvements but carry similar risks of customer attrition, which makes the allocation at the extremes optimal.

Moreover, Proposition 5 provides the tipping point for when this strategy is profitable, balancing the benefit of higher prices against the risk of losing customers. First, the strategy requires patient customers (small μ_1). Patient customers are likely to continue searching after encountering a poor recommendation. This provides a "safety net," allowing the platform to induce scarcity without excessive customer loss. Impatient customers, conversely, make the strategy too risky.

Second, the strategy is most profitable when suppliers are relatively myopic (high ρ). When suppliers are patient (low ρ), they can afford to wait for future sales, which inherently softens competition and leads to flatter effective prices across supplier types. Consequently, the price-increasing effect of induced scarcity is dampened since the effective prices customers encounter cannot be significantly altered by injecting scarcity. Conversely, when suppliers are myopic (high ρ), competition is more intense, which leads to high heterogeneity across effective prices. Hence, the strategic injection of scarcity is effective and yields a larger marginal increase in prices.

Third, the diversion is profitable only when the market is sufficiently thick ($m > m^\dagger$). If the market is thin (small m), the optimization problem is trivial: the platform must utilize all available suppliers to meet demand, leaving no scope for strategic diversion. Only when supply exceeds the critical threshold $m^\dagger$ does the platform have the flexibility to strategically allocate demand toward niche suppliers.

[6] *Definition of the critical capacity level $m^\dagger$.* Let $\delta_0(m)$ be the threshold type up to which capacity is fully utilized under the best-match policy given m. We define the *critical market size* $m^\dagger$ as the smallest supplier-arrival rate at which the capacity of best-match suppliers is sufficient to serve the entire customer base: $m^\dagger = \inf\left\{ m > 0 : \int_0^{\delta_0(m)} \frac{s^*(\delta; m)}{C(m)}\, d\delta = 1 \right\}$.

[7] The strict convexity of $p^*(\delta) + \delta$, shown in Appendix B, is sufficient but not necessary. The derivations in Appendix B reveal that our results would hold for a "loosely concave" effective price function $p^*(\delta) + \delta$.

5.2. Welfare Trade-offs and Policy Performance

The analysis in Section 5.1 reveals that the platform may optimally induce scarcity to boost revenue. We now examine the welfare implications of this strategy, demonstrating a fundamental trade-off between platform revenue/supplier profit and customer surplus, and we visualize the performance of the optimal policy.

The fundamental tension. We analyze the impact of diverting traffic toward niche suppliers using the D^B to D^ε framework defined earlier. The impact on customer welfare and market efficiency is unambiguously negative.

PROPOSITION 6 **(Scarcity reduces customer surplus and transaction volume)**. *Let $CS(\varepsilon)$ denote the equilibrium customer surplus and $Q(\varepsilon)$ the total transaction volume (mass of matches) under the shifted recommender D^ε. Then, for every baseline D^B and every feasible diversion (I_ℓ, I_h),*

$$\frac{dCS}{d\varepsilon}\Big|_{\varepsilon=0} < 0 \quad \textit{and} \quad \frac{dQ}{d\varepsilon}\Big|_{\varepsilon=0} < 0.$$

In other words, every shift toward high misfit suppliers strictly lowers *total customer surplus and the total volume of transactions.*

When the platform steers customers toward undesirable suppliers, customers face higher equilibrium prices (Proposition 4). This reduces both the expected utility per transaction and the overall probability of a successful match, leading to lower transaction volume and customer surplus.

In contrast, the interests of the platform and the suppliers are perfectly aligned regarding the strategic use of scarcity, due to the ad-valorem commission structure.

PROPOSITION 7 **(Alignment of supplier and platform incentives)**. *Let $R(\varepsilon)$ denote the platform revenue and $\Pi(\varepsilon)$ the total supplier profit under D^ε. Then,*

$$\operatorname{sign}\Big(\frac{dR}{d\varepsilon}\Big) = \operatorname{sign}\Big(\frac{d\Pi}{d\varepsilon}\Big) \quad \textit{at } \varepsilon = 0.$$

Scarcity injection increases platform revenue precisely when the gains from higher prices outweigh the losses from reduced transaction volume (Proposition 6). Propositions 6 and 7 establish a clear conflict: when inducing scarcity is profitable, it benefits the platform and suppliers at the expense of customers. This leads to the question: when do the interests of all parties align?

COROLLARY 1 **(Condition for incentive alignment)**. *Customer surplus, supplier profit, and platform revenue are simultaneously maximized by a best-match policy (i.e., the gap (δ_0, δ_1) in Theorem 3 is empty) if and only if:*

$$\mu_1 \geq \bar{\mu}(m, \rho) \quad \textit{or} \quad m \leq m^\dagger.$$

Alignment occurs when the conditions for profitable scarcity (Proposition 5) fail. This happens if high misfit recommendations are either too risky (customers are impatient, μ_1 high) or infeasible (market is thin, m low). However, when both customer patience and supplier availability are sufficiently high, the platform profits by strategically degrading the search experience.

5.3. Numerical Results and Illustrations of Main Findings

In this section, we use numerical experiments to illustrate the performance of the optimal recommendation policy characterized in Section 5.1 and to compare it with two natural benchmarks across a range of market conditions. The benchmark policies are: (i) *best-match (BM)*, which prioritizes the lowest misfit suppliers up to capacity saturation, formally $D_{BM}(\delta) = s^*(\delta, D_{BM})/C(D_{BM})$ for $\delta \le \delta_{BM}$ and 0 otherwise, where δ_{BM} ensures $\int D_{BM}(\delta)d\delta = 1$; and (ii) *uniform random (U)*, which assigns customers uniformly across the spectrum of active supplier types, such that $D_U(\delta) = 1/\bar{\delta}$ for $\delta \in [0, \bar{\delta}]$. [8] We first compare the revenues of these three policies as market thickness (the supply–demand ratio) and customer impatience $\mu_1 \in [0, 1]$ vary. We then compare the three policies by varying supplier myopia, captured by the discount factor ρ, together with customer impatience μ_1. We use R^*, R_{BM}, and R_U to denote the revenues under the optimal, best-match, and uniform random policies, respectively.

Note that this numerical analysis is made possible by our theoretical results that significantly simplify the computation required for this analysis. By Theorem 3, finding the optimal policy reduces to determining two cutoffs, δ_0 and δ_1. We compute δ_0 and δ_1 from a closed-form fixed-point equation that our equilibrium characterization yields (Appendix E).

Figure 1 compares revenues as market thickness (supply-demand ratio) and customer impatience (μ_1) vary, fixing $\rho = 1$. Panel (b) shows the revenue benefit of optimal recommender compared to Best-Match recommender (R^*/R_{BM}). The advantage depends on customer patience. When customers are patient (low μ_1), strategic diversion is optimal as shown in Proposition 5. The gains can exceed 25%. When customers are impatient (high μ_1), the risk of exit dominates, and the optimal policy converges to Best-Match ($R^*/R_{BM} = 1$), illustrating Corollary 1. The gains also require sufficient market thickness, as thin markets leave no room for diversion ($m \le m^\dagger$).

Panel (a) shows the revenue gain over not using a recommendation algorith (i.e, Uniformly Random), R^*/R_U. The optimal policy consistently outperforms Uniform Random, especially when customers are impatient (high μ_1), where the ratio often exceeds 2. Uniform Random performs poorly because it spreads recommendations too widely, leading to excessive customer exit. Interestingly, for very low μ_1, Uniform Random performs reasonably well relative to Best-Match (compare panels a

[8] The Uniform policy D_U randomizes uniformly over supplier *types*, independent of the mass of suppliers $s(\delta)$ at each type.

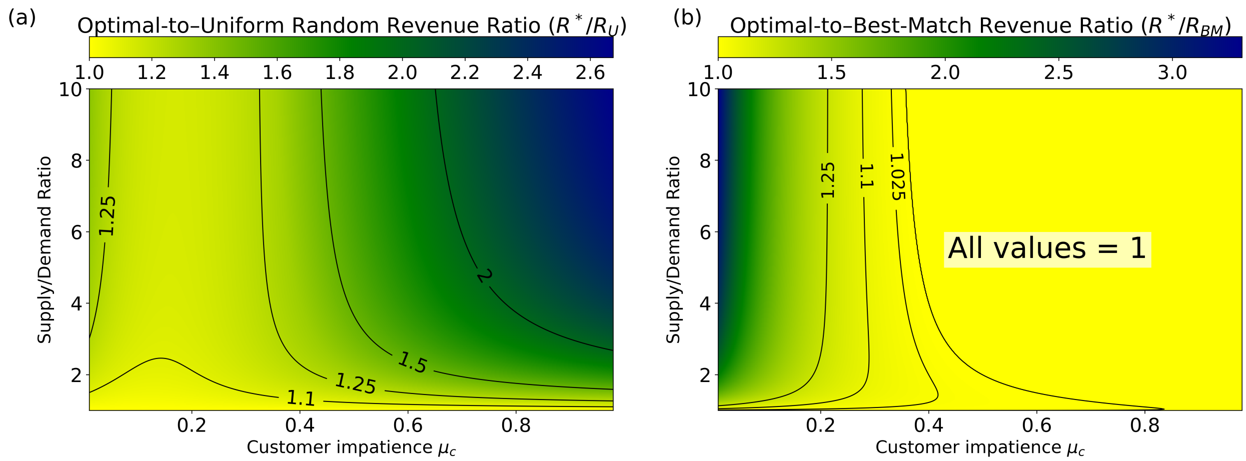

Figure 1 Revenue ratios across supply–demand ratio (market thickness) and customer impatience μ_1. Supplier myopia is fixed at $\rho = 1$. The x-axis varies customer impatience μ_1. The y-axis varies the supply-demand ratio by keeping $C(D)$ constant and varying $s(\delta, D) = s$ (assuming uniform supplier density for visualization).

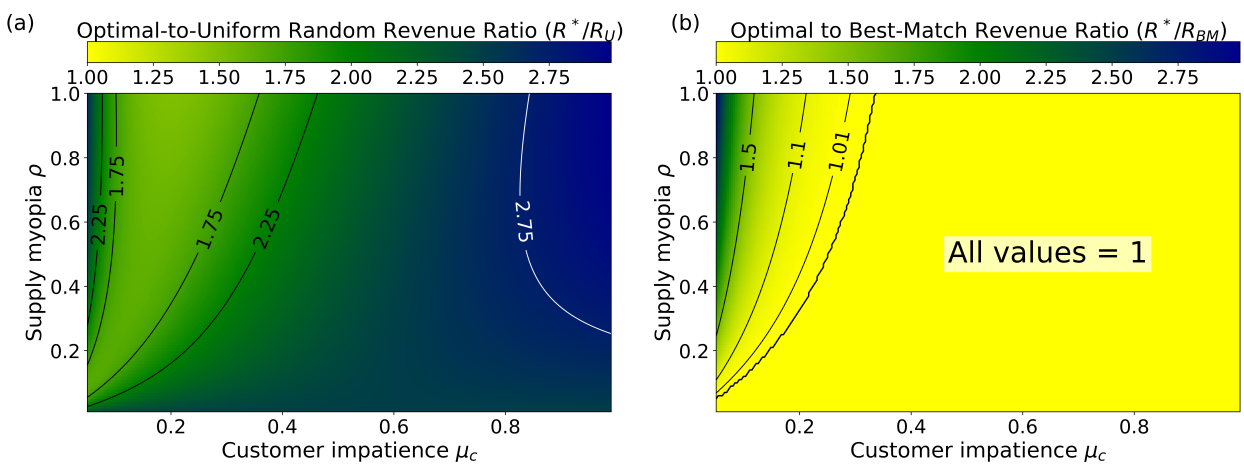

Figure 2 Revenue ratios across supplier myopia ρ and customer impatience μ_1, in the thick–market limit ($m \to \infty$). The theoretical boundary $\mu_1 = \rho/(1+2\rho)$ is visible in panel (b).

and b). This is because when μ_1 is low, injecting scarcity is beneficial (Proposition 5), and Uniform Random does inject scarcity, albeit "randomly" and suboptimally. However, for large μ_1, injecting any scarcity is suboptimal, which is where Uniform performs poorly because the very act of injecting scarcity is harmful.

Figure 2 examines the impact of supplier myopia (ρ) in the thick-market limit ($m \to \infty$), which isolates the interaction between ρ and μ_1 by removing capacity constraints. Panel (b) clearly illustrates the theoretical boundary from Proposition 5: $\bar{\mu}(\infty, \rho) = \rho/(1+2\rho)$. Below this boundary (patient customers, myopic suppliers), the platform optimally induces scarcity ($R^*/R_{BM} > 1$). Above it, the optimal policy coincides with Best-Match. The magnitude of the gain increases with supplier myopia (ρ), confirming the intuition that scarcity injection is more effective when competition is inherently intense (high ρ) and the effective price curves are steeper, rather than when competition is already softened by supplier patience (low ρ).

6. The Value of Personalization

As platforms collect increasingly granular data on customer preferences, patience, and valuations, the potential for personalization grows. While personalization intuitively improves allocative efficiency, its long-term impact on market equilibrium remains a question. In this section, we extend our analysis to a general equilibrium problem with J heterogeneous customer segments (OPT-G, as stated in Section 3.4) to investigate the structure of optimal personalized recommendations for each segment and quantify the Value of Information (VoI) that this personalization generates. To enable the analysis of this high-dimensional allocation problem, we establish a powerful structural property, named Segmental-Exclusivity (SE).

We uncover a nuanced picture where personalization serves a dual role. Beyond improving allocative efficiency, personalization enables the platform to strategically shape the competitive landscape, as in the single-segment case of Section 5. We show that the optimal personalized strategy retains the “Scarcity by Design" structure identified in the single-segment case, but applies it selectively, tailoring the degree of scarcity to each segment's characteristics.

We find that while personalization always increases platform revenue, its impact on customer surplus is ambiguous and often “distributional”. The balance between the efficiency gains and the exploitative potential of strategic scarcity is mediated by customer patience. Our analysis reveals that personalization can systematically disadvantage certain groups, such as patient bargain hunters or budget-constrained customers, highlighting significant equity concerns arising from algorithmically optimized markets.

6.1. The Structure of Optimal Personalized Recommendations

The optimization problem studied in this section is significantly more complex than the single-segment case studied in Section 5. In addition to the feedback loops where the recommendation policy shapes equilibrium prices and search behavior, the segments are now coupled through shared supplier capacity (constraint (2c)). This creates a high-dimensional, tightly coupled allocation problem embedded within an equilibrium framework.

Theorem 4 is one of our main results stating the key structural characterization that tames this complexity.

THEOREM 4 **(Structure of the Optimal Generalized Recommendation)**. *Let* $\mathbf{D}^*$ *be an optimal solution to OPT-G. Let* $[\underline{\delta}^*, \bar{\delta}^*]$ *be the convex hull of the active suppliers* $\mathcal{S}^*$*. The optimal* *policy exhibits the following structure (under standard regularity conditions, detailed in Appendix G):*

(i) ***Segmental-Exclusivity (SE):*** *Let $\mathcal{B} \subseteq [\underline{\delta}^*, \bar{\delta}^*]$ be the set where the capacity constraint* (2c) *binds. For almost every $\delta \in \mathcal{B}$, the capacity $s^*(\delta; \mathbf{D}^*)$ is allocated exclusively to a single customer segment j.*

(ii) ***Personalized Structure (Bind-Gap-Bind):*** *Let $[\underline{\delta}_j^*, \bar{\delta}_j^*]$ be the convex hull of the support of D_j^*. There exist cutoffs $\underline{\delta}_j^* \leq \delta_{j,0} \leq \delta_{j,1} \leq \bar{\delta}_j^*$ such that $D_j^*(\delta) = 0$ almost everywhere on the interior gap $(\delta_{j,0}, \delta_{j,1})$. On the extreme intervals $[\underline{\delta}_j^*, \delta_{j,0}]$ and $[\delta_{j,1}, \bar{\delta}_j^*]$, $D_j^*(\delta)$ saturates the available residual capacity (binding), i.e., $C_j(\mathbf{D}^*)D_j^*(\delta) = s^*(\delta; \mathbf{D}^*) - \sum_{k \neq j} C_k(\mathbf{D}^*)D_k^*(\delta)$ almost everywhere.*

(iii) ***Aggregate Structure (Bind-Gap-Bind):*** *There exist unique cutoffs $\underline{\delta}^* \leq \delta_0 \leq \delta_1 \leq \bar{\delta}^*$ such that the optimal aggregate exposure $y^*(\delta) = \sum_j C_j(\mathbf{D}^*)D_j^*(\delta)$ saturates the total capacity $s^*(\delta; \mathbf{D}^*)$ (binding) on $[\underline{\delta}^*, \delta_0] \cup [\delta_1, \bar{\delta}^*]$, and $y^*(\delta) = 0$ a.e. on the interior gap (δ_0, δ_1).*

The SE property (i) is fundamental for analytical tractability. Despite the complexity of J heterogeneous segments competing for shared capacity, the optimal solution results in *local specialization.* At almost any δ, the platform allocates it entirely to the single segment generating the highest marginal revenue. This result is powerful because it effectively decomposes the coupled optimization problem. Instead of a complex, multi-segment allocation decision for every unit of capacity, the problem elegantly decomposes into a collection of simpler, one-dimensional decisions (one for each supplier type). This allows equilibrium prices $p^*(\delta)$ to be determined locally.

The Personalized Structure (ii) shows that the “Scarcity by Design" strategy is applied at the most granular level. For each segment, the platform concentrates recommendations on the extremes—the very best and the very worst available suppliers. This relies on the strict convexity of the marginal revenue density $H_j(\delta)$ (Appendix G). The Aggregate Structure (iii) demonstrates that this strategic polarization persists at the aggregate level.

6.2. The Revenue Value of Information (VoI) and the Underlying Mechanism

In this section, we employ our general J segment analysis to understand the value of customizing recommendation algorithms to finer customer segments if better information is available to the platform. Note that in this richer information setting the platform can identify which segment the customer belongs to but not the private valuation. We analyze the increase in platform revenue from the utilization of more granular information. We define two optimization regimes:

- **Personalized** (P)**:** The platform utilizes all information (segments j). The maximum revenue is R^{P} (OPT-G). We let E^{P} be the induced equilibrium.

- **Non-Personalized (NP):** The platform pools segments (e.g., does not observe valuation differences), optimizing a common recommendation distribution for the pooled group. The maximum revenue is R^{NP}. We let E^{NP} be the induced equilibrium.

The Canonical Decomposition. The Total Value of Information (VoI) is: $\mathrm{VoI}^{\mathrm{Total}} = R^{\mathrm{P}} - R^{\mathrm{NP}}$. To disentangle the underlying drivers of VoI, we define the *optimal aggregate supplier exposure profile* under the NP regime:

$$y^{NP}(\delta) = \sum_{j=1}^{J} C_j(E^{NP})\, D_j^{NP}(\delta). \tag{3}$$

We introduce a counterfactual, denoted by $R^{\mathrm{P\,(NP\ Equil.)}}$, which is the maximum revenue achievable by optimally re-routing the fixed aggregate exposure $y^{\mathrm{NP}}(\delta)$ across all subgroups j, while keeping the equilibrium (prices, WTP distributions) *frozen* at E^{NP}. This maximizes immediate revenue given the fixed prices of E^{NP}. We decompose the total VoI:

$$\mathrm{VoI}^{\mathrm{Total}} = \underbrace{(R^{\mathrm{P\,(NP\ Equil.)}} - R^{\mathrm{NP}})}_{\mathrm{VoI}^{\mathrm{Frozen}}\ \text{(Matching Gain)}} + \underbrace{(R^{P} - R^{\mathrm{P\,(NP\ Equil.)}})}_{\mathrm{VoI}^{\mathrm{Adjustment}}\ \text{(Feedback Gain)}} . \tag{4}$$

$\mathrm{VoI}^{\mathrm{Frozen}}$ measures the gain from improved allocative efficiency given fixed prices. $\mathrm{VoI}^{\mathrm{Adjustment}}$ measures the revenue change due to the equilibrium shift. We apply the same decomposition to the total change in aggregate customer surplus (CS), $\Delta CS^{\mathrm{Total}} = CS^{\mathrm{P}} - CS^{\mathrm{NP}}$:

$$\Delta CS^{\mathrm{Total}} = \underbrace{(CS^{P\text{-Frozen}} - CS^{NP})}_{\Delta CS^{\mathrm{Matching}}\ \text{(Matching Effect)}} + \underbrace{(CS^{P} - CS^{P\text{-Frozen}})}_{\Delta CS^{\mathrm{Feedback}}\ \text{(Equilibrium Effect)}} .$$

PROPOSITION 8 **(Value of Information: Decomposition and Welfare Impact)**. *The Value of Information (VoI) is non-negative (*$\mathrm{VoI}^{Total} \geq 0$*). The decomposition in* (4) *reveals the following properties:*

1. ***Matching Gain (Allocative Efficiency):*** $\mathrm{VoI}^{Frozen} \geq 0$. *This gain results from improved allocative efficiency in revenue generation, leading to an increase in the number of successful transactions. However, the impact on customer surplus is ambiguous:* $\Delta CS^{Matching}$ *may be positive or negative.*
2. ***Feedback Gain (Strategic Exploitation):*** *The revenue impact* $\mathrm{VoI}^{Adjustment}$ *is ambiguous. However, the equilibrium adjustment strictly decreases customer surplus:* $\Delta CS^{Feedback} < 0$. *The magnitude of this negative effect increases as customer patience (*$1-\mu_j$*) increases.*
3. ***Total Impact:*** *Platform revenue increases (*$\mathrm{VoI}^{Total} \geq 0$*). The total impact on customer surplus (*ΔCS^{Total}*) is ambiguous. The potential harm to customer surplus is maximized when customers are highly patient (i.e., as* $\mu_j \to 0$ *for all* j*), because the negative Feedback Effect is strongest in this regime.*

Proposition 8 establishes that personalization always benefits the platform but creates a fundamental tension regarding customer welfare. Crucially, even the initial efficiency improvement (Matching Gain) does not guarantee an increase in customer surplus. We now detail the underlying mechanisms.

The mechanism behind the Matching Gain. The Matching Gain ($\mathrm{VoI}^{\mathrm{Frozen}}$) arises from the platform's ability to improve *allocative efficiency* (Proposition 8(i)). Under NP, the platform optimizes based on average characteristics, leading to misallocations (e.g., showing a premium supplier to a low-WTP segment). Personalization allows the platform to redirect the premium supplier to the high-WTP segment, which supports higher margins and conversion probabilities, thus increasing the number of matches and revenue.

However, this revenue-maximizing reallocation does not necessarily increase customer surplus ($\Delta CS^{\mathrm{Matching}}$ is ambiguous). While personalization increases the overall transaction volume, the impact on customer surplus depends on the distribution of that volume across segments with different valuation structures. Customer surplus is derived not just from making a purchase, but from the difference between valuation and price. If the platform shifts exposure from a segment with high valuation heterogeneity (where customers frequently capture significant surplus) to a segment that yields slightly higher revenue but allows the platform to extract nearly all the surplus, aggregate customer surplus can decrease even as revenue and transaction volume increase (see Appendix H for a formal derivation).

The mechanism behind the Feedback Gain. The Feedback Gain captures the equilibrium shift as the platform optimizes its strategy (Proposition 8(ii)). The platform strategically introduces scarcity tailored to each subgroup's search environment (Personalized Scarcity by Design, Theorem 4). This relies on the positive feedback loop: policies that increase expected future prices increase current willingness-to-pay, supporting higher equilibrium prices overall.

This strategic manipulation strictly harms customer surplus ($\Delta CS^{\mathrm{Feedback}} < 0$). By design, the strategy increases the mean effective prices ($\mathbf{M}$) faced by customers (Appendix H). Higher expected prices lower the continuation value of search, unambiguously reducing the expected lifetime surplus for every customer type. The strength of this mechanism is governed by customer patience; patient customers are more sensitive to changes in the future search environment, amplifying the impact of strategic manipulation.

Contrast with No-Personalization. The impact of personalization contrasts with the NP regime. Personalization introduces an efficiency gain ($\mathrm{VoI}^{\mathrm{Frozen}} \geq 0$). This efficiency gain potentially translates into customer surplus improvements. Furthermore, the strategic exploitation (Feedback Effect) remains a potent force that reduces customer surplus, potentially offsetting some of the allocative efficiency benefits.

6.3. Numerical Study: The Value of Information in Practice

We use numerical experiments to quantify the Value of Information (VoI) from personalization and analyze its underlying drivers—the Matching Gain and the Feedback Gain (Proposition 8). We examine how the platform leverages information regarding heterogeneity in customer patience (μ) and valuations (v).

We compare the *Personalized* policy (P), which optimizes segment-specific recommendations, against the *Non-Personalized* policy (NP), which optimizes a single, pooled recommendation distribution. In line with Theorems 3 and 4, we utilize a policy class, which concentrates recommendations on the best match ($\delta = 0$) and the worst active match ($\bar{\delta}$). This structure is parameterized by the scarcity level $\alpha_j \in [0,1]$, the probability that segment j is shown the worst match.

We assume myopic suppliers ($\rho = 1$) throughout. This represents a highly competitive environment where the potential for strategic price manipulation via scarcity is maximized (Proposition 5), allowing us to clearly illustrate the strategic effects and providing analytical tractability (see Appendix I for equilibrium computation details).

Experimental Design. To isolate the distinct mechanisms driving the VoI, our experiments utilize specific settings regarding market thickness (m):

1. **Experiment 1: Patience Heterogeneity (Strategic Value).** We study the value of information regarding the patience levels μ_j. To isolate the strategic Feedback Gain, we assume homogeneous valuations and a thick market ($m \to \infty$). This ensures capacity constraints are negligible and that policy differentiation is driven purely by varying sensitivity to scarcity based on patience.
2. **Experiment 2: Valuation Heterogeneity (Allocative Value).** We study the value of information regarding v_j. Crucially, we introduce a *finite capacity* constraint ($m < \infty$) on best-match suppliers. This is necessary because, in a thick market ($m \to \infty$) with ad-valorem fees and homogeneous patience, the optimal scarcity level depends primarily on μ, not valuation bounds (Appendix I). If capacity were infinite, P and NP policies would converge, yielding zero VoI. Finite capacity creates resource contention, activating the Matching Gain by forcing the platform to prioritize the allocation of scarce high-quality supply. We retain the barbell structure (allocation to the extremes) for analytical tractability, which captures the essential trade-off in allocation. Appendix J details the capacity-coupled optimization.

6.3.1. Experiment 1: The Strategic Value of Patience Information We analyze a thick market ($m \to \infty$) with two equal-sized segments ($\lambda_L = \lambda_H = 0.5$) sharing identical valuations $U[0,1]$, but differing in patience. We define a Patient (L) and an Impatient (H) segment. We fix the average impatience $\bar{\mu} = 0.5$ and vary the polarization gap $g \in [0,1)$, such that $\mu_H = \bar{\mu} + g/2$ and $\mu_L = \bar{\mu} - g/2$.

Figure 3 illustrates how the platform exploits information about patience heterogeneity, demonstrating that personalization enables targeted exploitation of patient customers.

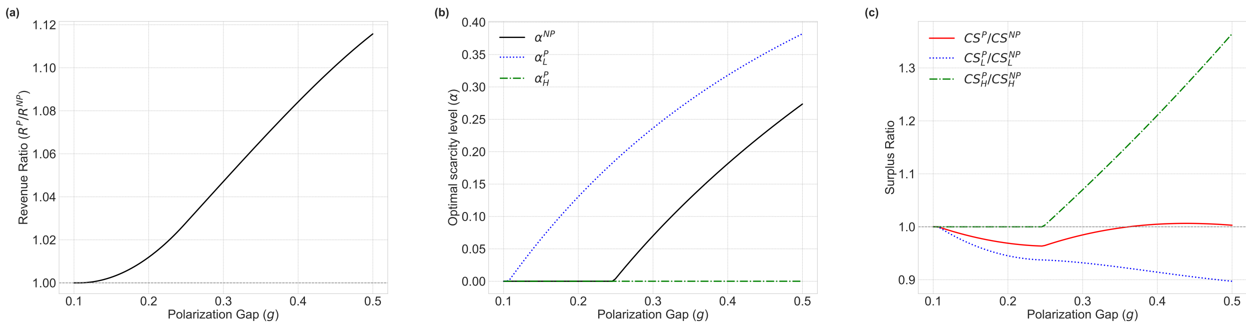

Figure 3 Experiment 1: Patience Heterogeneity and abundant suppliers ($m \to \infty$). Impact of polarization gap g on (a) Revenue Ratio $R^{\mathrm{P}}/R^{\mathrm{NP}}$; (b) Optimal Scarcity Levels; (c) Segment-level Surplus Ratios $CS^{\mathrm{P}}/CS^{\mathrm{NP}}$.

Revenue Gains (Panel a). The VoI increases significantly with the polarization gap g. As heterogeneity increases, the NP policy—which must apply a common scarcity level α^{NP}—becomes an increasingly poor compromise. Personalization allows the platform to tailor scarcity to each segment's tolerance, yielding revenue gains exceeding 8% when polarization is high.

Targeted Scarcity (Panel b). The mechanism behind these gains is the strategic deployment of scarcity. The NP policy must adopt a compromise (α^{NP}), lying between the ideal policies for the two segments. In contrast, the P policy tailors scarcity precisely. For the Impatient segment (H), the platform sets $\alpha^P_H = 0$. These customers exit quickly (high μ_H), making induced scarcity too risky (consistent with Proposition 5). For the Patient segment (L), the platform aggressively induces scarcity; α^P_L increases sharply with g. Patient customers are less likely to exit and are more sensitive to changes in the future search environment, making them ideal targets for the Feedback Gain mechanism to elevate prices.

Distributional Welfare Effects (Panel c). This strategy creates stark distributional consequences. Personalization effectively redistributes welfare from patient to impatient customers. The Impatient segment benefits significantly from personalization ($CS^P_H/CS^{NP}_H > 1$) as they are spared the scarcity imposed by the NP policy. However, the Patient segment is strictly worse off ($CS^P_L/CS^{NP}_L < 1$); they face higher scarcity and consequently higher equilibrium prices. While the aggregate surplus impact is slightly positive at high polarization, the primary outcome is the platform's ability to extract more revenue by exploiting the segment most tolerant to scarcity (e.g., bargain hunters).

6.3.2. Experiment 2: The Allocative Value of Valuation Information We now analyze a capacity-constrained market. We consider two segments with homogeneous patience (μ) but different valuations: a Budget segment (B) with $v_B \sim U[0, \bar{v}_B]$ and a General/Premium segment (G) with $v_G \sim U[0, \bar{v}_G]$, with $\bar{v}_G > \bar{v}_B$. We vary the market share of the General segment, s.

We analyze a representative scenario by setting the relative capacity $M_{rel} = 0.75$ (available capacity is 75% of the maximum potential demand, Appendix J) and $\mu = 0.5$ (an impatient market). This setting allows us to isolate the Matching Gain cleanly. Since $\mu = 0.5$ is relatively high, the incentive to strategically induce scarcity for the Feedback Gain is naturally muted (Proposition 5). Thus, the VoI in this setting is driven primarily by the efficient allocation of constrained capacity.

Before analyzing the numerical results, we establish a key theoretical result regarding how the platform optimally allocates scarce capacity when the capacity constraint binds.

PROPOSITION 9 **(Prioritization of High-Value Segments)**. *In a capacity-constrained market with homogeneous patience ($\mu_j = \mu$), the personalized policy (P) consistently imposes higher scarcity on the lower-valuation segment. That is, if $\bar{v}_G > \bar{v}_B$ and the capacity constraint binds, then $\alpha_G^P < \alpha_B^P$.*

(Proof in Appendix J). Proposition 9 formalizes the Matching Gain under constraints. Because the General segment generates higher expected revenue per transaction, the opportunity cost of allocating scarce best-match capacity to the Budget segment is high. The platform optimally responds by rationing the Budget segment more severely (higher α_B^P) to prioritize the General segment (lower α_G^P). This structural result confirms that the patterns observed in our representative scenario (Figure 4) are general properties of the optimal policy in this setting.

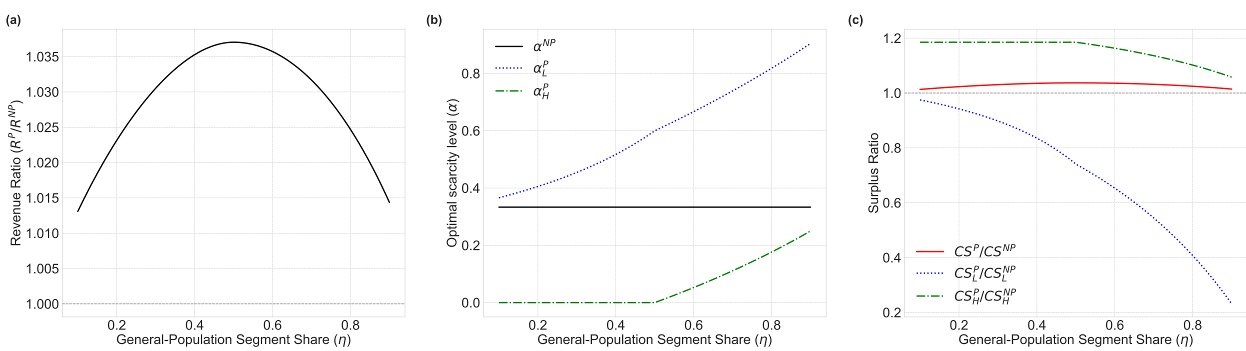

Figure 4 Experiment 2: Valuation Heterogeneity (Finite Capacity, $\mu = 0.50, M_{rel} = 0.75, \bar{v}_G = 1, \bar{v}_B = 0.5$). *(a)* Revenue ratio $R^{\mathrm{P}}/R^{\mathrm{NP}}$, *(b)* optimal scarcity levels, and *(c)* segment-level surplus ratios, plotted against the General/Premium segment share s.

Allocative Efficiency and Revenue Gains (Panel a). Personalization yields significant revenue gains, exceeding 12% when the General segment share (s) is intermediate. Even in this impatient market ($\mu = 0.5$) where the Feedback Gain is minimal, the ability to efficiently allocate constrained capacity (Matching Gain) drives substantial VoI.

Capacity Management and Policy Differentiation (Panel b). The NP policy imposes a constant scarcity level (α^{NP}) required to satisfy the capacity constraint, independent of the market composition (s). In contrast, the P policy dynamically adjusts scarcity based on s. Consistent with Proposition 9, the platform consistently imposes significantly higher scarcity on the Budget segment

($\alpha_B^P > \alpha_G^P$). Furthermore, as the share of the high-value General segment (s) increases, the shadow price of capacity rises (Appendix J). This increased contention forces the platform to increase scarcity for both segments to meet the constraint, although the relative prioritization of the General segment remains.

Welfare Impact and Prioritization (Panel c). The welfare impacts starkly highlight the consequences of revenue-maximizing allocation. The Budget segment is unambiguously harmed by personalization ($CS_B^P/CS_B^{NP} < 1$). They face higher scarcity under P than under NP across nearly all market compositions, as the platform reallocates best-match exposure toward the General segment. Conversely, the General segment benefits significantly, particularly when they are a minority (low s), as personalization shields them from the scarcity required under the pooled NP policy.

7. Conclusion and Discussion

This paper investigates the optimal design of recommendation systems in two-sided marketplaces, moving beyond immediate match quality to analyze long-term equilibrium effects. We develop a dynamic model that integrates strategic customer search, endogenous supplier pricing, and participation decisions in environments where evaluation is sequential. Our analysis reveals that maximizing short-term metrics does not necessarily align with long-term platform health.

The central insight is that a revenue-maximizing platform may strategically degrade recommendation quality—a strategy we term "Scarcity by Design." By occasionally steering customers toward less-preferred suppliers, the platform makes the search environment riskier, prompting customers to become less selective. Simultaneously, this strategy softens competition among top suppliers. Together, these equilibrium forces enable higher prices, boosting platform revenue and supplier profits, particularly when customers are patient and the market is thick. However, this strategy comes at the expense of customer surplus and overall transaction volume, highlighting a fundamental tension in platform objectives.

Our analysis of personalization reveals that granular customer data amplifies the platform's ability to manage this trade-off. We established the Segmental-Exclusivity (SE) property, which simplifies the complex allocation problem, and demonstrated that personalization yields value through two channels: improved allocative efficiency (Matching Gain) and enhanced strategic manipulation of the equilibrium (Feedback Gain). While personalization always increases revenue, it can have significant distributional consequences that harm vulnerable customer segments.

Managerial and Policy Implications. Our findings offer actionable insights for platform managers. In environments characterized by sequential search and customer patience (e.g., B2B services, freelance labor), strategically inducing scarcity can be a powerful lever for enhancing margins. Furthermore, when a high-quality supply is constrained, personalization is critical for optimizing allocative efficiency.

However, managers must balance these revenue gains against the potential long-term risks of reduced customer satisfaction and market inclusiveness. The distributional effects uncovered in our analysis raise significant policy considerations. Platforms can tailor scarcity to specific segments, often exploiting patient customers (bargain hunters) or disadvantaging budget-constrained users when capacity is limited. In the age of AI, as algorithmic personalization becomes more sophisticated and commonplace, the potential for exploiting vulnerable customer segments may significantly increase. This highlights the need for transparency in recommendation design and suggests a role for regulatory oversight to ensure fairness, concerning how platforms allocate visibility.

Limitations and Future Research. Our model relies on several assumptions that suggest avenues for future research. We consider a sequential search process, which is appropriate for high-value decisions. Exploring how recommendation design interacts with strategic behavior under alternative search protocols (e.g., simultaneous evaluation) would be valuable.

Our analysis focuses on a stationary equilibrium. This approach is best suited for evaluating the long-term impact of persistent platform policies, such as those updated periodically (e.g., monthly), rather than online policies that change with every action. Future work could explore learning dynamics, investigating how platforms and users adapt their strategies in non-stationary environments.

Finally, we consider revenue maximization via ad-valorem commissions, which is common for digital platforms. This commission structure directly links platform revenue to equilibrium prices, which is what creates the incentive for Scarcity by Design. Future research could explore how this incentive shifts under monetization strategies prioritizing transaction volume (e.g., per-transaction or subscription fees), or under hybrid monetization models, such as the interaction between organic recommendations and promoted listings (advertising), where platforms must balance organic scarcity with the incentive to sell visibility.

References

Abeliuk A, Berbeglia G, Cebrian M, Van Hentenryck P (2016) Assortment optimization under a multinomial logit model with position bias and social influence. *4OR* 14:57–75.

Adomavicius G, Bockstedt JC, Curley SP, Zhang J (2018) Effects of online recommendations on consumers' willingness to pay. *Information Systems Research* 29(1):84–102.

Asadpour A, Niazadeh R, Saberi A, Shameli A (2023) Sequential submodular maximization and applications to ranking an assortment of products. *Operations Research* 71(4):1154–1170.

Balseiro SR, Besbes O, Castro F (2024) Mechanism design under approximate incentive compatibility. *Operations Research* 72(1):355–372.

De los Santos B, Koulayev S (2017) Optimizing click-through in online rankings with endogenous search refinement. *Marketing Science* 36(4):542–564.

Dean M, Ravindran D, Stoye J (2025) A better test of choice overload. arXiv preprint arXiv:2212.03931 v3, URL https://arxiv.org/abs/2212.03931, version 3, 26 Jun 2025.

Ferreira KJ, Parthasarathy S, Sekar S (2022) Learning to rank an assortment of products. *Management Science* 68(3):1828–1848.

Fuller JB, Raman M, Palano J, Bailey A, Vaduganathan N, Kaufman E, Laverdière R, Lovett S (2020) Building the on-demand workforce. *Published by Harvard Business School and BCG* .

Gallego G, Li A, Truong VA, Wang X (2020) Approximation algorithms for product framing and pricing. *Operations Research* 68(1):134–160.

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

Hopenhayn HA (1992) Entry, exit, and firm dynamics in long-run equilibrium. *Econometrica* 60(5):1127–1150, URL http://dx.doi.org/10.2307/2951541.

Horton JJ (2017) The effects of algorithmic labor market recommendations: Evidence from a field experiment. *Journal of Labor Economics* 35(2):345–385.

Kim Y, Harburg E, Azria S, Shaw A, Gerber E, Gergle D, Zhang H (2016) Studying the effects of task notification policies on participation and outcomes in on-the-go crowdsourcing. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 4, 99–108.

Lee KH, Musolff L (2025) Two-sided markets shaped by platform-guided search.' .

Lei Y, Jasin S, Uichanco J, Vakhutinsky A (2022) Joint product framing (display, ranking, pricing) and order fulfillment under the multinomial logit model for e-commerce retailers. *Manufacturing & Service Operations Management* 24(3):1529–1546.

MacKenzie I, Meyer C, Noble S (2013) How retailers can keep up with consumers. Industry report, McKinsey & Company, URL https://www.mckinsey.com/industries/retail/our-insights/how-retailers-can-keep-up-with-consumers.

Papanastasiou Y, Sekar S (2024) Designing pandora's box: Product rankings and marketplace participation. Available at SSRN: https://ssrn.com/abstract=5006666 or http://dx.doi.org/10.2139/ssrn.5006666.

Pei C, Yang X, Cui Q, Lin X, Sun F, Jiang P, Ou W, Zhang Y (2019) Value-aware recommendation based on reinforcement profit maximization. *The World Wide Web Conference*, 3123–3129, WWW '19 (New York, NY, USA: Association for Computing Machinery), ISBN 9781450366748, URL http://dx.doi.org/10.1145/3308558.3313404.

Rios I, Torrico A (2023) Platform design in curated dating markets. *arXiv preprint arXiv:2308.02584* .

Shi P (2024) Optimal match recommendations in two-sided marketplaces with endogenous prices. *Management Science* .

Ursu RM (2018a) The power of rankings in online consumer search. *Marketing Science* 37(4):555–578.

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

Zhang M, Bockstedt JC (2016) Complements and substitutes in product recommendations: The differential effects on consumers' willingness-to-pay. *Proceedings of the 2016 Workshop on Interfaces and Human Decision Making for Recommender Systems (IntRS 2016)*, volume 1679 of *CEUR Workshop Proceedings*, 36–43 (Boston, MA, USA: CEUR-WS.org), URL `http://ceur-ws.org/Vol-1679/paper6.pdf`.

Zhou B, Zou T (2023) Competing for recommendations: The strategic impact of personalized product recommendations in online marketplaces. *Marketing Science* 42(2):360–376.

Appendix A: Proof of Theorem 2

In this section, we fix an arbitrary customer segment. Further, for ease of exposition, we drop any index indicating the customer segment. Consider a customer with valuation $v \in [\underline{v}, 1]$ who faces a sequence of suppliers drawn independently each period from the distribution D. The customer follows a stationary accept-reject threshold strategy, denoted by $\theta(v)$. Recall the instantaneous net utility from matching with a supplier of type δ is given by:

$$u_c(v,\delta) \,=\, v - p(\delta) - \delta.$$

Recall the acceptance probability:

$$q_c(v) \,=\, \Pr_{\delta\sim D}\big(u_c(v,\delta) \geq \theta(v)\big),$$

which is the probability that the customer meets a supplier providing at least the threshold utility $\theta(v)$. Then, we have the following result that characterize the customer's lifetime surplus.

Lemma 1 **(Customer lifetime surplus)**. *The expected lifetime surplus of the customer can be expressed as:*

$$\pi_c(v,\theta(v)) \,=\, \frac{\mathbb{E}_{\delta\sim D}\big[u_c(v,\delta)\,\mathbf{1}\{u_c(v,\delta)\geq\theta(v)\}\big]}{\mu_1 + (1-\mu_1)\,q_c(v)} \,=\, \frac{\theta(v)}{1-\mu_1}.$$

Furthermore, the threshold $\theta(v)$ satisfies the indifference condition:

$$\theta(v) \,=\, (1-\mu_1)\,\pi_c(v,\theta(v)).$$

Proof of Lemma 1 We start by clearly defining the customer's continuation value $\pi_c(v,\theta)$, representing the expected future surplus from following a threshold strategy θ. At any given period, the customer faces two possibilities upon meeting a supplier drawn independently from D:

- If $u_c(v,\delta) \geq \theta$, the customer accepts the offer immediately and obtains instantaneous surplus $u_c(v,\delta)$.
- If $u_c(v,\delta) < \theta$, the customer rejects the offer. With probability $(1-\mu_1)$, the customer continues searching, preserving the continuation value $\pi_c(v,\theta)$. Otherwise, with probability μ_1, the customer exits and obtains zero further surplus.

Therefore, the Bellman equation for the continuation value can be explicitly written as:

$$\pi_c(v,\theta) = \underbrace{\mathbb{E}_\delta\big[u_c(v,\delta)\,\mathbf{1}\{u_c(v,\delta)\geq\theta\}\big]}_{\text{immediate acceptance surplus}} + \underbrace{(1-\mu_1)\Pr_\delta(u_c(v,\delta)<\theta)\pi_c(v,\theta)}_{\text{continuation surplus after rejection}}.$$

Rearranging this Bellman equation clearly gives us the explicit expression for the lifetime surplus:

$$\pi_c(v,\theta) = \frac{\mathbb{E}_\delta\big[u_c(v,\delta)\,\mathbf{1}\{u_c(v,\delta)\geq\theta\}\big]}{\mu_1 + (1-\mu_1)q_c(v)}.$$

Next, we explicitly characterize the threshold condition. By definition, at the threshold utility $\theta(v)$, the customer is exactly indifferent between accepting the marginal offer and rejecting to continue. This implies the threshold must satisfy:

$$\theta(v) = (1-\mu_1)\pi_c(v,\theta(v)),$$

reflecting equality between immediate utility at the margin and the discounted expected future surplus.

Substituting this threshold condition directly into the previously derived expression for $\pi_c(v,\theta)$ yields:

$$\pi_c(v,\theta(v)) = \frac{\theta(v)}{1-\mu_1},$$

establishing precisely the second equality of the lemma.

Lastly, we verify that the denominator $\mu_1 + (1-\mu_1)q_c(v)$ is strictly positive. Since $\mu_1 \in (0,1)$, it immediately follows that this expression is strictly bounded away from zero, ensuring that the customer lifetime surplus π_c is well-defined and unique. This concludes the proof of Lemma 1. □

A.1. Proof of Proposition 1: Envelope derivative and slope bounds

In this subsection, we show that, in any stationary equilibrium, the threshold $\theta(v)$ is absolutely continuous, and its derivative satisfies, for almost every valuation v:

$$\theta'(v) = \frac{(1-\mu_1)\, q_c(v)}{\mu_1 + (1-\mu_1)\, q_c(v)} \in (0,1). \tag{5}$$

Consequently, the map $g(v) := v - \theta(v)$ is strictly increasing and invertible.

Proof. We structure the proof in several steps, and start by defining the key quantities.

Step 1: Defining Key Quantities and Equilibrium Condition.

We begin by defining two critical equilibrium-related terms. First, let

$$A(v) := \mathbb{E}_\delta \left[(v - p(\delta) - \delta) \cdot \mathbf{1}\{v - p(\delta) - \delta \geq \theta(v)\} \right],$$

represent the expected surplus conditional on acceptance. Second, define the customer's effective departure rate at valuation v by

$$B(v) := \mu_1 + (1-\mu_1) q_c(v),$$

accounting for both impatience (μ_1) and acceptance probability ($q_c(v)$). In equilibrium, the customer's value function satisfies the Bellman equation, yielding the lifetime surplus:

$$\pi_c(v) = \frac{A(v)}{B(v)}. \tag{6}$$

Step 2: Indifference Condition at Threshold.

At the acceptance threshold $\theta(v)$, the customer is precisely indifferent between accepting and continuing to wait. Thus, we must have:

$$A(v) = \frac{\theta(v)}{1-\mu_1}\, B(v). \tag{7}$$

Step 3: Derivative of Expected Surplus $A(v)$.

We next differentiate $A(v)$ with respect to valuation v. By changing variables and applying Leibniz's rule for differentiation under the integral sign, we obtain:

$$A'(v) = q_c(v) + \theta(v)\, q_c'(v). \tag{8}$$

Here, $q_c(v)$ appears naturally as the integrand evaluated at the boundary defined by $g(v) = v - \theta(v)$.

Step 4: Derivative of the Lifetime Surplus $\pi_c(v)$.

Using equations (6) and (8), along with $B'(v) = (1-\mu_1)q_c'(v)$, we apply the quotient rule to differentiate the lifetime surplus $\pi_c(v)$:

$$\pi_c'(v) = \frac{A'(v)B(v) - A(v)B'(v)}{B(v)^2}$$
$$= \frac{[q_c(v) + \theta(v)q_c'(v)]\, B(v) - A(v)(1-\mu_1)q_c'(v)}{B(v)^2}. \tag{9}$$

Substituting the equilibrium indifference condition from (7) simplifies this expression significantly. Indeed, after careful algebra, we obtain a clean, intuitive expression:

$$\pi_c'(v) = \frac{q_c(v)}{\mu_1 + (1-\mu_1)q_c(v)}. \tag{10}$$

Step 5: Deriving the Final Expression for $\theta'(v)$.

Lastly, recall that at equilibrium, the threshold $\theta(v)$ is directly related to the surplus $\pi_c(v)$ by:

$$\theta(v) = (1-\mu_1)\pi_c(v).$$

Differentiating this relationship with respect to v and substituting (10) yields:

$$\theta'(v) = (1-\mu_1)\pi_c'(v) = \frac{(1-\mu_1)q_c(v)}{\mu_1 + (1-\mu_1)q_c(v)}\ , \tag{11}$$

Step 6: Confirming Monotonicity and Bounds.

Since clearly $0 < \mu_1 < 1$ and $0 < q_c(v) < 1$, it follows directly that

$$0 < \theta'(v) < 1,$$

establishing absolute continuity and strict monotonicity of $g(v) = v - \theta(v)$, as required. This completes the proof. □

A.2. Steady-State Densities of Valuation and Willingness-to-Pay

We first derive the steady-state density of *valuations* among customers who are *present* in the market and then transform it into the density of *willingness-to-pay*.

Entry decision. A customer with valuation v enters the market only if her expected surplus is non-negative:

$$\pi_c(v)\ \geq\ 0.$$

Because $\pi_c(v)$ is strictly increasing in v (see (10)), there exists an endogenous cut-off $\underline{v}_{\mathrm{e}} \in [0,1]$ such that

$$\pi_c(v)\ <\ 0 \iff v < \underline{v}_{\mathrm{e}}, \qquad \pi_c(\underline{v}_{\mathrm{e}}) = 0.$$

We exclude the trivial no-entry case, so the entry set $[\underline{v}_{\mathrm{e}}, 1]$ is nonempty. Denote the entry indicator

$$I(v) := \mathbf{1}\{\, v \geq \underline{v}_{\mathrm{e}} \,\}.$$

Let $f_V^{\mathrm{arr}}(v)$ be the arrival density of valuations on $[0,1]$. The per-type inflow of entrants is then $\lambda_1\, f_V^{\mathrm{arr}}(v)\, I(v)$.

Step 1: Steady-state density of valuations. Each entrant of valuation v exits per period with hazard $\mu_1+(1-\mu_1)q_c(v)$ (impatience with prob. μ_1 plus acceptance with prob. $q_c(v)$ if not impatient). Flow balance implies the steady-state density among *present* customers at the start of a period:

$$c_V(v) \;=\; \frac{\lambda_1\, f_V^{\mathrm{arr}}(v)\, I(v)}{\mu_1+(1-\mu_1)q_c(v)}\,, \qquad 0\le v\le 1. \tag{12}$$

Step 2: Transformation to willingness-to-pay. Define $g(v):=v-\theta(v)$ and let $h=g^{-1}$. Since g is strictly increasing, the change-of-variables formula yields, for $\omega=g(v)$,

$$f_W(\omega) \;=\; \frac{c_V\big(h(\omega)\big)}{g'\big(h(\omega)\big)} \tag{13}$$

$$\overset{(12)}{=} \frac{\lambda_1\, f_V^{\mathrm{arr}}\big(h(\omega)\big)\, I\big(h(\omega)\big)}{\mu_1+(1-\mu_1)q_c\big(h(\omega)\big)}\,\frac{1}{g'\big(h(\omega)\big)}. \tag{14}$$

Using (11), $g'(v)=1-\theta'(v)=\frac{\mu_1}{\mu_1+(1-\mu_1)q_c(v)}$, so the denominator cancels and we obtain the general formula

$$f_W(\omega) \;=\; \frac{\lambda_1}{\mu_1}\, f_V^{\mathrm{arr}}\big(h(\omega)\big)\, I\big(h(\omega)\big). \tag{15}$$

COROLLARY 2 **(Uniform arrivals)**. *If arrival valuations are* uniform *on* $[0,1]$ *with density* λ_1*, i.e.* $f_V^{\mathrm{arr}}(v)\equiv 1$*, then*

$$f_W(\omega) \;=\; \frac{\lambda_1}{\mu_1}, \qquad \omega\in g\big([\underline{v}_{\mathrm{e}},1]\big).$$

Thus, at the *arrival stage* valuations may be non-uniform, and high-v customers tend to stay longer (because they accept more readily), so the *equilibrium valuation density* c_V is typically skewed toward high v. Remarkably, the map $g(v)=v-\theta(v)$ exactly removes this survival bias: the *equilibrium* willingness-to-pay density inherits the shape of the *arrival* valuation density and, for uniform arrivals, becomes uniform with constant density λ_1/μ_1 on its support.

Appendix B: Proof of Proposition 2 and Auxiliary Results on Optimal Prices

Notation. We allow the segment mix at a shown type to depend on the supplier type δ:

$$\beta_k(\delta)\in[0,1], \qquad \beta_1(\delta)+\beta_2(\delta)=1,$$

where $\beta_k(\delta)$ denotes the probability that the meeting customer belongs to segment j when a type-δ supplier is shown. The segment-specific misfit disutilitys are $d_1(\delta)=\delta$ and $d_2(\delta)=1-\delta$. Define the segment slacks

$$t_k(\delta):=\bar{\omega}_k-d_k(\delta)-p(\delta) \qquad (k=1,2).$$

The effective harmonic aggregation and the pooled slack are

$$\Delta_{\mathrm{eff}}(\delta) \;:=\; \Big(\tfrac{\beta_1(\delta)}{\Delta_1}+\tfrac{\beta_2(\delta)}{\Delta_2}\Big)^{-1}, \qquad t(\delta) \;:=\; \Delta_{\mathrm{eff}}(\delta)\Big(\tfrac{\beta_1(\delta)}{\Delta_1}\,t_1(\delta)+\tfrac{\beta_2(\delta)}{\Delta_2}\,t_2(\delta)\Big).$$

The per-meeting success probability and its price sensitivity are

$$\sigma(\delta)=\frac{t(\delta)}{\Delta_{\mathrm{eff}}(\delta)}, \qquad \frac{\partial\sigma}{\partial p}(\delta)=-\,\frac{1}{\Delta_{\mathrm{eff}}(\delta)}.$$

Localization and segmental-exclusivity. At the revenue-optimal recommendation $D^\star$, Appendix G (*segmental-exclusivity*) shows that at almost every δ where capacity binds, exposure is allocated to a single segment; hence $\beta_k(\delta)\in\{0,1\}$ on the occupied intervals. On each such interval $\Delta_{\mathrm{eff}}(\delta)$ and $\kappa(\delta):=\frac{1-\rho}{\rho\,\Delta_{\mathrm{eff}}(\delta)}$ are constant. Unless stated otherwise, all derivatives below are taken on such intervals; global conclusions then follow by concatenation across intervals.

B.1. Supplier profit with two customer segments

Throughout this appendix we omit the explicit dependence on the recommendation profile $\mathbf{D} = (D_1, D_2)$ and fix a supplier of mismfit type $\delta \in [0,1]$. Moreover, we factor out the ad–valorem commission and set $(1-\tau) = 1$ for notational simplicity. This normalization does not affect optimal pricing (FOCs are invariant).

Step 1: lifetime profit for a given price. In any meeting, the supplier earns the posted price p with probability $\sigma(\delta)$. With probability $(1-\sigma)$ the supplier cannot sell and a $(1-\rho)$ discount is implemented. Therefore, the lifetime profit is a simple geometric series:

$$\pi_s(p,\delta) = \frac{p\,\sigma(\delta)}{\rho + (1-\rho)\sigma(\delta)}. \tag{16}$$

Step 2: optimal lifetime profit. Differentiating (16) w.r.t. p gives the first-order condition

$$\big[\sigma + p\sigma'\big]\big[\rho + (1-\rho)\sigma\big] - p\sigma(1-\rho)\sigma' = 0. \tag{17}$$

Because $\sigma' < 0$, (17) solves uniquely for

$$p^*(\delta) = -\frac{\sigma(\delta)\,[\rho + (1-\rho)\sigma(\delta)]}{\rho\,\sigma'(\delta)}. \tag{18}$$

Lemma 2 **(Lifetime profit at the optimal price)**.

$$\pi_s^*(\delta) \;=\; -\frac{\sigma(\delta)^2}{\rho\,\sigma'(\delta)} \;=\; \frac{\Delta_{\text{eff}}\,\sigma(\delta)^2}{\rho}.$$

Proof of Lemma 2 From (16), $\pi_s(p,\delta) = \dfrac{p\,\sigma(\delta)}{\rho + (1-\rho)\sigma(\delta)}$. At the optimal price the FOC (18) holds: $p^*(\delta) = -\dfrac{\sigma(\delta)\,[\rho + (1-\rho)\sigma(\delta)]}{\rho\,\sigma'(\delta)}$. Substituting into (16) and simplifying,

$$\pi_s^*(\delta) \;=\; -\frac{\sigma(\delta)^2}{\rho\,\sigma'(\delta)} \;=\; \frac{\Delta_{\text{eff}}(\delta)\,\sigma(\delta)^2}{\rho},$$

because $\sigma'(\delta) = -1/\Delta_{\text{eff}}(\delta)$. □

Step 3: optimal per-period *profit.* A convenient way to get the per-period figure is to multiply the lifetime profit by

$$B(\delta) := \rho + (1-\rho)\sigma(\delta)\ .$$

Lemma 3 **(Per-period profit)**.

$$\pi^*_{\textit{per-period}}(\delta) \;=\; B(\delta)\,\pi_s^*(\delta) \;=\; \frac{t(\delta)^2}{\Delta_{\text{eff}}} \;+\; \frac{1-\rho}{\rho\,\Delta_{\text{eff}}^2}\,t(\delta)^3.$$

Proof of Lemma 3 Substitute $\sigma = t/\Delta_{\text{eff}}$ and $\sigma' = -1/\Delta_{\text{eff}}$ into $B\,\pi_s^*$ and expand. □

Step 4: expectation across supplier types. Let δ be distributed according to $\mathbf{D} = (D_1, D_2)$. Taking expectations term-wise in Lemma 3 yields

$$\mathbb{E}_D\big[\pi^*_{\text{per-period}}(\delta)\big] = \frac{\mathbb{E}[t^2]}{\Delta_{\text{eff}}} + \frac{1-\rho}{\rho\,\Delta_{\text{eff}}^2}\,\mathbb{E}[t^3]. \tag{19}$$

Corollary 3 **(Platform gross revenue)**. *Under segmental-exclusivity exposure, the per–period platform revenue can be written segment–wise as*

$$R(\mathbf{D}) \;=\; \tau \sum_{k\in\{1,2\}} C_j(\mathbf{D})\,\mathbb{E}_{\delta\sim D_j}\Big[t_k(\delta;\mathbf{D})^2 \;+\; \frac{1-\rho}{\rho\,\Delta_k}\,t_k(\delta;\mathbf{D})^3\Big],$$

where $\Delta_k = \bar{\omega}_k - \underline{\omega}_k$ *is the (segment–k) willingness–to–pay span, and* $t_k(\delta) = \bar{\omega}_k - d_k(\delta) - p^*(\delta)$.

Interpretation. On each segmental-exclusivity block (Appendix G), $\Delta_{\text{eff}}(\delta)$ takes a constant value equal to the active segment's Δ_k, so expectations in Corollary 3 can be read block–wise and then aggregated. The platform matches a total of $C_{\text{tot}} = \Delta_{\text{eff}}\left(\frac{\lambda_1}{\mu_1} + \frac{\lambda_2}{\mu_2}\right)$ active customers per period, so multiplying (19) with C_{tot} yields the desired result.

Remark (piecewise constancy). On any interval where $\beta_k(\delta)$ is constant (e.g., segmental-exclusivity blocks at $D^\star$), $\Delta_{\text{eff}}(\delta)$ and $\kappa(\delta)$ are constant, so (16)–(18) and Lemma 2 go through verbatim with $\sigma_p'(\delta) = -1/\Delta_{\text{eff}}(\delta)$. Where capacity is slack, the precise values of $\beta_k(\delta)$ do not affect the KKT conditions used in Appendices C and G, as the upper bound is not active.

B.2. Monotonicity of Supplier Optimal Price and Optimal Profit

Step 1 – closed form. Remember that $p^*(\delta) = \pi^*_{\text{per-period}}(\delta)/\sigma(\delta)$. Hence, Using $\sigma = t/\Delta_{\text{eff}}$ and equation (19) that provides an expression for $\pi^*_{\text{per-period}}$, the closed–form optimal price

$$p^*(\delta) = t(\delta) + \kappa(\delta)\, t(\delta)^2, \qquad \kappa(\delta) = \frac{1-\rho}{\rho\, \Delta_{\text{eff}}(\delta)}. \tag{20}$$

where

$$t(\delta) = \left(\tfrac{\beta_1}{\Delta_1} t_1(\delta) + \tfrac{\beta_2}{\Delta_2} t_2(\delta)\right)\Delta_{\text{eff}}, \qquad \Delta_{\text{eff}} := \left(\tfrac{\beta_1}{\Delta_1} + \tfrac{\beta_2}{\Delta_2}\right)^{-1}.$$

.

Step 2 – first derivative. On a block where the segment mix is constant, define

$$\Delta_B(\delta) := \frac{\beta_1(\delta)/\Delta_1 - \beta_2(\delta)/\Delta_2}{\beta_1(\delta)/\Delta_1 + \beta_2(\delta)/\Delta_2} \in [-1,1], \quad \text{so under segmental-exclusivity } \Delta_B(\delta) \equiv \pm 1.$$

Because $t_1 = \bar{\omega}_1 - \delta - p^*$ and $t_2 = \bar{\omega}_2 - (1-\delta) - p^*$, we obtain that $dt/d\delta = -\Delta B - dp^*/d\delta$. Differentiate $p^* = t + \kappa t^2$ once and solve for the slope:

$$\frac{dp^*}{d\delta} = -\frac{(1+2\kappa\, t)\,\Delta_B}{2+2\kappa\, t}. \tag{21}$$

segmental-exclusivity simplification. On a segmental-exclusivity block, $\Delta_B \in \{\pm 1\}$ and $2 + 2\kappa\, t(\delta) > 0$. Hence $\frac{dp^*}{d\delta} \in (-1, 0)$ when the occupying segment is $k = 1$ (so $\Delta_B = +1$), and $\frac{dp^*}{d\delta} \in (0,1)$ when it is $k = 2$ (so $\Delta_B = -1$). Consequently, for segment $k = 1$ the effective price $p^*(\delta) + \delta$ is strictly increasing in δ, whereas for segment $k = 2$ the effective price $p^*(\delta) + (1-\delta)$ is strictly decreasing in δ.

Step 3 – second derivative. Differentiate the expression just obtained:

$$\frac{d^2 p^*}{d\delta^2} = \frac{2\kappa\, \Delta_B^2}{(2+2\kappa\, t)^3} > 0. \tag{22}$$

because $\kappa > 0$ and the denominator is positive due to $t > 0$.

B.3. Auxiliary Results

LEMMA 4. *In any stationary equilibrium with* $M_k := \mathbb{E}_{\delta\sim D_j}[p^*(\delta)+d_k(\delta)]$*, we have the following:*

$$\underline{\omega}_k(M_k) = \frac{\mu_1+(1-\mu_1)\,M_k}{1+\rho}, \qquad \bar{\omega}_k(M_k) = \underline{\omega}_k(M_k)+\Delta_k(M_k), \quad \Delta_k(M_k) = \frac{\rho}{1+\rho}\big(\mu_k+(1-\mu_k)\,M_k\big).$$

Proof of Lemma 4. Fix a segment j and write $\mu_1=\mu_k$, $\lambda_1=\lambda_k$, $\underline{\omega}=\underline{\omega}_k$, $\bar{\omega}=\bar{\omega}_k$, $\Delta=\Delta_k=\bar{\omega}-\underline{\omega}$, and $M=M_k=\mathbb{E}_{\delta\sim D_j}\big[p^*(\delta)+d_k(\delta)\big]$. We proceed in three algebraic steps.

Step 1 (Uniform WTP ⇒ "rectangle" mass). By the uniformity of willingness-to-pay at equilibrium, the density of active customers over $\omega\in[\underline{\omega},\bar{\omega}]$ is constant and equals λ_1/μ_1. Hence, the steady-state mass of active customers in this segment is

$$C \;=\; \int_{\underline{\omega}}^{\bar{\omega}} \frac{\lambda_1}{\mu_1}\,d\omega \;=\; \frac{\lambda_1}{\mu_1}\,\Delta.$$

Step 2 (Structural equation for customer mass). Because arrivals of valuations are uniform on $[0,1]$ and the entry set is an interval, the measure that enters each period is

$$F \;=\; \lambda_1\,(1-\underline{\omega}),$$

since at the entry cutoff v_e the indifference condition gives $\theta_j(v_e)=0$ and thus $\underline{\omega}=\omega_k(v_e)=v_e$. A present customer exits within a period either by impatience or purchase, so the per-period exit hazard is $\mu_1+(1-\mu_1)\,q$, where q is the period acceptance probability for a random active customer of this segment. With uniform WTP and effective price $X=p^*(\delta)+d_k(\delta)$ shown according to D_j, we have for $X\in[\underline{\omega},\bar{\omega}]$:

$$\Pr\{\omega\ge X\} = \frac{\bar{\omega}-X}{\Delta} \quad\Longrightarrow\quad q \;=\; \mathbb{E}\left[\frac{\bar{\omega}-X}{\Delta}\right] \;=\; \frac{\bar{\omega}-M}{\Delta}.$$

The expected spell length equals $1/\big(\mu_1+(1-\mu_1)q\big)$, hence by geometric accumulation

$$C \;=\; \frac{F}{\mu_1+(1-\mu_1)q} \;=\; \frac{\lambda_1(1-\underline{\omega})}{\mu_1+(1-\mu_1)\frac{\bar{\omega}-M}{\Delta}}.$$

Equating this with Step 1's expression for C and simplifying gives

$$\frac{\lambda_1}{\mu_1}\Delta \;=\; \frac{\lambda_1(1-\underline{\omega})}{\mu_1+(1-\mu_1)\frac{\bar{\omega}-M}{\Delta}} \;\Longrightarrow\; \bar{\omega} \;=\; \mu_1+(1-\mu_1)\,M. \tag{23}$$

Step 3 ($p(0)=\underline{\omega}$ and the supplier FOC). Let $\delta_k^\star$ be the best-match type for segment j, so $d_k(\delta_k^\star)=0$. At the bottom of the market, the entry-margin condition implies that the lowest effective price equals the lowest WTP,

$$p^*(\delta_k^\star) \;=\; \underline{\omega}.$$

With uniform WTP on $[\underline{\omega},\bar{\omega}]$, a type-$\delta$ supplier's per-meeting success probability is $\sigma(\delta)=\big(\bar{\omega}-[p(\delta)+d_k(\delta)]\big)/\Delta$ and $\partial\sigma/\partial p=-1/\Delta$. Solving the (standard) dynamic FOC for the posted price under geometric discounting yields the quadratic best-response form (see Appendix B.2):

$$p^*(\delta) \;=\; t(\delta) \;+\; \frac{1-\rho}{\rho\,\Delta}\,t(\delta)^2, \qquad t(\delta) := \bar{\omega}-d_k(\delta)-p^*(\delta).$$

Evaluating at $\delta = \delta_k^\star$ gives $t(\delta_k^\star) = \bar{\omega} - p^*(\delta_k^\star) = \bar{\omega} - \underline{\omega} = \Delta$, hence

$$\underline{\omega} \;=\; \Delta + \frac{1-\rho}{\rho\,\Delta}\,\Delta^2 \;=\; \Delta\left(1+\frac{1-\rho}{\rho}\right) \;=\; \frac{\Delta}{\rho}.$$

Therefore $\Delta = \rho\underline{\omega}$ and $\bar{\omega} = (1+\rho)\,\underline{\omega}$.

Step 4 (Assemble). Combining ($\star$) with $\bar{\omega} = (1+\rho)\,\underline{\omega}$ yields

$$\underline{\omega} \;=\; \frac{\mu_1 + (1-\mu_1)\,M}{1+\rho}, \qquad \Delta \;=\; \rho\underline{\omega} \;=\; \frac{\rho}{1+\rho}\,\big(\mu_1 + (1-\mu_1)\,M\big),$$

and $\bar{\omega} = \underline{\omega} + \Delta$ follows immediately. Restoring the segment index k gives the stated formulas. □

LEMMA 5 **(Strict convexity results)**. *The per-period revenue of the platform stated in Corollary 3 is* strictly convex *in δ: $\frac{\partial^2 r}{\partial \delta^2} > 0$ for every $\delta \in [0, \bar{\delta}]$.*

Proof of Lemma 5. Write $p^*(\delta) = t(\delta) + \kappa t^2(\delta)$ with $\kappa = \dfrac{1-\rho}{\rho\Delta} > 0$ and $t(\delta) = \bar{\omega} - \delta - p^*(\delta)$.

First derivative of t. By the single-segment variant of equation in Appendix B where $\Delta B = 1$, we have the following:

$$\frac{dp^*}{d\delta} = -\frac{(1+2\kappa t)}{2+2\kappa t} \in (-1,0) \quad\Longrightarrow\quad t'(\delta) = -\frac{1}{2+2\kappa t(\delta)} \in (-1,0).$$

Second derivative of t. Similarly to the calculation of the first derivative, the single-segment equivalent of equation (21) where $\Delta B = 1$ yields the following:

$$\frac{d^2 p^*}{d\delta^2} \;=\; \frac{2\,\kappa\,\Delta_B^2}{\big(2+2\kappa\, t(\delta)\big)^3} \;>\; 0 \qquad\Longrightarrow\qquad t''(\delta) \;=\; \frac{\kappa\, t'(\delta)}{2\,\big(1+\kappa\, t(\delta)\big)^2} \;<\; 0. \tag{24}$$

because $t' < 0$ and $\kappa > 0$.

Convexity of the revenue $r(D)$ with respect to δ. Differentiating $r(\delta) = t^2 + \kappa t^3$ yields

$$\frac{\partial r}{\partial \delta} \;=\; (2t + 3\kappa t^2)\, t'(\delta), \qquad \frac{\partial^2 r}{\partial \delta^2} \;=\; (2+6\kappa t)\,\big[t'(\delta)\big]^2 \;+\; (2t+3\kappa t^2)\, t''(\delta).$$

Using $t'(\delta) = -1/(2+2\kappa t)$ and $t''(\delta) = \kappa t'(\delta)/\{2(1+\kappa t)^2\}$ gives

$$\frac{\partial^2 r}{\partial \delta^2} \;=\; \frac{3\kappa^2 t(\delta)^2 + 6\kappa\, t(\delta) + 2}{4\big(1+\kappa\, t(\delta)\big)^3} \;>\; 0,$$

establishing strict convexity. □

LEMMA 6 **(Comparative statics in M)**. *Fix a stationary single–segment equilibrium. For each seller type δ, let the effective price be $X(\delta) := p^*(\delta) + \delta$ and denote its mean by $M := \mathbb{E}[X]$. Define*

$$\bar{\omega}(M) = \mu_1 + (1-\mu_1)M, \qquad \Delta(M) = \bar{\omega}(M) - \omega(M) = \frac{\rho}{1+\rho}\big(\mu_1 + (1-\mu_1)M\big),$$

so that

$$\frac{\partial \Delta}{\partial M} = \frac{\rho}{1+\rho}\,(1-\mu_1).$$

Let

$$t := \bar{\omega}(M) - X, \qquad k := \frac{1-\rho}{\rho\,\Delta(M)}, \qquad r := t^2 + k\,t^3.$$

Then:

$$\frac{\partial t}{\partial M} > 0, \qquad 0 < \frac{\partial p^*}{\partial M} = \frac{\partial X}{\partial M} < 1, \qquad \frac{\partial r}{\partial M} = t\,(1-\mu_1)\left[1 + \frac{k\,t}{2+2\,k\,t}\left(1 + \frac{t^2}{\Delta^2}\,\frac{1-\rho}{1+\rho}\right)\right] > 0.$$

Proof of Lemma 6 By the seller FOC (derived earlier), the equilibrium price admits the closed form

$$p^* \,=\, t + k\,t^2, \qquad k = \frac{1-\rho}{\rho\,\Delta(M)}, \qquad t = \bar{\omega}(M) - X,$$

with $0 < t < \Delta$ (since $X \in [\omega, \bar{\omega}]$). Differentiating $r = t^2 + k\,t^3$ w.r.t. M gives

$$\frac{\partial r}{\partial M} = (2t + 3k\,t^2)\,\frac{\partial t}{\partial M} + t^3\,\frac{\partial k}{\partial M}.$$

From $\Delta(M) = \frac{\rho}{1+\rho}(\mu_1 + (1-\mu_1)M)$ we have

$$\frac{\partial k}{\partial M} = -\,\frac{1-\rho}{\rho}\,\frac{1}{\Delta^2}\,\frac{\partial \Delta}{\partial M} = -\,\frac{1-\rho}{\rho}\,\frac{1}{\Delta^2}\,\frac{\rho}{1+\rho}\,(1-\mu_1) = -\,\frac{1-\mu_1}{\Delta^2}\,\frac{1-\rho}{1+\rho}. \tag{25}$$

Next, differentiate $t = \bar{\omega}(M) - X$ and use $\bar{\omega}'(M) = 1 - \mu_1$:

$$\frac{\partial t}{\partial M} = (1-\mu_1) - \frac{\partial X}{\partial M} \quad \text{and} \quad \frac{\partial p^*}{\partial M} = \frac{\partial X}{\partial M} = (1 + 2k\,t)\frac{\partial t}{\partial M} + t^2\,\frac{\partial k}{\partial M}.$$

Eliminating $\partial X/\partial M$ yields

$$(2 + 2k\,t)\,\frac{\partial t}{\partial M} = (1-\mu_1) - t^2\,\frac{\partial k}{\partial M}.$$

Using (25), the right-hand side equals

$$(1-\mu_1) + \,(1-\mu_1)\,\frac{t^2}{\Delta^2}\,\frac{1-\rho}{1+\rho} \,>\, 0,$$

while the left-hand side has the positive factor $2 + 2k\,t > 0$. Hence

$$\frac{\partial t}{\partial M} = \frac{(1-\mu_1)\Big(1 + \frac{t^2}{\Delta^2}\,\frac{1-\rho}{1+\rho}\Big)}{2(1+k\,t)} \,>\, 0\ ,$$

proving the first inequality in the Lemma statement. Then,

$$\frac{\partial p^*}{\partial M} = (1-\mu_1) - \frac{\partial t}{\partial M} = (1-\mu_1)\left[1 - \frac{1 + \frac{t^2}{\Delta^2}\,\frac{1-\rho}{1+\rho}}{2(1+k\,t)}\right].$$

Because $0 < t/\Delta < 1$, $\frac{1-\rho}{1+\rho} < 1$, and $2(1 + k\,t) > 2$, the bracketed term lies in $(0, 1)$, so $0 < \partial p^*/\partial M < 1$, establishing the second inequality in the Lemma statement.

Finally, substitute the expressions for $\partial t/\partial M$ and $\partial k/\partial M$ into $\partial r/\partial M$ and simplify:

$$\begin{aligned}
\frac{\partial r}{\partial M} &= (2t + 3k\,t^2)\frac{\partial t}{\partial M} + t^3\,\frac{\partial k}{\partial M} \\
&= t(1-\mu_1)\frac{2 + 3k\,t}{2(1+k\,t)} \,+\, t^3\left[\frac{k\,t}{2(1+k\,t)}\right]\frac{(1-\mu_1)}{\Delta^2}\,\frac{1-\rho}{1+\rho} \\
&= t(1-\mu_1)\left[1 + \frac{k\,t}{2 + 2\,k\,t}\Big(1 + \frac{t^2}{\Delta^2}\,\frac{1-\rho}{1+\rho}\Big)\right],
\end{aligned}$$

which is strictly positive since all multiplicative factors are positive. This proves last inequality in the Lemma statement. This completes the proof.

Appendix C: Proof of Theorem 3

Before we provide the fully technical arguments, we start by providing a high-level proof sketch.

C.1. Proof sketch

We first define the marginal-revenue density (Gateaux derivative) $H(\delta; D)$,[9] which captures how total platform revenue changes when a small amount of recommendation mass is shifted towards supplier type δ, given that the current recommendation distribution is D. Next, we show that $H(\delta; D)$ has a U-shaped pattern for any D: it decreases up to a unique interior point $\delta^\star$, and then increases thereafter. Intuitively, the U-shaped pattern of the marginal-revenue density explains why it can be beneficial for the platform to skip the "middle-δ-region" and divert customers to their least-preferred suppliers. With this property established, we then write down the Karush–Kuhn–Tucker (KKT) conditions associated with the platform's optimisation problem under the endogenous capacity constraint. These conditions yield clear implications for the optimal solution: whenever the upper bound binds, we must have $H(\delta) \geq \nu$; whenever no probability is assigned, we must have $H(\delta) \leq \nu$; and equality $H(\delta) = \nu$ holds only at points where the constraint is slack (where ν is the Lagrange multiplier for the mass constraint). Now that we have established this crucial relationship, we conclude that the U-shape of H implies a *bind-or-zero* structure for the optimal recommendation density. Specifically, the U-shape implies that we need to find two cut-offs δ_0 and δ_1 with $0 \leq \delta_0 \leq \delta^\star \leq \delta_1 \leq \bar{\delta}$ such that the upper bound is saturated on the intervals $[0, \delta_0]$ and $[\delta_1, \bar{\delta}]$. At the same time, no probability is placed in the middle region (δ_0, δ_1) due to the U-shape (the marginal benefit is highest at the extremes). Since the ceiling ratio $s(\delta, D)/C(D)$ is strictly decreasing, these two thresholds are uniquely determined by mass-balance equations. Finally, we observe how the location of (δ_0, δ_1) depends on the supplier arrival rate m that directly influences the capacity levels. When m is small, the probability concentrates exclusively on low-misfit suppliers, leaving the upper block empty ($\delta_1 = \bar{\delta}$). As m increases, an additional block of high-misfit suppliers emerges, and eventually, if capacity becomes completely slack as $m \to \infty$, the solution simplifies to the two-point distribution supported at the extremes $\{0, \bar{\delta}\}$.

C.2. The Main Challenges in Analysis

The proof sketch above outlines the core strategy: analyzing the Gateaux derivative $H(\delta; D)$ and leveraging its convexity. However, the optimization problem analyzed here is complex and presents additional challenges beyond showing the convexity of the Gateaux derivative, as it involves maximization over a space of distributions (D), where both the objective function and the constraints depend endogenously on D through the equilibrium fixed point $M(D)$. The subsequent detailed sections address three fundamental mathematical challenges necessary for a complete proof.

Challenge 1: Handling State-Dependent (Endogenous) Constraints (Appendix C.5 and C.7). Perhaps the most significant technical challenge arises because the capacity constraints are endogenous. The capacity ceiling, $\overline{D}(\delta; D) = s(\delta; D)/C(D)$, depends on the policy D itself.

- Intuitively, the platform is optimizing against a moving target. Changing the recommendation strategy D alters the market equilibrium (prices, customer mass $C(D)$, supplier mass $s(\delta; D)$), which, in turn, changes the very capacity constraints the platform faces. Standard optimization techniques assume constraints are fixed. We must rigorously verify that the feedback effects on the constraints do not invalidate the intuition derived from analyzing a simplified "frozen" problem.

[9] The Gateaux derivative $H(\delta; D)$ used throughout this analysis rigorously accounts for the feedback effect of a change in D on the equilibrium (prices, willingness-to-pay, and total customer mass $C(D)$).

- Standard KKT theory does not directly apply when constraints depend on the decision variable. We ensure that the optimality conditions derived for the simplified problem correctly characterize the optimum of the original problem (Appendix C.7). This validation step, referred to as “Lifting" is crucial for rigor.

Challenge 2: Ensuring Existence and Uniqueness of the Solution (Appendix C.6). While the U-shape of $H(\delta)$ suggests prioritizing the extremes, this must be formalized to ensure the resulting policy is feasible and unique.

- Intuitively, the KKT conditions imply that the platform uses capacity where the marginal revenue $H(\delta)$ exceeds the shadow price of mass, ν. This shadow price ν (the Lagrange multiplier for the mass constraint) is endogenous; it acts as an equilibrium price for recommendation visibility. We must ensure there is a unique price ν^* that exactly "clears the market"—meaning the total mass allocated equals 1 ($\int D^* = 1$).
- Appendix C.6 provides the rigorous mechanism. It defines the mass map $\Phi(\nu)$ (total mass allocated at price ν) and proves it is continuous and strictly decreasing. The Intermediate Value Theorem then guarantees the existence of a unique ν^*, uniquely defining the cutoffs (δ_0, δ_1).

Challenge 3: Characterizing Comparative Statics and Market Thickness (Appendix C.8). Theorem 3 characterizes both the structure of D^* and its behavior as $m \to \infty$. This analysis is vital for managerial insights regarding market thickness. Appendix C.8 rigorously proves, using monotone comparative statics, how the optimal policy changes as supplier capacity (m) increases. This leads to the crucial insight that in thick markets, the optimal strategy converges to the polarized two-point distribution.

We now proceed with the detailed mathematical arguments. We start by describing two key notation conventions:

- $K(\delta; M)$: The *revenue kernel* (expected revenue per unit of recommendation mass on type δ, given equilibrium scalar M).
- $H(\delta; D)$: The *Gateaux derivative* (marginal-revenue density) of $R(D)$, defined as $H(\delta; D) \equiv \frac{\delta R}{\delta D}(\delta \mid D)$.

C.3. Gateaux Derivative with Equilibrium Feedback

The platform's per-period revenue under density D is

$$R(D) = \int_0^{\bar{\delta}} K(\delta; M(D))\, D(\delta)\, d\delta.$$

The scalar $M(D)$ summarizes the mean effective price, defined by the equilibrium fixed point

$$M(D) = \int_0^{\bar{\delta}} X(\delta; M(D))\, D(\delta)\, d\delta,$$

where $X(\delta; M) = p^*(\delta; M) + \delta$ is the effective price. We use the notation $\mathbb{E}_D[\cdot] = \int_0^{\bar{\delta}} (\cdot)\, D(\delta)\, d\delta$.

To find the marginal revenue density $H(\delta; D)$, we analyze how $R(D)$ changes under a small, mass-preserving perturbation. Consider $D_\varepsilon = D + \varepsilon h$ with $\int h = 0$, and let $M_\varepsilon = M(D_\varepsilon)$. We calculate the derivative $\frac{d}{d\varepsilon} R(D_\varepsilon)\big|_{\varepsilon=0}$.

Differentiating $R(D_\varepsilon)$ with respect to ε at $\varepsilon = 0$ yields two terms: the direct effect of changing D, and the indirect effect via the change in M.

$$\frac{d}{d\varepsilon}R(D_\varepsilon)\Big|_{\varepsilon=0} = \underbrace{\int K(\delta;M)\,h(\delta)\,d\delta}_{\text{Direct effect}} + \underbrace{\frac{dM_\varepsilon}{d\varepsilon}\Big|_{\varepsilon=0}\cdot \mathbb{E}_D\big[\partial K/\partial M\big]}_{\text{Indirect (feedback) effect}}. \tag{26}$$

We differentiate the fixed-point condition for M_ε w.r.t. ε and solve for $dM/d\varepsilon$:

$$\frac{dM_\varepsilon}{d\varepsilon}\Big|_{\varepsilon=0} = \frac{\int X(\delta;M)\,h(\delta)\,d\delta}{1-\mathbb{E}_D\big[\partial X/\partial M\big]}. \tag{27}$$

By Lemma 6, we have $0 < \partial X/\partial M < 1$, ensuring the denominator is strictly positive.

We substitute (27) into (26). We define the *feedback multiplier* $\Gamma(D)$:

$$\Gamma(D) \equiv \frac{\mathbb{E}_D[\partial K/\partial M]}{1-\mathbb{E}_D[\partial X/\partial M]}.$$

By Lemma 6, $\partial K/\partial M > 0$ (via $\partial r/\partial M > 0$), hence $\Gamma(D) > 0$. Collecting terms, we obtain:

$$\frac{d}{d\varepsilon}R(D_\varepsilon)\Big|_{\varepsilon=0} = \int \underbrace{\Big(K(\delta;M(D)) + \Gamma(D)\,X(\delta;M(D))\Big)}_{=:\,H(\delta;D)}\,h(\delta)\,d\delta. \tag{28}$$

The integrand is the Gateaux derivative, the true marginal-revenue density:

$$H(\delta;D) = K(\delta;M(D)) + \Gamma(D)\,X(\delta;M(D))$$

H ranks types by their total contribution to revenue, including equilibrium adjustments.

C.4. Strict Convexity and the “U-Shape”

We now show that the marginal revenue density $\delta \mapsto H(\delta;D)$ is strictly convex on $[0,\bar{\delta}]$.

(i) *Kernel convexity.* Lemma 5 proves that the underlying revenue component r is strictly convex in δ. Since $K(\delta;M)$ is a positive scaling of $r(\cdot)$, the kernel $K(\delta;M)$ is strictly convex in δ.

(ii) *Effective-price convexity.* Equation (22)) shows $p^{*\prime\prime}(\delta) > 0$. Therefore, the effective price $X(\delta;M) = p^*(\delta;M) + \delta$ is strictly convex in δ.

Since the feedback multiplier $\Gamma(D) > 0$ (established in Appendix C.3), $H(\cdot;D)$ is a positive weighted sum of two strictly convex functions. Thus, $H(\cdot;D)$ is strictly convex in δ. Strict convexity implies a unique minimizer of $H(\cdot;D)$. This corresponds to a "U-shape" (decreasing then increasing) provided the minimizer is interior (verified later in this section).

C.5. KKT Conditions and the Bind–Gap–Bind Rule

The platform maximizes $R(D)$ subject to constraints. The primary challenge is that the capacity ceiling $\overline{D}(\delta;D) := s(\delta;D)/C(D)$ is state-dependent (it depends on D). To address this, we employ a *freeze–then–lift* approach. We first assume an optimum D^* exists. We "freeze" the equilibrium objects at their values corresponding to D^*, and analyze the resulting simplified optimization problem. In Section C.7, we provide the rigorous justification (the "Lifting" argument), proving that the necessary conditions for this simplified problem indeed characterize the optimum of the original, state-dependent problem.

Define the fixed ceiling $\overline{D}^*(\delta) := s(\delta; D^*)/C(D^*)$. The *frozen* optimization problem is:

$$(\mathrm{P_{fix}}) \qquad \max_{D} \quad \int_0^{\bar\delta} H(\delta; D^*)\, D(\delta)\, d\delta$$
$$\text{s.t.} \quad \int_0^{\bar\delta} D(\delta)\, d\delta = 1, \qquad 0 \le D(\delta) \le \overline{D}^*(\delta) \quad \text{a.e.}$$

This is a linear program. We introduce multipliers: ν for the mass constraint, $\zeta(\delta) \ge 0$ for the lower bound $D(\delta) \ge 0$, and $\eta(\delta) \ge 0$ for the upper bound $\overline{D}^*(\delta) - D(\delta) \ge 0$. The Lagrangian is:

$$\mathcal{L}(D, \nu, \eta, \zeta) = \int H(\delta; D^*) D(\delta)\, d\delta - \nu\Big(\int D - 1\Big)$$
$$+ \int \zeta(\delta)\, D(\delta)\, d\delta + \int \eta(\delta)\big(\overline{D}^*(\delta) - D(\delta)\big)\, d\delta.$$

The KKT conditions are:

$$\text{Stationarity:} \quad H(\delta; D^*) - \nu + \zeta(\delta) - \eta(\delta) = 0 \quad \text{a.e.,} \tag{29}$$
$$\text{Comp. Slackness:} \quad \zeta(\delta)\, D(\delta) = 0, \quad \eta(\delta)\left(\overline{D}^*(\delta) - D(\delta)\right) = 0. \tag{30}$$

These conditions imply the standard *bind–or–zero* rule. Rearranging (29) gives $H(\delta; D^*) - \nu = \eta(\delta) - \zeta(\delta)$.

- If $H(\delta; D^*) > \nu$, then $\eta(\delta) > 0$. By (30), $D^*(\delta) = \overline{D}^*(\delta)$ (upper bound binds) and $\zeta(\delta) = 0$.
- If $H(\delta; D^*) < \nu$, then $\zeta(\delta) > 0$. By (30), $D^*(\delta) = 0$ (lower bound binds) and $\eta(\delta) = 0$.

Because $H(\cdot; D^*)$ is strictly convex (U-shaped), the set where $H > \nu$ consists of intervals at the extremes, and the set where $H < \nu$ is the interval in the middle. This confirms the bind-gap-bind structure.

C.6. Existence and Uniqueness of the Two Cutoffs

We now rigorously establish the existence and uniqueness of the thresholds (δ_0, δ_1) that satisfy the KKT conditions (C.5) and the mass constraint $\int D^* = 1$. Let $H^*(\delta) := H(\delta; D^*)$.

Step 1: Defining the Cutoff Functions. By C.4, $H^*(\delta)$ is strictly convex and continuously differentiable on $[0, \bar\delta]$. Let $\delta^\star = \arg\min_\delta H^*(\delta)$. Define the range of relevant shadow prices as $\mathcal{N} = [\min H^*, \max H^*]$.

Due to strict convexity, H^* is strictly decreasing on $[0, \delta^\star]$ (where $H^{*\prime}(\delta) \le 0$) and strictly increasing on $[\delta^\star, \bar\delta]$ (where $H^{*\prime}(\delta) \ge 0$). We define the cutoff functions $\alpha(\nu)$ and $\beta(\nu)$.

- $\alpha(\nu)$ is the unique solution to $H^*(\delta) = \nu$ in $[0, \delta^\star]$, if $\nu \le H^*(0)$; otherwise $\alpha(\nu) = 0$.
- $\beta(\nu)$ is the unique solution to $H^*(\delta) = \nu$ in $[\delta^\star, \bar\delta]$, if $\nu \le H^*(\bar\delta)$; otherwise $\beta(\nu) = \bar\delta$.

The superlevel set $\{\delta : H^*(\delta) \ge \nu\}$ is $[0, \alpha(\nu)] \cup [\beta(\nu), \bar\delta]$.

Step 2: Properties of the Cutoff Functions. We establish continuity and strict monotonicity. Where $\alpha(\nu) \in (0, \delta^\star)$, we have $H^{*\prime}(\alpha(\nu)) < 0$. We apply the Implicit Function Theorem (IFT) to $G(\delta, \nu) = H^*(\delta) - \nu = 0$. Since $\partial G/\partial\delta = H^{*\prime}(\delta) \ne 0$, the IFT guarantees that $\alpha(\nu)$ is continuously differentiable and:

$$\alpha'(\nu) = -\frac{\partial G/\partial \nu}{\partial G/\partial \delta} = \frac{1}{H^{*\prime}(\alpha(\nu))} < 0.$$

Similarly, where $\beta(\nu) \in (\delta^\star, \bar\delta)$, we have $H^{*\prime}(\beta(\nu)) > 0$, so:

$$\beta'(\nu) = \frac{1}{H^{*\prime}(\beta(\nu))} > 0.$$

Thus, α and β are continuous and strictly monotone on their respective domains.

Step 3: The Mass Map and its Monotonicity. We define the mass map $\Phi : \mathcal{N} \to \mathbb{R}_{\geq 0}$, using the frozen ceiling $\overline{D}^*(\delta)$:

$$\Phi(\nu) := \int_0^{\alpha(\nu)} \overline{D}^*(\delta)\, d\delta + \int_{\beta(\nu)}^{\bar{\delta}} \overline{D}^*(\delta)\, d\delta.$$

Continuity of $\Phi(\nu)$ follows from the continuity of the integration limits (α, β).

To show strict monotonicity, we differentiate $\Phi(\nu)$ using the Leibniz Integral Rule (where α and β are interior):

$$\Phi'(\nu) = \overline{D}^*(\alpha(\nu)) \cdot \alpha'(\nu) - \overline{D}^*(\beta(\nu)) \cdot \beta'(\nu).$$

Since the supplier density is positive ($\overline{D}^*(\delta) > 0$), and given $\alpha'(\nu) < 0$ and $\beta'(\nu) > 0$, we conclude $\Phi'(\nu) < 0$. Therefore, $\Phi(\nu)$ is strictly decreasing over $\mathcal{N}$.

Step 4: Existence and Uniqueness of ν^*. We analyze the range of the mass map $\Phi(\nu)$ to establish the existence of a unique shadow price ν^* such that $\Phi(\nu^*) = 1$. We examine the behavior of $\Phi(\nu)$ at the boundaries of the domain $\mathcal{N} = [\min H^*, \max H^*]$.

Step 4.1: Analyzing the Boundaries of $\mathcal{N}$. Let $\nu_{\min} = \min H^*$ and $\nu_{\max} = \max H^*$.

- **Lower Boundary ($\nu = \nu_{\min}$):** When the shadow price is set to the minimum marginal revenue, the superlevel set $\{\delta : H^*(\delta) \geq \nu_{\min}\}$ covers the entire domain $[0, \bar{\delta}]$. The mass map evaluates to the total normalized capacity available at the equilibrium D^*:

$$\Phi(\nu_{\min}) = \int_0^{\bar{\delta}} \overline{D}^*(\delta)\, d\delta.$$

- **Upper Boundary ($\nu = \nu_{\max}$):** When the shadow price is set to the maximum marginal revenue, the superlevel set $\{\delta : H^*(\delta) \geq \nu_{\max}\}$ contains only the maximizers of H^*. Since H^* is strictly convex, it cannot be constant on any interval of positive measure. Therefore, the set of maximizers has measure zero (it consists only of the boundaries $\{0, \bar{\delta}\}$ if the maximum is attained there). Thus, $\Phi(\nu_{\max}) = 0$.

Step 4.2: Rigorous Justification of Sufficient Capacity ($\Phi(\nu_{\min}) \geq 1$). We must rigorously establish that the total capacity at the optimum is sufficient to accommodate a total mass of 1. This is a direct consequence of the feasibility of the optimal solution D^*. By definition, D^* must be feasible for the original problem (P). Feasibility requires satisfaction of two constraints:

1. Mass constraint: $\int_0^{\bar{\delta}} D^*(\delta) d\delta = 1$.
2. Capacity constraint: $D^*(\delta) \leq \overline{D}(\delta; D^*)$ almost everywhere.

Recall that the frozen ceiling is defined as $\overline{D}^*(\delta) = \overline{D}(\delta; D^*)$. We integrate the capacity constraint (2) over the domain $[0, \bar{\delta}]$:

$$\int_0^{\bar{\delta}} D^*(\delta) d\delta \leq \int_0^{\bar{\delta}} \overline{D}^*(\delta) d\delta.$$

We substitute the mass constraint (1) into the left-hand side, and the definition of $\Phi(\nu_{\min})$ into the right-hand side:

$$1 \leq \Phi(\nu_{\min}).$$

This rigorously confirms that the total capacity is sufficient. (Alternatively, suppose for contradiction that $\Phi(\nu_{\min}) < 1$. Then $\int_0^{\bar{\delta}} D^*(\delta)d\delta \leq \Phi(\nu_{\min}) < 1$, which contradicts the requirement that D^* is a probability distribution.)

Step 4.3: Applying the Intermediate Value Theorem. We have established that $\Phi(\nu)$ is continuous and strictly decreasing (Step 3). Its range spans from $\Phi(\nu_{\min}) \geq 1$ down to $\Phi(\nu_{\max}) = 0$.

The Intermediate Value Theorem states that a continuous function on a closed interval takes on every value between its minimum and maximum. Applying this theorem to $\Phi(\nu)$ on the interval $\mathcal{N} = [\nu_{\min}, \nu_{\max}]$, since $0 \leq 1 \leq \Phi(\nu_{\min})$, there must exist a $\nu^* \in \mathcal{N}$ such that $\Phi(\nu^*) = 1$. Furthermore, because $\Phi(\nu)$ is strictly decreasing, this ν^* is unique. This unique shadow price defines the unique cutoffs $\delta_0 := \alpha(\nu^*)$ and $\delta_1 := \beta(\nu^*)$, establishing the unique bind–gap–bind structure.

C.7. Lifting to the State-Dependent Ceiling

The analysis in Appendix C.5 characterized the optimal solution D^* based on the "frozen" optimization problem (P_{fix}). We must now rigorously verify that this characterization holds for the original problem (P), where the ceiling $\overline{D}(\delta; D)$ is state-dependent (endogenous).

High-level ideas: We justify the "Freeze-then-Lift" approach by demonstrating that the first-order necessary conditions (FONC) of the original problem (P) imply that the optimum D^* must also solve the specific linear program (P_{fix}). We rely on the established regularity of the revenue functional $R(D)$. As shown in Section C.3, $R(D)$ is Gateaux differentiable, and its derivative at D^* in the direction h is given by $\delta R[D^*; h] = \int H(\delta; D^*)h(\delta)d\delta$.

Lemma 7 **(FONC Implies Optimality in the Frozen LP)**. *If D^* is an optimal solution to the original problem (P), then D^* is also an optimal solution to the frozen linear program (P_{fix}) defined in Section C.5.*

Proof of Lemma 7 Let D^* be optimal for (P). We analyze the First-Order Necessary Conditions (FONC) for optimality.

Step 1: The Feasible Set and Tangent Cone. Let $\mathcal{F}$ be the feasible set for the original problem (P). Let $\mathcal{F}_{\text{fix}}$ be the feasible set for the frozen problem (P_{fix}). Since D^* is feasible for (P), it satisfies $0 \leq D^*(\delta) \leq \overline{D}(\delta; D^*)$. By definition of the frozen ceiling $\overline{D}^*(\delta) = \overline{D}(\delta; D^*)$, we have $D^* \in \mathcal{F}_{\text{fix}}$. The Tangent Cone $TC(D^*, \mathcal{F})$ is the set of directions h such that $D^* + \varepsilon h$ approximates a feasible path in $\mathcal{F}$ starting at D^*.

Step 2: The First-Order Necessary Condition (FONC). If D^* maximizes $R(D)$ over $\mathcal{F}$, and $R(D)$ is Gateaux differentiable, then the derivative at D^* must be non-positive along any direction h in the Tangent Cone:

$$\delta R[D^*; h] = \int_0^{\bar{\delta}} H(\delta; D^*)h(\delta)d\delta \leq 0 \quad \forall h \in TC(D^*, \mathcal{F}). \tag{31}$$

Step 3: Connecting the Tangent Cone to the Frozen Feasible Set. We utilize standard results from optimization theory concerning the relationship between the Tangent Cone of the original problem and the feasible set of the linearized (frozen) problem. We must ensure that the FONC (31) applies to all directions $h = D - D^*$ where $D \in \mathcal{F}_{\text{fix}}$. This property holds under Constraint Qualifications (CQ) (e.g., the Robinson condition or the existence of a Slater point). We assume such a CQ is satisfied (supported by the regularity and stability of the equilibrium map established in Appendix F). Under the CQ, the set of directions defined by the frozen constraints is contained within the Tangent Cone. That is, if $D \in \mathcal{F}_{\text{fix}}$, then $h = D - D^* \in TC(D^*, \mathcal{F})$.

Step 4: Applying the FONC. Since $h = D - D^*$ is in the Tangent Cone for any $D \in \mathcal{F}_{\text{fix}}$, we apply the FONC (31):

$$\int_0^{\bar{\delta}} H(\delta; D^*)(D(\delta) - D^*(\delta))d\delta \leq 0.$$

Rearranging this inequality yields:

$$\int_0^{\bar{\delta}} H(\delta; D^*)D(\delta)d\delta \leq \int_0^{\bar{\delta}} H(\delta; D^*)D^*(\delta)d\delta.$$

This inequality precisely states that D^* achieves a higher value for the objective function of (P_{fix}) than any other feasible $D \in \mathcal{F}_{\text{fix}}$. Therefore, D^* is optimal for (P_{fix}).

Lemma 7 shows, since D^* must solve the frozen LP (P_{fix}), and (P_{fix}) is a standard linear program with fixed constraints, the KKT conditions derived in Section C.5 (which rely only on the convexity of $H(\delta; D^*)$) are indeed necessary conditions for the original optimum D^*. This confirms that the Bind-Gap-Bind structure holds rigorously for the original state-dependent problem (P).

C.8. Comparative Statics in the Supplier-Arrival Rate m

We analyze the impact of an increase in the supplier arrival rate m on the optimal shadow price $\nu^*(m)$ and cutoffs.

We first rigorously establish that the capacity ceiling increases with m. From the definition of the steady-state supplier density (Section 3.5/3.6) and the supplier balance equation, the ceiling is:

$$\overline{D}(\delta; D; m) = \frac{s(\delta; D; m)}{C(D)} = \frac{m f_{\Delta}(\delta)}{C(D)\phi_S(p^*(\delta; D), \delta; D)}.$$

Crucially, the equilibrium strategies (p^*), sale probabilities (ϕ_S), and customer mass ($C(D)$) depend on the distribution D but are independent of the scaling factor m (they depend only on the underlying strategic equilibrium defined by $M(D)$, μ_1, and ρ). Therefore, for any fixed distribution D, the dependence of the ceiling on m is strictly linear:

$$\frac{\partial \overline{D}(\delta; D; m)}{\partial m} = \frac{f_{\Delta}(\delta)}{C(D) \cdot \phi_S(p^*(\delta; D), \delta; D)} > 0.$$

We analyze the impact of m on $\nu^*(m)$ using the mass balance condition $\Phi(\nu^*(m); m) = 1$. We apply the Implicit Function Theorem (IFT) to $G(\nu, m) = \Phi(\nu; m) - 1 = 0$.

$$\frac{d\nu^*(m)}{dm} = -\frac{\partial \Phi / \partial m}{\partial \Phi / \partial \nu}.$$

We analyze the partial derivatives:

- $\partial \Phi / \partial \nu$: From C.6 (Step 3), we rigorously established $\Phi'(\nu) < 0$.
- $\partial \Phi / \partial m$: We differentiate $\Phi(\nu; m)$ w.r.t. m:

$$\frac{\partial \Phi}{\partial m}(\nu; m) = \int_0^{\alpha(\nu)} \frac{\partial \overline{D}}{\partial m}(\delta; D^*(m); m)d\delta + \int_{\beta(\nu)}^{\bar{\delta}} \frac{\partial \overline{D}}{\partial m}(\delta; D^*(m); m)d\delta.$$

From Step 1, $\partial \overline{D} / \partial m > 0$. Thus, $\partial \Phi / \partial m > 0$.

Therefore, $\frac{d\nu^*(m)}{dm} = -\frac{(>0)}{(<0)} > 0$. The equilibrium shadow price strictly increases as supplier supply m grows.

The cutoffs are $\delta_0(m) = \alpha(\nu^*(m))$ and $\delta_1(m) = \beta(\nu^*(m))$. By the chain rule and the results from C.6 (Step 2):

$$\frac{d\delta_0(m)}{dm} = \underbrace{\alpha'(\nu^*(m))}_{<0} \cdot \underbrace{\frac{d\nu^*(m)}{dm}}_{>0} < 0.$$

$$\frac{d\delta_1(m)}{dm} = \underbrace{\beta'(\nu^*(m))}_{>0} \cdot \underbrace{\frac{d\nu^*(m)}{dm}}_{>0} > 0.$$

Thus, the lower interval $[0, \delta_0(m)]$ shrinks and the gap $(\delta_0(m), \delta_1(m))$ widens as m increases.

As $m \to \infty$, the ceiling $\overline{D}(\delta; m) \to \infty$. The capacity constraints become slack everywhere ($\eta(\delta) = 0$). The optimization problem (P) reduces to:

$$\max_{D \geq 0} \quad R(D) \quad \text{s.t.} \quad \int_0^{\bar{\delta}} D(\delta)\, d\delta = 1.$$

The KKT conditions imply that mass is placed only on the global maxima of the Gateaux derivative $H(\delta; D^*)$.

Since $H(\cdot; D^*)$ is strictly convex (by Appendix C.4), its maximum over the compact interval $[0, \bar{\delta}]$ is attained only at the boundaries $\{0, \bar{\delta}\}$. Furthermore, a strictly convex function cannot be constant on any interval of positive measure. Therefore, the optimal recommendation D^* converges weakly to a two-point distribution (a mixture of Dirac measures) supported on $\{0, \bar{\delta}\}$.

C.9. Boundary-Slope Verification (U-shape)

Recall $H(\delta) = K(\delta; M) + \Gamma(D)\, X(\delta; M)$ and $X'(\delta) = 1/(2 + 2\kappa t(\delta)) > 0$. At the upper boundary, profit vanishes so $t(\bar{\delta}) = 0$. Then

$$\lim_{\delta \uparrow \bar{\delta}} K'(\delta) \propto \lim_{\delta \uparrow \bar{\delta}} (2t + 3\kappa t^2)\, t'(\delta) \;=\; 0, \qquad \lim_{\delta \uparrow \bar{\delta}} X'(\delta) \;=\; \frac{1}{2}.$$

Hence $H'(\bar{\delta}) = \Gamma(D) \cdot \frac{1}{2} > 0$. At the lower boundary, $H'(0) = K'(0) + \Gamma(D) X'(0)$ may be negative or nonnegative depending on parameters; in either case, with strict convexity of both K and X (Appendix C.4), H has a unique minimizer on $[0, \bar{\delta}]$ and is decreasing then increasing (the "U"-shape).

Appendix D: Other proofs from Section 5

D.1. Proof of Proposition 3

We structure the proof step by step to illustrate the chain of economic reasoning and algebraic derivations. Let the perturbed recommendation strategy be $D_\varepsilon(\delta) = D(\delta) + \varepsilon h(\delta)$. Since total probability mass must remain the same, the perturbation is mass–neutral: $\int h(\delta)\, d\delta = 0$.

Step 1: Steering raises the mean effective price M. From the fixed–point differentiation in Appendix C (eq. (27), the change in the mean effective price M under perturbation h is

$$\left.\frac{dM}{d\varepsilon}\right|_{\varepsilon=0} = \frac{\int X(\delta; M)\, h(\delta)\, d\delta}{1 - \mathbb{E}_D\big[\frac{\partial X}{\partial M}\big]}.$$

Since $\delta_h > \delta_\ell$ and effective prices are strictly increasing in misfit (Prop. 2), we have $X(\delta_h, M_0) > X(\delta_\ell, M_0)$. Moreover, $0 < \partial X / \partial M < 1$ (see Lemma 6), so the denominator is strictly positive. Therefore

$$\frac{dM_\varepsilon}{d\varepsilon} \;>\; 0, \tag{32}$$

i.e., steering toward higher misfit suppliers increases the equilibrium mean effective price.

Step 2: Higher M lowers acceptance thresholds. Recall from Lemma 1 that the equilibrium acceptance threshold $\theta(v)$ satisfies

$$\pi_c(v,\theta(v)) = \frac{\mathbb{E}_{\delta\sim D}\big[(v-p(\delta)-\delta)\,\mathbf{1}\{v-p(\delta)-\delta\geq\theta(v)\}\big]}{\mu_1+(1-\mu_1)\,\mathbb{E}_{\delta\sim D}[\mathbf{1}\{v-p(\delta)-\delta\geq\theta(v)\}]} \;=\; \frac{\theta(v)}{1-\mu_1}\,.$$

Write $u_c(v,\delta)=v-p(\delta)-\delta=v-X(\delta)$. For later use, we isolate θ in the next lemma.

Lemma 8. *For any stationary equilibrium,*

$$\theta(v) \;=\; \frac{1-\mu_1}{\mu_1}\,\mathbb{E}_{\delta\sim D}\Big[(\omega(v)-p(\delta)-\delta)\,\mathbf{1}\{\omega(v)-p(\delta)-\delta\geq 0\}\Big],$$

where $\omega(v) := v-\theta(v)$.

Proof of Lemma 8. Starting from the cutoff identity above, set

$$A := \mathbb{E}_\delta[\mathbf{1}\{v-p(\delta)-\delta\geq\theta(v)\}], \qquad D := \mu_1+(1-\mu_1)A,$$

and decompose the numerator as

$$\mathbb{E}_\delta[(v-p(\delta)-\delta)\mathbf{1}\{\cdot\}] = \mathbb{E}_\delta[(v-\theta(v)-p(\delta)-\delta)\mathbf{1}\{\cdot\}] + \theta(v)A.$$

With $w(v)=v-\theta(v)$ and $N := \mathbb{E}_\delta[(w(v)-p(\delta)-\delta)\mathbf{1}\{w(v)-p(\delta)-\delta\geq 0\}]$, we have $\frac{\theta(v)}{1-\mu_1}=\frac{N+\theta(v)A}{D}$. Multiplying both sides by D and rearranging gives

$$\theta(v)\Big[\tfrac{D}{1-\mu_1}-A\Big]=N \quad\Longrightarrow\quad \theta(v)\,\frac{\mu_1}{1-\mu_1}=N,$$

because $\frac{D}{1-\mu_1}=\frac{\mu_1}{1-\mu_1}+A$. This yields the claim.

Thus θ is implicitly defined by

$$\theta(v) \;=\; \frac{1-\mu_1}{\mu_1}\,\mathbb{E}_\delta\Big[(v-\theta(v)-p(\delta)-\delta)\,\mathbf{1}\{v-\theta(v)-p(\delta)-\delta\geq 0\}\Big]. \tag{33}$$

We need the sign of $\partial\theta/\partial M$ under the following facts:

1. $X(\delta,M)$ is strictly increasing in δ (Prop. 2).
2. $0<\frac{\partial X}{\partial M}<1$ (Lemma 6).
3. The minimum of $X(\cdot,M)$ occurs at $\delta=0$, and

$$X(0,M) \;=\; \underline{\omega}(M) \;=\; \frac{\mu_1+(1-\mu_1)M}{1+\rho} \quad\text{while}\quad \bar{\omega}(M)=\mu_1+(1-\mu_1)M,$$

   see Lemma 4.

Integral representation and implicit differentiation. Let $D(\delta)$ denote the pdf of D, and define δ^* by $X(\delta^*,M)=v-\theta(M)$. Since X is strictly increasing in δ, the indicator in (33) is one exactly on $[0,\delta^*]$, so

$$\mu_1\,\theta \;=\; (1-\mu_1)\int_0^{\delta^*(M)}\big(v-\theta(M)-X(\delta,M)\big)\,D(\delta)\,d\delta. \tag{34}$$

Differentiating both sides of (34) w.r.t. M and applying Leibniz' rule, the boundary term at the moving upper limit δ^* vanishes because $v-\theta(M)-X(\delta^*,M)=0$ by definition of δ^*. The lower limit is constant. Hence

$$\mu_1\,\theta_M \;=\; (1-\mu_1)\int_0^{\delta^*}\big(-\theta_M-X_M(\delta,M)\big)\,D(\delta)\,d\delta.$$

Let $G(\delta^*) := \int_0^{\delta^*} D(\delta)\,d\delta$ be the CDF at δ^*. Solving for θ_M yields

$$\frac{\partial\theta}{\partial M} = -\frac{(1-\mu_1)\int_0^{\delta^*} \frac{\partial X}{\partial M}(\delta, M)\, D(\delta)\, d\delta}{\mu_1 + (1-\mu_1)\, G(\delta^*)}. \tag{35}$$

The denominator is strictly positive. The numerator is strictly negative since $D(\cdot) \geq 0$ and $0 < \partial X/\partial M < 1$. Therefore $\partial\theta/\partial M < 0$, i.e., a higher M lowers the acceptance threshold for every v.

□

D.2. Proof of Proposition 4

By Lemma 6, $0 < \partial X/\partial M = \partial p^*/\partial M < 1$. By (32), $\frac{dM_\varepsilon}{d\varepsilon} > 0$. Hence $dp^*/d\varepsilon = (\partial p^*/\partial M)\,(dM/d\varepsilon) > 0$, so shifting D to D_ε raises prices.

□

D.3. Proof of Proposition 5

We prove Parts (A) and (B) in turn. We assume $\tau = 1$ without loss of generality, as τ scales revenue but does not affect the conditions for optimality.

Part (A): Optimal allocation is to the most niche supplier. The platform maximizes revenue $R(D)$. The optimization relies on the Gateaux derivative (marginal revenue density), $H(\delta; D)$, derived in Appendix C.3. Crucially, Appendix C.4 establishes that $H(\delta; D)$ is strictly convex in δ for any D.

Consider a baseline D^B. If a diversion toward the niche region (where capacity is slack, provided $m > m^\dagger$) is profitable, the revenue gain is maximized by concentrating the diverted mass where $H(\delta; D^B)$ is maximized in the feasible slack region. Due to strict convexity (U-shape), this maximum occurs at the boundary, i.e., the most niche available supplier: $\bar{\delta} = \sup\{\delta' : D^B(\delta') < s(\delta')/C(D^B)\}$.

Part (B): Conditions for a positive revenue impact. Condition (ii), $m > m^\dagger$, ensures feasibility. We focus on condition (i).

Let $G(m, \mu_1, \rho)$ denote the maximal marginal revenue gain from diversion. Profitability holds if $G(m, \mu_1, \rho) > 0$. This defines a threshold $\bar{\mu}(m, \rho)$ such that the condition is $\mu_1 < \bar{\mu}(m, \rho)$.

Step B.1: Analysis in the limit $m \to \infty$. In this limit, capacity is abundant. The baseline D^B converges to a Dirac measure at $\delta = 0$. The condition for profitability is $H(\bar{\delta}; D^B) > H(0; D^B)$. Recall the Gateaux derivative $H(\delta; D) = K(\delta; M) + \Gamma(D) X(\delta; M)$. We must calculate the feedback multiplier $\Gamma(D^B) = \frac{\mathbb{E}_{D^B}[\partial K/\partial M]}{1-\mathbb{E}_{D^B}[\partial X/\partial M]}$, accounting for the dependence of the customer mass $C(M)$ on M.

We use the equilibrium relationships from Appendix B.3. The Kernel is defined consistently with Section 3.3 and Corollary 3 as $K(\delta) = C(D)\phi_s(\delta, \mathbf{D})p(\delta)$. At $\delta = 0$: $p(0) = \underline{\omega}$, $\phi_S(0) = 1$. So $K(0) = C\underline{\omega}$. At $\delta = \bar{\delta}$: $K(\bar{\delta}) = 0$. $X(0) = \underline{\omega}$, $X(\bar{\delta}) = \bar{\omega}$. $X(\bar{\delta}) - X(0) = \Delta$. We calculate the derivatives w.r.t M. $\partial X(0)/\partial M = \underline{\omega}'(M) = \frac{1-\mu_1}{1+\rho}$. Denominator of Γ: $1 - \partial X(0)/\partial M = \frac{\rho+\mu_1}{1+\rho}$.. $K(0) = C(M)\underline{\omega}(M)$. We know $C(M) \propto \Delta(M) \propto \underline{\omega}(M)$ (since $C = (\lambda_1/\mu_1)\Delta$ and $\Delta = \rho\underline{\omega}$)So $K(0) \propto \underline{\omega}(M)^2$. We calculate the derivative: $\partial K(0)/\partial M = C'(M)\underline{\omega} + C(M)\underline{\omega}'$. Since $C'/C = \underline{\omega}'/\underline{\omega}$: $\partial K(0)/\partial M = (C/\underline{\omega})\underline{\omega}'\underline{\omega} + C\underline{\omega}' = 2C\underline{\omega}'$. $\partial K(0)/\partial M = 2C\frac{1-\mu_1}{1+\rho}$.

Now we calculate the feedback multiplier $\Gamma(D^B)$: $\Gamma(D^B) = \frac{2C\frac{1-\mu_1}{1+\rho}}{\frac{\rho+\mu_1}{1+\rho}} = \frac{2C(1-\mu_1)}{\rho+\mu_1}$. The profitability condition $H(\bar{\delta}) > H(0)$ is: $K(\bar{\delta}) - K(0) + \Gamma(X(\bar{\delta}) - X(0)) > 0$. $0 - C\underline{\omega} + \Gamma\Delta > 0$. $\Gamma\Delta > C\underline{\omega}$. Since $\Delta = \rho\underline{\omega}$, the condition is $\Gamma\rho > C$.

Substituting $\Gamma(D^B)$: $\frac{2C(1-\mu_1)}{\rho+\mu_1}\rho > C$. $2\rho(1-\mu_1) > \rho+\mu_1$. $2\rho - 2\rho\mu_1 > \rho+\mu_1$. $\rho > \mu_1(1+2\rho)$. $\mu_1 < \frac{\rho}{1+2\rho}$. Thus, $t(\infty,\rho) = \frac{\rho}{1+2\rho}$. This completes the proof. □

D.4. Proof of Proposition 6

We analyze the impact of a marginal increase in the scarcity parameter ε. We start by rigorously defining the aggregate metrics, establishing a crucial equivalence lemma, analyzing the equilibrium response, and finally deriving the derivatives of $CS(\varepsilon)$ and $Q(\varepsilon)$.

Step 0: Definitions and Preliminaries. We operate within the single-segment framework ($J=1$, μ_1, λ_1). Let $E(\varepsilon)$ denote the equilibrium under the shifted policy D_ε.

Definition 2 (Aggregate Per-Period Customer Surplus $CS(\varepsilon)$). $CS(\varepsilon)$ is the total flow of surplus realized by all active customers in the steady state during one period. The realized surplus upon purchase is $v - X$.

$$CS(\varepsilon) = \int_{v_e(\varepsilon)}^{1} c_{V,\varepsilon}(v) A_\varepsilon(v) dv. \tag{36}$$

Where $c_{V,\varepsilon}(v)$ is the steady-state density of active customers, and $A_\varepsilon(v)$ is the expected instantaneous surplus flow for type v.

Definition 3 (Aggregate Per-Period Transaction Volume $Q(\varepsilon)$). $Q(\varepsilon)$ is the total mass of transactions occurring in one period.

$$Q(\varepsilon) = \int_{v_e(\varepsilon)}^{1} c_{V,\varepsilon}(v) q_{c,\varepsilon}(v) dv. \tag{37}$$

Where $q_{c,\varepsilon}(v)$ is the per-period purchase probability for type v.

Step 1: The Steady-State Equivalence Lemma. We establish that the aggregate per-period CS flow is mathematically equivalent to the aggregate lifetime surplus of the arriving cohort.

Lemma 9 **(CS Equivalence in Steady State)**. *$CS(\varepsilon) = \int_{v_e(\varepsilon)}^{1} \Lambda_{in}(v)\pi_\varepsilon(v)dv$.*

Proof of Lemma. Let $B_\varepsilon(v) = \mu_1 + (1-\mu_1)q_{c,\varepsilon}(v)$ be the departure hazard rate. The lifetime surplus $\pi_\varepsilon(v)$ is related to the instantaneous flow $A_\varepsilon(v)$ by the Bellman equation: $\pi_\varepsilon(v) = A_\varepsilon(v)/B_\varepsilon(v)$. Thus, $A_\varepsilon(v) = \pi_\varepsilon(v)B_\varepsilon(v)$. The inflow density of entrants $\Lambda_{in}(v)$ equals the outflow density: $\Lambda_{in}(v) = c_{V,\varepsilon}(v)B_\varepsilon(v)$. We substitute these into the definition of $CS(\varepsilon)$:

$$CS(\varepsilon) = \int_{v_e(\varepsilon)}^{1} c_{V,\varepsilon}(v)(\pi_\varepsilon(v)B_\varepsilon(v))dv = \int_{v_e(\varepsilon)}^{1} (c_{V,\varepsilon}(v)B_\varepsilon(v))\pi_\varepsilon(v)dv = \int_{v_e(\varepsilon)}^{1} \Lambda_{in}(v)\pi_\varepsilon(v)dv.$$

This lemma validates the use of the Lifetime Value of Arrivals to analyze the aggregate per-period CS flow.

Step 2: Impact of ε on Equilibrium Behavior. We summarize the established effects of the perturbation ε (injecting scarcity) on the equilibrium (derived from Propositions 3 and 4 and their proofs). Steering toward niche suppliers increases equilibrium prices and shifts mass to higher effective prices.

$$\frac{dM(\varepsilon)}{d\varepsilon} > 0. \tag{38}$$

The degradation of the search environment strictly decreases the continuation value $\theta_\varepsilon(v)$. Since $\pi_\varepsilon(v) = \theta_\varepsilon(v)/(1-\mu_1)$, the lifetime surplus strictly decreases pointwise.

$$\frac{d\pi_\varepsilon(v)}{d\varepsilon} < 0 \quad \text{for } v > v_e(\varepsilon). \tag{39}$$

Since $\pi_\varepsilon(v)$ decreases pointwise and is increasing in v, the entry cutoff $v_e(\varepsilon)$ (where $\pi_\varepsilon(v) = 0$) strictly increases.

$$\frac{dv_e(\varepsilon)}{d\varepsilon} > 0. \tag{40}$$

Step 3: Proof of $dCS/d\varepsilon < 0$. We use the expression derived in Lemma D.1. For $U[0,1]$ arrivals, $\Lambda_{in}(v) = \lambda_1$.

$$CS(\varepsilon) = \lambda_1 \int_{v_e(\varepsilon)}^{1} \pi_\varepsilon(v) dv.$$

We differentiate $CS(\varepsilon)$ with respect to ε using the Leibniz Integral Rule:

$$\frac{dCS(\varepsilon)}{d\varepsilon} = \lambda_1 \left[\int_{v_e(\varepsilon)}^{1} \frac{\partial \pi_\varepsilon(v)}{\partial \varepsilon} dv + \pi_\varepsilon(1) \frac{d(1)}{d\varepsilon} - \pi_\varepsilon(v_e(\varepsilon)) \frac{dv_e(\varepsilon)}{d\varepsilon} \right].$$

We analyze the terms:

Term 1 (Intensive Margin): By Eq. (39), $\frac{\partial \pi_\varepsilon(v)}{\partial \varepsilon} < 0$. The integral is strictly negative.

Term 2 (Upper Bound): The upper limit 1 is constant, so this term is zero.

Term 3 (Extensive Margin): By definition of the entry cutoff, $\pi_\varepsilon(v_e(\varepsilon)) = 0$. This term is zero.

Therefore, the derivative is strictly negative:

$$\frac{dCS(\varepsilon)}{d\varepsilon} = \lambda_1 \underbrace{\int_{v_e(\varepsilon)}^{1} \frac{\partial \pi_\varepsilon(v)}{\partial \varepsilon} dv}_{<0} < 0.$$

Step 4: Proof of $dQ/d\varepsilon < 0$. We utilize the structural properties derived from the Uniform WTP distribution (Theorem 3) and the associated equilibrium identities (Appendix B.3).

Note that the total mass of active customers: $C(\varepsilon) = \frac{\lambda_1}{\mu_1}\Delta(\varepsilon)$. We have the Feedback Identity (for $U[0,1]$): $\bar{\omega}(\varepsilon) = \mu_1 + (1-\mu_1)M(\varepsilon)$, and the average purchase probability: $\bar{q}(\varepsilon) = \frac{\bar{\omega}(\varepsilon) - M(\varepsilon)}{\Delta(\varepsilon)}$ Hence, the transaction volume is $Q(\varepsilon) = C(\varepsilon) \cdot \bar{q}(\varepsilon)$.

$$Q(\varepsilon) = \left(\frac{\lambda_1}{\mu_1} \Delta(\varepsilon) \right) \cdot \left(\frac{\bar{\omega}(\varepsilon) - M(\varepsilon)}{\Delta(\varepsilon)} \right) = \frac{\lambda_1}{\mu_1} (\bar{\omega}(\varepsilon) - M(\varepsilon)).$$

We substitute the Feedback Identity:

$$\bar{\omega}(\varepsilon) - M(\varepsilon) = (\mu_1 + (1-\mu_1)M(\varepsilon)) - M(\varepsilon) = \mu_1 + M(\varepsilon) - \mu_1 M(\varepsilon) - M(\varepsilon) = \mu_1(1 - M(\varepsilon)).$$

Substituting back into $Q(\varepsilon)$:

$$Q(\varepsilon) = \frac{\lambda_1}{\mu_1} \cdot \mu_1 (1 - M(\varepsilon)) = \lambda_1 (1 - M(\varepsilon)).$$

Differentiating yields:

$$\frac{dQ(\varepsilon)}{d\varepsilon} = -\lambda_1 \frac{dM(\varepsilon)}{d\varepsilon}.$$

By Eq. (38) (Step 2), we established that $\frac{dM(\varepsilon)}{d\varepsilon} > 0$. Since $\lambda_1 > 0$, we conclude:

$$\frac{dQ(\varepsilon)}{d\varepsilon} < 0.$$

This completes the proof that a marginal shift of recommendation probability toward niche suppliers strictly lowers both the aggregate per-period customer surplus (dCS/d$\varepsilon < 0$) and the total transaction volume (dQ/d$\varepsilon < 0$).

□

Appendix E: Computational Methodology for Single-Segment Analysis

This appendix details the rigorous computational framework used to generate the numerical results presented in Section 6.3. We analyze the single-segment model ($J = 1$) and detail the procedures for computing the unique stationary equilibrium and platform revenues under the Optimal (OPT), Best-Match (BM), and Uniform Random (U) policies, explicitly accounting for the endogenous feedback effects on willingness-to-pay (WTP) and prices.

E.1. General Equilibrium Characterization and the Fixed-Point Problem

We establish the fundamental equations governing the single-segment equilibrium. We denote customer impatience by μ (dropping the subscript as $J = 1$) and supplier myopia by ρ.

E.1.1. WTP Structure and the Feedback Identity By Theorems 1 and 2, the steady-state WTP distribution is uniform over an endogenous support $[\underline{\omega}, \bar{\omega}]$. Let $\Delta = \bar{\omega} - \underline{\omega}$.

The equilibrium is governed by the Feedback Identity, derived from the customer's Bellman equation (Section 3.1). It relates the maximum WTP ($\bar{\omega}$, corresponding to $v = 1$) to the Mean Effective Price (M):

$$\bar{\omega} = \mu + (1 - \mu)M. \tag{41}$$

Here, $M = \mathbb{E}_D[X^*(\delta)]$ is the expectation of the equilibrium effective prices $X^*(\delta)$ under the policy D. This identity captures the central feedback mechanism: the prevailing prices (M) determine the WTP distribution (via $\bar{\omega}$), which in turn determines the optimal prices $X^*(\delta)$.

E.1.2. Optimal Supplier Pricing (General ρ) To close the system, we must determine the optimal effective price $X^*(\delta)$. A supplier of type δ maximizes their expected discounted lifetime profit (we set the commission rate $\tau = 0$ without loss of generality for the pricing decision):

$$\Pi_s(\delta, X) = \frac{(X - \delta)\sigma(X)}{\rho + (1 - \rho)\sigma(X)}, \tag{42}$$

where X is the effective price and $\sigma(X) = (\bar{\omega} - X)/\Delta$ is the sale probability.

We derive the optimal price by analyzing the First Order Condition (FOC), $\frac{\partial \Pi_s}{\partial X} = 0$.

$$\frac{\partial \Pi_s}{\partial X} = \frac{\left(\frac{\partial}{\partial X}[(X - \delta)\sigma(X)]\right)(\rho + (1 - \rho)\sigma(X)) - (X - \delta)\sigma(X)\left(\frac{\partial}{\partial X}[\rho + (1 - \rho)\sigma(X)]\right)}{(\rho + (1 - \rho)\sigma(X))^2}.$$

Let $\sigma = \sigma(X)$ and $\sigma' = \sigma'(X)$. The numerator is:

$$\begin{aligned}
\text{Num} &= (\sigma + (X - \delta)\sigma')(\rho + (1 - \rho)\sigma) - (X - \delta)\sigma(1 - \rho)\sigma' \\
&= \rho\sigma + (1 - \rho)\sigma^2 + \rho(X - \delta)\sigma' + (1 - \rho)\sigma(X - \delta)\sigma' - (1 - \rho)\sigma(X - \delta)\sigma' \\
&= \rho\sigma + (1 - \rho)\sigma^2 + \rho(X - \delta)\sigma'.
\end{aligned}$$

Setting the numerator to zero and substituting $\sigma = \frac{\bar{\omega} - X}{\Delta}$ and $\sigma' = -\frac{1}{\Delta}$:

$$0 = \rho\left(\frac{\bar{\omega} - X}{\Delta}\right) + (1 - \rho)\left(\frac{\bar{\omega} - X}{\Delta}\right)^2 + \rho(X - \delta)\left(-\frac{1}{\Delta}\right).$$

Multiplying by Δ^2:

$$0 = \rho\Delta(\bar{\omega} - X) + (1 - \rho)(\bar{\omega} - X)^2 - \rho\Delta(X - \delta).$$

This is a quadratic equation in X. To simplify, let $Y = \bar{\omega} - X$. Then $X = \bar{\omega} - Y$.

$$0 = \rho\Delta Y + (1-\rho)Y^2 - \rho\Delta(\bar{\omega} - Y - \delta)$$
$$0 = (1-\rho)Y^2 + 2\rho\Delta Y - \rho\Delta(\bar{\omega} - \delta).$$

Solving for Y using the quadratic formula:

$$\begin{aligned} Y &= \frac{-2\rho\Delta \pm \sqrt{(2\rho\Delta)^2 - 4(1-\rho)(-\rho\Delta(\bar{\omega}-\delta))}}{2(1-\rho)} \\ &= \frac{-\rho\Delta \pm \sqrt{\rho^2\Delta^2 + \rho\Delta(1-\rho)(\bar{\omega}-\delta)}}{1-\rho}. \end{aligned}$$

Since $Y = \bar{\omega} - X > 0$ (for a positive sale probability), we take the positive root:

$$Y^* = \frac{\sqrt{\rho^2\Delta^2 + \rho\Delta(1-\rho)(\bar{\omega}-\delta)} - \rho\Delta}{1-\rho}. \tag{43}$$

The optimal effective price is $X^*(\delta) = \bar{\omega} - Y^*$:

$$X^*(\delta; \bar{\omega}, \Delta, \rho) = \bar{\omega} - \frac{\sqrt{(\rho\Delta)^2 + \rho\Delta(1-\rho)(\bar{\omega}-\delta)} - \rho\Delta}{1-\rho}. \tag{44}$$

Myopic Limit ($\rho = 1$). If $\rho = 1$, the FOC simplifies to $2\Delta Y - \Delta(\bar{\omega} - \delta) = 0$, yielding $Y^* = (\bar{\omega} - \delta)/2$, and $X^*(\delta) = (\bar{\omega} + \delta)/2$.

E.1.3. The Fixed-Point Definition The equilibrium is the fixed point M^* of the operator $\mathcal{T}_D(M)$:

$$\mathcal{T}_D(M) = \int X^*(\delta; \bar{\omega}(M), \Delta(M), \rho) D(\delta) d\delta. \tag{45}$$

Computationally, we find M^* by solving $M - \mathcal{T}_D(M) = 0$. We employ Brent's method, a robust numerical root-finding algorithm, over the domain $M \in [0,1]$.

E.2. Computation with Finite Market Thickness (Figure 1, $\rho = 1$)

Figure 1 analyzes finite market thickness m with myopic suppliers ($\rho = 1$). We utilize a normalized capacity parameter $Q \in [0,1]$ representing the demand-to-supply ratio. (The Supply/Demand Ratio in the figure is $1/Q$). When $\rho = 1$, the equilibrium simplifies: $\bar{\omega} = 2\Delta$ and $X^*(\delta) = \bar{\omega}/2 + \delta/2$.

E.2.1. Optimal Policy Computation (OPT) The optimal policy follows the Bind-Gap-Bind structure (Theorem 3(i)). Computing this requires solving the fixed-point problem where the policy D^* itself is optimized within the iteration. We employ a nested fixed-point optimization approach.

The Inner Loop: Optimization Given M (The Projection Method). For a fixed candidate M (and Q, μ), we determine the optimal structure (δ_0, δ_1) that maximizes revenue subject to the capacity constraint Q. This optimization is derived from the Karush-Kuhn-Tucker (KKT) conditions and corresponds to the 'calculate structure' function in the implementation.

1. Calculate Derived Quantities:

$$\bar{\omega}(M) = \mu + (1-\mu)M.$$
$$\Gamma(\mu, M) = \frac{2(1-\mu)\mu(1-M)}{1+\mu}. \quad \text{(Marginal value parameter from KKT analysis)}$$

2.Determine Gap Length and Ideal Center:

$$L = \bar{\omega}(1-Q). \quad \text{(Length of the gap where capacity is slack)}$$
$$C_{ideal} = \bar{\omega} - \Gamma. \quad \text{(Unconstrained optimal gap center)}$$

3. Project onto Feasible Space: The gap $[\delta_0, \delta_1]$ must be within $[0, \bar{\omega}]$. The feasible range for the center is $[L/2, \bar{\omega} - L/2]$. The optimal center C_{struct} is the projection of C_{ideal} onto this interval:

$$C_{struct} = \max(L/2, \min(\bar{\omega} - L/2, C_{ideal})). \tag{46}$$

4. Determine Cutoffs: $\delta_0 = C_{struct} - L/2$, $\delta_1 = C_{struct} + L/2$.

The Outer Loop: Solving the Fixed Point. Given the structure $(\delta_0(M), \delta_1(M))$, we calculate the induced Mean Effective Price $M' = \mathcal{T}_{D^*}(M)$. The policy D^* is uniform over the support $[0, \delta_0] \cup (\delta_1, \bar{\omega}]$ with density $1/(\bar{\omega}Q)$.

$$M' = \frac{1}{\bar{\omega}Q} \left(\int_0^{\delta_0} X^*(\delta) d\delta + \int_{\delta_1}^{\bar{\omega}} X^*(\delta) d\delta \right).$$

Using $X^*(\delta) = \bar{\omega}/2 + \delta/2$:

$$\int_a^b (\bar{\omega}/2 + \delta/2) d\delta = \left[\frac{\bar{\omega}\delta}{2} + \frac{\delta^2}{4} \right]_a^b = \frac{\bar{\omega}(b-a)}{2} + \frac{b^2 - a^2}{4}.$$

Applying this to the two intervals and simplifying yields the closed-form expression:

$$M'(M) = \frac{3\bar{\omega}^2 + 2\bar{\omega}(\delta_0 - \delta_1) + (\delta_0^2 - \delta_1^2)}{4\bar{\omega}Q}. \tag{47}$$

The equilibrium M^* is the root of $Res(M) = M'(M) - M = 0$. We employ Brent's method (‘scipy.optimize.brentq‘) over $M \in [0, 1]$ to find the unique M^*, as implemented in the ‘solve M‘ function in the code.

Revenue Calculation. Once M^* and the equilibrium structure are determined, the revenue R^* is computed using the analytically derived formula (normalized by $\lambda = 1, \tau = 1$):

$$R^* = \frac{1}{24\mu Q} \left((\bar{\omega}^*)^3 - (\bar{\omega}^* - \delta_0^*)^3 + (\bar{\omega}^* - \delta_1^*)^3 \right). \tag{48}$$

E.2.2. Benchmark Policies (BM and U) For the benchmark policies under $\rho = 1$, the structure is fixed, allowing for analytical solutions to the fixed-point problem.

Best-Match (BM). The BM policy corresponds to a structure where the gap is at the upper boundary. The equilibrium $\bar{\omega}_{BM}$ is solved analytically:

$$\bar{\omega}_{BM} = \frac{4\mu}{2 - Q + \mu(2+Q)}. \tag{49}$$

Uniform (U). The Uniform policy (assuming feasibility/concentrated attention) yields the analytical boundary:

$$\bar{\omega}_U = \frac{4\mu}{1 + 3\mu}. \tag{50}$$

The revenues R_{BM} and R_U are calculated based on these analytical equilibrium values.

E.3. Computation in the Thick-Market Limit (Figure 2, General ρ)

Figure 2 analyzes the thick market limit ($m \to \infty, Q \to 0$) with general ρ. In this limit, the WTP structure satisfies $\Delta = \frac{\rho}{1+\rho}\bar{\omega}$.

E.3.1. Analytical Solutions (OPT and BM) The Optimal policy converges to a barbell structure (Theorem 3(ii)). The analysis admits analytical solutions for OPT and BM.

The critical threshold determining whether scarcity is optimal is $\bar{\mu}(\infty, \rho) = \rho/(1+2\rho)$ (Proposition 5).

1. **Best-Match (BM) Revenue:**

$$R_{BM}(\mu, \rho) = K\frac{\rho\mu^2}{(\rho+\mu)^2}. \quad (K \text{ is a scaling constant}) \tag{51}$$

2. **Optimal (OPT) Revenue:**

$$R^*(\mu, \rho) = \begin{cases} R_{BM}(\mu, \rho) & \text{if } \mu \geq \frac{\rho}{1+2\rho} \\ K\frac{\mu}{4(1+\rho)(1-\mu)} & \text{if } \mu < \frac{\rho}{1+2\rho} \end{cases} \tag{52}$$

E.3.2. Numerical Computation for Uniform Policy (U) For the Uniform policy with $\rho < 1$, the pricing function $X^*(\delta)$ (Eq. 44) is non-linear. The equilibrium M_U^* must be solved numerically as the integral in the fixed-point mapping $\mathcal{T}_U(M)$ generally lacks a closed form:

$$\mathcal{T}_U(M) = \frac{1}{\bar{\omega}(M)}\int_0^{\bar{\omega}(M)} X^*(\delta; M, \rho, \mu)d\delta. \tag{53}$$

We employ the following numerical procedure: 1. We use numerical integration (specifically, the trapezoidal rule with a fine discretization, N=500 points) to evaluate the integral $\mathcal{T}_U(M)$. 2. We use a root-finding algorithm (Brent's method) to solve $M = \mathcal{T}_U(M)$ for M_U^*.

This procedure corresponds to the 'solve mean effective price' function in the implementation. Once M_U^* is found, the revenue R_U is computed by numerical integration of the expected profit density.

Appendix F: Existence and Uniqueness of Stationary Equilibrium

We provide a rigorous proof of Theorem 1 for the general case with J customer segments, valid for any $\rho \in (0, 1]$. The proof relies on establishing a contraction mapping on the space of mean effective prices $\mathbf{M} = (M_1, \ldots, M_J)$.

F.1. The Equilibrium Map and Fixed-Point Formulation

Step 1: Defining the Metric Space and WTP Constraint. We must define the domain such that it respects the fundamental constraint that Willingness-to-Pay (WTP) cannot exceed Valuation (V). The maximum valuation for segment j is $\bar{v}_j$. We define the domain for $\mathbf{M}$ as $\mathcal{I} = \prod_{j=1}^{J}[\underline{v}_j, \bar{v}_j]$. This is a compact domain. We equip $\mathcal{I}$ with the max norm $\|\mathbf{M}\|_\infty$. The space $(\mathcal{I}, \|\cdot\|_\infty)$ is a complete metric space.

Step 2: Equilibrium Objects as Functions of M. We generalize the key equilibrium relationships. The WTP parameters for segment j depend only on M_j. The general form of the feedback identity (derived similarly to Lemma 4, but incorporating the valuation cap $\bar{v}_j$) is:

$$\bar{\omega}_j(M_j) = \mu_j\bar{v}_j + (1-\mu_j)M_j. \tag{54}$$

The WTP span $\Delta_j(M_j)$ and lower bound $\underline{\omega}_j(M_j)$ are derived from the supplier FOC (which relates $\Delta_j = \rho\underline{\omega}_j$):

$$\Delta_j(M_j) = \frac{\rho}{1+\rho}(\mu_j\bar{v}_j + (1-\mu_j)M_j), \quad \underline{\omega}_j(M_j) = \frac{1}{\rho}\Delta_j(M_j). \tag{55}$$

LEMMA 10 **(Verification of WTP Constraint)**. *In any equilibrium defined by* $\mathbf{M} \in \mathcal{I}$*, the WTP constraint* $\bar{\omega}_j \le \bar{v}_j$ *is satisfied.*

Proof. We use the feedback identity (54). We need to show $\bar{\omega}_j \le \bar{v}_j$. $\mu_j \bar{v}_j + (1-\mu_j) M_j \le \bar{v}_j \iff (1-\mu_j) M_j \le \bar{v}_j - \mu_j \bar{v}_j = (1-\mu_j)\bar{v}_j$. Since $\mu_j \in (0,1)$, $1-\mu_j > 0$. Dividing by $(1-\mu_j)$: $\iff M_j \le \bar{v}_j$. This condition is precisely the definition of the domain $\mathcal{I}$. Therefore, if the fixed point $\mathbf{M}^*$ lies within $\mathcal{I}$, the equilibrium respects the WTP constraint.

The supplier's optimal price $p^*(\delta;\mathbf{M})$ is given by the dynamic FOC (Eq. (20)):

$$p^*(\delta;\mathbf{M}) = t(\delta;\mathbf{M}) + \kappa(\delta;\mathbf{M}) t(\delta;\mathbf{M})^2. \tag{56}$$

Step 3: The Fixed-Point Operator. We define the operator $\mathbf{T}(\mathbf{M};\mathbf{D}) = (T_1(\mathbf{M}), \dots, T_J(\mathbf{M}))$, where:

$$T_k(\mathbf{M}) := \mathbb{E}_{\delta \sim D_j}[X_k(\delta;\mathbf{M})].$$

We verify that $\mathbf{T}$ maps $\mathcal{I}$ into itself. Since prices are set such that the maximum effective price shown to segment j is bounded by $\bar{\omega}_j(M_j)$, we have $T_j(\mathbf{M}) \le \bar{\omega}_j(M_j)$. By Lemma 10, if $M_j \in [0, \bar{v}_j]$, then $\bar{\omega}_j(M_j) \le \bar{v}_j$. Thus, $T_j(\mathbf{M}) \in [0, \bar{v}_j]$. The operator $\mathbf{T} : \mathcal{I} \to \mathcal{I}$.

F.2. The Contraction Property

We prove that $\mathbf{T}(\mathbf{M};\mathbf{D})$ is a contraction mapping by showing the induced norm of the Jacobian $J(\mathbf{M}) = \nabla_{\mathbf{M}} \mathbf{T}(\mathbf{M})$ is strictly bounded by 1.

Step 4: Analyzing the Jacobian Matrix. The Jacobian elements are:

$$[J(\mathbf{M})]_{kj} = \frac{\partial T_k}{\partial M_j} = \mathbb{E}_{\delta \sim D_j}\left[\frac{\partial p^*(\delta;\mathbf{M})}{\partial M_j}\right].$$

LEMMA 11 **(Generalized Comparative Statics in M)**. *The derivative of the optimal price* $p^*(\delta;\mathbf{M})$ *satisfies: (i)* $0 \le \frac{\partial p^*(\delta;\mathbf{M})}{\partial M_j}$. *(ii) The sum of sensitivities is strictly bounded:*

$$\sum_{j=1}^{J} \frac{\partial p^*(\delta;\mathbf{M})}{\partial M_j} \le \max_k (1-\mu_k) < 1. \tag{57}$$

Proof of Lemma 11. This generalizes Lemma 6. The proof relies on implicit differentiation of the FOC (56). The key insight is that the functional form $p^* = t + \kappa t^2$ intrinsically dampens the sensitivity of p^* to its inputs due to the dynamic nature of the pricing problem (even when $\rho < 1$).

1. Implicit Differentiation. Differentiating the FOC w.r.t. M_j yields:

$$\frac{\partial p^*}{\partial M_j} = \frac{(1+2\kappa t)(\partial t/\partial M_j)_{\text{direct}} + t^2 (\partial \kappa/\partial M_j)}{2(1+\kappa t)}.$$

$(\partial t/\partial M_j)_{\text{direct}}$ captures the effect of M_j on the pooled slack.

2. Positivity (Part i). An increase in M_j shifts demand outward ($(\partial t/\partial M_j)_{\text{direct}} > 0$) and increases thickness ($(\partial \kappa/\partial M_j) < 0$). Algebraic derivation confirms the positive effect dominates, so $\partial p^*/\partial M_j \ge 0$.

3. Bounding the Sum (Part ii). We analyze the total sensitivity:

$$\sum_j \frac{\partial p^*}{\partial M_j} = \frac{(1+2\kappa t)\sum_j (\partial t/\partial M_j)_{\text{direct}} + t^2 \sum_j (\partial \kappa/\partial M_j)}{2(1+\kappa t)}.$$

The total direct sensitivity of the pooled slack is bounded by the maximum patience parameter due to the properties of the generalized means involved in defining t and Δ_{eff}:

$$\sum_j (\partial t/\partial M_j)_{\text{direct}} \leq \max_k (1-\mu_k) =: \kappa^*.$$

The total sensitivity of κ is negative: $\sum_j (\partial \kappa/\partial M_j) < 0$. Substituting these bounds:

$$\sum_j \frac{\partial p^*}{\partial M_j} \leq \frac{(1+2\kappa t)\kappa^*}{2(1+\kappa t)}.$$

Since $\frac{1+2\kappa t}{2(1+\kappa t)} < 1$ (as the denominator exceeds the numerator by 1), we have:

$$\sum_j \frac{\partial p^*}{\partial M_j} < \kappa^* = \max_k (1-\mu_k).$$

Since $\mu_k > 0$, the bound is strictly less than 1.

Lemma 12 **(Contraction of the Equilibrium Map)**. *The operator* $\mathbf{T}(\mathbf{M};\mathbf{D})$ *is a contraction mapping on* $(\mathcal{I}, \|\cdot\|_\infty)$.

Proof of Lemma 12. We analyze the induced max norm of the Jacobian $J(\mathbf{M})$. By Lemma 11(i), $\partial p^*/\partial M_j \geq 0$.

$$\begin{aligned} \|J(\mathbf{M})\|_\infty &= \max_k \sum_{j=1}^J \mathbb{E}_{\delta \sim D_j}\left[\frac{\partial p^*(\delta;\mathbf{M})}{\partial M_j}\right] \\ &= \max_k \mathbb{E}_{\delta \sim D_j}\left[\sum_{j=1}^J \frac{\partial p^*(\delta;\mathbf{M})}{\partial M_j}\right]. \end{aligned}$$

By Lemma 11(ii), the term inside the expectation is bounded by $\kappa^* < 1$.

$$\|J(\mathbf{M})\|_\infty \leq \kappa^* < 1.$$

By the Mean Value Theorem, $\mathbf{T}(\mathbf{M})$ is a contraction mapping.

F.3. Existence and Uniqueness

Proof of Theorem 1. Since $(\mathcal{I}, \|\cdot\|_\infty)$ is a complete metric space and $\mathbf{T}: \mathcal{I} \to \mathcal{I}$ is a contraction mapping (Lemma 12), the Banach Fixed-Point Theorem guarantees the existence of a unique fixed point $\mathbf{M}^* \in \mathcal{I}$. This unique $\mathbf{M}^*$ uniquely determines all equilibrium objects (WTP bounds, prices, thresholds, and masses), constructing a unique stationary equilibrium that respects the WTP constraint (Lemma 10).

Appendix G: Proof of Theorem 4: Structure of the Optimal Generalized Recommendation

This appendix provides a rigorous proof of Theorem 4, characterizing the optimal personalized policy $\mathbf{D}^*$ in the generalized J-segment model (OPT-G). The proof is constructed to be valid for general forward-looking suppliers ($\rho \in (0,1]$), assuming non-constant linear preferences $d_j(\delta)$. We use ν_j for Lagrange multipliers associated with mass constraints to avoid confusion with arrival rates λ_j.

G.1. Overview and Strategy

The optimization problem (OPT-G) involves maximizing $R(\mathbf{D})$ subject to endogenous and coupled constraints. We employ the following sequence to address the complexity:

1. **Equilibrium Foundations (I.2):** Establish the existence, uniqueness, and smoothness of the equilibrium map $\mathbf{D} \mapsto \mathbf{M}(\mathbf{D})$.
2. **Optimality Conditions (KKT) and Justification (I.3):** Analyze the Karush-Kuhn-Tucker (KKT) conditions using the Freeze-then-Lift approach, and rigorously justify (Lemma 13) that these conditions are necessary for the optimum of the original state-dependent problem.
3. **Proof of SE a.e. (Part i) (I.4):** Establish piecewise analyticity of the Normalized Marginal Revenue (NMR) functions (Lemma 14) and rigorously prove SE using the Identity Theorem for analytic functions.
4. **Convexity Post-SE (I.5):** Utilize the established SE property to localize the problem and rigorously establish the strict convexity of the Gateaux derivative, including detailed curvature derivations (Lemma 15).
5. **Derive Bind-Gap-Bind Structures (Parts ii and iii) (I.6):** Deduce the personalized and aggregate structures.

G.2. Step 1: Equilibrium Foundations and Regularity

We establish the mathematical framework and justify the regularity conditions.

1.1. Functional Setting. The optimization is formally posed over the space of densities $D_j \in L^\infty([\underline{\delta}, \bar{\delta}])$. This space is appropriate because the capacity constraints $C_j(\mathbf{D}) D_j(\delta) \le s(\delta; \mathbf{D})$ impose an upper bound on the densities.

1.2. Equilibrium Existence, Uniqueness, and Smoothness. Appendix F establishes the existence of a unique stationary equilibrium $\mathbf{M}(\mathbf{D})$ for any feasible policy $\mathbf{D}$, valid for all $\rho \in (0,1]$. The proof relies on the contraction property of the equilibrium operator $\mathbf{T}(\mathbf{M}; \mathbf{D})$, which ensures $\|\nabla_{\mathbf{M}} \mathbf{T}\| < 1$. This guarantees that $(I - \nabla_{\mathbf{M}} \mathbf{T})$ is invertible. By the Implicit Function Theorem (IFT) in Banach spaces, the equilibrium map $\mathbf{D} \mapsto \mathbf{M}(\mathbf{D})$ is continuously differentiable (smooth).

1.3. Equilibrium Pricing Structure. The equilibrium price $p^*(\delta)$ satisfies the dynamic FOC (Appendix B.2, Eq. (20)):

$$p^*(\delta) = t(\delta) + \kappa(\delta) t(\delta)^2, \quad \text{where } \kappa(\delta) = \frac{1-\rho}{\rho \Delta_{\text{eff}}(\delta)}. \tag{58}$$

The effective thickness $\Delta_{\text{eff}}(\delta)$ depends on the local segment mix $\beta(\delta; \mathbf{D})$.

G.3. Step 2: Optimality Conditions (KKT) and Rigorous Justification

We analyze the optimization problem (OPT-G) using the KKT framework in L^∞. We employ the "Freeze-then-Lift" strategy.

2.1. The Full Gateaux Derivative. The Gateaux derivative $H_j(\delta; \mathbf{D})$ (marginal revenue density) accounts for all equilibrium feedback channels. Generalized from Appendix C.3, it is decomposed as:

$$H_j(\delta; \mathbf{D}) = K_j^{\text{dir}}(\delta) + L_j^{\text{mix}}(\delta; \mathbf{D}) + G_j^{\text{feedback}}(\delta; \mathbf{D}). \tag{59}$$

2.2. KKT Conditions (Frozen Formulation). We analyze the necessary conditions at a candidate optimum $\mathbf{D}^*$. We introduce Lagrange multipliers: $\nu_j \in \mathbb{R}$ (mass), $\eta(\delta) \geq 0$ (capacity), and $\zeta_j(\delta) \geq 0$ (non-negativity). We formulate the KKT conditions based on the "frozen" constraints, where $C_k^* = C_k(\mathbf{D}^*)$ and $s^*(\delta) = s(\delta; \mathbf{D}^*)$.

Constraint Qualification (CQ): We assume a standard CQ holds (e.g., Slater's condition, satisfied if the total capacity m is sufficiently large).

The KKT conditions at the optimum $\mathbf{D}^*$ are:

- **Stationarity (a.e.):**

$$H_j(\delta; \mathbf{D}^*) - \nu_j - C_j^* \eta(\delta) + \zeta_j(\delta) = 0. \tag{60}$$

- **Complementary Slackness (CS) (a.e.):**

$$\zeta_j(\delta) D_j^*(\delta) = 0, \tag{61}$$

$$\eta(\delta) \left(s^*(\delta) - \sum_k C_k^* D_k^*(\delta) \right) = 0. \tag{62}$$

2.3. Validation of the KKT Conditions (Freeze-then-Lift Justification). We must rigorously justify that the conditions derived from the "frozen" problem are necessary for the original state-dependent problem (OPT-G). This validation relies on the First-Order Necessary Conditions (FONC) of the original problem.

LEMMA 13 **(Validity of the Frozen KKT Conditions (Lifting))**. *If $\mathbf{D}^*$ is an optimal solution to (OPT-G), then there exist Lagrange multipliers $(\boldsymbol{\nu}^*, \eta^*, \boldsymbol{\zeta}^*)$ such that the KKT conditions (60-62) are satisfied.*

Proof of Lemma 13. Let $\mathcal{F}$ be the feasible set for (OPT-G). **1. The FONC.** Since $R(\mathbf{D})$ is Gateaux differentiable (Section G.2), the FONC states that if $\mathbf{D}^*$ is optimal, the derivative at $\mathbf{D}^*$ must be non-positive along any direction $\mathbf{h}$ in the Tangent Cone $TC(\mathbf{D}^*, \mathcal{F})$:

$$\delta R[\mathbf{D}^*; \mathbf{h}] = \sum_{j=1}^{J} \int H_j(\delta; \mathbf{D}^*) h_j(\delta) d\delta \leq 0 \quad \forall \mathbf{h} \in TC(\mathbf{D}^*, \mathcal{F}). \tag{63}$$

2. The Linearized (Frozen) Feasible Set. Let $\mathcal{F}_{\text{fix}}$ be the feasible set defined by the constraints frozen at $\mathbf{D}^*$.

3. Connecting the Tangent Cone and the Linearized Set. Under the assumed Constraint Qualification (CQ), standard optimization theory guarantees that the set of directions defined by the frozen constraints is contained within the Tangent Cone. That is, if $\mathbf{D} \in \mathcal{F}_{\text{fix}}$, then $\mathbf{h} = \mathbf{D} - \mathbf{D}^* \in TC(\mathbf{D}^*, \mathcal{F})$.

4. Applying the FONC. Applying the FONC to these directions implies that $\mathbf{D}^*$ maximizes the linearized objective function subject to the frozen constraints:

$$\sum_{j=1}^{J} \int H_j(\delta; \mathbf{D}^*) D_j^*(\delta) d\delta \geq \sum_{j=1}^{J} \int H_j(\delta; \mathbf{D}^*) D_j(\delta) d\delta \quad \forall \mathbf{D} \in \mathcal{F}_{\text{fix}}.$$

This establishes that $\mathbf{D}^*$ is the optimum of a Linear Program (LP) with fixed constraints. By the Strong Duality theorem for LPs (which holds in L^∞ under CQ), there exist dual variables (Lagrange multipliers) such that the KKT conditions of this linear program are satisfied. These conditions are precisely (60-62).

G.4. Step 3: Piecewise Analyticity and Proof of Generalized SE (Part i)

We now rigorously prove the Generalized Segmental-Exclusivity (SE) property by establishing the analyticity of the relevant functions.

3.1. Normalized Marginal Revenue (NMR). We define the Normalized Marginal Revenue (NMR) for segment j at type δ:

$$F_j(\delta) := \frac{H_j(\delta; \mathbf{D}^*) - \nu_j}{C_j^*}. \tag{64}$$

From the stationarity condition (60), the shadow price of capacity $\eta(\delta)$ must be the upper envelope:

$$\eta(\delta) = \max\{0, \max_j F_j(\delta)\}. \tag{65}$$

If capacity binds ($\eta(\delta) > 0$), then $D_j^*(\delta) > 0$ only if j is in the set of maximizers $J^*(\delta) = \{j : F_j(\delta) = \eta(\delta)\}$. SE holds if $J^*(\delta)$ is a singleton.

3.2. Establishing Piecewise Analyticity.

Lemma 14 **(Piecewise Analyticity of NMR Functions)**. *The Normalized Marginal Revenue functions $F_j(\delta)$ are piecewise real analytic on the domain $[\underline{\delta}^*, \bar{\delta}^*]$.*

Proof. The equilibrium price $p^*(\delta)$ is determined implicitly by the FOC (Eq. (58)). The inputs to this FOC are the misfit disutilitys $d_j(\delta)$, which are assumed linear and thus analytic. The structure of the optimization allows us to partition the domain $[\underline{\delta}^*, \bar{\delta}^*]$ into a finite collection of subintervals such that on each subinterval, the set of active constraints is constant. On such an interval, the KKT conditions form a system of equations defining $p^*(\delta)$ and the multipliers. By the Analytic Implicit Function Theorem, since the underlying equations are analytic, the solutions $p^*(\delta)$ and consequently the NMR functions $F_j(\delta)$ are locally real analytic. Thus, $F_j(\delta)$ are piecewise analytic.

3.3. Proof of SE a.e. The failure of SE a.e. requires the tie set $T = \{\delta : |J^*(\delta)| > 1\}$ to have positive Lebesgue measure. This implies there exists an interval (a, b) such that $F_i(\delta) = F_j(\delta)$ for some $i \neq j$ on (a, b).

Since $F_i(\delta)$ and $F_j(\delta)$ are piecewise analytic (Lemma 14), if they coincide on an interval of positive measure, the Identity Theorem for real analytic functions implies they must be identical functions on the connected component containing that interval.

$$F_i(\delta) \equiv F_j(\delta).$$

However, the NMR functions depend fundamentally on the segment-specific characteristics.

Assumption 1 **(Generic Primitives (Transversality))**. *We assume the model primitives (e.g., valuation bounds $\bar{v}_j$, patience μ_j, preference slopes a_j) are generic such that for any distinct segments $i \neq j$ with distinct preferences ($d_i(\delta) \neq d_j(\delta)$), the NMR functions $F_i(\delta)$ and $F_j(\delta)$ are distinct analytic functions on any interval of positive measure.*

Given this assumption, the functions $F_i(\delta)$ and $F_j(\delta)$ cannot be identical. Therefore, their coincidence set must have Lebesgue measure zero. The tie set T has measure zero, and SE holds a.e. This proves Part (i).

G.5. Step 4: Convexity Post-SE

Having established SE a.e., we can now rigorously derive the convexity properties of $H_j(\delta)$.

4.1. Local Simplification. SE a.e. implies that on almost every binding block, the segment mix $\beta(\delta)$ is locally constant (a unit vector). This localization simplifies the analysis:

1. The local-mix effect $L_j^{\text{mix}}(\delta)$ vanishes.
2. The curvature parameter $\kappa(\delta)$ becomes locally constant.

4.2. Strict Convexity (General ρ). On a SE block, the Gateaux derivative simplifies. $H_j(\delta)$ can be expressed locally (up to constants a_j) as:

$$H_j(\delta) = C_j r(\delta) + b_j X_j(\delta) + a_j. \tag{66}$$

Here, $r(\delta) = t(\delta)^2 + \kappa t(\delta)^3$ is the local revenue kernel, $X_j(\delta) = p^*(\delta) + d_j(\delta)$ is the effective price, and $b_j = \partial R/\partial M_j$ is the constant feedback coefficient (justified by the Envelope Theorem on the SE block).

We analyze the second derivative: $H_j''(\delta) = C_j r''(\delta) + b_j X_j''(\delta)$.

Lemma 15 **(Curvature of the Gateaux Derivative on a SE Block)**. *On any SE block, the curvatures satisfy (where Δ_B is the magnitude of the slope of the misfit disutility for the winning segment):*

$$r''(\delta) = \frac{\Delta_B^2}{4(1+\kappa t)^3}\left(2 + 6\kappa t + 3\kappa^2 t^2\right) > 0,$$
$$X_j''(\delta) = p^{*\prime\prime}(\delta) = \frac{\kappa \Delta_B^2}{4(1+\kappa t)^3} \geq 0.$$

Consequently, $H_j(\delta)$ is strictly convex under mild regularity conditions (guaranteed by the contraction property of the equilibrium map, Appendix H, which bounds the feedback coefficient b_j).

Proof of Lemma 15. We provide the detailed algebraic derivation, generalizing the results from Appendix B.2. We analyze the derivatives on a SE block where κ is constant and preferences are linear. Let $d'(\delta)$ be the slope of the misfit disutility for the winning segment. Let $A = 2 + 2\kappa t$.

First Derivatives. From $t + p = \bar{\omega} - d(\delta)$ and $p = t + \kappa t^2$. $t' + p' = -d'$. $p' = t'(1 + 2\kappa t)$. $t'(2 + 2\kappa t) = -d'$. $t' = -d'/A$.

Second Derivatives. $d''(\delta) = 0$. $A' = 2\kappa t'$. We differentiate t': $t'' = \frac{d'A'}{A^2} = \frac{d'(2\kappa t')}{A^2} = \frac{2\kappa d'(-d'/A)}{A^2} = -\frac{2\kappa (d')^2}{A^3}$. We differentiate p':

$$\begin{aligned} p'' &= t''(1 + 2\kappa t) + 2\kappa (t')^2 \\ &= \left(-\frac{2\kappa (d')^2}{A^3}\right)(A-1) + 2\kappa\left(\frac{(d')^2}{A^2}\right) \\ &= \frac{2\kappa (d')^2}{A^3}[-(A-1) + A] = \frac{2\kappa (d')^2}{A^3}. \end{aligned}$$

Let $\Delta_B^2 = (d')^2$. Substituting $A = 2(1+\kappa t)$: $X_j''(\delta) = p''(\delta) = \frac{2\kappa \Delta_B^2}{8(1+\kappa t)^3} = \frac{\kappa \Delta_B^2}{4(1+\kappa t)^3} \geq 0$.

Curvature of $r(\delta)$. $r = t^2 + \kappa t^3$. $r'' = (2 + 6\kappa t)(t')^2 + (2t + 3\kappa t^2)t''$. Substitute t' and t'':

$$\begin{aligned} r'' &= (2 + 6\kappa t)\frac{(d')^2}{A^2} + (2t + 3\kappa t^2)\left(-\frac{2\kappa (d')^2}{A^3}\right) \\ &= \frac{(d')^2}{A^3}\left[(2 + 6\kappa t)A - 2\kappa(2t + 3\kappa t^2)\right]. \end{aligned}$$

Expanding the bracketed term: $(2 + 6\kappa t)(2 + 2\kappa t) - (4\kappa t + 6\kappa^2 t^2) = (4 + 16\kappa t + 12\kappa^2 t^2) - (4\kappa t + 6\kappa^2 t^2) = 4 + 12\kappa t + 6\kappa^2 t^2$. $r''(\delta) = \frac{\Delta_B^2}{4(1+\kappa t)^3}(2 + 6\kappa t + 3\kappa^2 t^2) > 0$.

Lemma 16 **(Strict Convexity Post-SE)**. *Under the regularity implied by the contraction property, the Gateaux derivative $H_j(\delta; \mathbf{D}^*)$ is strictly convex on SE blocks.*

G.6. Step 5: Proof of Bind-Gap-Bind Structures (Parts ii and iii)

5.1. Decoupling and the Frozen LPs. We utilize the Freeze-then-Lift Separability (established implicitly by Lemma 13) to analyze the structure segment by segment. We define the residual capacity $U_j(\delta;\mathbf{D}^*)$ available to segment j. Lemma 13 confirms that D_j^* must solve the frozen LP ($P_{\text{fix},j}$), optimizing D_j against the residual capacity U_j.

5.2. Personalized Structure (Part ii). We analyze the frozen LP for segment j. The KKT conditions imply the standard optimality rule:

- Bind: $D_j^*(\delta) = U_j(\delta;\mathbf{D}^*)$ if $H_j(\delta) > \nu_j^*$.
- Gap: $D_j^*(\delta) = 0$ if $H_j(\delta) < \nu_j^*$.

Since $H_j(\delta)$ is strictly convex (Lemma 16), it is "U-shaped". The superlevel set $\{\delta : H_j(\delta) \geq \nu_j^*\}$ consists of at most two closed intervals at the extremes of the segment's support. This yields the Personalized Bind-Gap-Bind structure.

5.3. Aggregate Structure (Part iii). The shadow price of capacity $\eta^*(\delta)$ is the upper envelope of the NMR functions $F_j(\delta)$ (Eq. (65)). Since each $F_j(\delta)$ is strictly convex, the upper envelope $\eta^*(\delta)$ is also convex.

The Aggregate Gap $[\delta_0, \delta_1]$ is defined where capacity is slack, i.e., $\eta^*(\delta) = 0$. The binding regions are the extremes: $[\underline{\delta}^*, \delta_0] \cup [\delta_1, \bar{\delta}^*]$.

5.4. Zero Exposure in the Gap. In the interior of the gap (δ_0, δ_1), $\eta^*(\delta) = 0$. By stationarity (60), $H_j(\delta) \leq \nu_j^*$. Since $H_j(\delta)$ is strictly convex (and piecewise analytic), it cannot equal a constant ν_j^* over any interval of positive measure. Thus, $H_j(\delta) < \nu_j^*$ a.e. in the gap. By the KKT conditions, $D_j^*(\delta) = 0$ a.e. for all j. Therefore, the aggregate exposure $y^*(\delta) = 0$ a.e. in the gap (δ_0, δ_1).

This completes the rigorous proof of Theorem 4 for general $\rho \in (0, 1]$.

Appendix H: Proof of Proposition 8: Value of Information Decomposition and Welfare Impact

This appendix provides a proof of Proposition 8.

H.1. Mathematical Preliminaries: Defining Per-Period Customer Surplus

We establish the definition and calculation of the aggregate per-period Customer Surplus (CS_j) in a stationary equilibrium.

Definition 4 (Aggregate Per-Period Customer Surplus). CS_j is the total surplus realized by all active customers in segment j during a single period in the steady state. It is calculated as the integral of the expected instantaneous surplus flow $A_j(v)$ over the steady-state distribution of active customers $c_{V_j}(v)$.

$$CS_j = \int_{\nu_j^*} c_{V_j}(v) A_j(v) dv. \tag{67}$$

Where $A_j(v)$ is the expected immediate realized surplus $(v - X)$ conditional on acceptance:

$$A_j(v) = \mathbb{E}_{\delta \sim D_j}[(v - X(\delta)) \cdot \mathbf{1}\{v - X(\delta) \geq \theta_j(v)\}]. \tag{68}$$

Lemma 17 **(Equivalence of CS Definitions in Steady State)**. *The aggregate per-period surplus (Instantaneous Flow) equals the total expected lifetime surplus of the arriving cohort of entrants (Lifetime Value of Arrivals).*

Proof. We use the fundamental relationships derived from the model dynamics (Appendices A and B). Let $\pi_{c,j}(v)$ be the lifetime expected surplus. Let $B_j(v) = \mu_j + (1-\mu_j)q_{c,j}(v)$ be the effective departure rate (hazard rate).

1. Bellman Equation Relationship: The definition of $\pi_{c,j}(v)$ via the Bellman equation implies: $\pi_{c,j}(v) = A_j(v)/B_j(v)$. Rearranging gives the relationship between instantaneous flow and lifetime value: $A_j(v) = \pi_{c,j}(v)B_j(v)$.

2. Steady-State Balance Equation: The inflow density of entrants $\Lambda_{in,j}(v)$ equals the outflow density: $\Lambda_{in,j}(v) = c_{V_j}(v)B_j(v)$.

3. Substitution: Substitute these into the definition of CS_j: $CS_j = \int c_{V_j}(v)A_j(v)dv = \int c_{V_j}(v)\left(\pi_{c,j}(v)B_j(v)\right)dv$. $CS_j = \int \left(c_{V_j}(v)B_j(v)\right)\pi_{c,j}(v)dv = \int \Lambda_{in,j}(v)\pi_{c,j}(v)dv$. The final expression is the Lifetime Value of Arrivals. The equivalence holds rigorously.

The Decomposition and the Frozen Step. We analyze the decomposition $\Delta CS^{\text{Total}} = \Delta CS^{\text{Matching}} + \Delta CS^{\text{Feedback}}$.

The "Frozen" step (OPT-F) involves reallocating the aggregate exposure $y^{NP}(\delta)$ from D^{NP} to D^F to maximize revenue, while holding the equilibrium behavior (prices $X^{NP}(\delta)$ and thresholds $\theta_j^{NP}(v)$) fixed at E^{NP}.

We define the Matching Gain based on the immediate change in the instantaneous surplus flow, holding the population masses fixed at C_j^{NP}. $\Delta CS^{\text{Matching}} = CS^{P\text{-Frozen}} - CS^{NP}$.

H.2. Proof of Part (i): Matching Gain (Efficiency)

Step 1.1: Proving $\text{VoI}^{Frozen} \geq 0$. By definition of the optimization (OPT-F), $R^{\text{P (NP Equil.)}} \geq R^{NP}$. Thus, $\text{VoI}^{\text{Frozen}} \geq 0$.

Step 1.2: Analyzing $\Delta CS^{Matching}$ (Ambiguity). We must rigorously derive the Surplus Kernel $S_j(\delta)$ and compare it with the Revenue Kernel $K_j(\delta)$.

LEMMA 18 **(Derivation of the True Surplus Kernel $S_j(\delta)$).** *The Surplus Kernel $S_j(\delta)$—the expected realized surplus $(v - X)$ when a random active customer meets supplier δ—is composed of two parts: the surplus relative to WTP, and the continuation value.*

$$S_j(\delta) = \underbrace{\frac{\Delta_j}{2}\phi_j(\delta)^2}_{\text{Part A: Relative to WTP}} + \underbrace{Avg(\theta_j(v)|Sale)\cdot\phi_j(\delta)}_{\text{Part B: Continuation Value}}. \tag{69}$$

Proof. $S_j(\delta)$ is the expected value of $v - X(\delta)$ conditional on a match and acceptance. We utilize the crucial result that the distribution of WTP ω among active customers is Uniform on $[\underline{\omega}_j, \bar{\omega}_j]$ (Theorem 2). The normalized density is $1/\Delta_j$.

A sale occurs if $\omega \geq X(\delta)$. The realized surplus is $v - X(\delta)$. We use the identity $v = \omega + \theta(v)$. Let $h_j(\omega)$ be the inverse map $v = h_j(\omega)$.

$S_j(\delta) = \frac{1}{\Delta_j}\int_{X(\delta)}^{\bar{\omega}_j}(h_j(\omega) - X(\delta))d\omega$. $S_j(\delta) = \frac{1}{\Delta_j}\int_{X(\delta)}^{\bar{\omega}_j}(\omega - X(\delta) + \theta_j(h_j(\omega)))d\omega$. We decompose the integral:

Part A (Relative to WTP): $\frac{1}{\Delta_j}\int_{X(\delta)}^{\bar{\omega}_j}(\omega - X(\delta))d\omega = \frac{1}{\Delta_j}\left[\frac{(\omega - X(\delta))^2}{2}\right]_{X(\delta)}^{\bar{\omega}_j} = \frac{(\bar{\omega}_j - X(\delta))^2}{2\Delta_j}$. Let $t_j(\delta) = \bar{\omega}_j - X(\delta)$. The sale probability is $\phi_j(\delta) = t_j(\delta)/\Delta_j$. Part A $= \frac{t_j(\delta)^2}{2\Delta_j} = \frac{\Delta_j}{2}\phi_j(\delta)^2$.

Part B (Continuation Value): $\frac{1}{\Delta_j}\int_{X(\delta)}^{\bar{\omega}_j}\theta_j(h_j(\omega))d\omega$. We can rewrite this as: $\frac{\bar{\omega}_j - X(\delta)}{\Delta_j} \cdot \left(\frac{1}{\bar{\omega}_j - X(\delta)}\int_{X(\delta)}^{\bar{\omega}_j}\theta_j(h_j(\omega))d\omega\right)$. This is $\phi_j(\delta)\cdot\text{Avg}(\theta_j(v)|\text{Sale})$.

Lemma 19 **(Ambiguity of Matching Effect on CS)**. *The reallocation that maximizes revenue (*$\text{VoI}^{Frozen} \geq 0$*) may decrease customer surplus (*$\Delta CS^{Matching} < 0$*).*

Proof. (OPT-F) maximizes the Revenue Kernel $K_j(\delta) = p^*(\delta)\phi_j(\delta)$. We analyze the impact on CS based on the Surplus Kernel $S_j(\delta)$ (Lemma 18).

$S_j(\delta)$ depends on Δ_j (Part A) and the structure of $\theta_j(v)$ (Part B). When heterogeneity exists (e.g., in patience μ_j or valuations $\bar{v}_j$), these components vary across segments.

The relationship between K_j and S_j is complex: $S_j(\delta) = \frac{\Delta_j}{2(p^*)^2} K_j(\delta)^2 + \text{Avg}(\theta_j(v)|\text{Sale}) \cdot \frac{K_j(\delta)}{p^*(\delta)}$.

The structure of $S_j(\delta)$ is fundamentally different from $K_j(\delta)$ due to the presence of the continuation value term (Part B) and the non-linear dependence on K_j in Part A (weighted by Δ_j).

Demonstration via Heterogeneity: Consider heterogeneity in patience μ_j. Let $\mu_H > \mu_L$. Impatient segments (H) tend to have larger WTP spans ($\Delta_H > \Delta_L$). Patient segments (L) have higher continuation values ($\theta_L(v) > \theta_H(v)$). Part A favors H. Part B favors L. If the platform shifts allocation from L to H (because $K_H > K_L$), revenue increases. However, this shifts surplus realization from a segment with high continuation values (L) to one with low continuation values (H). If the loss in Part B (Continuation Value) dominates the potential gain in Part A (WTP Slack), the overall surplus decreases.

Since the optimization objective (revenue) does not align with the welfare measure (true surplus), $\Delta CS^{\text{Matching}}$ is ambiguous.

H.3. Proof of Part (ii): Feedback Gain (Strategic Exploitation)

We analyze the impact of the equilibrium adjustment from the Frozen state towards E^P.

Step 2.1: Analyzing $\text{VoI}^{Adjustment}$. The sign of $\text{VoI}^{\text{Adjustment}} = R^P - R^{\text{P (NP Equil.)}}$ is mathematically ambiguous.

Step 2.2: Proving $\Delta CS^{Feedback} < 0$.

Lemma 20 **(Strategic Optimization Increases Mean Effective Prices)**. *The optimal personalized policy* $\mathbf{D}^P$ *(BGB structure) strategically increases the mean effective prices* M_j.

Proof. 1. **Convexity:** The effective price $X_j(\delta)$ is convex in δ on SE blocks (Appendix I, Lemma I.1), for any $\rho \in (0,1]$. 2. **BGB Strategy:** The BGB structure shifts mass from the interior to the boundaries (a mean-preserving spread). 3. **Impact:** By Jensen's inequality applied to the convex function $X_j(\delta)$, this reallocation increases $M_j = \mathbb{E}_{D_j}[X_j(\delta)]$. This is amplified by the positive equilibrium feedback ($\partial X / \partial M > 0$, Appendix H).

Lemma 21 **(Strategic Optimization Harms CS)**. *The equilibrium adjustment driven by the revenue-maximizing strategy strictly decreases aggregate customer surplus (*$\Delta CS^{Feedback} < 0$*).*

Proof. 1. **Mechanism.** By Lemma 20, the strategic optimization towards E^P involves an increase in $\mathbf{M}$.

2. **Impact on Continuation Values.** The customer's acceptance threshold $\theta_j(v)$ is strictly decreasing in M_j. This is a rigorous comparative static derived from the Bellman equation (Generalization of Appendix D.1): $\frac{\partial \theta_j(v)}{\partial M_j} < 0$.

3. **Impact on Lifetime Surplus.** The customer's expected lifetime surplus is $\pi_{c,j}(v) = \theta_j(v)/(1-\mu_j)$. The increase in $\mathbf{M}$ strictly decreases $\pi_{c,j}(v)$ pointwise.

4. **Impact on Aggregate Per-Period CS.** By Lemma 17, the aggregate per-period CS equals the Lifetime Value of Arrivals: $CS_j = \int \Lambda_{in,j}(v)\pi_{c,j}(v)dv$. Since $\pi_{c,j}(v)$ decreases pointwise, and the entry set $\mathcal{V}_j^*$ also shrinks (as the entry cutoff v_e increases), the aggregate CS strictly decreases.

H.4. Proof of Part (iii): Total Impact

Step 3.1: Total Revenue Impact. $\text{VoI}^{\text{Total}} = R^P - R^{NP} \geq 0$.

Step 3.2: Total Customer Surplus Impact.

$$\Delta CS^{\text{Total}} = \underbrace{\Delta CS^{\text{Matching}}}_{\text{Ambiguous Sign (Lemma 19)}} + \underbrace{\Delta CS^{\text{Feedback}}}_{<0\text{ (Lemma 21)}}. \tag{70}$$

The total impact is ambiguous and can be negative. [Omer: Analyse the impact of μ values] This completes the proof of Proposition 8.

Appendix I: Equilibrium Analysis for Experiment 1 (Thick Market, Patience Heterogeneity)

This appendix details the equilibrium derivation for the model used in Experiment 1. This analysis corresponds to the setting with a thick market ($m \to \infty$), myopic suppliers ($\rho = 1$), heterogeneous patience (μ_j), and homogeneous valuations ($\bar{v}_j = 1$). This setting features a coupled equilibrium system, where segments interact through the endogenous market boundary $\bar{\delta}$. This analysis aligns with the provided reference material "Personalization with Heterogeneous Patience".

I.1. X.1. Environment and Equilibrium Structure

We analyze the stationary equilibrium under a personalized barbell policy $\boldsymbol{\alpha} = \{\alpha_j\}$.

WTP Geometry and Pricing. In the unique stationary equilibrium (Theorems 1 and 2), the Willingness-to-Pay (WTP) distribution for segment j is uniform over an endogenous support $[\underline{\omega}_j, \bar{\omega}_j]$. Let $\Delta_j = \bar{\omega}_j - \underline{\omega}_j$ be the WTP span.

A myopic supplier ($\rho = 1$) maximizes immediate profit. The optimal effective price $X_j^*(\delta)$ (posted price plus misfit δ) is found by maximizing the expected profit:

$$\max_X \quad \pi_j(X) = (X-\delta)\cdot P(\omega_j \geq X) = (X-\delta)\frac{\bar{\omega}_j - X}{\Delta_j}.$$

The First Order Condition (FOC) is:

$$\frac{\partial \pi_j}{\partial X} = \frac{1}{\Delta_j}\left[(1)(\bar{\omega}_j - X) + (X-\delta)(-1)\right] = 0$$

$$\bar{\omega}_j - 2X + \delta = 0 \implies X_j^*(\delta) = \frac{\bar{\omega}_j + \delta}{2}.$$

In equilibrium, the lowest WTP must equal the lowest effective price (at $\delta = 0$): $\underline{\omega}_j = X_j^*(0) = \bar{\omega}_j/2$. This establishes the equilibrium WTP structure: $U[\Delta_j, 2\Delta_j]$. The optimal effective price is $X_j^*(\delta) = \Delta_j + \delta/2$.

The Feedback Identity. The WTP span Δ_j is determined endogenously by the Feedback Identity, derived from the customer's Bellman equation (Section 3.1). Since $\bar{v}_j = 1$ in Experiment 1:

$$2\Delta_j = \mu_j + (1-\mu_j)M_j, \tag{71}$$

where M_j is the mean effective price faced by segment j.

Market Coupling via $\bar{\delta}$. The market boundary $\bar{\delta}$ is endogenous and set by the segment with the largest WTP span, the "Top Segment" $j^* = \arg\max_k \Delta_k$. Thus, $\bar{\delta} = 2\Delta_{j^*}$. When valuations are homogeneous, the top segment is the most impatient one: $j^* = \arg\max_k \mu_k$.

I.2. Solving the Coupled Equilibrium System

We solve the fixed-point system defined by the Feedback Identity and the definition of M_j under the barbell policy. Let $c_j = 1 + \mu_j$ and $d_j = 1 - \mu_j$.

Top Segment (j^*). The top segment faces the best match ($\delta = 0$) with price $X^*_{j^*}(0) = \Delta_{j^*}$, and the worst match ($\delta = \bar{\delta} = 2\Delta_{j^*}$) with price $X^*_{j^*}(\bar{\delta}) = \Delta_{j^*} + (2\Delta_{j^*})/2 = 2\Delta_{j^*}$. The mean effective price is:

$$
\begin{aligned}
M_{j^*} &= (1-\alpha_{j^*})X^*_{j^*}(0) + \alpha_{j^*}X^*_{j^*}(\bar{\delta}) \\
&= (1-\alpha_{j^*})\Delta_{j^*} + \alpha_{j^*}(2\Delta_{j^*}) = \Delta_{j^*}(1+\alpha_{j^*}).
\end{aligned}
$$

Substituting M_{j^*} into the Feedback Identity (71):

$$
\begin{aligned}
2\Delta_{j^*} &= \mu_{j^*} + d_{j^*}\Delta_{j^*}(1+\alpha_{j^*}) \\
\Delta_{j^*}(2 - d_{j^*}(1+\alpha_{j^*})) &= \mu_{j^*} \\
\Delta_{j^*}(2 - (1 - \mu_{j^*} + \alpha_{j^*}(1-\mu_{j^*}))) &= \mu_{j^*} \\
\Delta_{j^*}(1 + \mu_{j^*} - d_{j^*}\alpha_{j^*}) &= \mu_{j^*} \\
\Delta_{j^*}(c_{j^*} - d_{j^*}\alpha_{j^*}) &= \mu_{j^*}. \\
\Delta_{j^*}(\alpha_{j^*}) &= \frac{\mu_{j^*}}{c_{j^*} - d_{j^*}\alpha_{j^*}}. \qquad (72)
\end{aligned}
$$

Other Segments ($j \neq j^*$). For these segments, $\bar{\delta} > 2\Delta_j$. We assume that the effective price observed at the bad arm is $\bar{\delta}$ (which results in no sale).

$$
M_j = (1-\alpha_j)\Delta_j + \alpha_j\bar{\delta}.
$$

Substituting M_j into the Feedback Identity (71):

$$
\begin{aligned}
2\Delta_j &= \mu_j + d_j((1-\alpha_j)\Delta_j + \alpha_j\bar{\delta}) \\
\Delta_j(2 - d_j(1-\alpha_j)) &= \mu_j + d_j\alpha_j\bar{\delta} \\
\Delta_j(2 - (1 - \mu_j - \alpha_j(1-\mu_j))) &= \mu_j + d_j\alpha_j\bar{\delta} \\
\Delta_j(1 + \mu_j + d_j\alpha_j) &= \mu_j + d_j\alpha_j\bar{\delta} \\
\Delta_j(c_j + d_j\alpha_j) &= \mu_j + d_j\alpha_j\bar{\delta}. \\
\Delta_j(\alpha_j \mid \bar{\delta}) &= \frac{\mu_j + d_j\alpha_j\bar{\delta}}{c_j + d_j\alpha_j}. \qquad (73)
\end{aligned}
$$

Equations (72) and (73), coupled via $\bar{\delta} = 2\Delta_{j^*}$, define the equilibrium system solved numerically in Experiment 1.

I.3. Revenue Optimization

The platform maximizes total revenue $R = \sum_j R_j$. We first derive the revenue function.

Stationary Mass and Revenue. According to Theorem 2, the steady-state density of active customers is λ_j/μ_j over the WTP support $[\Delta_j, 2\Delta_j]$. The total mass is:

$$C_j = \int_{\Delta_j}^{2\Delta_j} \frac{\lambda_j}{\mu_j} d\omega = \frac{\lambda_j \Delta_j}{\mu_j}. \tag{74}$$

The expected commission per interaction for segment j is $\mathbb{E}[K_j]$. Revenue is generated only at the good arm ($\delta = 0$), where the price is $p_j^*(0) = \Delta_j$.

$$\mathbb{E}[K_j] = (1-\alpha_j)(\tau p_j^*(0)) + \alpha_j(0) = \tau(1-\alpha_j)\Delta_j. \tag{75}$$

The total revenue flow from segment j is:

$$R_j(\boldsymbol{\alpha}) = C_j \cdot \mathbb{E}[K_j] = \left(\frac{\lambda_j \Delta_j}{\mu_j}\right) \cdot (\tau(1-\alpha_j)\Delta_j) = \frac{\tau\lambda_j}{\mu_j}(1-\alpha_j)\Delta_j(\boldsymbol{\alpha})^2. \tag{76}$$

Personalized (P) Optimization. The platform optimizes $\{\alpha_j\}$ to maximize $\sum_j R_j(\boldsymbol{\alpha})$. This is a coupled optimization problem. We analyze the optimization for the top segment, which determines $\bar{\delta}$.

We maximize $R_{j^*}(\alpha_{j^*})$. Using (72):

$$R_{j^*}(\alpha_{j^*}) \propto (1-\alpha_{j^*}) \left(\frac{\mu_{j^*}}{c_{j^*} - d_{j^*}\alpha_{j^*}}\right)^2. \tag{77}$$

We maximize $f(\alpha) = \frac{1-\alpha}{(c-d\alpha)^2}$. The FOC (detailed derivation in the provided reference material) yields the maximizer $\alpha^* = (2d-c)/d$. Substituting c_{j^*}, d_{j^*}:

$$\alpha_{j^*}^* = \frac{2(1-\mu_{j^*}) - (1+\mu_{j^*})}{1-\mu_{j^*}} = \frac{1-3\mu_{j^*}}{1-\mu_{j^*}}.$$

The optimal policy is $\alpha_{j^*}^P = \max\{0, \alpha_{j^*}^*\}$. This determines $\bar{\delta}^P$. For $j \neq j^*$, the platform maximizes $R_j(\alpha_j \mid \bar{\delta}^P)$. This involves solving a quadratic FOC (detailed in the reference material). The system is solved iteratively or via multivariate optimization, as implemented in the code for Experiment 1.

Non-Personalized (NP) Optimization. The platform chooses a single α. The equilibrium spans $\Delta_j(\alpha)$ are calculated using this common α, identifying the top segment j^* (highest μ_j) and the resulting $\bar{\delta}(\alpha) = 2\Delta_{j^*}(\alpha)$. The optimal α^{NP} maximizes $R^{NP}(\alpha) = \sum_j R_j(\alpha)$. This requires solving the coupled system and optimizing numerically.

Appendix J: Capacity-Constrained Analysis (Experiment 2) and Proof of Proposition 9

This appendix details the optimization framework for Experiment 2, where a finite capacity constraint $\bar{M}$ is active. As justified in the main text (Section 6.3) and Appendix I, this analysis utilizes a stylized barbell model to ensure tractability when analyzing the Matching Gain under capacity constraints. This analysis aligns with the provided reference material "Capacity-Coupled Toolkit". We assume homogeneous patience ($\mu_j = \mu$) and heterogeneous valuations ($\bar{v}_j$).

J.1. Stylized Model Structure under Capacity Constraints

We adopt the stylized model structure consistent with the "Capacity-Coupled Toolkit". Let $D(\alpha) = 1 + \mu - (1-\mu)\alpha$.

Model Primitives. In this framework, the equilibrium relationships are defined locally (decoupled WTP spans):

- WTP Span: $\Delta_j(\alpha_j) = \frac{\mu \bar{v}_j}{D(\alpha_j)}$.
- Stationary Mass (Derived in the Toolkit): $C_j(\alpha_j) = \frac{\lambda_j}{D(\alpha_j)}$.
- Exposure (Best Match): $Exp_j(\alpha_j) = (1-\alpha_j)C_j = \lambda_j \frac{1-\alpha_j}{D(\alpha_j)}$.
- Revenue (Ad-Valorem, Derived in the Toolkit):

$$\begin{aligned} R_j(\alpha_j) &= \tau \cdot p_j^* \cdot \text{Sales}_j = \tau \cdot \Delta_j \cdot (1-\mu) Exp_j(\alpha_j) \\ &= \tau \left(\frac{\mu \bar{v}_j}{D(\alpha_j)} \right) (1-\mu) \left(\lambda_j \frac{1-\alpha_j}{D(\alpha_j)} \right) \\ &= \tau \lambda_j \mu (1-\mu) \bar{v}_j \frac{1-\alpha_j}{D(\alpha_j)^2}. \end{aligned}$$

Note the distinction in the revenue structure compared to the model in Appendix I. Here, $R_j \propto \bar{v}_j$, whereas in Appendix I, $R_j \propto \bar{v}_j^2$. This difference arises from the distinct derivations of the stationary mass C_j in the respective underlying models. We proceed using the definitions consistent with the Experiment 2 implementation.

J.2. KKT Analysis of the Constrained Problem

The platform solves:

$$\max_{\{\alpha_j\}} \quad \sum_j R_j(\alpha_j) \quad \text{s.t.} \quad \sum_j Exp_j(\alpha_j) \le \bar{M}.$$

We form the Lagrangian with shadow price $\eta \ge 0$:

$$\mathcal{L} = \sum_j R_j(\alpha_j) + \eta \left(\bar{M} - \sum_j Exp_j(\alpha_j) \right). \tag{78}$$

Derivatives. We calculate the derivatives required for the FOC.

1. Derivative of Exposure: $\frac{\partial Exp_j}{\partial \alpha_j} = \lambda_j \frac{d}{d\alpha_j} \left(\frac{1-\alpha_j}{D(\alpha_j)} \right)$. Using the quotient rule ($u = 1-\alpha, v = D(\alpha), v' = -(1-\mu)$):

$$\begin{aligned} \frac{u'v - uv'}{v^2} &= \frac{(-1)D(\alpha) - (1-\alpha)(-(1-\mu))}{D(\alpha)^2} \\ &= \frac{-(1+\mu-(1-\mu)\alpha) + (1-\mu)(1-\alpha)}{D(\alpha)^2} \\ &= \frac{-1-\mu+(1-\mu)\alpha + 1-\mu-(1-\mu)\alpha}{D(\alpha)^2} = \frac{-2\mu}{D(\alpha)^2}. \end{aligned}$$

2. Derivative of Revenue: Let $K_j = \tau \lambda_j \mu (1-\mu) \bar{v}_j$.

$$\frac{\partial R_j}{\partial \alpha_j} = K_j \frac{d}{d\alpha_j} \left(\frac{1-\alpha_j}{D(\alpha_j)^2} \right).$$

Using the quotient rule ($u = 1-\alpha, v = D(\alpha)^2, v' = 2D(\alpha)(-(1-\mu))$):

$$\begin{aligned} \frac{u'v - uv'}{v^2} &= \frac{(-1)D(\alpha)^2 - (1-\alpha) \cdot 2D(\alpha)(-(1-\mu))}{D(\alpha)^4} \\ &= \frac{-D(\alpha) + 2(1-\mu)(1-\alpha)}{D(\alpha)^3}. \end{aligned}$$

Expanding the numerator:

$$\begin{aligned} \text{Num} &= -(1+\mu-(1-\mu)\alpha) + 2(1-\mu) - 2(1-\mu)\alpha \\ &= (1-3\mu) - (1-\mu)\alpha. \end{aligned}$$

First-Order Condition (Interior Solution). Assuming $0 < \alpha_j < 1$, the FOC is $\frac{\partial R_j}{\partial \alpha_j} - \eta \frac{\partial Exp_j}{\partial \alpha_j} = 0$.

$$K_j \frac{(1-3\mu)-(1-\mu)\alpha_j}{D(\alpha_j)^3} - \eta \left(\lambda_j \frac{-2\mu}{D(\alpha_j)^2} \right) = 0.$$

Substitute K_j and multiply by $D(\alpha_j)^3/(\lambda_j \mu)$:

$$\tau(1-\mu)\bar{v}_j\left((1-3\mu)-(1-\mu)\alpha_j\right) + 2\eta D(\alpha_j) = 0.$$

We solve this linear equation for α_j. Substitute $D(\alpha_j) = (1+\mu) - (1-\mu)\alpha_j$.

$$\tau(1-\mu)\bar{v}_j(1-3\mu) - \tau(1-\mu)^2\bar{v}_j\alpha_j + 2\eta(1+\mu) - 2\eta(1-\mu)\alpha_j = 0.$$

Isolating α_j:

$$\alpha_j\left[\tau(1-\mu)^2\bar{v}_j + 2\eta(1-\mu)\right] = \tau(1-\mu)\bar{v}_j(1-3\mu) + 2\eta(1+\mu).$$

The optimal interior policy as a function of the shadow price η is:

$$\alpha_j(\eta) = \frac{\tau(1-\mu)\bar{v}_j(1-3\mu) + 2\eta(1+\mu)}{(1-\mu)\left(\tau(1-\mu)\bar{v}_j + 2\eta\right)}. \tag{79}$$

The optimal policy α_j^P is this expression clipped to $[0,1]$. The equilibrium η^* is determined by complementary slackness, solving the capacity equation if the constraint binds.

J.3. Proof of Proposition 9

We must demonstrate that $\alpha_j(\eta)$ is strictly decreasing in the valuation $\bar{v}_j$ when the shadow price $\eta > 0$. We analyze the derivative of the interior solution (79) with respect to $\bar{v}_j$.

Let $V = \bar{v}_j$. We define the following constants, which are independent of V:

$$C_1 = \tau(1-\mu)(1-3\mu)$$
$$C_2 = 2\eta(1+\mu)$$
$$C_3 = \tau(1-\mu)^2$$
$$C_4 = 2\eta(1-\mu)$$

The function is structured as $\alpha(V) = \frac{C_1 V + C_2}{C_3 V + C_4}$.

We apply the quotient rule to find the derivative $\frac{\partial \alpha}{\partial V}$:

$$\begin{aligned}\frac{\partial \alpha}{\partial V} &= \frac{C_1(C_3 V + C_4) - C_3(C_1 V + C_2)}{(C_3 V + C_4)^2} \\ &= \frac{C_1 C_4 - C_3 C_2}{(C_3 V + C_4)^2}.\end{aligned}$$

The sign of the derivative is determined by the numerator $N = C_1 C_4 - C_3 C_2$. We substitute the constants:

$$N = \left(\tau(1-\mu)(1-3\mu)\right)\cdot\left(2\eta(1-\mu)\right) - \left(\tau(1-\mu)^2\right)\cdot\left(2\eta(1+\mu)\right).$$

We factor out the common terms $2\eta\tau(1-\mu)^2$:

$$\begin{aligned}N &= 2\eta\tau(1-\mu)^2\left[(1-3\mu)-(1+\mu)\right] \\ &= 2\eta\tau(1-\mu)^2\left[1-3\mu-1-\mu\right] \\ &= 2\eta\tau(1-\mu)^2[-4\mu] \\ &= -8\eta\tau\mu(1-\mu)^2.\end{aligned}$$

Since $\tau > 0, \mu \in (0,1)$, and we assume the constraint binds ($\eta > 0$), the numerator N is strictly negative.

$$\frac{\partial \alpha_j(\eta)}{\partial \bar{v}_j} < 0. \tag{80}$$

The optimal scarcity level is strictly decreasing in the segment's valuation cap. Given that $\bar{v}_G > \bar{v}_B$, it follows that $\alpha_G(\eta^*) < \alpha_B(\eta^*)$. This holds strictly even if one policy hits a boundary (0 or 1), as they cannot hit the same boundary simultaneously due to different $\bar{v}_j$ (unless η is sufficiently large or small to force both to the same boundary, in which case the relationship holds weakly; however, strict inequality holds over the range where differentiation occurs).

Appendix K: Robustness Check: Endogenous Customer Departure

In the main model, we assume the customer departure probability μ (impatience) is exogenous. This assumption is critical for the tractability of the dynamic equilibrium model, specifically enabling the Uniform WTP result (Theorem 2). In this appendix, we investigate the implications of endogenizing this departure decision. We analyze two distinct mechanisms for endogeneity. The first mechanism, where departure depends on individual valuations, leads to intractability. The second mechanism, where departure depends on aggregate market conditions, preserves tractability and, surprisingly, reinforces the core "Scarcity by Design" mechanism. This analysis demonstrates that our baseline model provides a conservative estimate of the strategic benefits of scarcity.

We focus on the single-segment case ($J = 1$) with valuations $v \sim U[0,1]$. We analyze a model where the departure rate depends on the aggregate quality of the marketplace, proxied by the Mean Effective Price M. This captures the idea that customers become more likely to abandon the platform when the expected search environment deteriorates through poorer recommendations (higher M).

Model Setup. $\mu = \mu(M)$, where $M = \mathbb{E}_D[p^*(\delta) + \delta]$. We assume $\mu'(M) = \eta > 0$.

Preservation of Tractability. Since $\mu(M)$ is independent of v, the structure of the individual customer's optimization problem remains identical to the exogenous case analyzed in Appendix A. The indifference condition remains linear:

$$\theta(v) = (1 - \mu(M))\pi_c(v).$$

The derivation of the Uniform WTP property relies only on this linearity and the v-independence of μ. Therefore, Theorem 2 still holds, and all structural results regarding the existence, uniqueness, and structure of the optimal recommendations (Theorems 1 and 3) are preserved.

K.1. Impact on the Profitability of Scarcity ($m \to \infty$)

We now analyze how this tractable endogeneity affects the condition under which "Scarcity by Design" is profitable. We focus on the thick-market limit ($m \to \infty$), where the feasibility constraint is relaxed. We compare the new threshold for profitability with the exogenous threshold derived in Proposition 5 ($\mu < \rho/(1+2\rho)$).

This mechanism introduces two competing effects when the platform induces scarcity (increases M):

1. **Attrition Effect (Negative):** Higher M increases $\mu(M)$, reducing the steady-state customer mass $C(M)$.
2. **Desperation Effect (Positive):** Higher $\mu(M)$ means customers are more impatient. This increased impatience raises their WTP, allowing suppliers to increase prices further.

The Profitability Condition. In the thick market limit, the baseline policy D^B is a Dirac measure at $\delta = 0$. Scarcity (diverting mass to $\bar{\delta}$) is profitable if the Gateaux derivative at $\bar{\delta}$ exceeds the marginal revenue at 0: $H(\bar{\delta}) > H(0)$. As derived in Appendix D.3, this condition simplifies to:

$$\Gamma\rho > C,$$

where Γ is the feedback multiplier and C is the steady-state mass of customers.

The Generalized Feedback Multiplier Γ. The feedback multiplier captures how a change in the recommendation policy (which changes M) feeds back into the revenue through equilibrium adjustments. It is defined as:

$$\Gamma = \frac{dK(0)/dM}{1 - dX(0)/dM}.$$

Here, $K(0)$ is the revenue kernel at $\delta = 0$, and $X(0)$ is the effective price at $\delta = 0$. Crucially, we must use *total derivatives* with respect to M, accounting for the endogenous relationship $M \to \mu(M)$.

Equilibrium Identities. We use the fundamental equilibrium relationships for $U[0,1]$ valuations (derived in Appendix B.3), adapted for the dependence on $\mu(M)$.

1. Feedback Identity: $\bar{\omega}(M) = \mu(M) + (1 - \mu(M))M$.
2. WTP Bounds: $\underline{\omega}(M) = \frac{1}{1+\rho}\bar{\omega}(M)$; $\Delta(M) = \rho\underline{\omega}(M)$.
3. Customer Mass: $C(M) = \frac{\lambda}{\mu(M)}\Delta(M) = \frac{\lambda\rho}{\mu(M)}\underline{\omega}(M)$.
4. Equilibrium Objects at $\delta = 0$: $X(0) = \underline{\omega}(M)$; $K(0) = C(M)\underline{\omega}(M)$.

Step 1: Calculating the Total Derivative of Price $dX(0)/dM$ (The Desperation Effect). We calculate the total derivative of $X(0) = \underline{\omega}(M)$. First, we differentiate $\bar{\omega}(M)$:

$$\begin{aligned}
\frac{d\bar{\omega}}{dM} &= \frac{d}{dM}[\mu(M) + M - \mu(M)M] \\
&= \mu'(M) + 1 - (\mu'(M)M + \mu(M)) \\
&= (1 - \mu(M)) + \mu'(M)(1 - M).
\end{aligned}$$

Now, we differentiate $\underline{\omega}(M)$:

$$\frac{d\underline{\omega}}{dM} = \frac{1}{1+\rho}\frac{d\bar{\omega}}{dM} = \frac{1}{1+\rho}\left[(1 - \mu(M)) + \mu'(M)(1 - M)\right]. \tag{81}$$

Interpretation (Desperation Effect): In the exogenous case ($\mu' = 0$), the derivative is $\frac{1-\mu}{1+\rho}$. Since $\mu' > 0$ and $M < 1$, the second term $\mu'(1 - M)$ is positive. This means that $d\underline{\omega}/dM$ is strictly larger under endogenous departure. An increase in M not only directly increases prices (the first term) but also increases the departure rate $\mu(M)$. Higher impatience makes customers less selective (more "desperate"), which allows suppliers to further increase prices (the second term).

Step 2: Calculating the Total Derivative of Revenue $dK(0)/dM$ (The Attrition Effect). We calculate the total derivative of $K(0)$. We express $K(0)$ explicitly:

$$K(0) = C(M)\underline{\omega}(M) = \frac{\lambda\rho}{\mu(M)}\underline{\omega}(M)^2.$$

We use logarithmic differentiation for tractability. $\ln K(0) = \ln(\lambda\rho) - \ln\mu(M) + 2\ln\underline{\omega}(M)$. Differentiating with respect to M:

$$\frac{1}{K(0)}\frac{dK(0)}{dM} = -\frac{\mu'(M)}{\mu(M)} + \frac{2}{\underline{\omega}(M)}\frac{d\underline{\omega}}{dM}.$$

Rearranging:

$$\frac{dK(0)}{dM} = K(0)\left[\frac{2}{\underline{\omega}(M)}\frac{d\underline{\omega}}{dM} - \frac{\mu'(M)}{\mu(M)}\right]. \tag{82}$$

Interpretation (Attrition Effect): The first term represents the revenue gain from higher prices (amplified by the Desperation Effect). The second term, $-\mu'/\mu$, is strictly negative and represents the revenue loss due to increased attrition (a smaller customer base $C(M)$).

Step 3: The Generalized Profitability Condition. We return to the profitability condition $\Gamma\rho > C$. We substitute the definition of Γ:

$$\frac{dK(0)/dM}{1 - dX(0)/dM}\rho > C.$$

We substitute $K(0) = C\underline{\omega}$ and $X(0) = \underline{\omega}$.

$$\frac{dK(0)/dM}{1 - d\underline{\omega}/dM}\rho > C.$$

Substitute the expression for $dK(0)/dM$ from Eq. (82):

$$\frac{C\underline{\omega}\left[\frac{2}{\underline{\omega}}\frac{d\underline{\omega}}{dM} - \frac{\mu'}{\mu}\right]}{1 - d\underline{\omega}/dM}\rho > C.$$

We cancel C (which is positive) from both sides:

$$\frac{\underline{\omega}\left[\frac{2}{\underline{\omega}}\frac{d\underline{\omega}}{dM} - \frac{\mu'}{\mu}\right]}{1 - d\underline{\omega}/dM}\rho > 1.$$

Multiply both sides by $1 - d\underline{\omega}/dM$. (We must verify this term is positive. From Eq. (81), $d\underline{\omega}/dM$ is positive. Since $1 - \mu < 1$ and $\mu'(1 - M)$ is typically small in a stable equilibrium, $d\underline{\omega}/dM < 1$. We assume this holds).

$$\rho\underline{\omega}\left[\frac{2}{\underline{\omega}}\frac{d\underline{\omega}}{dM} - \frac{\mu'}{\mu}\right] > 1 - \frac{d\underline{\omega}}{dM}.$$

Distribute $\rho\underline{\omega}$ on the LHS:

$$2\rho\frac{d\underline{\omega}}{dM} - \rho\underline{\omega}\frac{\mu'}{\mu} > 1 - \frac{d\underline{\omega}}{dM}.$$

Rearrange to collect terms involving $d\underline{\omega}/dM$:

$$2\rho\frac{d\underline{\omega}}{dM} + \frac{d\underline{\omega}}{dM} > 1 + \rho\underline{\omega}\frac{\mu'}{\mu}.$$

Factor out $d\underline{\omega}/dM$:

$$\underbrace{\frac{d\underline{\omega}}{dM}(1 + 2\rho)}_{\text{Marginal Benefit (MB)}} > \underbrace{1 + \rho\underline{\omega}\frac{\mu'}{\mu}}_{\text{Marginal Cost (MC)}}. \tag{83}$$

Interpretation: The LHS represents the total marginal benefit of inducing scarcity (increasing M). The RHS represents the total marginal cost.

Step 4: Dominance Analysis (Desperation vs. Attrition). We analyze how endogeneity ($\mu' > 0$) affects the profitability condition compared to the exogenous case ($\mu' = 0$). We compare the change in the Marginal Benefit (LHS) and the Marginal Cost (RHS).

Change in Marginal Benefit (Δ_{MB}): This captures the increased benefit due to the Desperation Effect.

$$\Delta_{MB} = MB_{\text{Endo}} - MB_{\text{Exo}} = (1+2\rho)\left(\frac{d\underline{\omega}}{dM}_{Endo} - \frac{d\underline{\omega}}{dM}_{Exo}\right).$$

We calculate the difference in the derivatives using Eq. (81):

$$\begin{aligned}\frac{d\underline{\omega}}{dM}_{Endo} - \frac{d\underline{\omega}}{dM}_{Exo} &= \frac{1}{1+\rho}[(1-\mu)+\mu'(1-M)] - \frac{1}{1+\rho}[1-\mu] \\ &= \frac{\mu'(1-M)}{1+\rho}.\end{aligned}$$

Substituting this back into Δ_{MB}:

$$\Delta_{MB} = (1+2\rho)\frac{\mu'(1-M)}{1+\rho} = \frac{1+2\rho}{1+\rho}\mu'(1-M).$$

Change in Marginal Cost (Δ_{MC}): This captures the increased cost due to the Attrition Effect.

$$\Delta_{MC} = MC_{Endo} - MC_{Exo} = \left(1+\rho\underline{\omega}\frac{\mu'}{\mu}\right) - (1) = \rho\underline{\omega}\frac{\mu'}{\mu}.$$

Comparison. We compare Δ_{MB} and Δ_{MC}. We must evaluate these at the baseline equilibrium (D^B).

Baseline Equilibrium Values. We need the equilibrium value of M (and $\underline{\omega}$) in the thick market under D^B. At D^B, $M = X(0) = \underline{\omega}$. We use the definition of $\underline{\omega}$ (assuming μ is fixed at the equilibrium level for this calculation):

$$\underline{\omega} = \frac{\mu + (1-\mu)M}{1+\rho}.$$

Setting $M = \underline{\omega}$:

$$\begin{aligned}M(1+\rho) &= \mu + (1-\mu)M \\ M(1+\rho) &= \mu + M - \mu M \\ M\rho + \mu M &= \mu \\ M(\rho+\mu) &= \mu \implies M = \underline{\omega} = \frac{\mu}{\rho+\mu}.\end{aligned}$$

We also need $1-M$:

$$1 - M = 1 - \frac{\mu}{\rho+\mu} = \frac{\rho+\mu-\mu}{\rho+\mu} = \frac{\rho}{\rho+\mu}.$$

And the ratio $\underline{\omega}/\mu$:

$$\frac{\underline{\omega}}{\mu} = \frac{1}{\rho+\mu}.$$

Evaluating the Changes at Equilibrium. We substitute the equilibrium values into Δ_{MB} and Δ_{MC}.

Δ_{MB} evaluation:

$$\Delta_{MB} = \frac{1+2\rho}{1+\rho}\mu'\left(\frac{\rho}{\rho+\mu}\right).$$

Δ_{MC} evaluation:

$$\Delta_{MC} = \rho\mu'\left(\frac{1}{\rho+\mu}\right).$$

The Comparison Inequality. We check if $\Delta_{MB} > \Delta_{MC}$.

$$\frac{1+2\rho}{1+\rho}\mu'\left(\frac{\rho}{\rho+\mu}\right) \overset{?}{>} \rho\mu'\left(\frac{1}{\rho+\mu}\right).$$

We cancel the common positive factor $\mu'\frac{\rho}{\rho+\mu}$ from both sides:

$$\frac{1+2\rho}{1+\rho} \overset{?}{>} 1.$$

We analyze the term on the LHS:

$$\frac{1+2\rho}{1+\rho} = \frac{1+\rho+\rho}{1+\rho} = 1 + \frac{\rho}{1+\rho}.$$

Since the supplier discount rate $\rho > 0$, the term $\frac{\rho}{1+\rho}$ is strictly positive. Therefore:

$$\frac{1+2\rho}{1+\rho} > 1.$$

Conclusion on Mechanism 2. We have rigorously demonstrated that $\Delta_{MB} > \Delta_{MC}$. The increased benefit from the Desperation Effect (higher prices due to increased impatience) strictly dominates the increased cost from the Attrition Effect (lost customers due to increased departure).

This implies that the generalized profitability condition (Eq. (83)) is easier to satisfy when $\mu' > 0$ compared to the exogenous case. The critical threshold for μ below which scarcity is profitable is strictly higher than the exogenous benchmark $\rho/(1+2\rho)$.

K.2. Implications for the Main Model

The analysis in this appendix provides strong support for the robustness of the paper's main findings and justifies the modeling choices:

1. **Tractability Justification:** Endogenizing departure based on individual valuations (Mechanism 1) breaks the Uniform WTP property and renders the model intractable. This justifies the use of the exogenous assumption in the main paper to maintain analytical tractability in a complex equilibrium setting.
2. **Robustness of "Scarcity by Design":** A tractable model of endogenous departure based on aggregate market conditions (Mechanism 2) preserves the structural results. Crucially, it reveals that endogeneity reinforces the strategic feedback loop. The Desperation Effect dominates the Attrition Effect.
3. **Conservative Estimates:** Because endogenous departure strengthens the incentive to induce scarcity, the main model (with exogenous departure) provides a conservative *lower bound* on the potential revenue gains from "Scarcity by Design."