

Accuracy Discrimination in Personalized Product Recommendation

Wei Peiyu*

January 20, 2026

Abstract

I examine how a platform designs personalized product recommendations with discriminatory accuracy based on consumers' preference intensities to maximize the probability of trade. The platform can increase recommendation accuracy for a consumer at the cost of the consumer's privacy and higher endogenous prices set by third-party firms. When consumers' types are observable, the platform's optimal mechanism is perfectly misaligned with consumer preferences: it offers accurate recommendations only to "picky" enough consumers who benefit the least. In contrast, when the platform must screen consumers, it behaves as if maximizing ex post consumer surplus by allowing consumers to self-select any accuracy levels. The results highlight the importance of requiring platforms to obtain consumers' consent before providing personalized services, even when the platform's interests are partly aligned with those of consumers.

Keywords— Personalized Product Recommendations, Platform Design, Online Privacy, Accuracy Discrimination

*Faculty of Business and Economics, The University of Hong Kong. Email: pywei@connect.hku.hk

1 Introduction

Online shopping platforms collect vast amounts of data from consumers and have developed a somewhat uncanny ability to predict consumers' tastes for a variety of products, which has grown even further with recent rapid advancements in artificial intelligence technology. With this ability, platforms are able to offer personalized recommendations of products that match consumers' tastes. However, it may not always be profitable for platforms to fully exploit this ability for two reasons. First, consumers may feel uncomfortable knowing that their data are used for the platform's profit. Excessive personalization can thus raise privacy concerns and reduce user engagement. Second, if firms understand that their products are being recommended to consumers with higher valuations through personalized recommendations, they may increase prices in response to a flatter demand curve. This will deter purchases and again reduce consumer engagement. If platforms find it optimal to strategically restrict the extent of personalization, a natural question arises: Is it optimal to treat different consumers with different levels of recommendation accuracy? If so, which consumers should receive more or less accurate recommendations?

To address these questions, this paper studies how a platform optimally provides personalized product recommendations to consumers who are differentiated both horizontally and vertically. Horizontally, consumers have different tastes for the products available for sale on the platform, and the platform can match each consumer to her ideal product with varying degrees of accuracy. Vertically, consumers differ in the intensity of their preferences for the product under consideration, measured by the value of their outside options from not buying the product at all. In other words, consumers with higher outside options are "pickier", in the sense that they are willing to buy only if the match quality is very high. Consumers with lower outside options are "less picky" and more likely to purchase the recommended product with higher probability compared to the pickier ones *ex ante*. The objective is to determine whether pickier or less picky consumers are more profitable for the platform to serve with accurate recommendations.

To be specific, the platform serves as an intermediary between infinitely many *ex-ante* homogeneous firms with pricing power and a continuum of *ex-ante* heterogeneous consumers. The platform aims to promote trade and commits to a mechanism that decides the recommendation accuracy for each consumer, which is defined by the probability that

its recommended product is the consumer's most preferred one that brings the maximum valuation among all available products on the platform. Personalization also imposes privacy costs on consumers, the magnitude of which increases with the recommendation accuracy. Consumers either pay the privacy cost to receive recommendations or exit and take their outside options.

As a preliminary result, I show that the absence of firm commitment to a fixed price leads to a misalignment of interests between the platform and the firms, beyond merely the difference in how they benefit from trades. To see this, note that in equilibrium, a product delivering the highest possible valuation to all consumers will be purchased by any consumer who chooses to receive a recommendation. Although firms prefer the platform to provide perfectly accurate recommendations to all consumers, the platform always finds this suboptimal because consumers would anticipate that firms extract all the surplus from trades and therefore would not pay the privacy cost in the first place. Consequently, the platform must strategically allocate a limited amount of accuracy across heterogeneous consumers, creating an incentive for a discriminatory mechanism.

The main results show that the optimal allocation of accuracy to maximize trade probability varies significantly depending on whether consumers' outside options are observable to the platform. When outside options are observable (Section 4.3), I show that although the platform provides recommendations to all consumers who purchase with positive probability, only pickier consumers receive positive accuracy, while less picky ones receive random recommendations. Since pickier consumers are ex-ante less likely to purchase a given product, they have greater potential to promote trade from the platform's perspective. Under the assumption that the distribution of consumers' valuations is log-concave, this tendency to favor pickier consumers persists even when endogenous pricing effects are considered. A notable implication is that such an allocation of accuracy is actually extremely inefficient in terms of consumer surplus (or total welfare), given that the platform allocates recommendation accuracy to consumers who benefit from it the least.

When the platform cannot observe consumers' outside options (Section 4.4), it must screen consumers whose private "types" are the scales of their outside options, and the magnitudes of privacy costs imposed on consumers serve as a useful screening device. This screening problem departs from standard mechanism design problem since the allocation of accuracy

inevitably affects the endogenous price set by firms, which in turn influences consumers' incentives to report their types truthfully. To address this challenge, I demonstrate that, by the envelope theorem, the trade probability depends only on the payoff to the least picky type in the market as well as the equilibrium price. Consequently, the accuracy provided to other types is relevant to the platform's objective only through its impact on the endogenous price. In other words, for an allocation rule to be optimal, given the payoff to the least picky type, it must minimize the accuracy assigned to other types to induce the lowest possible price. The question is how this minimization complies with incentive compatibility. I show that the optimal mechanism features a decreasing allocation of accuracy with respect to type, and each consumer receives her most preferred level of accuracy, subject to an upper bound. This result not only differs from, but also reverses, the outcome under observable types.

Although endogenous pricing makes standard welfare analysis analytically intractable, comparing outcomes across informational regimes yields meaningful insights (Section 4.5) nonetheless yields meaningful insights. Interpreting recommendation accuracy as a limited resource and the recommendation mechanism as its allocation, the platform's optimal mechanism under observable types behaves as if it minimizes consumer surplus among allocations with the same total supply of accuracy, even though the platform's interests are partly aligned with those of consumers. By contrast, under private types, the optimal mechanism behaves as if it maximizes consumer surplus subject to the same upper bound on accuracy. Notably, these contrasting welfare implications are driven by the same underlying force: the platform's tendency to allocate more recommendation accuracy to consumers with higher outside option values. This comparison highlights the importance of limiting mandatory personalization and requiring consumer consent, even when the platform's objectives are partly aligned with those of consumers.

The analysis can be extended to settings in which the platform earns a commission per transaction that depends on price (Section 5.1) and in which the platform can provide a bundle of multiple recommendations to each consumer (Section 5.2). I show that, in both cases, the form of the optimal mechanism remains largely unchanged. I also extend the analysis to a Hotelling framework in which consumers and products differ in their locations (Section 5.3). While the platform continues to allocate recommendation accuracy primarily to pickier consumers, the welfare implications are reversed: in this environment, pickier

consumers benefit more from accurate recommendations than others. This result highlights the importance of distinguishing between different notions of consumer pickiness.

A key departure of this paper from existing research on discriminatory practices is that the platform always has an incentive to misallocate the “good”, that is, to provide as much accuracy as possible to those who need it the least. In standard mechanism design, the principal tends to allocate the good to the agent with the highest valuation to extract greater profit. Viewing recommendation accuracy as a good allocated by the platform, the platform’s incentive is precisely the opposite. This contrast arises from the assumption of a trade-maximizing platform, under which the platform’s notion of “efficiency” is opposite to that of consumers, despite their partially aligned interests.

The remainder of the paper is organized as follows. Section 2 reviews related literature. Section 3 introduces the model. Section 4 presents the optimal recommendation mechanisms under observable and private types, respectively, and corresponding welfare implications. Section 5 relaxes some model assumptions and examines the robustness of the main results. Section 6 concludes.

2 Related Literature

In the model, a key trade-off of the platform is between higher recommendation accuracy and lower price. Recent works find that personalization may change firms’ price decision. Although most of these works claim a positive effect, some differ from this paper in that they rely on price maker’s ability to price discriminate (Baik and Larson (2023), Belleflamme and Vergote (2016) and Braghieri (2019)). Ichihashi (2020) examines the consumer’s optimal disclosure of preference to the firm who can make use of the disclosure to offer personalized price. He proves that the firm always commits to no price discrimination in order to encourage more disclosure. In my paper, while the firm does not have commitment power, since the platform ensures the anonymity of consumer, the consumer’s concern is no longer that disclosure increases the price. Instead, consumer worries about privacy loss brought by targeting. The concern about price effect of disclosure is then shifted to the platform.

With the assumption that firms are not allowed to price discriminate, the present paper explains the positive impact of personalization on price by the logic that the firm knows that

his product is displayed to the consumer with higher valuation, which reduces the marginal loss in demand incurred by increasing the price. This logic is similar to that of [De Corniere and De Nijs \(2016\)](#), where winning an auction of advertising convey information advantages to the winner firm, and [Hagiu and Jullien \(2011\)](#), where personalization expands the expected demand from the consumer. In terms of modeling, the most similar one to the present paper is [Hidir and Vellodi \(2021\)](#). They also involve the same trade-off, but they assume that the outside option of consumer is not persistent and only available before entering, thus consumers have identical vertical preference and there is no issue of accuracy discrimination. Instead, they focus on market segmentation with respect to horizontal tastes.

The present paper also contributes to literature on how platforms rank products and consumer search. In [Teh and Wright \(2022\)](#), similar to the present paper, the platform knows how much a product matches a consumer but does not know the consumer's willingness to pay. A key deviation is that they assume the seller can also propose a commission fees in addition to the product price and the platform cannot commit to a recommendation rule in the sense that it depends on the sellers' proposal. The reason why steering increases the equilibrium price comes from the increase of commission fees in equilibrium. Their model does not involve discriminatory recommendation mechanism since consumers are ex-ante homogeneous. In [Rhodes and Zhou \(2024\)](#), ex-ante homogeneous consumers can choose to share their data with the platform at the cost of privacy. Sharing data has negative externality to the total surplus of consumers since it increases the equilibrium price. The present paper features a similar setting, but considers a more general form of privacy cost varying with the levels of personalization and introduces different types of consumers to explore discriminatory behaviors.

There is also literature that discusses the possibility "match discrimination". [Condorelli and Szentes \(2023\)](#) considers a platform which can match sellers with different qualities and consumers with different values. To maximize consumers surplus, it is optimal for the platform to match low value consumers with high quality sellers, which incentivizes sellers to offer lower price for consumers with higher values. In contrast, my model features ex-ante homogeneous sellers and excludes the possibility of price discrimination, which makes the trade-probability maximized mechanism differ from welfare-maximized one. This difference is not featured in their model.

3 The Model

3.1 Settings

The model involves a platform (it), infinite but countable number of firms (he) with horizontally differentiated products and a unit continuum measure of consumers (she). The platform serves as an intermediary to recommend a product from firms to consumers.

Consumers and Products A representative consumer has one unit demand for the product, which can be matched or not for her. All products are ex-ante homogeneous with production costs normalized to zero, and each product can be either a match or a generic product for the consumer ex post. The consumer's valuation of a product is denoted by v . If the product is a match, $v = 1$ for sure¹; otherwise, the product is generic and v is randomly distributed on $[0, 1]$ and drawn from a cumulative distribution function (CDF) $F(\cdot)$. Matches and valuations are independent and identically distributed (i.i.d.) across consumers and firms. Whether a product has a match to a product is only observable by the platform to reflect the platform's power of analyzing consumer data.

The consumers are ex ante heterogeneous in their persistent outside options. I assume that firms can never observe the consumer types to model the platform's commitment to keep the data from third-party firms. The scale of a consumer's outside option is denoted by $\theta \in [0, 1]$, which is drawn from a CDF $G(\cdot)$. I assume that the CDFs $F(\cdot)$ and $G(\cdot)$ have full support and admit differentiable probability density functions (PDFs) $f(\cdot)$ and $g(\cdot)$.

Personalized Product Recommendation Each consumer can choose whether to receive a personalized recommendation from the platform. If she chooses to receive it, given her taste, the platform commits to a recommendation rule which is specified by the *recommendation accuracy* defined as the probability that a matched product is recommended, or

$$\lambda \in [0, 1].$$

The usage of consumer data in selection of recommended product intrudes the consumers' privacy, and therefore incurs privacy costs for the consumers, which is denoted by $c \in [0, +\infty)$

¹A more natural and commonly used interpretation of the assumption that a matched product yields valuation equal to 1 is to define the matched product as the product that delivers the highest valuation among all available products. In this case, the probability that the highest valuation is weakly less than 1 is given by $\lim_{N \rightarrow \infty} F(v)^N = 0$. Hence, the highest valuation equals 1 almost surely.

and paid by the consumers before they seeing the recommended product. To capture the endogenous trade-off between prediction accuracy and consumer privacy, I assume that any improvement in the platform's predictive accuracy leads to an increase in privacy loss for the consumer. Specifically, for any given level of accuracy λ , the privacy costs c must satisfy

$$c = \gamma \cdot e(\lambda), \quad (1)$$

where $\gamma > 0$ are constants and $e(\cdot) \geq 0$ is a strictly increasing, strictly convex and differentiable function with $e(0) = e'(0) = 0$. In other words, both the level of privacy cost and its marginal increase weakly grow with recommendation accuracy.

Platform Design The platform commits to a *mechanism* allocating recommendation accuracy λ and the corresponding scale of privacy cost c to each consumer. By the Revelation Principle, regardless of whether targeted strategy is available, it is without loss to focus on direct mechanisms, where consumers truly report their types. Thereby the platform's mechanism of recommendation can be represented by

$$\lambda \equiv (\lambda_\theta)_{\theta \in [0,1]}.$$

A representative firm n offers price P_n for his product after observing the platform's mechanism. I assume that firms do not observe the platform's identification of product-consumer matches, and thus cannot engage in price discrimination. The firm n 's pricing strategies is a function of platform's mechanism:

$$p_n : \lambda \longmapsto P_n.$$

Since I focus on symmetric equilibrium, it is without loss to drop the index n and assume that all firms use the same strategy p .

Payoffs A consumer derives utility equal to the value of the product she purchases minus its price. A firm's expected profit equals its revenue from selling its product. The platform's profit comes from a constant commission fee per transaction and thus maximizes the trade probability.

Timeline I summarize the game in chronological order in the following:

1. The platform announces a recommendation mechanism λ which is observed by all players.
2. Each consumer chooses whether to receive a recommendation of product. A representative consumer with type θ and report θ' pays the privacy cost equal to $c_{\theta'} = \gamma \cdot e(\lambda_{\theta'})$ if she chooses to receive the recommendation.
3. The platform selects the recommended product k for the consumer according to the recommendation accuracy $\lambda_{\theta'}$. Each firm n decides his price P_n simultaneously.
4. The consumer observes the valuation of the recommended product v_k and its price P_k . If the consumer accepts the recommendation and make a purchase, she get $v_k - P_k$; otherwise she gets θ .

3.2 Interpretation of Model Assumptions

Here I further illustrate and justify some assumptions I made in the model:

Personalization and Privacy Loss Platforms such as Google, Amazon, and Meta develop algorithms that learn consumers' tastes from their purchase records and browsing histories. In this paper, the platform's ability to correctly identify a consumer's matched product is summarized by a single parameter, λ . This formulation is based on a simplifying assumption that the platform either learns nothing about a consumer's taste or learns it perfectly. I further assume that the platform can improve its predictive accuracy at the cost of consumer privacy. The parameter γ captures the efficiency with which the algorithm extracts information about consumers' tastes from collected data: a lower value of γ corresponds to a more efficient algorithm, allowing the platform to infer consumer preferences more precisely for a given level of data collection and privacy loss. In practice, platforms often obtain such data through mandatory information disclosure during user registration or usage, which serves as a primary input for learning consumers' tastes.

Third-Party Firms' Pricing Power The model assumes that firms retain pricing power that is outside the platform's control. While individual-level price discrimination is infeasible because consumers are typically anonymous to firms, price discrimination at the platform level may still arise. This is because Most Favored Nation (MFN) provisions, which require firms to offer prices no higher than those on other distribution channels, are not universally

available to platforms. Indeed, MFN clauses imposed by platforms such as Amazon have been subject to investigation by European Union antitrust authorities.² As a result, an overly accurate recommendation algorithm may place all consumers on the platform at a disadvantage by inducing higher equilibrium prices, thereby reducing the probability of trade. I also assume that the platform cannot condition its recommendations on prices. On e-commerce platforms such as Amazon, organic recommendations are primarily driven by relevance and user behavior, whereas price considerations are largely confined to search filters or sponsored listings.

Conflict of Interests between Platform and Firms One might conjecture that the only source of misalignment between the platform and firms arises from differences in their payoff functions. In fact, there is an additional channel. Because each firm knows that its price will not be observed by consumers until after they have incurred the privacy cost required to receive a recommendation, the firm fails to internalize the possibility that consumers may opt out of the recommendation. Consequently, even when the platform and firms share the same payoff function, they may disagree on the optimal allocation of recommendation accuracy (see Section 5.1). The underlying assumption is that firms lack commitment power to promise not to charge a high price when provided with a highly accurate recommendation.³ As a result, the price that is optimal from the platform’s perspective is strictly lower than the price preferred by firms. This distinction reflects an important feature of real-world platforms: concerns about user traffic and engagement are typically borne by the platform, whereas third-party firms are largely insulated from these considerations.

4 Analysis

4.1 The Platform’s Problem

In this subsection I define the platform’s problem formally. To examine whether a representative consumer is willing to receive the recommendation from platform, in the following I derive the net benefit for a consumer with type θ from reporting θ' and

²See [Baker and Scott Morton \(2017\)](#) for an overview of enforcement actions against platform MFNs, and [Peitz and Sobolev \(2024\)](#) for a discussion of their effects.

³[Hidir and Vellodi \(2021\)](#) adopts a similar assumption, arguing that platforms may refrain from fully segmenting the market for related reasons.

receiving the recommendation given the platform's mechanism λ and the common price P set symmetrically by firms:

$$\begin{aligned} U(\lambda_{\theta'}, P; \theta) &= \lambda_{\theta'} \cdot (1 - P) + (1 - \lambda_{\theta'}) \cdot \int_{P+\theta}^1 (v - P) dF(v) + \theta \cdot (1 - \lambda_{\theta'}) \cdot F(P + \theta) - \gamma \cdot e(\lambda_{\theta'}) - \theta \\ &= \int_{P+\theta}^1 [1 - (1 - \lambda_{\theta'}) F(v)] dv - \gamma \cdot e(\lambda_{\theta'}). \end{aligned} \quad (2)$$

If the consumer chooses to receive the recommendation, she will always buy the recommended product and gets $1 - P$ if it is a match⁴, which is the best one she expects to get from a recommendation. This happens with probability λ_{θ} by the definition of recommendation accuracy. With probability $1 - \lambda_{\theta}$, the product is not a match and the consumer purchases if and only if $v \geq P + \theta$ and gets expected surplus equal to $\int_{P+\theta}^1 (v - P) dF(v)$. By (2), the consumer chooses to receive the recommendation if and only if

$$U(\lambda_{\theta}, P; \theta) \geq 0 \implies \int_{P+\theta}^1 [1 - (1 - \lambda_{\theta'}) F(v)] dv \geq \gamma \cdot e(\lambda_{\theta'}).$$

The firm's payoff from offering P given the recommendation mechanism λ and the set of consumers that choose to receive the recommendation $\Theta \subset [0, 1]$ is

$$V(P; \lambda, \Theta) \propto P \cdot \int_{\Theta} [1 - (1 - \lambda_{\theta}) F(P + \theta)] dG(\theta). \quad (3)$$

The platform's profit is proportional to the trade probability, or

$$W(\lambda, \Theta; P) = \int_{\Theta} [1 - (1 - \lambda_{\theta}) F(P + \theta)] dG(\theta).$$

Note that although the platform only decides the mechanism, it is without loss to assume that it can also select which consumer to enter the market and which to continue searching beyond the recommended product, as long as the selected consumer finds it profitable to do so under the mechanism λ ; and can also decide the equilibrium price, as long as all firms find it optimal to do so. Throughout the paper, when I refer to the platform's *strategy*, I mean the

⁴Although allowing consumers with outside options $\theta > 1 - P$ to enter the platform without ever making a purchase does not violate individual rationality, I instead assume that only consumers with a strictly positive probability of purchase choose to enter. This can be ensured by assuming $e(0) = \varepsilon \downarrow 0$; that is, the entry cost of receiving a recommendation is always strictly positive but arbitrarily small.

triple $\{\lambda, \Theta, P\}$. To summarize, the platform's problem is to find a strategy to maximize its profit:

$$\max_{\substack{\lambda \in [0,1]^{[0,1]}, \\ \Theta \subset [0,1] \\ P \geq 0}} W(\lambda, \Theta; P)$$

subject to the consumers' individual rationality constraints

$$U(\lambda_\theta, P; \theta) \geq 0 \text{ for any } \theta \in \Theta, \quad (\text{IR})$$

and the constraint to ensure equilibrium price equal to P :

$$P = p(\lambda, \Theta) \in \min\{\arg \max_{P'} V(P'; \lambda, \Theta)\}. \quad (\text{EP})$$

For evaluation of optimal mechanism under private types, we also need the consumer's incentive compatibility constraint to make sure all entering consumers chooses the assigned outcomes. Specifically, for any $\theta, \theta' \in \Theta$, we have

$$U(\lambda_{\theta'}, P; \theta) \leq U(\lambda_\theta, P; \theta). \quad (\text{IC})$$

4.2 The Platform's Trade-off

To understand the platform's trade-off in determining discriminatory accuracy levels in recommended slots, it is crucial to see how firms react in pricing to the platform's strategy. First I claim that the first-order condition described in the following is a necessary condition of (EP):

Lemma 1. *If a price P satisfies (EP) given platform's policy λ and set of consumers Θ who choose to receive the recommendation, then it must be*

$$\frac{\partial}{\partial P'} V(P'; \lambda, \Theta, P) \Big|_{P'=P} = \int_{\Theta} [1 - (1 - \lambda_\theta) (F(P + \theta) + P f(P + \theta))] dG(\theta) = 0. \quad (\text{FOC-S})$$

With Lemma 1, I show that higher recommendation accuracy always increases the equilibrium price and it is always suboptimal for the platform to assign the highest accuracy, or $\lambda_\theta = 1$ to all the consumers:

Lemma 2. Let $\{\lambda, \Theta, P\}$ and $\{\tilde{\lambda}, \Theta, \tilde{P}\}$ be two different platform's strategies satisfying (EP). If there exists a set of types $\theta \in \Theta$ of positive measure such that $\lambda_\theta > \tilde{\lambda}_\theta$, and $\lambda_\theta = \tilde{\lambda}_\theta$ elsewhere, then it must be

$$P > \tilde{P}.$$

Furthermore, if λ features $\lambda_\theta = 1$ for any $\theta \in \Theta$, then there exists a unique equilibrium in which all firms charge price $P = 1$ and the platform's profit is equal to 0.

Just like the platform, the firm also refrains from too high a price for fear of losing too many consumers. When the recommendation accuracy λ is set to 1, this fear disappears because consumer is always served with a perfectly matched product and will buy regardless of the price upon entering. Consequently, there is no horizontal differentiation in consumers' valuations, and the firms just offer price equal to 1. However, anticipating that the equilibrium price is 1 the consumers know that they never get strictly positive payoff from transactions. Given non-zero privacy costs is required to enter the platform, the consumers choose not to enter in the first place, leaving the platform with zero profit.

Although firm always benefits from high accuracy, Lemma 2 indicates that this need not be the case from the platform and the consumer's views. The increase of price not only directly decreases the consumer's willingness to buy, but also decreases the number of entering consumers by making them less optimistic that they can gain from the recommendation of product, for which Lemma 2 provides an extreme example. This logic is similar to [Hidir and Vellodi \(2021\)](#), where the platform controlled market segmentation effectively increases the firm's market power and brings misaligned interests. In practice, while firms only care about their profits given the demand curve in a market, the platforms have to further worry about the level of customer engagement. In the following analysis we will see how this misalignment of interests makes the accuracy of recommendation become a scarce resource that the platform needs to allocate it carefully.

In the rest of paper, I assume the regularity condition for the firm's price strategy defined by the following:

Condition 1 (Regularity). $\Phi(x|P) \equiv F(P+x) + P \cdot f(P+x)$ is strictly increasing in x on $[0, 1-P]$.

Basically, $\Phi(x|P) - 1$ is the product of virtual valuation $P - \frac{1-F}{f}$ and PDF f . Therefore the effect of Condition 1 is similar to regularity condition in other mechanism design literature.

4.3 Optimal Mechanism under Observable Types

In this section, I discuss the platform's optimal recommendation mechanism when the platform is able to observe the consumers' outside options and enforce the allocation of the outcomes personalization λ_θ freely. Denote the optimal strategy of platform under observable types as $\{\hat{\lambda}^{\text{ot}}, \hat{\Theta}^{\text{ot}}, \hat{P}^{\text{ot}}\}$ with "ot" standing for "observable types".

As a preliminary, consider a relaxed problem, where the equilibrium price $\hat{P}$ is taken as given. Furthermore, γ is set to 0 and there is no privacy cost for consumers, therefore we can drop the consumers' individual rationality constraints (IR) and assume that any consumers with outside options $\theta \leq 1 - P$ would choose to receive the recommendation. Then it remains to determine the allocation of accuracy λ_θ . The platform aims to maximize

$$\int_0^{1-P} [1 - (1 - \lambda_\theta) F(P + \theta)] dG(\theta).$$

subject to (FOC-S)

$$\int_0^{1-P} (1 - \lambda_\theta) [F(P + \theta) + P f(P + \theta)] dG(\theta) = \int_0^{1-P} dG(\theta)$$

It is obvious that this problem is linear with respect to $\hat{\lambda}_\theta$, the optimal mechanism depends on the ratio of marginal benefit to the shadow price to give more accuracy to every type θ , or

$$s(\theta; P) \equiv \frac{F(P + \theta)}{F(P + \theta) + P f(P + \theta)}. \quad (4)$$

This is why imposing the following condition on the distribution of generic product values is helpful:

Condition 2 (Strict Log-concavity). *The PDF $f(\cdot)$ satisfies that $\frac{f(v)}{f'(v)}$ is strictly increasing in v .*

Note that the above condition implies that $F(\cdot)$ is log-concave. Then for $\theta > \theta'$ we have $s(\theta; P) > s(\theta'; P)$, which indicates that it is more cost-efficient for the platform to give higher accuracy λ to type θ than to θ' . As a result, for any given targeted price, the solution to the

relaxed problem is to give as much accuracy as possible to the consumers with higher outside options:

Lemma 3. *Given a targeted equilibrium price P , if $\gamma = 0$, there exists $\theta^+ \in [0, 1 - P]$ such that the platform's optimal allocation of accuracy under observable types satisfies*

$$\hat{\lambda}_\theta^0 = \mathbf{1}_{\theta \geq \theta^+}.$$




Figure 1: First-Best Accuracy Allocation for the Platform

The platform finds it optimal to provide accurate recommendation to consumers with high outside options (or the “pickier” ones), while leaving those with low outside options with no personalization. To illustrate the intuition behind, recall that the platform knows if it predicts the consumer’s taste correctly, the consumer will buy the recommended product with probability 1⁵. As a result, the marginal benefit from giving higher level of accuracy to a consumer with type θ is exactly the probability that the consumer would not buy the product if the recommendation is completely random, or $F(P + \theta)$, which is strictly increasing in θ . A practical explanation is that pickier consumers can be regarded as those with “larger

⁵This result follows from the assumption that a consumer’s valuation for a matched product is deterministically equal to 1. Alternatively, one can verify that the same result continues to hold if the model is modified so that the valuation equals $v \sim F^{\text{match}}(\cdot)$ with probability λ , where $F^{\text{match}}(\cdot)$ is a cumulative distribution function on $[0, 1]$ such that $\frac{F(v) - F^{\text{match}}(v)}{f(v) - f^{\text{match}}(v)}$ is strictly increasing in v .

potentials”, which is profitable to serve with very accurate product recommendations. In contrast, since less picky consumers already tend to make a purchase with higher probability, it is a waste to provide them with the scarce accuracy to further increase their chances to buy. This benefit from giving pickier consumers more accurate recommendation is then proven to outweigh the cost caused by giving the firm a stronger incentive to increase the price thanks to the log-concavity of $F(\cdot)$ ⁶.

The optimal allocation rule $\hat{\lambda}_\theta^0$ is depicted in Figure 1, which illustrates the meaning of “accuracy discrimination”: The platform tries its best to match the picky consumers’ tastes while completely ignores the eager consumers’. It seems to imply that pickier consumers gain an advantage from this discrimination. However, this is not necessarily the case since accurate recommendations are costly to consumers when $\gamma > 0$, and a consumer’s optimal levels of accuracy are strictly decreasing in her outside option θ . Technically,

$$\frac{dU(L; P, \theta)}{dL} = \frac{d}{dL} \left[\int_{P+\theta}^1 1 - (1 - L) F(v) dv - \gamma \cdot e(L) \right] = \int_{P+\theta}^1 F(v) dv - \gamma \cdot e'(L), \quad (5)$$

where a consumer’s marginal benefit from higher accuracy L shrinks as her outside option θ increases. Therefore, although pickier consumers are more likely to get their matched products, the privacy costs may be too high for them to bear compared to the benefits of accurate recommendations.

When the participation constraints (IR) are considered, some consumers find it unprofitable to enter the platform with accuracy level $\hat{\lambda}_\theta^0$. This is because although those very picky consumers get the most accurate recommendation, their gains from the trade are very low relative to the privacy loss caused by personalization. To address this problem and describe the platform’s first-best policy considering consumers’ participation constraint, let $\Lambda(\theta; P)$ be a function representing the level of accuracy at which a consumer with effective outside option value t is indifferent to entering the platform, given that privacy cost $\gamma \cdot e(\Lambda(\theta; P))$ equals the lowest one to satisfy the feasibility constraint and the equilibrium

⁶If the firm increases price, on one hand, he earns more whenever trade happens, with which his profits increases by $1 - (1 - \lambda) F$; on the other hand, the trade probability decreases marginally by $(1 - \lambda) f$, and thus he makes $P \cdot (1 - \lambda) f$ less profit. The first term is exactly the same as type θ ’s probability to buy. That is, both are equally sensitive to the change of λ (with sensitivity equal to F). Therefore, the first term can be canceled out by the platform’s incentive to increase the trade probability, and there remains the second term with sensitivity to λ equal to f . In the end, to compare the benefit-cost ratios between giving accuracy to different types is equivalent to comparing $F(P + \theta) / f(P + \theta)$, which is strictly increasing in θ by log-concavity.

price equals P :

$$\int_{P+\theta}^1 1 - (1 - \Lambda(\theta; P)) F(v) dv = \gamma \cdot e(\Lambda(\theta; P)).$$

The uniqueness of the value of $\Lambda(\cdot; P)$ is guaranteed by the convexity of $e(\cdot)$. One can verify that with the convexity of $e(\cdot)$, $\Lambda(\cdot; P)$ represents the *highest* accuracy that a consumer is willing to afford to receive the recommendation. Now it is ready to present the first main result:

Theorem 1. *Under observable types, the platform's optimal strategy $\{\hat{\lambda}^{ot}, \hat{\Theta}^{ot}, \hat{P}^{ot}\}$ satisfies the following properties:*

(a) *The set of consumers who choose to receive the recommendation is*

$$\hat{\Theta}^{ot} = [0, 1 - \hat{P}^{ot}].$$

(b) *The accuracy allocation function $\hat{\lambda}_{\theta}^{ot}$ is an increasing two-step function, with a higher step at 1 and a lower step at 0, followed by a strictly decreasing tail. Or, there exists $\theta_{ot}^{\dagger}, \theta_{ot}^{\ddagger} \in \hat{\Theta}^{ot}$ with $\theta_{ot}^{\dagger} \leq \theta_{ot}^{\ddagger}$ and $\Lambda(\theta_{ot}^{\ddagger}; \hat{P}^{ot}) = 1$ such that*

$$\hat{\lambda}_{\theta}^{ot} = \begin{cases} \mathbf{1}_{\theta \geq \theta_{ot}^{\dagger}} & , \theta \in [0, \theta_{ot}^{\ddagger}); \\ \Lambda(\theta; \hat{P}^{ot}), & \theta \in [\theta_{ot}^{\ddagger}, 1 - \hat{P}^{ot}); \\ 0 & , \theta \in [1 - \hat{P}^{ot}, 1]. \end{cases} \quad (6)$$

Theorem 1 presents the optimal strategy of the platform, which is depicted in Figure 2 where the gray area stands for the set of consumers that choose to receive the recommendation. First, the platform is optimal to engage as many consumers as possible. This is intuitive since more consumers always contribute to the platform's profit, and it is always possible to let a consumer receive personalized recommendation by setting the accuracy equal to 0 so the privacy cost paid by the consumer to enter is also 0. Secondly, the optimal allocation of accuracy $\hat{\lambda}_{\theta}^{ot}$ is still a step function for $\theta \leq \theta_{ot}^{\ddagger}$, which shares the same logic behind the optimality of $\hat{\lambda}_{\theta}^0$. When $\theta > \theta_{ot}^{\ddagger}$, setting $\hat{\lambda}_{\theta}^{ot} = 1$ definitely violates the constraint of individual rationality by the definition of $\theta_{ot}^{\ddagger}$. Instead, the accuracy is set to the highest one that guarantees the consumer's entry, which is exactly the definition of $\Lambda(\cdot; \hat{P}^{ot})$. This is optimal considering the individual rationality constraint (IR) because further

increasing the accuracy for the consumer will always force her to exit and the platform gets no profit.




Figure 2: Optimal Policy under Observable Types

Note that the optimal mechanism under observable types, or $\hat{\lambda}_\theta^{\text{ot}}$, is actually the counterpart of the classic first-degree price discrimination in the context of accuracy allocation. To see this, recall that a consumer with type θ has a payoff given by $\int_{p+\theta}^1 1 - (1 - \lambda_\theta) \cdot F(v) dv - c_\theta$. If we interpret the accuracy λ_θ as the “probability of getting the good (recommendation accuracy)” and the privacy cost c_θ as a “transfer” to the platform, then the platform’s optimal strategy is to sell the good only to consumers with higher enough valuations and to require a transfer equal to each consumer’s valuation. This contrasts to $\hat{\lambda}_\theta^{\text{ot}}$, which gives the “good” only to consumers with lower enough “valuations”. The key difference arises from how the platform benefits from trades. In the price discrimination framework, the platform directly extracts consumer surplus through transfers and therefore favors consumers that have higher profit to extract, while in the accuracy discrimination case, the platform benefits from transaction and therefore favors those has lower “valuations” but greater potential to promote trade.

Although accuracy discrimination does not aim to extract all the surplus from consumers, it is worth noting that the platform’s optimal policy under observable types, $\hat{\lambda}_\theta^{\text{ot}}$, still allocates accuracy in a way that is significantly misaligned with the preferences of consumers: although

less picky consumers benefit more from accurate recommendations than pickier ones do, it only gives positive accuracy to the latter. This is due to the misalignment of interests between the platform and the consumers. It is obvious that the platform fails to take the consumers' privacy cost into consideration, but more importantly, although both parties benefit from transactions, they benefit in different ways. The platform simply aims to promote trade regardless of who made the purchase, while consumers care how much they can get from the recommendation, which is determined by their outside options. Later we will see how this misalignment of interests is resolved by not allowing the platform to observe consumers' types.

4.4 Optimal Mechanism under Private Types

In this subsection, I derive the platform's optimal mechanism when consumers' outside options are private information. Denote the optimal strategy of platform under private types as $\{\hat{\lambda}^{\text{pt}}, \hat{\Theta}^{\text{pt}}, \hat{P}^{\text{pt}}\}$ with "pt" standing for "private types". Unlike the observable types case, the platform must now screen consumers and design the mechanism subject to incentive compatibility constraints (IC).

First I describe the behavior of accuracy allocation λ_θ and set of consumers willing to receive the recommendation Θ if only considering incentive compatibility. Now altering the recommendation accuracy designed for one consumer may incentivize her or other consumers to deviate. Besides, the platform is no longer able to decide who to receive the product recommendation: a consumer opts in the personalized recommendation if and only if she finds receiving the recommendation makes her better off. In the following I show that, any mechanism that features an allocation rule similar to $\hat{\lambda}_\theta^{\text{ot}}$ (that is, pickier consumers get strictly higher level of accuracy) always violates incentive compatibility constraint; while the personalized recommendation set is still downward-closed, just like that under observable types:

Lemma 4. *If a mechanism λ satisfies incentive compatibility (IC) and individual rationality (IR), then*

- (a) *the allocation rule of accuracy λ_θ is weakly decreasing in θ ;*
- (b) *there exists $\hat{\theta} \in \Theta$ such that the set of consumers who choose to enter is given by $\Theta = [0, \hat{\theta}]$.*

Lemma 4 is intuitive: for (a), less picky consumers always benefit more than those pickier ones from the trade by the expression of $U(\lambda, P; \theta)$, thus tends to choose higher accuracy to multiply this benefit; for (b), now that the platform cannot observe consumers' types, a consumer with lower type θ is always free to pretend to be a higher type θ' . If she does so, her payoff is strictly higher than that of a truth-telling consumer with θ' . By incentive compatibility, her equilibrium payoff is weakly higher than misreporting one, which in turn means that lower types always have strictly higher payoff than the higher types. That is, whenever consumer with θ' enters the market, θ does so, too. This implies that it is always without loss to use $\hat{\theta}$ to represent the participation set as $\hat{\Theta} = [0, \hat{\theta}]$. Note that by (a), it is never incentive compatible to replicate a similar allocation rule to $\hat{\lambda}^{\text{ot}}$ with private types: higher accuracy is only possible to be chosen by less picky consumers.

Although the platform's problem of allocating recommendation accuracy shares a similar form to a classical auction design, as illustrated in Section 4.3, a key difficulty arises from endogenous pricing. In particular, it is *impossible* to analyze consumers' incentive compatibility constraints in isolation, because firms' equilibrium prices respond to the allocation of accuracy. These prices, in turn, shape consumers' incentives to choose their recommendation accuracy. To address this difficulty, I introduce a method to simplify the platform's problem before deriving the optimal mechanism. By the Envelope Principle, we always have

$$\begin{aligned}
\frac{dU(\lambda_\theta, P; \theta)}{d\theta} &= \frac{\partial U(\lambda_{\theta'}, P; \theta)}{\partial \theta} \Big|_{\theta'=\theta} = \left\{ \frac{d}{d\theta} \int_{P+\lceil \hat{\theta} \rceil}^1 (v - P - \lceil \hat{\theta} \rceil) d[(1 - \lambda_{\theta'}) F(v) + \lambda_{\theta'} \mathbb{1}_{v \geq 1}(v)] \right\} \Big|_{\theta'=\theta} \\
&\quad \begin{array}{l} \text{i) decrease} \\ \text{trade prob.} \end{array} \quad \begin{array}{l} \text{ii) decrease} \\ \text{net surplus} \end{array} \\
&= -0 - [1 - (1 - \lambda_\theta) F(P + \theta)] \\
&= -[1 - (1 - \lambda_\theta) F(P + \theta)],
\end{aligned} \tag{7}$$

that is, under any incentive compatible mechanism, the marginal change of consumer's value function with respect to her type is exactly the opposite of the trade probability. This is not surprising. Since the consumer's payoff is defined as the net gain from the recommendation relative to outside option, the marginal change of consumer's payoff with respect to type comes from i) the decrease of trade probability and ii) the decrease of net trade surplus (represented by the first and the second θ in (7) respectively). The former only affects

consumer when she barely purchases and thus decreases her payoff by 0, while the latter matters only when trade happens and the consumer gives up her outside option and therefore decreases her payoff by an amount equal to the trade probability. Then by the Envelope Theorem, the marginal decrease of a consumer's value function with respect to her type is equal to the trade probability. This connection makes it possible to reconsider the platform's optimization problem with

$$\begin{aligned}
W(\lambda, [0, \hat{\theta}], P) &= \int_0^{\hat{\theta}} [1 - (1 - \lambda_{\theta}) F(P + \theta)] dG(\theta) \\
&= - \int_0^{\hat{\theta}} \frac{dU(\lambda_{\theta}, P; \theta)}{d\theta} \cdot g(\theta) d\theta \\
&= g(0) \cdot (U(\lambda_{\theta=0}, P; 0) - U(\lambda_{\hat{\theta}}, P; \hat{\theta})) + \int_0^{\hat{\theta}} (U(\lambda_{\theta}, P; \theta) - U(\lambda_{\hat{\theta}}, P; \hat{\theta})) \cdot g'(\theta) d\theta \\
&= \int_0^{\hat{\theta}} U(\lambda_{\theta}, P; \theta) \cdot g'(\theta) d\theta + g(0) \cdot U(\lambda_{\theta=0}, P; \theta = 0),
\end{aligned}$$

where $[0, \hat{\theta}]$ is the set of consumers who choose to accept the platform's recommendations and $U(\lambda_{\hat{\theta}}, P; \hat{\theta}) = 0$ by individual rationality. Specifically, if the distribution of consumers' outside option values is uniform, that is, $\theta \sim U[0, 1]$ and $G(\theta) = \theta$ and $g(\theta) = 1$ for any $\theta \in [0, 1]$, we have

$$W(\lambda, [0, \hat{\theta}], P) = U(\lambda_{\theta=0}, P; \theta = 0). \quad (8)$$

In other words, if $\lambda_{\theta=0}$ is given, to maximize the platform's payoff, it remains to minimize the equilibrium price $P = p(\lambda, \Theta)$ while ensuring that λ satisfies (IC) and (IR) for any $\theta \in [0, \hat{\theta}]$.

Now I start to derive the optimal mechanism $\hat{\lambda}^{\text{pt}}$. Consider the following auxiliary optimization problem given a pair $(\bar{L}, P)$:

$$\begin{aligned}
&\min_{\substack{\lambda \in [0,1]^{[0,1]} \\ \hat{\theta} \in [0,1]}} \int_0^{\hat{\theta}} \{1 - (1 - \lambda_{\theta}) \cdot [F(P + \theta) + P f(P + \theta)]\} d\theta \quad (\text{AP}) \\
&\text{s.t. } \int_{P+\theta}^1 [1 - (1 - \lambda_{\theta}) \cdot F(v)] dv - \gamma \cdot e(\lambda_{\theta}) = \int_{\theta}^{\hat{\theta}} [1 - (1 - \lambda_{\theta'}) \cdot F(P + \theta')] d\theta' \\
&\quad \text{for any } \theta \in [0, \hat{\theta}], \quad (\text{AP-ICa}) \\
&\int_P^1 [1 - (1 - \bar{L}) F(v)] dv - \gamma \cdot e(\bar{L}) = \int_0^{\hat{\theta}} [1 - (1 - \lambda_{\theta'}) \cdot F(P + \theta')] d\theta'. \quad (\text{AP-ICb})
\end{aligned}$$

Basically, (AP) provides necessary conditions for an allocation rule λ to be optimal when the

accuracy level for consumer with type 0 is $\lambda_0 = \bar{L}$ and equilibrium price is P : it must minimize the LHS of (FOC-S) so that the endogenous price is also minimized; besides, it must satisfy two constraints of incentive compatibility including i) (AP-ICa) restricting for type $\theta \in (0, \hat{\theta}]$, the LHS of which represents the actual utility of type θ given that the accuracy level equal to λ_θ and the RHS of which represents the utility level required by the Envelope theorem (7); ii) (AP-ICb) restricting for type $\theta = 0$, the logic of which is similar to (AP-ICa). Otherwise, if an allocation rule is not the solution to (AP), there exists a profitable deviation for the platform to decrease the price without decreasing the trade probability. Therefore, any allocation mechanism $\hat{\lambda}^{pt}$ is optimal only if it is a solution of the auxiliary problem (AP) with $\bar{L} = \hat{\lambda}_0^{pt}$ and $P = \hat{P}^{pt}$, which I described in the following:

Theorem 2. *If $G(\cdot) = U[0,1]$, then the platform's optimal strategy under private types, $\{\hat{\lambda}^{pt}, \hat{\Theta}^{pt}, \hat{P}^{pt}\}$, satisfies the following properties:*

- (a) $\hat{\lambda}_\theta^{pt}$ is continuous in θ . Moreover, there exists $\bar{L}^{pt} \in [0,1]$ and $\hat{P}^{pt} \in (0,1)$ such that $\lambda = \hat{\lambda}^{pt}$ and $\hat{\theta} = 1 - \hat{P}^{pt}$ solve the problem (AP) with $\bar{L} = \bar{L}^{pt}$ and $P = \hat{P}^{pt}$.
- (b) $\hat{\Theta}^{pt} = [0, 1 - \hat{P}^{pt}]$ and there exists $\theta^* \in \hat{\Theta}^{pt}$ such that $\hat{\lambda}_\theta^{pt}$ is strictly decreasing and satisfies

$$\gamma \cdot e'(\hat{\lambda}_\theta^{pt}) = \int_{\hat{P}^{pt} + \theta}^1 F(v) dv \quad (9)$$

on $[\theta^*, 1 - \hat{P}^{pt}]$ and equals $\bar{L}^{pt}$ on $[0, \theta^*)$.

Now I explain the structure of $\hat{\lambda}_\theta^{pt}$, which consists of a constant segments and a strictly decreasing segment satisfying (9), separated by θ^* , as described in Theorem 2. To see what this structure represents, suppose that λ_θ is differentiable with respect to θ . Take the derivative of binding (AP-ICa) with respect to θ , we have

$$\frac{d\lambda_\theta}{d\theta} \cdot \left[\int_{P+\theta}^1 F(v) dv - \gamma \cdot e'(\lambda_\theta) \right] = 0.$$

That is, the optimal accuracy allocation should either i) remain constant across neighboring types, expressed as $\frac{d\lambda_\theta}{d\theta} = 0$, or ii) ensure that the consumer is indifferent to a marginal change in accuracy, formalized by

$$\frac{d}{dL} U(L, P; \theta) = \frac{d}{dL} \left[\int_{P+\theta}^1 1 - (1-L) F(v) dv - \gamma \cdot e(L) \right] = \int_{P+\theta}^1 F(v) dv - \gamma \cdot e'(L) = 0.$$





Figure 3: Optimal Mechanism under Private Types

In either case, the consumer has no incentive to misreport their type as a neighboring one, thereby satisfying incentive compatibility. In other words, if we only focus on continuous accuracy allocation function λ_θ , the only incentive compatible possibility that is optimal for the platform is the one described in Theorem 2. I will explain why λ_θ is optimal to be continuous later.

When $\theta \in [0, \theta^*)$, that is, when $\hat{\lambda}_\theta^{\text{pt}} = \bar{L}^{\text{pt}}$ is a constant, determining the optimal choice of accuracy level is less straightforward compared to the case where $\theta \in [\theta^*, 1 - \hat{p}^{\text{pt}}]$. It is hard to determine whether $\bar{L}^{\text{pt}} = 1$ or $\bar{L}^{\text{pt}} < 1$, which depends on the optimal choice of the equilibrium price, or by (FOC-S), $\int_{\hat{\theta}^{\text{pt}}} (1 - \hat{\lambda}_\theta^{\text{pt}}) \cdot \Phi(\theta | \hat{p}^{\text{pt}}) d\theta$. When it is relatively high, which means the platform is willing to have higher accuracy of recommendation to increase trade possibility at the cost of higher price, we have $\bar{L}^{\text{pt}} = 1$; when it is relatively low, the platform's trade-off reverses and $\bar{L}^{\text{pt}} < 1$ just like what Figure 3 shows.

The accuracy allocation under private types $\hat{\lambda}^{\text{pt}}$ described in Figure 3 contrasts significantly with its counterpart under observable types $\hat{\lambda}_\theta^{\text{ot}}$ described in Figure 2 in terms of its monotonic behavior and continuity, in the sense that part of $\hat{\lambda}_\theta^{\text{ot}}$ has a strictly increasing part with a single discontinuity. While it is evident from incentive compatibility that $\hat{\lambda}_\theta^{\text{pt}}$ is decreasing, its continuity is less straightforward and carries a non-trivial implication: intermediate options are always available for consumers to choose, which benefits those with medium outside

options. Note that both (IC) and (IR) do not forbid discontinuity of λ_θ . For example, it is incentive compatible to have λ_θ that is discontinuous at θ_d and constant on $[\theta_l, \theta_d)$ and $(\theta_d, \theta_h]$, as presented by the blue line in the top graph of Figure 4.



It is always profitable for the platform to “spread” the decrease of λ_θ equal to $L_h - L_l$ at θ_d (the vertical blue line) to $\tilde{\theta}_1$ with decrease $L_h - \tilde{L}$ and to $\tilde{\theta}_2$ with decrease $\tilde{L} - L_l$ (the red dotted lines). This is because that $\int_{\theta}^{\theta_d} F(P + \theta') d\theta'$ is more concave than $\int_{\theta}^{\theta_d} [F(P + \theta') + P f(P + \theta')] d\theta'$, given that $F(\cdot)$ is log-concave. When the spreading keeps $(L_h - \tilde{L}) \cdot \int_{\tilde{\theta}_1}^{\theta_d} F(P + \theta) d\theta + (\tilde{L} - L_l) \cdot \int_{\tilde{\theta}_2}^{\theta_d} F(P + \theta) d\theta$ unchanged (the two big dots at the same height), it effectively reduces the LHS of (FOC-S), which decreases $(L_h - \tilde{L}) \cdot \int_{\tilde{\theta}_1}^{\theta_d} [F(P + \theta) + P f(P + \theta)] d\theta + (\tilde{L} - L_l) \cdot \int_{\tilde{\theta}_2}^{\theta_d} [F(P + \theta) + P f(P + \theta)] d\theta$ (the two big dots of which the left one is lower).

Figure 4: A Profitable Deviation from Discontinuous λ_θ

To understand why a discontinuous allocation of accuracy is suboptimal, consider the

following deviating allocation $\tilde{\lambda}$, with $\tilde{\lambda}_\theta = \tilde{L} \in (L_l, L_h)$ for any $\theta \in [\tilde{\theta}_1, \tilde{\theta}_2)$ and remains unchanged elsewhere. It is illustrated by the red dotted line in the top graph of Figure 4. To satisfy incentive compatibility, $\tilde{\theta}_1$ and $\tilde{\theta}_2$ are uniquely pinned down by

$$(L_h - \tilde{L}) \cdot \int_{P+\tilde{\theta}_1}^1 F(v) dv = \gamma \cdot e(L_h) - \gamma \cdot e(\tilde{L})$$

and

$$(\tilde{L} - L_l) \cdot \int_{P+\tilde{\theta}_2}^1 F(v) dv = \gamma \cdot e(\tilde{L}) - \gamma \cdot e(L_l)$$

using the expression of consumer's payoff (2). Similarly, the original policy λ also satisfies incentive compatibility at θ_d , therefore we have

$$(L_h - L_l) \cdot \int_{P+\theta_d}^1 F(v) dv = \gamma \cdot e(L_h) - \gamma \cdot e(L_l).$$

Combine the three equations, we have

$$\begin{aligned} L_h \cdot \int_{\tilde{\theta}_1}^{\theta_d} F(P + \theta) d\theta + L_l \cdot \int_{\theta_d}^{\tilde{\theta}_2} F(P + \theta) d\theta &= \tilde{L} \cdot \int_{\tilde{\theta}_1}^{\tilde{\theta}_2} F(P + \theta) d\theta \\ \implies (L_h - \tilde{L}) \cdot \int_{\tilde{\theta}_1}^{\theta_d} F(P + \theta) d\theta &= (\tilde{L} - L_l) \cdot \int_{\theta_d}^{\tilde{\theta}_2} F(P + \theta) d\theta. \end{aligned}$$

To prove the profitability of this deviation, note that it does not change the accuracy level of type $\theta = 0$, therefore by (8), it is sufficient to prove that the deviation decreases the LHS of (FOC-S), or $\int_{\tilde{\theta}_1}^{\tilde{\theta}_2} [1 - (1 - \lambda_\theta) F(P + \theta)] d\theta$. That is,

$$\begin{aligned} &L_h \cdot \int_{\tilde{\theta}_1}^{\theta_d} (F(P + \theta) + P f(P + \theta)) d\theta + L_l \cdot \int_{\theta_d}^{\tilde{\theta}_2} (F(P + \theta) + P f(P + \theta)) d\theta \\ &< \tilde{L} \cdot \int_{\tilde{\theta}_1}^{\tilde{\theta}_2} (F(P + \theta) + P f(P + \theta)) d\theta \\ \iff &(L_h - \tilde{L}) \cdot \int_{\tilde{\theta}_1}^{\theta_d} (F(P + \theta) + P f(P + \theta)) d\theta < (\tilde{L} - L_l) \cdot \int_{\theta_d}^{\tilde{\theta}_2} (F(P + \theta) + P f(P + \theta)) d\theta, \end{aligned}$$

which is obvious because by the log-concavity of $F(\cdot)$ we have

$$\frac{\int_{\tilde{\theta}_1}^{\theta_d} F(P + \theta) d\theta}{\int_{\tilde{\theta}_1}^{\theta_d} (F(P + \theta) + P f(P + \theta)) d\theta} < \frac{\int_{\theta_d}^{\tilde{\theta}_2} F(P + \theta) d\theta}{\int_{\theta_d}^{\tilde{\theta}_2} (F(P + \theta) + P f(P + \theta)) d\theta}. \quad (10)$$

Intuitively, the bottom panel of Figure 4 provides a more direct illustration of the profitability of the deviation. Instead of directly choosing the level of accuracy for each type, it is useful to consider the optimal allocation of the *decrease* in λ_θ . Under this interpretation, a discontinuity of λ_θ at θ_d can be viewed as concentrating the entire decrease in λ_θ over the interval $[\theta_l, \theta_h]$ at a single point θ_d . The profitability of spreading this decrease becomes apparent once we note that $\int_{\theta}^{\theta_d} [F(P + \theta') + P f(P + \theta')] d\theta'$ is more concave than $\int_{\theta}^{\theta_d} F(P + \theta') d\theta'$. As a result, spreading the decrease in λ_θ yields a profitable deviation for the platform. Iterating this spreading argument leads to a continuous allocation of accuracy, as characterized by $\hat{\lambda}_\theta^{\text{pt}}$ ⁷. Further spreading would violate the incentive compatibility constraint and is therefore infeasible.

An important implication of the optimal mechanism under private types is that there is no distortion in recommendation accuracy for consumers with $\theta > \theta^*$, as formalized by (9). This contrasts with the classic second-degree price discrimination result in [Mussa and Rosen \(1978\)](#), where quality is distorted downward for all but the highest type due to the principal's rent-extraction motive under monopoly. As explained above, the absence of distortion here is driven by inequality (10). One may note that it reflects the same underlying force as the property that the benefit–cost ratio of marginally increasing recommendation accuracy (4) is strictly increasing in the consumer's type. Interestingly, this very same force leads the platform to generate extreme distortion in recommendation accuracy when types are observable, as discussed in Section 4.3. Thus, the introduction of incentive compatibility *reverses* the effect of the platform's underlying motive.

4.5 Welfare Implication

When it comes to welfare analysis, it may seem natural to conclude that consumers are better off under private types than under observable types, given that the platform must provide information rents to consumers due to information asymmetry. In fact, this conclusion is ambiguous and difficult to establish, since it is unclear whether the equilibrium price under private types, $\hat{P}^{\text{pt}}$, is higher or lower than that under observable types, $\hat{P}^{\text{ot}}$. More importantly, as mentioned earlier, the platform does not extract profits directly from consumers and shares

⁷The formal proof of continuity of $\hat{\lambda}_\theta^{\text{pt}}$ also rules out additional forms of discontinuity beyond those illustrated in Figure 4. These cases are addressed in Step 2 of the proof of Theorem 2 in Appendix A.

some common interests with them. Therefore, private information does not necessarily make consumers better off.

In light of this challenge, this section instead discusses the welfare implications by viewing the recommendation mechanism as an allocation of recommendation accuracy, taking the equilibrium price as given, and asking whether this allocation is efficient from an ex post perspective. The following proposition addresses this question by identifying the opposite welfare implications of the optimal mechanisms in the two cases.

Proposition 1. For $(\hat{\lambda}^{ot}, \hat{\Theta}^{ot}, \hat{P}^{ot})$ defined in Theorem 1 and $(\hat{\lambda}^{pt}, \hat{\Theta}^{pt}, \hat{P}^{pt})$ in Theorem 2 we have

(a)

$$\begin{aligned} \hat{\lambda}^{ot} = & \arg \min_{\lambda \in [0,1]^{[0,1]}} \int_{\hat{\Theta}^{ot}} U(\lambda_{\theta}, \hat{P}^{ot}; \theta) dG(\theta); \\ \text{s.t. } & \int_{\hat{\Theta}^{ot}} (\lambda - \hat{\lambda}_{\theta}^{ot}) dG(\theta) = 0; \\ & \forall \theta \in \hat{\Theta}^{ot}, U(\lambda_{\theta}, \hat{P}^{ot}; \theta) \geq 0 \end{aligned}$$

(b)

$$\hat{\lambda}^{pt} = \arg \max_{\lambda \in [0, \bar{L}^{pt}]^{[0,1]}} \int_{\hat{\Theta}^{pt}} U(\lambda_{\theta}, \hat{P}^{pt}; \theta) dG(\theta),$$

In the case where consumers' types are observable to the platform, although the platform does not benefit from directly extracting surplus from consumers, it nevertheless behaves as if it were playing a zero-sum game with them. That is, it behaves as if it aims to minimize consumer surplus subject to a given total supply of recommendation accuracy and the individual rationality constraint. This is because the optimal mechanism, $\hat{\lambda}^{ot}$, allocates the highest accuracy to those who value it the least and the lowest accuracy to those who value it the most. One can verify that the consumer's payoff function, $U(\lambda, P; \theta)$, is submodular with respect to λ and θ , which explains why the surplus-minimizing allocation should be positive assortative.

Conversely, when consumers' types are private information, the optimal mechanism, $\hat{\lambda}^{pt}$, behaves as if the platform's interests are fully aligned with those of consumers, subject to a common upper bound on accuracy. This property is evident for $\theta \in [\theta^*, \hat{\theta}^{pt}]$, since the first-order condition for maximizing a consumer's payoff is always satisfied, that is,

$$\int_{P+\theta}^1 F(v) dv = \gamma \cdot e'(\hat{\lambda}_{\theta}^{pt}),$$

and each consumer receives her most desirable level of accuracy. However, for lower-type consumers, $\theta \in [0, \theta^*)$, it is possible that they are served with excessively low accuracy if

$$\int_{P+\theta}^1 F(v) dv > \gamma \cdot e'(\bar{L}^{\text{pt}}),$$

so that their most desirable level of accuracy cannot be achieved. This explains why the maximization is taken over all accuracy allocations with a common upper bound equal to $\bar{L}^{\text{pt}}$.

The contrasting properties of allocation rules $\hat{\lambda}^{\text{ot}}$ and $\hat{\lambda}^{\text{pt}}$ cannot be simply solely by the standard notion of information rents accruing to consumers. In classical auction and mechanism design settings, the principal benefits from higher trade prices, which increases profits but necessarily harms agents. In the present model, the principal's (i.e., the platform) interest is partially aligned with the agents' (i.e., the consumers), in the sense that both desire a lower price. When consumer types are observable, the platform can fully exploit consumers' propensity to trade, which directs the accuracy to the picky consumers. By contrast, when types are private and the platform has to leave the choices to the consumers, the same underlying incentive instead aligns the platform's objectives with consumer welfare, as illustrated at the end of Section 4.4.

Proposition 1 highlights the importance of requiring consumers' consent for personalization, even when price discrimination is unavailable and the platform's interests are partly aligned with those of consumers. Although such a requirement does not necessarily increase consumer surplus or overall welfare, it ensures the "most" envy-free allocation of recommendation accuracy from an ex post perspective. While price discrimination has widely been criticized by researchers and prohibited on many e-commerce platforms, this regulation does not fully eliminate the inefficiency arising from allowing platforms to unilaterally determine the extent of personalization. To conclude, it appears fair and welfare-improving for regulations to require online platforms to obtain consumers' consent before providing personalized services.

5 Extensions

I relax several model assumptions and examine whether the main results change.

5.1 Price-Dependent Commission Fees

In the baseline model, I assume that the platform cares only about the probability of trade. I now consider a more general environment in which the platform earns a commission fee that depends on the price charged by the firm. Its payoff is given by

$$Y(\lambda, \hat{\Theta}, P) = y(P) \cdot \int_{\hat{\Theta}} [1 - (1 - \lambda_{\theta}) F(P + \theta)] d\theta,$$

where $y(\cdot)$ is a positive and increasing function. In contrast to the original setting, the platform may now induce a higher price than in the benchmark even though this reduces the trade probability.

One may think that if $y(P) \propto P$, the problem becomes trivial because that the platform's interest and the firm's are completely aligned, eliminating any trade-off between accurate recommendation and low price. That would suggest the optimal mechanism sets $\lambda_{\theta} = 1$ for any $\theta \in \Theta$. However, as discussed in Section 4.2, the misalignment of interests still persists: the firm fails to fully internalize the consequences of setting a too-high price. Even when the platform and the firm share exactly the same payoff function, the optimal mechanism remains non-trivial:

Proposition 2. *When the platform maximizes $Y(\lambda, \hat{\Theta}, P)$, under the assumptions imposed in Theorem 1 and 2, we have*

- (a) *Under observable types, the optimal mechanism $\hat{\lambda}^{yot}$ features an increasing step function followed by a decreasing tail, similar in structure to $\hat{\lambda}^{ot}$ in (6) but with different cutoff values $\theta^{\dagger}$ and $\theta^{\ddagger}$;*
- (b) *Under private types, an incentive compatible mechanism $\hat{\lambda}^{ypt}$ is optimal if and only if the associated equilibrium price, pinned down by $\hat{P}^{ypt} = p(\hat{\lambda}^{ypt}, [0, 1 - \hat{P}^{ypt}])$, solves the following equation with respect to $\hat{P}$:*

$$\left. \frac{dy(P)}{dP} \right|_{P=\hat{P}} \cdot \left[\int_{\hat{P}}^1 1 - (1 - \hat{\lambda}_0^{ypt}) F(v) dv - \gamma \cdot e(\hat{\lambda}_0^{ypt}) \right] = y(\hat{P}) \cdot \left[1 - (1 - \hat{\lambda}_0^{ypt}) F(\hat{P}) \right]. \quad (11)$$

Furthermore, if $\hat{P}^{ypt} > \hat{P}^{pt}$, the optimal mechanism $\hat{\lambda}^{ypt}$ is not unique; if instead $\hat{P}^{ypt} = \hat{P}^{pt}$, we have $\hat{\lambda}^{ypt} = \hat{\lambda}^{pt}$.

When the platform benefits from a higher price, it is natural to expect that the equilibrium

price under the platform's optimal mechanism is weakly higher than in the baseline model, in which the platform benefits only from trade itself. Conditional on a given price, the platform's objective reduces to maximizing the probability of trade, which coincides with the baseline problem. As a result, the structure of the optimal policy remains very similar to that of the benchmark. The key difference lies in the platform's optimal choice of the equilibrium price $\hat{p}^{ypt}$. If $\hat{p}^{ypt} = \hat{p}^{pt}$, which can occur when the marginal benefit from a higher price, $\frac{dM}{dP}/M$, is sufficiently small, the platform must implement the same mechanism $\hat{\lambda}^{pt}$ as in the baseline model in order to keep prices low. By contrast, if $\hat{p}^{ypt} > \hat{p}^{pt}$, the platform is no longer constrained to mechanisms that induce the lowest possible price. In this case, by (8), the trade probability equals the payoff of the lowest type, $\int_{\hat{p}^{ypt}}^1 [1 - (1 - \hat{\lambda}_0^{ypt})F(v)] dv - \gamma \cdot e(\hat{\lambda}_0^{ypt})$, and the platform can choose any mechanism that optimally trades off this value against the induced equilibrium price. Overall, when the platform charges a price-dependent commission fee, the optimal mechanism retains key structural properties of the baseline model, highlighting the robustness of the main results.

5.2 Multiple Recommendations

Consider an environment in which the platform can provide a bundle of multiple recommendations to each consumer, while each consumer still has unit demand. The consumer observes the values and prices of all recommended products simultaneously and then decides whether to purchase one of them or to exit. Specifically, let $n > 1$ denote the number of recommendations. Since the consumer has unit demand and always would purchase a perfect match, it is without loss of generality to redefine the recommendation accuracy λ as the probability that at least one of the n recommended products perfectly matches the consumer's taste. I show that this change of settings does not significantly change main results:

Proposition 3. *If the number of recommended products satisfies $n < \infty$, then the optimal recommendation mechanisms share the same structure as those described in Theorems 1 and 2, differing only in parameter values.*

Compared with the baseline model, allowing for multiple recommendations is nontrivial, as it introduces price competition among the recommended firms. Nevertheless, the

platform's trade-off between higher recommendation accuracy and lower equilibrium prices persists, and the benefit–cost ratio of marginally increasing recommendation accuracy remains increasing in a consumer's outside option value. In the Appendix, I show that the platform's problem retains essentially the same structure, and therefore the main results of the baseline model continue to hold even when price competition is taken into account.

5.3 Hotelling Model

In the baseline model, the only relevant ex-ante heterogeneity of consumers is the difference in their outside options. One may wonder whether the result of accuracy misallocation survives under more common forms of consumer preference heterogeneity, such as a Hotelling framework in which both consumers and products are differentiated by their locations on a line. In this setting, consumers locating closer to the end are “pickier,” in the sense that their expected valuation for a randomly recommended product is lower.

Specifically, consider consumers' preferences described by θ distributed on $[-1/2, 1/2]$ and characters of sellers' products are denoted by $b \in [-1/2, 1/2]$. The products distribution is denoted by $F(\cdot)$ with $f(-b) = f(b)$ to reflect the symmetry of the distribution function. A consumer located at θ has valuation for a product located at b equal to

$$v(\theta, b) = 1 - |\theta - b|.$$

A consumer has type with larger absolute value $|\theta|$ should be one that is “pickier”, in the sense that her expected valuation for a random product is lower. This is corresponding to that with larger outside option value in the original model. Furthermore, a “matched” product for type θ now is defined as the product with location $b = \theta$ and gives the consumer with the highest possible valuation $v(\theta, \theta) = 1$.

With these changes of settings and others unchanged, the platform still tends to allocate as much recommendation accuracy as possible to pickier consumers since otherwise they won't make purchases with high possibility, similar to the baseline model. However, these pickier consumers now derive a higher marginal benefit from more accurate recommendations,

because random recommendations are especially costly for them. The results are summarized in the following⁸:

Condition 3. *The CDF $F(\cdot)$ satisfies that*

$$\frac{1 - F(t + \Delta) + F(t - \Delta)}{f(t + \Delta) + f(t - \Delta)}$$

is increasing in $|t|$ for any $t \in [-1/2, 1/2]$ and $\Delta \in (0, 1)$ with $f(t + \Delta) + f(t - \Delta) > 0$ ⁹.

Proposition 4. *With the Hotelling model described above and under Condition 3, if $\gamma = 0$, there exists $t^\dagger > 0$ such that the platform's optimal allocation of recommendation accuracy under observable types, denoted by $\hat{\lambda}^H$, satisfies*

$$\hat{\lambda}_\theta^H = \mathbf{1}_{|\theta| \geq t^\dagger}.$$

Moreover, the consumer's marginal benefit from receiving higher recommendation accuracy is increasing in $|\theta|$ for a given price P , that is,

$$\frac{\partial^2}{\partial |\theta| \partial L} U(L, P; \theta) > 0.$$

The key difference between the two settings lies in how pickiness is defined. With heterogeneity in outside option values, pickier consumers are those who are intrinsically more indifferent across products. While more accurate recommendations increase the likelihood of trade for these consumers more, they derive less utility. By contrast, under heterogeneity in locations, pickier consumers value match quality just as much as others but face a more restricted set of suitable matches. Consequently, they benefit more from increased recommendation accuracy, aligning their interests with those of the platform. This distinction in the definition of pickiness also corresponds to different real-world environments. In the

⁸Under the Hotelling model with private types, the platform's problem can no longer be simplified in the same way as in the baseline model. As a result, the presence of endogenous pricing makes it difficult to characterize the optimal recommendation mechanism. For this reason, I restrict attention here to the case of observable types.

⁹This condition is obviously satisfied for $t + \Delta < 0$ or $t - \Delta > 0$ by any single-peak distribution functions, that is, $f(b) = f(-b)$ is increasing in b for $b < 0$. For other cases, this condition can still be satisfied by some common single-peak distributions. For example, let $f(b) = 4b + 2$ for $b \in [-1/2, 0]$ and $f(b) = -4b + 2$ for $b \in (0, 1/2]$. Also we have $F(b) = 2b^2 + 2b + \frac{1}{2}$ for $b \in [-1/2, 0]$ and $F(b) = -2b^2 + 2b + \frac{1}{2}$. Then when $t - \Delta < 0 < t + \Delta$ the condition becomes

$$\frac{1 + 4(t^2 + \Delta^2) - 4\Delta}{4 - 4\Delta}$$

increasing in $|t|$, which is obviously satisfied.

former case, consumers differ primarily in their willingness to pay or their attractiveness of outside options, such as casual shoppers versus loyal users. In the latter case, consumers share similar valuations for match quality but differ in how difficult it is to find suitable options, for example, shoppers with very specific size, style, or compatibility constraints compared with more flexible shoppers.

6 Conclusions

I conclude the paper by discussing the platform's role in the model. On the one hand, since firms cannot commit to refraining from setting high prices in response to highly accurate recommendations, the platform acts as a commitment device that ensures reasonable pricing and sustained consumer engagement. This mechanism benefits not only consumers but also firms by enabling a market that would otherwise fail to exist. On the other hand, because the platform's payoff directly depends on the probability of trade, it has an inherent incentive to allocate as much recommendation accuracy as possible to pickier consumers—an outcome that ultimately harms consumers when their types are observable. By prohibiting mandatory personalization and allowing consumers to freely choose their preferred level of personalization, the platform's interests become largely more aligned with those of consumers. The results thus provide an explanation for why a seemingly neutral platform may end up with implementing an extremely inefficient recommendation mechanism, as well as a regulatory insight into how such inefficiencies can be mitigated.

References

- Baik, Alicia, Simon Anderson and Nathan Larson (2023) "Price discrimination in the information age: Prices, poaching, and privacy with personalized targeted discounts," *The Review of Economic Studies*, 90 (5), 2085–2115.
- Baker, Jonathan B and Fiona Scott Morton (2017) "Antitrust enforcement against platform MFNs," *Yale LJ*, 127, 2176.
- Belleflamme, Paul and Wouter Vergote (2016) "Monopoly price discrimination and privacy: The hidden cost of hiding," *Economics Letters*, 149, 141–144.

- Braghieri, Luca (2019) "Targeted advertising and price discrimination in intermediated online markets," *Available at SSRN 3072692*.
- Condorelli, Daniele and Balazs Szentes (2023) "Buyer-Optimal Platform Design," *The RAND Journal of Economics*.
- De Corniere, Alexandre and Romain De Nijs (2016) "Online advertising and privacy," *The RAND Journal of Economics*, 47 (1), 48–72.
- Hagiu, Andrei and Bruno Jullien (2011) "Why do intermediaries divert search?" *The RAND Journal of Economics*, 42 (2), 337–362.
- Hidir, Sinem and Nikhil Vellodi (2021) "Privacy, personalization, and price discrimination," *Journal of the European Economic Association*, 19 (2), 1342–1363.
- Ichihashi, Shota (2020) "Online privacy and information disclosure by consumers," *American Economic Review*, 110 (2), 569–595.
- Mussa, Michael and Sherwin Rosen (1978) "Monopoly and product quality," *Journal of Economic theory*, 18 (2), 301–317.
- Peitz, Martin and Anton Sobolev (2024) "Product recommendations and price parity clauses," Technical report, University of Bonn and University of Mannheim, Germany.
- Rhodes, Andrew and Jidong Zhou (2024) "Personalization and privacy choice," *Available at SSRN 4795118*.
- Teh, Tat-How and Julian Wright (2022) "Intermediation and steering: Competition in prices and commissions," *American Economic Journal: Microeconomics*, 14 (2), 281–321.

Appendix A Omitted Proofs

Proof of Lemma 1

Denote the set of consumers types that choose to receive the recommendation and purchase the recommended product with positive probability as

$$\overline{\Theta}(P'|\lambda, \Theta, P) \equiv \{\theta \in \Theta | P' + \theta \leq 1\}$$

given that the recommended firm's price of recommended product is P' , and the platform's strategy is $\{\lambda, \Theta, P\}$. First I prove that (FOC-S) is a necessary condition for P to be optimal for the firm. Let

$$V_0(P'_1, P'_2|\lambda, \Theta, P) = P'_1 \cdot \int_{\overline{\Theta}(P'_2|\lambda, \Theta, P)} 1 - (1 - \lambda_\theta) F(P'_1 + \theta) dG(\theta).$$

Obviously, we have $V(P; \lambda, \Theta) = V_0(P', P'|\lambda, \Theta, P)$. Since the LHS of (FOC-S) is $\frac{\partial V_0(P'_1, P'_2|\lambda, \Theta, P)}{\partial P'_1} \Big|_{P'_1=P', P'_2=P}$, to prove that the FOC is expressed by (FOC-S), I need only prove that

$$\frac{\partial V_0(P'_1, P'_2|\lambda, \Theta, P)}{\partial P'_2} \Big|_{P'_1=P', P'_2=P} = 0.$$

It is sufficient to prove that for any $\theta \in \overline{\Theta}(P|\lambda, \Theta, P)$, we have

$$P' + \theta = 1 \implies \lambda_\theta = 0.$$

Suppose not, that is, $\lambda_\theta > 0$. Note that we have

$$\begin{aligned} U(\lambda_\theta; P, \theta) &= \int_{P'+\theta}^1 1 - (1 - \lambda_\theta) F(v) dv - \gamma \cdot e(\lambda_\theta) \\ &= 0 - \gamma \cdot e(\lambda_\theta) < 0, \end{aligned}$$

which contradicts to that θ chooses to receive the recommendation and completes the proof. $\square$

Proof of Lemma 2

First I prove that higher accuracy increases the equilibrium price. Note that the LHS of (FOC-S) is strictly increasing in λ_θ , we always have $\frac{\partial}{\partial P'} V(P'; \tilde{\lambda}, \Theta) \Big|_{P'=P} > 0$. It then remains to prove that for any $P' < \tilde{P}$, we have

$$V(P'; \lambda, \Theta, \tilde{P}) - V(P'; \lambda, \Theta, \tilde{P}) < V(P'; \tilde{\lambda}, \Theta, \tilde{P}) - V(P'; \tilde{\lambda}, \Theta, \tilde{P})$$

where the RHS is definitely negative by definition of equilibrium price. It always holds and completes the proof.

Second I prove that when $\lambda_\theta = 1$ for any $\theta \in [0, 1]$, then there is a unique equilibrium which features $P = 1$ and $\Theta = \emptyset$, and therefore platform's profit $W(\lambda, \Theta, P) = 0$. Suppose not, that is, $P \in (0, 1 - \gamma \cdot e(1))$ and there exists $\theta \in [0, 1]$ such that $U(\lambda_\theta, P; \theta) > 0$. Note that since $v = 1$ with probability 1, any consumers who choose to receive the recommendations will always buy the product. Therefore, the consumer's profit becomes

$$V(P'; \lambda, \Theta, P) = P' \cdot \int_{\Theta} dG(\theta)$$

and the FOC becomes

$$\frac{\partial}{\partial P'} V(P'; \lambda, \Theta, P) \Big|_{P'=P} = \int_{\Theta} dG(\theta) = 0,$$

which is a contradiction to $\Theta \neq \emptyset$ and completes the proof. $\square$

Proof of Lemma 3

Suppose not, that is, there exists sets of types $\Theta_l, \Theta_h \subsetneq \hat{\Theta}^{\text{ot}}$, such that i) for any $\theta_l \in \Theta_l$ and $\theta_h \in \Theta_h$, we have $\theta_h > \theta_l$; ii) for any $\theta \in \Theta_l$, $\hat{\lambda}_\theta > 0$; iii) for any $\theta \in \Theta_h$, either $U(\hat{\lambda}_\theta, P; \theta) > 0$ and $\hat{\lambda}_\theta < 1$.

Consider the following new policy $\tilde{\lambda}$: i) for $\theta \in \Theta_l$, let $1 - \tilde{\lambda}_\theta = (1 + \varepsilon_l) \cdot (1 - \hat{\lambda}_\theta)$; ii) for $\theta \in \Theta_h$, let $1 - \tilde{\lambda}_\theta = (1 - \varepsilon_h) \cdot (1 - \hat{\lambda}_\theta)$ where $\varepsilon_l, \varepsilon_h \in (0, 1)$ satisfy

$$\varepsilon_l \cdot \int_{\Theta_l} (1 - \hat{\lambda}_\theta) \cdot \Phi(\theta|\hat{P}) dG(\theta) = \varepsilon_h \cdot \int_{\Theta_h} (1 - \hat{\lambda}_\theta) \cdot \Phi(\theta|\hat{P}) dG(\theta)$$

and sufficiently small so that $U(\hat{\lambda}_\theta, P; \theta) \geq 0$. Note that if the set of consumers who choose

to receive the recommendation does not change, $\tilde{\lambda}$ does not change FOC since the LHS of (FOC-S) remains the same:

$$\begin{aligned}
& \int_{\Theta_l} (1 - \tilde{\lambda}_\theta) \cdot \Phi(\theta|\hat{P}) d\theta + \int_{\Theta_h} (1 - \tilde{\lambda}_\theta) \cdot \Phi(\theta|\hat{P}) d\theta \\
&= \int_{\Theta_l} (1 - \hat{\lambda}_\theta) \cdot \Phi(\theta|\hat{P}) d\theta + \int_{\Theta_h} (1 - \hat{\lambda}_\theta) \cdot \Phi(\theta|\hat{P}) d\theta \\
&\quad + \varepsilon_l \cdot \int_{\Theta_l} (1 - \hat{\lambda}_\theta) \cdot \Phi(\theta|\hat{P}) d\theta - \varepsilon_h \cdot \int_{\Theta_h} (1 - \hat{\lambda}_\theta) \cdot \Phi(\theta|\hat{P}) d\theta \\
&= \int_{\Theta_l} (1 - \hat{\lambda}_\theta) \cdot \Phi(\theta|\hat{P}) d\theta + \int_{\Theta_h} (1 - \hat{\lambda}_\theta) \cdot \Phi(\theta|\hat{P}) d\theta.
\end{aligned}$$

Then the profitability is given by i) the trade probability increases by

$$\begin{aligned}
& W(\tilde{\lambda}, \hat{\Theta}, \hat{P}) - W(\hat{\lambda}, \hat{\Theta}, \hat{P}) \\
&= - \int_{\Theta_l} (1 - \tilde{\lambda}_\theta) \cdot F(\hat{P} + \theta) dG(\theta) - \int_{\Theta_h} (1 - \tilde{\lambda}_\theta) \cdot F(\hat{P} + \theta) dG(\theta) \\
&\quad + \int_{\Theta_l} (1 - \hat{\lambda}_\theta) \cdot F(\hat{P} + \theta) dG(\theta) + \int_{\Theta_h} (1 - \hat{\lambda}_\theta) \cdot F(\hat{P} + \theta) dG(\theta) \\
&= - \varepsilon_l \cdot \int_{\Theta_l} (1 - \hat{\lambda}_\theta) \cdot F(\hat{P} + \theta) dG(\theta) + \varepsilon_h \cdot \int_{\Theta_h} (1 - \hat{\lambda}_\theta) \cdot F(\hat{P} + \theta) dG(\theta) \\
&\propto - \frac{\int_{\Theta_l} (1 - \hat{\lambda}_\theta) \cdot F(\hat{P} + \theta) dG(\theta)}{\int_{\Theta_l} (1 - \hat{\lambda}_\theta) \cdot \Phi(\theta|\hat{P}) dG(\theta)} + \frac{\int_{\Theta_h} (1 - \hat{\lambda}_\theta) \cdot F(\hat{P} + \theta) dG(\theta)}{\int_{\Theta_h} (1 - \hat{\lambda}_\theta) \cdot \Phi(\theta|\hat{P}) dG(\theta)} > 0
\end{aligned}$$

where the inequality comes from the log-concavity of $F(\cdot)$.

Now I prove that $\hat{P}$ remains the platform's optimal price. to this end, I prove that for any $P' > \hat{P}$, we have $V(P'; \tilde{\lambda}, \bar{\Theta}(P'|\tilde{\lambda}, \hat{\Theta}, \hat{P})) \leq V(\hat{P}; \tilde{\lambda}, \bar{\Theta}(\hat{P}|\tilde{\lambda}, \hat{\Theta}, \hat{P}) = \hat{\Theta})$. Alternatively, I prove that

$$\begin{aligned}
& \left. \frac{d}{d\varepsilon_l} (V(P'; \tilde{\lambda}, \bar{\Theta}(P'|\tilde{\lambda}, \hat{\Theta}, \hat{P})) - V(\hat{P}; \tilde{\lambda}, \hat{\Theta})) \right|_{\varepsilon_l=0} \\
&\leq \left. \frac{d}{d\varepsilon_l} (V(P'; \tilde{\lambda}, \Theta_l \cup \Theta_h) - V(\hat{P}; \tilde{\lambda}, \Theta_l \cup \Theta_h)) \right|_{\varepsilon_l=0} \\
&= \left. \frac{d}{d\varepsilon_l} \left\{ \int_{\hat{P}}^{P'} \int_{\Theta_l \cup \Theta_h} [1 - (1 - \tilde{\lambda}_\theta) \Phi(\theta; P'')] dG(\theta) dP'' \right\} \right|_{\varepsilon=0} \\
&= - \int_{\hat{P}}^{P'} \left[\int_{\Theta_l} (1 - \hat{\lambda}_\theta^{\text{ot}}) \cdot \Phi(\theta; P'') dG(\theta) - \left. \frac{d\varepsilon_h}{d\varepsilon_l} \right|_{\varepsilon_l=0} \cdot \int_{\Theta_h} (1 - \hat{\lambda}_\theta^{\text{ot}}) \cdot \Phi(\theta; P'') dG(\theta) \right] dP'' \leq 0.
\end{aligned}$$

It is sufficient to prove that $\frac{\Phi(\theta'; P'')}{\Phi(\theta'; \hat{P})} \geq \frac{\Phi(\theta; P'')}{\Phi(\theta; \hat{P})}$ for any $\theta' \leq \theta$, which is obvious given that

$$\begin{aligned} \frac{d}{d\theta} \frac{\Phi(\theta; P'')}{\Phi(\theta; \hat{P})} &= \frac{d}{d\theta} \frac{F(P'' + \theta) + P'' f(P'' + \theta)}{F(\hat{P} + \theta) + \hat{P} f(\hat{P} + \theta)} \\ &\propto \frac{f(P'' + \theta) + P'' f'(P'' + \theta)}{F(P'' + \theta) + P'' f(P'' + \theta)} - \frac{f(\hat{P} + \theta) + \hat{P} f'(\hat{P} + \theta)}{F(\hat{P} + \theta) + \hat{P} f(\hat{P} + \theta)} < 0, \end{aligned}$$

where the negativity comes from i) $\frac{f(P'' + \theta)}{F(P'' + \theta)} < \frac{f(\hat{P} + \theta)}{F(\hat{P} + \theta)}$ and $\frac{f'(P'' + \theta)}{f(P'' + \theta)} < \frac{f'(\hat{P} + \theta)}{f(\hat{P} + \theta)}$ and ii) $\frac{f(v)}{F(v)} > \frac{f'(v)}{f(v)}$ by the log-concavity of $f(\cdot)$. $\square$

Proof of Theorem 1

(a) First I prove that it is without loss to assume that $\hat{\Theta}^{\text{ot}}$ is downward-closed. Suppose not, let $\hat{\theta}^{\text{ot}} = \sup \hat{\Theta}^{\text{ot}}$. Consider $\varepsilon > 0$ such that there exists $\tilde{\theta} \in (0, \hat{\theta}^{\text{ot}} - \varepsilon)$ such that

$$\int_{\hat{\Theta}^{\text{ot}}} dG(\theta) = \int_{[0, \tilde{\theta}] \cup ([\tilde{\theta}, \hat{\theta}^{\text{ot}} - \varepsilon] \cap \hat{\Theta}^{\text{ot}})} dG(\theta).$$

Note that $\frac{d\tilde{\theta}}{d\varepsilon} = \frac{g(\hat{\theta}^{\text{ot}} - \varepsilon)}{g(\tilde{\theta})}$. Now let $\Theta_l = [0, \tilde{\theta}] \setminus \hat{\Theta}^{\text{ot}}$ and $\Theta_h = \hat{\Theta}^{\text{ot}} \setminus [0, \hat{\theta}^{\text{ot}} - \varepsilon]$ so that

$$\Theta_l \cup ([0, \hat{\theta}^{\text{ot}} - \varepsilon] \cap \hat{\Theta}^{\text{ot}}) = [0, \tilde{\theta}] \cup ([\tilde{\theta}, \hat{\theta}^{\text{ot}} - \varepsilon] \cap \hat{\Theta}^{\text{ot}})$$

and

$$([0, \hat{\theta}^{\text{ot}} - \varepsilon] \cap \hat{\Theta}^{\text{ot}}) \cup \Theta_h = \hat{\Theta}^{\text{ot}}.$$

Since $\Theta_l \cap \Theta_h = \emptyset$ we have $\int_{\Theta_l} dG(\theta) = \int_{\Theta_h} dG(\theta)$. It is obvious that $\sup \Theta_l \leq \inf \Theta_h$. Define a function, $O : \Theta_l \rightarrow \Theta_h$, such that for any $\theta \in \Theta_l$,

$$g(O(\theta)) \cdot O'(\theta) = g(\theta)$$

and

$$O(\inf \Theta_l) = \inf \Theta_h$$

are satisfied. The existence of solution to this ODE together with the boundary condition is ensured by Picard's existence theorem and that $g(\cdot)$ is continuous. Note that i) $O(\cdot)$ is strictly increasing and thus allows inverse $O^{-1}(\cdot)$; ii) for any $\theta \in \Theta_l$ we have $O(\theta) > \theta$. Consider

the following new platform strategy $\{\tilde{\lambda}, \tilde{\Theta}, \tilde{P}\}$ such that i) $\tilde{\Theta} = [0, \tilde{\theta}]$; ii) for any $\theta \in \Theta_l$, $\delta \cdot (1 - \tilde{\lambda}_\theta) = 1 - \hat{\lambda}_{O(\theta)}^{\text{ot}}$ with

$$\delta = \frac{\int_{\Theta_l} (1 - \hat{\lambda}_\theta^{\text{ot}}) \cdot \Phi(\theta|\hat{P}) d\theta}{\int_{\Theta_l} (1 - \hat{\lambda}_{O(\theta)}^{\text{ot}}) \cdot \Phi(O(\theta)|\hat{P}) d\theta} = \frac{\int_{\Theta_l} (1 - \hat{\lambda}_\theta^{\text{ot}}) \cdot \Phi(\theta|\hat{P}) d\theta}{\int_{\Theta_h} (1 - \hat{\lambda}_\theta^{\text{ot}}) \cdot \Phi(\theta|\hat{P}) d\theta} \in (0, 1);$$

iii) everything else remains the same. This strategy would always satisfy feasibility and result in $\tilde{P} = \hat{P}^{\text{ot}}$ because $\hat{P}^{\text{ot}}$ satisfies (FOC-S):

$$\begin{aligned} & \int_0^{\tilde{\theta}} [1 - (1 - \tilde{\lambda}_\theta) \cdot \Phi(\theta|\hat{P})] dG(\theta) \\ &= \int_{\Theta_l} [1 - (1 - \tilde{\lambda}_\theta) \cdot \Phi(\theta|\hat{P})] \cdot g(O^{-1}(O(\theta))) d\theta + \int_{\hat{\Theta}^{\text{ot}} \cap [0, \tilde{\theta}]} [1 - (1 - \hat{\lambda}_\theta^{\text{ot}}) \cdot \Phi(\theta|\hat{P})] dG(\theta) \\ &= \int_{\Theta_h} [1 - (1 - \tilde{\lambda}_{O^{-1}(\theta)}) \cdot \Phi(O^{-1}(\theta)|\hat{P})] \cdot \frac{g(O^{-1}(\theta))}{O'(O^{-1}(\theta))} d\theta + \int_{\hat{\Theta}^{\text{ot}} \cap [0, \tilde{\theta}]} [1 - (1 - \hat{\lambda}_\theta^{\text{ot}}) \cdot \Phi(\theta|\hat{P})] dG(\theta) \\ &= \int_{\Theta_h} [1 - (1 - \hat{\lambda}_\theta^{\text{ot}}) \cdot \Phi(\theta|\hat{P})] \cdot g(\theta) d\theta + \int_{\hat{\Theta}^{\text{ot}} \cap [0, \tilde{\theta}]} [1 - (1 - \hat{\lambda}_\theta^{\text{ot}}) \cdot \Phi(\theta|\hat{P})] dG(\theta) \\ &= \int_{\hat{\Theta}^{\text{ot}}} [1 - (1 - \hat{\lambda}_\theta^{\text{ot}}) \cdot \Phi(\theta|\hat{P})] dG(\theta) = 0. \end{aligned}$$

To prove that the equilibrium price remains the same, it is sufficient to further prove that

$$\begin{aligned} & \frac{d}{d\varepsilon} (V(P'; \tilde{\lambda}, \tilde{\Theta}(P'|\tilde{\lambda}, \tilde{\Theta}, \hat{P})) - V(\hat{P}; \tilde{\lambda}, \tilde{\Theta})) \Big|_{\varepsilon_l=0} \\ & \leq \frac{d}{d\varepsilon} [(V(P'; \tilde{\lambda}, \Theta_l) - V(\hat{P}; \tilde{\lambda}, \Theta_l) - (V(P'; \tilde{\lambda}, \Theta_h) - V(\hat{P}; \tilde{\lambda}, \Theta_h))] \Big|_{\varepsilon_l=0} \\ & = \frac{d}{d\varepsilon} \left\{ \int_{\hat{P}}^{P'} \int_{\Theta_l} [1 - (1 - \tilde{\lambda}_\theta) \Phi(\theta; P'')] dG(\theta) dP'' - \int_{\Theta_h} [1 - (1 - \tilde{\lambda}_\theta) \Phi(\theta; P'')] dG(\theta) dP'' \right\} \Big|_{\varepsilon=0} \\ & = - \int_{\hat{P}}^{P'} \left[\frac{d\tilde{\theta}}{d\varepsilon} \Big|_{\varepsilon=0} \cdot (1 - \tilde{\lambda}_{\tilde{\theta}(0)}) \cdot \Phi(\tilde{\theta}(0); P'') \cdot g(\tilde{\theta}(0)) - (1 - \hat{\lambda}_{\hat{\theta}^{\text{ot}}}^{\text{ot}}) \cdot \Phi(\hat{\theta}^{\text{ot}}; P'') \cdot g(\hat{\theta}^{\text{ot}}) \right] dP'' \\ & = - g(\hat{\theta}^{\text{ot}}) \cdot \int_{\hat{P}}^{P'} \left[1/\delta(0) \cdot (1 - \hat{\lambda}_{\tilde{\theta}(0)}^{\text{ot}}) \cdot \Phi(\tilde{\theta}(0); P'') - (1 - \hat{\lambda}_{\hat{\theta}^{\text{ot}}}^{\text{ot}}) \cdot \Phi(\hat{\theta}^{\text{ot}}; P'') \right] dP'' \leq 0. \end{aligned}$$

It is sufficient to prove that

$$\begin{aligned} \delta(0) &= \frac{(1 - \hat{\lambda}_{\tilde{\theta}(0)}^{\text{ot}}) \cdot \Phi(\tilde{\theta}(0)|\hat{P})}{(1 - \hat{\lambda}_{\hat{\theta}^{\text{ot}}}^{\text{ot}}) \cdot \Phi(\hat{\theta}^{\text{ot}}|\hat{P})} \leq \frac{(1 - \hat{\lambda}_{\tilde{\theta}(0)}^{\text{ot}}) \cdot \Phi(\tilde{\theta}(0); P'')}{(1 - \hat{\lambda}_{\hat{\theta}^{\text{ot}}}^{\text{ot}}) \cdot \Phi(\hat{\theta}^{\text{ot}}; P'')} \\ &\Longleftarrow \frac{\Phi(\hat{\theta}^{\text{ot}}; P'')}{\Phi(\hat{\theta}^{\text{ot}}; \hat{P})} \leq \frac{\Phi(\tilde{\theta}(0); P'')}{\Phi(\tilde{\theta}(0); \hat{P})}, \end{aligned}$$

which is satisfied by $\tilde{\theta}(0) < \hat{\theta}^{\text{ot}}$.

Now I prove the profitability of the new strategy. It is sufficient to prove that

$$\begin{aligned} & \left. \frac{d}{d\varepsilon} \left(W(\tilde{\lambda}, \tilde{\Theta}, \tilde{P} = \hat{P}) - W(\hat{\lambda}^{\text{ot}}, \hat{\Theta}^{\text{ot}}, \hat{P}^{\text{ot}}) \right) \right|_{\varepsilon=0} \\ &= \left. \frac{d\tilde{\theta}}{d\varepsilon} \right|_{\varepsilon=0} \cdot \left[1 - (1 - \tilde{\lambda}_{\tilde{\theta}(0)}) F(\hat{P}^{\text{ot}} + \tilde{\theta}(0)) \right] \cdot g(\tilde{\theta}(0)) - \left[1 - (1 - \hat{\lambda}_{\hat{\theta}^{\text{ot}}}^{\text{ot}}) F(\hat{P}^{\text{ot}} + \hat{\theta}^{\text{ot}}) \right] \cdot g(\hat{\theta}^{\text{ot}}) \\ &= g(\hat{\theta}^{\text{ot}}) \cdot \left\{ -1/\delta(0) \cdot (1 - \tilde{\lambda}_{\tilde{\theta}(0)}) F(\hat{P}^{\text{ot}} + \tilde{\theta}(0)) + (1 - \hat{\lambda}_{\hat{\theta}^{\text{ot}}}^{\text{ot}}) F(\hat{P}^{\text{ot}} + \hat{\theta}^{\text{ot}}) \right\} \geq 0. \end{aligned}$$

to this end, note that

$$\begin{aligned} \delta(0) &= \frac{(1 - \hat{\lambda}_{\tilde{\theta}(0)}^{\text{ot}}) \cdot (F(\hat{P} + \tilde{\theta}(0)) + P f(\hat{P} + \tilde{\theta}(0)))}{(1 - \hat{\lambda}_{\hat{\theta}^{\text{ot}}}^{\text{ot}}) \cdot (F(\hat{P} + \hat{\theta}^{\text{ot}}) + P f(\hat{P} + \hat{\theta}^{\text{ot}}))} \leq \frac{(1 - \tilde{\lambda}_{\tilde{\theta}(0)}) F(\hat{P}^{\text{ot}} + \tilde{\theta}(0))}{(1 - \hat{\lambda}_{\hat{\theta}^{\text{ot}}}^{\text{ot}}) F(\hat{P}^{\text{ot}} + \hat{\theta}^{\text{ot}})} \\ &\Longleftarrow \frac{F(\hat{P} + \tilde{\theta}(0)) + P f(\hat{P} + \tilde{\theta}(0))}{F(\hat{P}^{\text{ot}} + \tilde{\theta}(0))} \geq \frac{F(\hat{P} + \hat{\theta}^{\text{ot}}) + P f(\hat{P} + \hat{\theta}^{\text{ot}})}{F(\hat{P}^{\text{ot}} + \hat{\theta}^{\text{ot}})}, \end{aligned}$$

which is always satisfied by the log-concavity of $F(\cdot)$. Therefore, the new platform strategy is always profitable for the platform, which serves as a contradiction and completes the proof.

Now it is always without loss to focus on a downward-closed $\hat{\Theta}^{\text{ot}} = [0, \hat{\theta}^{\text{ot}}]$. Now I prove that $\hat{\theta}^{\text{ot}} = 1 - \hat{P}^{\text{ot}}$. Suppose not, that is, $\hat{\theta}^{\text{ot}} < 1 - \hat{P}^{\text{ot}}$. Consider the following new platform strategy $\{\tilde{\lambda}, \tilde{\Theta}, \tilde{P}\}$ such that i) $\tilde{\Theta} = [0, \hat{\theta}^{\text{ot}} + \varepsilon]$; ii) for any $\theta \in (\hat{\theta}^{\text{ot}}, \hat{\theta}^{\text{ot}} + \varepsilon]$, $\tilde{\lambda}_{\theta} = 0$; iii) everything else is the same. I prove the profitability of this new strategy by prove that $\tilde{P} < \hat{P}^{\text{ot}}$. To see this, note that if the equilibrium price remains unchanged, or $\tilde{P} = \hat{P}^{\text{ot}}$, if a firm deviates to a higher price $P' \in (\hat{P}^{\text{ot}}, 1 - \hat{\theta}^{\text{ot}} - \varepsilon]$, the change of profit is

$$\begin{aligned} & V(P'; \tilde{\lambda}, \tilde{\Theta}, \hat{P}^{\text{ot}}) - V(\hat{P}^{\text{ot}}; \tilde{\lambda}, \tilde{\Theta}, \hat{P}^{\text{ot}}) \\ &= P' \cdot \int_0^{\hat{\theta}^{\text{ot}} + \varepsilon} D_{\theta}(P'; \tilde{\lambda}_{\theta}, \hat{P}^{\text{ot}}) dG(\theta) - \hat{P}^{\text{ot}} \cdot \int_0^{\hat{\theta}^{\text{ot}} + \varepsilon} D_{\theta}(\hat{P}^{\text{ot}}; \tilde{\lambda}_{\theta}, \hat{P}^{\text{ot}}) dG(\theta) \\ &= P' \cdot \int_0^{\hat{\theta}^{\text{ot}}} D_{\theta}(P'; \hat{\lambda}_{\theta}^{\text{ot}}, \hat{P}^{\text{ot}}) dG(\theta) - \hat{P}^{\text{ot}} \cdot \int_0^{\hat{\theta}^{\text{ot}}} D_{\theta}(\hat{P}^{\text{ot}}; \hat{\lambda}_{\theta}^{\text{ot}}, \hat{P}^{\text{ot}}) dG(\theta) \\ &\quad + \int_{\hat{\theta}^{\text{ot}}}^{\hat{\theta}^{\text{ot}} + \varepsilon} \int_{\hat{P}^{\text{ot}}}^{P'} [1 - \Phi(\theta|P'')] dP'' dG(\theta) \\ &< P' \cdot \int_0^{\hat{\theta}^{\text{ot}}} D_{\theta}(P'; \hat{\lambda}_{\theta}^{\text{ot}}, \hat{P}^{\text{ot}}) dG(\theta) - \hat{P}^{\text{ot}} \cdot \int_0^{\hat{\theta}^{\text{ot}}} D_{\theta}(\hat{P}^{\text{ot}}; \hat{\lambda}_{\theta}^{\text{ot}}, \hat{P}^{\text{ot}}) dG(\theta) \leq 0, \end{aligned}$$

where the first inequality comes from that

$$1 - \Phi(\theta|P'') < 1 - \Phi(\theta|\hat{P}^{\text{ot}}) < 1 - (1 - \lambda_{\theta'}) \Phi(\theta'|\hat{P}^{\text{ot}})$$

for any $P'' > \hat{P}^{\text{ot}}$, $\theta \in (\hat{\theta}^{\text{ot}}, \hat{\theta}^{\text{ot}} + \varepsilon]$ and $\theta' \in [0, \hat{\theta}^{\text{ot}}]$; the second inequality comes from that the firm finds it optimal to offer $\hat{P}^{\text{ot}}$ under $\hat{\lambda}^{\text{ot}}$. If the firm deviates to $P' > 1 - \hat{\theta}^{\text{ot}} - \varepsilon$, we have the similar logic to the above. Therefore, the new equilibrium price $\tilde{P}$ is weakly lower than $\hat{P}^{\text{ot}}$, which completes the proof.

Proof of Lemma 4

(a) For any $\theta' < \theta$, by (IC) we have

$$\int_{P+\theta}^1 1 - (1 - \lambda_{\theta}) F(v) dv - c_{\theta} \geq \int_{P+\theta}^1 1 - (1 - \lambda_{\theta'}) F(v) dv - c_{\theta'}$$

and

$$\int_{P+\theta'}^1 1 - (1 - \lambda_{\theta'}) F(v) dv - c_{\theta'} \geq \int_{P+\theta'}^1 1 - (1 - \lambda_{\theta}) F(v) dv - c_{\theta}.$$

Combine the two inequalities, we have

$$(\lambda_{\theta} - \lambda_{\theta'}) \cdot \int_{P+\theta'}^{P+\theta} F(v) dv \leq 0 \implies \lambda_{\theta} \leq \lambda_{\theta'},$$

given $\theta > \theta'$.

(b) Suppose not, that is, there exists $\theta' < \theta$ such that $U(\lambda_{\theta}, P; \theta) \geq 0$ but $U(\lambda_{\theta'}, P; \theta') < 0$.

Note that by incentive compatibility, we have

$$U(\lambda_{\theta'}, P; \theta') \geq U(\lambda_{\theta}, P; \theta') = U(\lambda_{\theta}, P; \theta) + \int_{P+\theta'}^{P+\theta} 1 - (1 - \lambda_{\theta}) F(v) dv \geq 0,$$

which is a contradiction. □

Proof of Theorem 2

The reason why the optimal mechanism should be the solution to (AP) is obvious: otherwise there exists a new mechanism so that the equilibrium price is strictly lower while

i) the accuracy level for type $\theta = 0$ remains $\bar{L}$ and ii) the payoff of type 0 remains the same, which strictly increases the trade probability by (8), a contradiction.

Now I solve (AP). By the method of Karush-Kuhn-Tucker conditions, there exists $\mu(\cdot), \kappa \in \mathbb{R}$ and $\eta_1(\cdot), \eta_2(\cdot) \geq 0$, such that for any $\theta \in [0, \hat{\theta}]$, for a mechanism λ to be optimal, we have

$$\mu(\theta) \cdot \left(\int_{P+\theta}^1 F(v) dv - \gamma \cdot e'(\lambda_\theta) \right) + \eta_1(\theta) - \eta_2(\theta) = \left(1 + \int_0^\theta \mu(\theta') d\theta' - \kappa \right) \cdot F(P+\theta) + P f(P+\theta). \quad (\text{A.1})$$

Similarly, we can evaluate KKT condition with respect to the choice of $\hat{\theta}$, or

$$\begin{aligned} & \eta_1(\hat{\theta}) \cdot \lambda_{\hat{\theta}} + \eta_2(\hat{\theta}) \cdot (1 - \lambda_{\hat{\theta}}) + \mu(\hat{\theta}) \cdot \left[\int_{P+\hat{\theta}}^1 1 - (1 - \lambda_{\hat{\theta}}) F(v) dv - \gamma \cdot e'(\lambda_{\hat{\theta}}) \right] \\ &= \left(1 + \int_0^{\hat{\theta}} \mu(\theta') d\theta' - \kappa \right) \cdot [1 - (1 - \lambda_{\hat{\theta}}) F(P + \hat{\theta})] - (1 - \lambda_{\hat{\theta}}) \cdot P f(P + \hat{\theta}). \end{aligned} \quad (\text{A.2})$$

By complementary slackness, we have $\eta_1(\theta) \cdot \lambda_\theta = 0$ and $\eta_2(\theta) \cdot (1 - \lambda_\theta) = 0$. I derive the solution to (AP), or, $\hat{\lambda}$, in several steps:

Step 1. If there exists $\theta^c, \varepsilon > 0$ such that $\hat{\lambda}_\theta^{\text{pt}}$ is continuous on interval $[\theta^c, \theta^c + \varepsilon] \subset \hat{\Theta}^{\text{pt}}$, then either $\frac{d\lambda_\theta}{d\theta} = 0$ or $\int_{P+\theta}^1 F(v) dv = \gamma \cdot e'(\lambda_\theta)$. To see this, take the derivative of (AP-ICa) w.r.t. θ , we have

$$\begin{aligned} -[1 - (1 - \lambda_\theta) \cdot F(P + \theta)] &= -[1 - (1 - \lambda_\theta) \cdot F(P + \theta)] + \frac{d\lambda_\theta}{d\theta} \cdot \int_{P+\theta}^1 F(v) dv - \frac{d\lambda_\theta}{d\theta} \cdot \gamma \cdot e'(\lambda_\theta) \\ &\implies \lambda'(\theta) \cdot \left(\int_{P+\theta}^1 F(v) dv - \gamma \cdot e'(\lambda_\theta) \right) = 0, \end{aligned}$$

which completes Step 1. In the rest of proof I made the following definition: let $\lambda_\theta^- \equiv \lim_{\theta' \uparrow \theta} \hat{\lambda}_{\theta'}^{\text{pt}}$ and $\lambda_\theta^+ \equiv \lim_{\theta' \downarrow \theta} \hat{\lambda}_{\theta'}^{\text{pt}}$.

Step 2. I prove that $\hat{\lambda}_\theta^{\text{pt}}$ is continuous everywhere. Suppose not, by Steps 1 it is sufficient to consider three cases:

Case A There exists $\theta^a \in (0, \hat{\theta})$ and $\varepsilon_-^a, \varepsilon_+^a > 0$ such that $\hat{\lambda}_\theta^{\text{pt}}$ is i) constant on $[\theta^a - \varepsilon_-^a, \theta^a]$; ii) continuous but not constant (and henceforth satisfies $\int_{P+\theta}^1 F(v) dv = \gamma \cdot e'(\hat{\lambda}_\theta^{\text{pt}})$ by Step 2) on $[\theta^a, \theta^a + \varepsilon_+^a]$; iii) $\lambda^-(\theta^a) > \lambda^+(\theta^a)$. Note that since $\int_{P+\theta^a}^1 F(v) dv = \gamma \cdot e'(\lambda^+(\theta^a))$ and $e(\cdot)$ is strictly convex, we have $\int_{P+\theta^a}^1 F(v) dv < \gamma \cdot e'(L)$ for any $L \in (\lambda^+(\theta^a), \lambda^-(\theta^a))$. Therefore we

have

$$\begin{aligned}
& \int_{P+\theta^a}^1 1 - (1 - \lambda^+(\theta^a)) \cdot F(v) dv - \gamma \cdot e(\lambda^+(\theta^a)) \\
&= \int_{P+\theta^a}^1 1 - (1 - \lambda^-(\theta^a)) \cdot F(v) dv - \gamma \cdot e(\lambda^-(\theta^a)) - \int_{\lambda^+(\theta^a)}^{\lambda^-(\theta^a)} \left[\int_{P+\theta^a}^1 F(v) dv - \gamma \cdot e'(L) \right] dL \\
&> \int_{P+\theta^a}^1 1 - (1 - \lambda^-(\theta^a)) \cdot F(v) dv - \gamma \cdot e(\lambda^-(\theta^a)) \\
&= \int_{\theta^a}^{\hat{\theta}} 1 - (1 - \hat{\lambda}_{\theta'}^{\text{pt}}) \cdot F(P + \theta') d\theta'
\end{aligned}$$

which contradicts to (AP-ICa).

Case B There exists $\theta^b \in (0, \hat{\theta})$ and $\varepsilon_-^b, \varepsilon_+^b > 0$ such that $\hat{\lambda}_{\theta}^{\text{pt}}$ is i) constant on $[\theta^b, \theta^b + \varepsilon_+^b)$; ii) continuous (and henceforth satisfies $\int_{P+\theta}^1 F(v) dv = \gamma \cdot e'(\hat{\lambda}_{\theta}^{\text{pt}})$ by Step 1) on $[\theta^b - \varepsilon_-^b, \theta^b)$; iii) $\lambda^+(\theta^b) < \lambda^-(\theta^b)$. Note that since $\int_{P+\theta^b}^1 F(v) dv = \gamma \cdot e'(\lambda^-(\theta^b))$ and $e(\cdot)$ is strictly convex, we have $\int_{P+\theta^b}^1 F(v) dv > e'(L)$ for any $L \in (\lambda^+(\theta^b), \lambda^-(\theta^b))$. Similarly, we have

$$\begin{aligned}
& \int_{P+\theta^b}^1 1 - (1 - \lambda^-(\theta^b)) \cdot F(v) dv - \gamma \cdot e(\lambda^-(\theta^b)) \\
&= \int_{P+\theta^b}^1 1 - (1 - \lambda^+(\theta^b)) \cdot F(v) dv - \gamma \cdot e(\lambda^+(\theta^b)) + \int_{\lambda^+(\theta^b)}^{\lambda^-(\theta^b)} \left[\int_{P+\theta^b}^1 F(v) dv - \gamma \cdot e'(L) \right] dL \\
&> \int_{P+\theta^b}^1 1 - (1 - \lambda^+(\theta^b)) \cdot F(v) dv - \gamma \cdot e(\lambda^+(\theta^b)) \\
&= \int_{\theta^b}^{\hat{\theta}} 1 - (1 - \hat{\lambda}_{\theta'}^{\text{pt}}) \cdot F(P + \theta') d\theta'.
\end{aligned}$$

Therefore, we still have the same contradiction.

Case C¹⁰ There exists $0 \leq \theta_l < \theta_h \leq \hat{\theta}$ and $\varepsilon_-^l, \varepsilon_+^h \geq 0$ satisfying $\theta_l > 0 \iff \varepsilon_-^l > 0$ and $\theta_h < \hat{\theta} \iff \varepsilon_+^h > 0$ such that $\hat{\lambda}_{\theta}^{\text{pt}}$ is i) strictly decreasing and continuous on $[\theta_l - \varepsilon_-^l, \theta_l)$ and $[\theta_h, \theta_h + \varepsilon_+^h)$; ii) either constant or discontinuous on $[\theta_l, \theta_h)$. In other words, for $\theta \in (\theta_l, \theta_h)$, there exists $n \in \mathbb{N}^+$ with $n > 1$, $\lambda^-(\theta_l) \geq L_1 > L_2 > \dots > L_n \geq \lambda^+(\theta_h)$ and $\theta_l = \theta_0 < \theta_1 < \theta_2 < \dots < \theta_n = \theta_h$, such that $\theta \in [\theta_{i-1}, \theta_i) \implies \hat{\lambda}_{\theta}^{\text{pt}} = L_i$.

¹⁰The spirit of this subcase is the same as what Figure 4 shows.

Consider a deviation from λ_θ on $[\theta_{j-1}, \theta_{j+1})$. Let $\tilde{L}_{j+1} \in (L_{j+1}, L_j)$ and

$$\tilde{\lambda}_\theta = \begin{cases} L_j, & \theta \in [\theta_{j-1}, \tilde{\theta}_j); \\ \tilde{L}_{j+1}, & \theta \in [\tilde{\theta}_j, \tilde{\theta}_{j+1}); \\ L_{j+1}, & \theta \in [\tilde{\theta}_{j+1}, \theta_{j+1}), \end{cases}$$

where $\tilde{\theta}_j, \tilde{\theta}_{j+1}$ are pinned down by

$$\int_{P+\tilde{\theta}_j}^1 F(v) dv = \frac{e(L_j) - e(\tilde{L}_{j+1})}{L_j - \tilde{L}_{j+1}} \quad (\text{A.3})$$

and

$$\int_{P+\tilde{\theta}_{j+1}}^1 F(v) dv = \frac{e(\tilde{L}_{j+1}) - e(L_{j+1})}{\tilde{L}_{j+1} - L_{j+1}}. \quad (\text{A.4})$$

The existence of $\tilde{\theta}_j$ and $\tilde{\theta}_{j+1}$ satisfying $\theta_j \in (\tilde{\theta}_j, \tilde{\theta}_{j+1})$ is guaranteed by:

$$\begin{aligned} \int_{P+\theta_{j-1}}^1 F(v) dv &\geq \gamma \cdot e'(L_j) > \gamma \cdot \frac{e(L_j) - e(\tilde{L}_{j+1})}{L_j - \tilde{L}_{j+1}} \\ &> \gamma \cdot \frac{e(L_j) - e(L_{j+1})}{L_j - L_{j+1}} = \int_{P+\theta_j}^1 F(v) dv \\ &> \gamma \cdot \frac{e(\tilde{L}_{j+1}) - e(L_{j+1})}{\tilde{L}_{j+1} - L_{j+1}} > \gamma \cdot e'(L_{j+1}) \geq \int_{P+\theta_{j+1}}^1 F(v) dv, \end{aligned}$$

where the first and the last inequalities is from $\int_{P+\theta_l}^1 F(v) dv - \gamma \cdot e'(\lambda^-(\theta_l)) = \int_{P+\theta_h}^1 F(v) dv - \gamma \cdot e'(\lambda^+(\theta_h)) = 0$ and the convexity of $e(\cdot)$; other inequalities is from the convexity of $e(\cdot)$; the equality is by (AP-ICa) at θ_j^- and θ_j^+ .

It turns out that by construction (AP-ICa) remains satisfied for any $\theta \in [\theta_{j-1}, \theta_{j+1})$. To see

this, note that for $\theta \in [\tilde{\theta}_j, \tilde{\theta}_{j+1})$, we have

$$\begin{aligned}
& \int_{P+\theta}^1 1 - (1 - \tilde{\lambda}_\theta) \cdot F(v) dv - \int_\theta^{\hat{\theta}} 1 - (1 - \tilde{\lambda}_{\theta'}) \cdot F(P + \theta') d\theta' - \gamma \cdot e(\tilde{\lambda}_\theta) \\
&= \int_{P+\theta}^1 1 - (1 - \tilde{L}_{j+1}) \cdot F(v) dv - \int_{\tilde{\theta}_{j+1}}^{\hat{\theta}} 1 - (1 - \hat{\lambda}_{\theta'}^{\text{pt}}) \cdot F(P + \theta') d\theta' \\
&\quad - \int_\theta^{\tilde{\theta}_{j+1}} 1 - (1 - \tilde{L}_{j+1}) \cdot F(P + \theta') d\theta' - \gamma \cdot e(\tilde{L}_{j+1}) \\
&= \int_{P+\tilde{\theta}_{j+1}}^1 1 - (1 - \tilde{L}_{j+1}) \cdot F(v) dv - \int_{\tilde{\theta}_{j+1}}^{\hat{\theta}} 1 - (1 - \hat{\lambda}_{\theta'}^{\text{pt}}) \cdot F(P + \theta') d\theta' - \gamma \cdot e(\tilde{L}_{j+1}) \\
&= \int_{P+\tilde{\theta}_{j+1}}^1 1 - (1 - L_{j+1}) \cdot F(v) dv - \int_{\tilde{\theta}_{j+1}}^{\hat{\theta}} 1 - (1 - \hat{\lambda}_{\theta'}^{\text{pt}}) \cdot F(P + \theta') d\theta' - \gamma \cdot e(L_{j+1}) = 0,
\end{aligned}$$

where the first equality is by the definition of $\tilde{\lambda}_\theta$, the third equality is by (A.4) and the last one comes from (AP-ICa) at $\theta = \tilde{\theta}_{j+1}$. That is, if $\hat{\lambda}^{\text{pt}}$ satisfies (AP-ICa) then so does $\tilde{\lambda}$. Similarly, for $\theta \in [\theta_{j-1}, \tilde{\theta}_j)$ we have

$$\begin{aligned}
& \int_{P+\theta}^1 1 - (1 - \tilde{\lambda}_\theta) \cdot F(v) dv - \int_\theta^{\hat{\theta}} 1 - (1 - \tilde{\lambda}_{\theta'}) \cdot F(P + \theta') d\theta' - \gamma \cdot e(\tilde{\lambda}_\theta) \\
&= \int_{P+\theta}^1 1 - (1 - L_j) \cdot F(v) dv - \int_{\tilde{\theta}_{j+1}}^{\hat{\theta}} 1 - (1 - \hat{\lambda}_{\theta'}^{\text{pt}}) \cdot F(P + \theta') d\theta' - \int_{\tilde{\theta}_j}^{\tilde{\theta}_{j+1}} 1 - (1 - \tilde{L}_{j+1}) \cdot F(P + \theta') d\theta' \\
&\quad - \int_\theta^{\tilde{\theta}_j} 1 - (1 - L_j) \cdot F(P + \theta') d\theta' - \gamma \cdot e(L_j) \\
&= \int_{P+\tilde{\theta}_j}^1 1 - (1 - L_j) \cdot F(v) dv - \int_{\tilde{\theta}_{j+1}}^{\hat{\theta}} 1 - (1 - \hat{\lambda}_{\theta'}^{\text{pt}}) \cdot F(P + \theta') d\theta' - \int_{\theta_j}^{\tilde{\theta}_{j+1}} 1 - (1 - L_{j+1}) \cdot F(P + \theta') d\theta' \\
&\quad - \int_{\tilde{\theta}_j}^{\theta_j} 1 - (1 - L_j) \cdot F(P + \theta') d\theta' - \gamma \cdot e(L_j) \\
&= \int_{P+\tilde{\theta}_j}^1 1 - (1 - L_j) \cdot F(v) dv - \int_{\tilde{\theta}_j}^{\hat{\theta}} 1 - (1 - \hat{\lambda}_{\theta'}^{\text{pt}}) \cdot F(P + \theta') d\theta' - \gamma \cdot e(L_j) = 0,
\end{aligned}$$

where the second equality comes from

$$\begin{aligned}
& (L_j - \tilde{L}_{j+1}) \cdot \int_{P+\tilde{\theta}_j}^1 F(v) dv + (\tilde{L}_{j+1} - L_{j+1}) \cdot \int_{P+\tilde{\theta}_{j+1}}^1 F(v) dv \\
&= \gamma \cdot (e(L_j) - e(L_{j+1})) = (L_j - L_{j+1}) \cdot \int_{P+\theta_j}^1 F(v) dv \\
&\implies \int_{\tilde{\theta}_j}^{\theta_j} 1 - (1 - L_j) \cdot F(P + \theta') d\theta' + \int_{\theta_j}^{\tilde{\theta}_{j+1}} 1 - (1 - L_{j+1}) \cdot F(P + \theta') d\theta' = \int_{\tilde{\theta}_j}^{\tilde{\theta}_{j+1}} 1 - (1 - \tilde{L}_{j+1}) \cdot F(P + \theta') d\theta'
\end{aligned} \tag{A.5}$$

by (A.3), (A.4) and (AP-ICa).

Now it remains to prove the profitability. By (A.5) we have

$$(L_j - \tilde{L}_{j+1}) \cdot \int_{\tilde{\theta}_j}^{\theta_j} F(P + \theta) d\theta = (\tilde{L}_{j+1} - L_{j+1}) \cdot \int_{\theta_j}^{\tilde{\theta}_{j+1}} F(P + \theta) d\theta$$

and thus

$$\begin{aligned} & \int_{\theta_{j-1}}^{\theta_{j+1}} (1 - \hat{\lambda}_{\theta}^{\text{pt}}) \cdot [F(P + \theta) + P f(P + \theta)] d\theta - \int_{\theta_{j-1}}^{\theta_{j+1}} (1 - \tilde{\lambda}_{\theta}) \cdot [F(P + \theta) + P f(P + \theta)] d\theta \\ &= -(L_j - \tilde{L}_{j+1}) \cdot \int_{\tilde{\theta}_j}^{\theta_j} [F(P + \theta) + P f(P + \theta)] d\theta + (\tilde{L}_{j+1} - L_{j+1}) \cdot \int_{\theta_j}^{\tilde{\theta}_{j+1}} [F(P + \theta) + P f(p + \theta)] d\theta < 0, \end{aligned} \quad (\text{A.6})$$

where the inequality comes from $\frac{\int_{\tilde{\theta}_j}^{\tilde{\theta}_{j+1}} [F(P + \theta) + P f(P + \theta)] d\theta}{\int_{\tilde{\theta}_j}^{\tilde{\theta}_{j+1}} F(P + \theta) d\theta} < \frac{\int_{\tilde{\theta}_j}^{\theta_j} [F(P + \theta) + P f(P + \theta)] d\theta}{\int_{\tilde{\theta}_j}^{\theta_j} F(P + \theta) d\theta}$ by the log-concavity of $F(\cdot)$, therefore $\tilde{\lambda}$ is a profitable deviation for the platform, a contradiction.

Step 3. I prove that there is at most one connected interval on which $\hat{\lambda}_{\theta}^{\text{pt}}$ is strictly decreasing and continuous. Suppose not, then by Step 3, it must be that there exists $0 < \theta_l < \theta_h < \hat{\theta}$ and $\varepsilon_-^l, \varepsilon_+^h \geq 0$ such that $\hat{\lambda}_{\theta}$ i) satisfies $\int_{P+\theta}^1 F(v) dv = \gamma \cdot e'(\hat{\lambda}_{\theta})$ for $\theta \in [\theta_l - \varepsilon_-^l, \theta_l]$ and $[\theta_h, \theta_h + \varepsilon_+^h]$; ii) is constant on $[\theta_l, \theta_h]$, or $\lambda^+(\theta_l) = \lambda^-(\theta_h)$; iii) $\lambda^-(\theta_l) = \lambda^+(\theta_l) = \hat{\lambda}_{\theta_l}^{\text{pt}}$ and $\lambda^-(\theta_h) = \lambda^+(\theta_h) = \hat{\lambda}_{\theta_h}^{\text{pt}}$. However, this means that $\int_{P+\theta_l}^1 F(v) dv = \int_{P+\theta_h}^1 F(v) dv$, which is a contradiction to $F(v) > 0$.

Step 4. Denote $[\theta^a, \theta^b]$ with $0 \leq \theta^a \leq \theta^b \leq \hat{\theta}$ as the unique connected interval (could be empty) on which $\hat{\lambda}_{\theta}^{\text{pt}}$ is strictly decreasing and continuous. In this step, I prove that if $\lambda(0) > \lambda_{\hat{\theta}}$, then $\hat{\theta}$, $\hat{\lambda}_{\hat{\theta}}^{\text{pt}}$, θ^a and θ^b are collectively pinned down by

$$\gamma \cdot e(\hat{\lambda}_{\hat{\theta}}^{\text{pt}}) = 0, \quad (\text{A.7})$$

$$\int_{P+\theta^a}^1 F(v) dv = \gamma \cdot e'(\bar{L}) \quad (\text{A.8})$$

and

$$\theta^b = \hat{\theta} = 1 - P. \quad (\text{A.9})$$

First I prove that $\gamma \cdot e(\hat{\lambda}_{\hat{\theta}}^{\text{pt}}) = 0$. Suppose not, then it must be $\gamma \cdot e(\hat{\lambda}_{\hat{\theta}}^{\text{pt}}) > 0$. Therefore we have

$$\int_{P+\hat{\theta}}^1 F(v) dv < \gamma \cdot e'(\hat{\lambda}_{\hat{\theta}}^{\text{pt}})$$

by $U(\hat{\lambda}_{\hat{\theta}}^{\text{pt}}, P; \hat{\theta}) = \int_{P+\hat{\theta}}^1 1 - (1 - \hat{\lambda}_{\hat{\theta}}^{\text{pt}}) F(v) dv - \gamma \cdot e(\hat{\lambda}_{\hat{\theta}}^{\text{pt}}) = 0$. Furthermore, let $\theta = \hat{\theta}$ and (A.1) becomes

$$\mu(\hat{\theta}) \cdot \left(\int_{P+\hat{\theta}}^1 F(v) dv - \gamma \cdot e'(\hat{\lambda}_{\hat{\theta}}^{\text{pt}}) \right) = \left(1 + \int_0^{\hat{\theta}} \mu(\theta') d\theta' - \kappa \right) \cdot F(P + \hat{\theta}) + P f(P + \hat{\theta}) + \eta_2(\hat{\theta}),$$

which implies that $1 + \int_0^{\hat{\theta}} \mu(\theta') d\theta' - \kappa < 0$. And (A.2) becomes

$$0 = \left(1 + \int_0^{\hat{\theta}} \mu(\theta') d\theta' - \kappa \right) \cdot \left[1 - (1 - \hat{\lambda}_{\hat{\theta}}^{\text{pt}}) F(P + \hat{\theta}) \right] - (1 - \hat{\lambda}_{\hat{\theta}}^{\text{pt}}) \cdot P f(P + \hat{\theta}) - \eta_2(\hat{\theta}) \cdot (1 - \hat{\lambda}_{\hat{\theta}}^{\text{pt}}).$$

Combine both equations, we have

$$\begin{aligned} 0 &< -\lambda_{\hat{\theta}} \cdot \mu(\hat{\theta}) \cdot \left(\int_{P+\hat{\theta}}^1 F(v) dv - \gamma \cdot e'(\hat{\lambda}_{\hat{\theta}}^{\text{pt}}) \right) \\ &= \left(1 + \int_0^{\hat{\theta}} \mu(\theta') d\theta' - \kappa \right) \cdot (1 - F(P + \hat{\theta})) - P f(P + \hat{\theta}) - \eta_2(\hat{\theta}) < 0, \end{aligned}$$

which is a contradiction.

Now I have proved that $\lambda_{\hat{\theta}} = 0$, it is obvious that $\hat{\theta} = 1 - P$. Besides, by definition we have $\hat{\lambda}_{\theta_b}^{\text{pt}} = \hat{\lambda}_{\hat{\theta}}^{\text{pt}} = 0$ and $\int_{P+\theta_b}^1 F(v) dv = \gamma \cdot e'(\hat{\lambda}_{\theta_b}^{\text{pt}}) = 0$, which completes the proof of (A.9). The same logic applies to θ_a and we also have (A.8). Let $\theta^* = \theta_a$, we get exactly the same result as stated in Theorem 2. $\square$

For other omitted proofs, please refer to the Online Appendix.