

PERSONALIZED DISCOUNTS AND
CONSUMER SEARCH

By

Zikun Liu, Jiwoong Shin and Jidong Zhou

April 2025

COWLES FOUNDATION DISCUSSION PAPER NO. 2440



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

Personalized Discounts and Consumer Search

Zikun Liu*

Jiwoong Shin[†]

Jidong Zhou[‡]

April 2025

Abstract

The growing availability of big data enables firms to predict consumer search outcomes and outside options more accurately than consumers themselves. This paper examines how a firm can utilize such superior information to offer personalized buy-now discounts intended to deter consumer search. However, discounts can also serve as signals of attractive outside options, potentially encouraging rather than discouraging consumer search. We show that, despite the firm's ability to tailor discounts across a continuum of consumer valuations, the firm-optimal equilibrium features a simple two-tier discount scheme, comprising a uniform positive discount when the consumer outside option is intermediate and no discount when the outside option is low or high. Furthermore, compared to a scenario where the firm lacks superior information, we find that the firm earns lower profits, consumers search more while their welfare remains unchanged, and total welfare declines.

Keywords: Consumer search, superior information, personalized prices, signaling, buy-now discounts, search deterrence

*School of Management, Yale University. zikun.liu@yale.edu

[†]School of Management, Yale University. jiwoong.shin@yale.edu

[‡]School of Management, Yale University. jidong.zhou@yale.edu

1 Introduction

Consumers often face imperfect information about product qualities and prices, relying heavily on search to inform their purchasing decisions. Firms seeking to influence consumer search behaviors increasingly utilize detailed information about consumers’ preferences to predict their search outcomes and outside options. They then deploy targeted incentives—such as personalized “buy-now” discounts—to discourage further search by making an immediate purchase more attractive. However, such tactics can backfire: substantial discounts may arouse consumer suspicion about the firm’s motives and prompt even more extensive exploration of alternatives.

Consider a consumer deliberating between the BMW 3-series, Mercedes C-class, and other models. Unsure about their specific attributes, she visits various dealerships. At the BMW dealership, her interest in the 3-series becomes evident. In response, the salesperson offers a \$5,000 discount to encourage an immediate purchase. Although tempting, this discount prompts her to suspect that the salesperson might be anticipating her potential preference for the Mercedes—so she decides to extend her search. Similarly, a real estate agent with deep market knowledge and detailed property information—often unavailable to average buyers—can leverage this information advantage. For instance, the agent may discern from prior interactions that a family prioritizes good schools. Knowing that other properties in highly rated school districts may soon come onto the market, she might offer an early-bird discount contingent on a commitment before the public open house, aiming to prevent the family from exploring potentially more attractive alternatives. However, such an offer can inadvertently signal that even better options remain undiscovered, prompting buyers to keep looking. Insurance companies sometimes use detailed demographic and claims data to offer targeted discounts designed to discourage customers from shopping around. Yet, such offers can imply that even better deals exist, potentially spurring additional consumer search.

These examples highlight how personalized discounts intended to discourage search can instead encourage it, underscoring the intricate interplay between firms’ informational edge and consumers’ inferences. Firms typically possess deeper insights into consumer preferences and match values—both for their own offerings and for competitors’—gained through direct sales interactions (Wernerfelt, 1994; Rogers, 2013) and reinforced by modern analytics and predictive algorithms (Choi et al., 2024; Rafieian and Zuo, 2024). Although consumers may discover their preferences through hands-on experience, high search costs and lack of complete information about all available alternatives preserve the firm’s informational edge. Consumers, aware of the firm’s informational advantage, may interpret personalized discounts as signals that superior alternatives exist. This

signaling effect has a dual effect: it can steer consumers toward products that better match their preferences—boosting consumer surplus—or, conversely, stimulate unnecessary searches, ultimately harming both market efficiency and firm’s profit.

This paper explores the optimal design of personalized discounts when consumers engage in search and make strategic inferences in the presence of firm’s superior information. Specifically, we consider a search model where a firm (e.g., a retailer) interacts with a consumer who is uncertain about the value of her outside option (e.g., a competing product). The firm, by contrast, knows how attractive the outside option is for the consumer and can tailor a buy-now discount accordingly to induce an immediate purchase. The consumer must then decide whether to accept the offer or continue searching for the outside option, weighing the immediate discount against the potential benefit of finding a better alternative. We characterize the perfect Bayesian equilibrium of this game and compare it to a benchmark in which the firm lacks superior information and cannot condition discounts on the outside option. This framework captures a central tension: while better information enables more precise targeting, it also makes the consumer wonder why the offer is being targeted in the first place.

We show that, perhaps surprisingly, the firm’s profit-maximizing strategy, despite the ability to fine-tune offers across a continuum of consumer types, has a simple two-tier structure: a uniform positive discount for consumers with intermediate outside option values, and no discount for those with low or high outside option values. This coarse two-tier policy is profit-maximizing because it avoids unnecessary discounts for consumers with low outside options (who would return to buy even after searching) and unprofitable discounts for those with high outside options (who would require a rather high discount to stop searching). This finding challenges the usual intuition that richer consumer data would lead to more complex pricing strategies. It contributes to the pricing literature by demonstrating that simple pricing strategies can outperform more granular, data-intensive ones. It also helps explain why, in practice, firms often favor streamlined discounting schemes despite having access to extensive consumer data.

Equally striking are the impacts of superior information on market outcomes. In particular, we find that the firm’s ability to predict consumer outside options backfires: instead of boosting profits, having superior information leaves the firm worse off compared to the benchmark where it lacks such information. The reason is that personalized discounts influence the consumer’s inference: a discount signals that the consumer likely has a good outside option, which in turn encourages the consumer to search more. Anticipating this, the firm must either offer deeper discounts or allow more customers to walk away, both of which reduce its profit.

Moreover, even though consumers search more under the superior information regime, aggregate consumer surplus remains unchanged compared to the benchmark with symmetric information. Intuitively, the additional search benefits those consumers who under-searched in the benchmark but harms those who now search excessively. Since consumer surplus stays constant while firm profit declines, total welfare is lower under superior information. This efficiency loss is primarily driven by consumers who engage in excessive search in response to personalized discounts.

In the next section, we position our work in the context of the related literature before proceeding to the model and results.

Literature review

Our paper is related to several strands of research in economics and marketing. At its core, our paper builds upon consumer search models with product differentiation (Wolinsky, 1986; Anderson and Renault, 1999). Most relevant to our paper is the literature on consumer search deterrence. Extensive consumer search can intensify market competition, prompting firms to devise strategies that deter such behavior by increasing the cost of acquiring product information (Ellison and Wolitzky, 2012) or making exploding offers or offering buy-now discounts (Armstrong and Zhou, 2016). In particular, buy-now discount is widely used as a sales tactic allowing firms to price discriminate based on consumer search history.¹ Armstrong and Zhou (2016) show that such a sales technique can discourage consumers from searching for alternatives, resulting in sub-optimal product-to-consumer matches and higher prices paid by consumers. In contrast, our paper evaluates the effectiveness of such tactics when the firm knows more information about consumer preferences than themselves. We argue that with firm superior information, personalized offers can lead to consumer suspicion and then encourage consumer search, rather than discourage it.²

The key mechanism in our model is the signaling effect of personalized discounts offered by a firm that has better knowledge of market alternatives than consumers themselves. The implications of firms' superior information have been extensively explored across various domains, including sales assistance (Wernerfelt, 1994; Lee et al., 2024), product line design (Xu and Dukes, 2019), pricing (Li and Xu, 2022; Xu and Dukes, 2022), product recommendation (Dzyabura and Hauser, 2019) and advertising (Rafeian and Zuo, 2024). This body of work, however, largely addresses the

¹More broadly, there is a burgeoning literature on price discrimination in search markets. See, e.g., Preuss (2023), Groh (2024), and Mauring (2025).

²On page 50 of Armstrong and Zhou (2016), the authors mention the possibility that a seller might know a consumer's outside option: "When a seller's choice of buy-later policy can be made contingent on the outside option, a savvy buyer might then use the seller's policy as a signal of her outside option." However, they do not explore the equilibrium implications of the firm having superior information about the consumer's outside option.

context in which firms have superior information on consumers' valuations for their own products. Our model extends this framework by incorporating scenarios where firms also possess superior information about consumers' valuations of competing products (Wernerfelt, 1994), aligning more closely with recent developments in targeted advertising and personalized recommendations (Shin and Yu, 2021; Choi et al., 2024; Rafieian and Zuo, 2024). Shin and Yu (2021) models how targeted advertising utilizes consumer data to predict consumers' preferences not only for the firm's own products but also for competitors' products within the same category, emphasizing firms' ability to identify potential consumer interest before consumers themselves become aware. Also, Rafieian and Zuo (2024) and Choi et al. (2024) examine personalized recommendation systems, noting that algorithms have an advantage in accurately predicting consumers' preferences across various product options available on the platform because the vast number of options in online marketplaces imposes significant search costs, thereby effectively precluding consumers from fully understanding all available market alternatives.

Li and Xu (2022) is closest to our paper in studying the implications of superior knowledge on pricing strategies. In their model, the firm has private information about a consumer's valuation for its product and engages in personalized pricing. The consumer can infer her valuation from the offered price and, additionally, can incur a cost to inspect the product and learn her true valuation. Li and Xu (2022) also find that due to the signaling effect of personalized prices the firm may suffer from having superior information. Our model considers a broader definition of superior knowledge that includes not only a firm's own product but also its competitive products. More importantly, a fundamental distinction lies in the role of consumer learning: in their model product inspection can be beneficial to the firm by making it more likely for high-valuation consumers to purchase; in contrast, in our model consumer search for alternatives is always detrimental to the firm, as it only raises the chance that consumers will buy from elsewhere. This difference is crucial in explaining why the firm in our model always suffers from superior information, whereas in their model the firm can sometimes benefit. Also, unlike Li and Xu, who restrict their analysis to the case with binary consumer valuations and so binary discount levels, we study a more general setup with a continuum of valuations. This more generalized approach allows us to demonstrate that simplicity in pricing, employing a two-tier discount, can outperform more complex, fine-tuned discount strategies.

More broadly, our paper is also related to the literature on the informed principal problem (see, e.g., Myerson, 1983, and Maskin and Tirole, 1992 for seminar works). In this framework, a principal with private information offers a contract/mechanism to an agent, from which the agent can potentially infer the principal's private information. In our main model the seller acts as the

principal, as she has private information and commits to a selling mechanism which consists of a discounted buy-now price and a regular buy-later price. That literature, however, generally does not consider the possibility that the agent can costly verify the principal’s private information. It also typically does not examine the impact of private information relative to a setting where the principal is uninformed. A few exceptions include, for instance, Beaudry (1994), Silvers (2012), and Bedard (2017) which identify cases in various contexts—different from ours—where the principal may actually prefer to remain uninformed.

2 The Model

Our setup is built on the monopoly model in Armstrong and Zhou (2016). There is a monopoly firm (e.g., a local car dealership) in the market. Consumers value its product at u , which is assumed to be i.i.d. across consumers and follows a distribution on $[\underline{u}, \bar{u}]$ with a cumulative distribution function (CDF) $F(u)$. Consumers also face an outside option. The surplus of the outside option is denoted by v and it is also i.i.d. across consumers and follows a distribution on $[\underline{v}, \bar{v}]$ with CDF $G(v)$. The value of v is initially unknown to a consumer but can be discovered through search, which incurs a cost $s > 0$. We assume that (i) consumers cannot take the outside option without searching it first, (ii) they have free recall and so can return to buy the firm’s product after inspecting the outside option without paying any extra search cost, and (iii) u and v are independent of each other, $0 \leq \underline{v} < \bar{u} - p$, and $\mathbb{E}[v - \underline{v}] > s$ (i.e., the value of the outside option is high and dispersed enough relative to the search cost).

The firm sells its product at a regular list price p . This list price is assumed to be exogenous (e.g., it is a nationwide price predetermined by the company’s headquarter).³ However, when a consumer visits the firm, it can offer her a “buy-now” discount $\tau \in [0, p]$ aiming to incentivize the consumer to buy its product immediately without searching for the outside option. After receiving the offer, the consumer then decides whether to search for the outside option. If she chooses not to search, she purchases the firm’s product immediately at the discounted price $p - \tau$. Otherwise, she will incur the search cost and discover v , and then decide which option to take. Also, if the consumer returns to buy the firm’s product, the initial discount is no longer available and she must pay the regular price p . In other words, the firm has commitment power when offering a buy-now discount.

Both the firm and consumers are risk neutral. Consumers search optimally to maximize their

³As we will discuss in Section 5.3, considering an endogenous list price does not affect our main insights.

expected net surplus, and the firm sets the discount to maximize its expected profit based on its rational expectation about consumers’ search behavior. If offering no discount yields the same profit as offering a positive discount, we assume the tie-break rule that the firm prefers the former. Without loss of generality, we normalize the firm’s unit production cost to zero.

Two information regimes

During the sales process, sellers often gain detailed information about consumers’ preferences and needs, enabling them to assess a consumer’s match with their product (Wernerfelt, 1994; Rogers, 2013). We therefore assume throughout this paper that the firm observes the consumer’s match value with their product u . Modern data analytics and predictive algorithms further enhance sellers’ ability to predict not only a consumer’s preferences for their own products but also their potential match with competitors’ offerings. This may give firms “superior information” about consumers’ future search outcomes that consumers do not fully possess due to the significant effort required to explore alternative options and a lack of comprehensive knowledge about the supply side (Wernerfelt, 1994; Shin and Yu, 2021; Rafeian and Zuo, 2024).

In our model, we consider two information regimes, depending on whether the firm also observes consumers’ valuation for their outside option. In the benchmark case, the firm learns a consumer’s valuation u but remains unaware of her valuation v for the outside option. This is referred to as the “symmetric information” regime where the firm and the consumer possess the same amount of information. In this regime, the firm’s buy-now discount τ can depend only on u but not on v .

In the main case, the firm learns *both* u and v when a consumer visits. This is referred to as the “superior information” regime where the firm knows more about the consumer’s preferences than the consumer herself. This information advantage allows the firm to tailor “buy-now” discount based on both u and v . However, a rational consumer, recognizing this information asymmetry, may infer information about v from the discount offered and adjust her search strategy accordingly. This potential signaling channel will play a crucial role in our analysis. In particular, buy-now discounts may actually encourage consumer search, counteracting the firm’s intention to secure an immediate sale.

Timing of the game

As summarized on Figure 1 below, our game unfolds as follows: At the first stage, a consumer visits the firm, during which both she and the firm learn her valuation u for the firm’s product. In the regime of superior information, the firm also learns the consumer’s valuation v for the outside

option. At the second stage, the firm offers a personalized buy-now discount $\tau(u, v) \in [0, p]$ to the consumer. Note that in the regime of symmetric information, $\tau(u, v)$ must be invariant with respect to v . At the third stage, the consumer decides whether to buy the firm's product immediately at the discounted price, or continue to search for the outside option. If she chooses to search, at the fourth stage, she observes her v and decides whether to take the outside option or return to buy the firm's product at the regular price p .

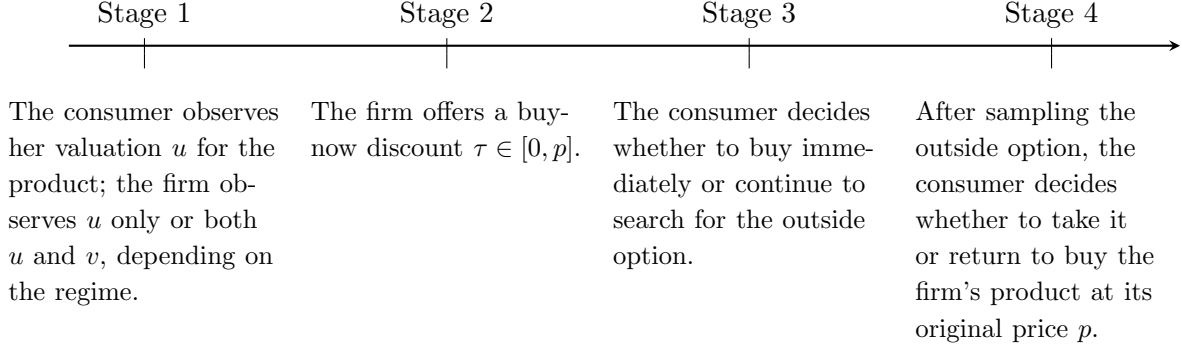


Figure 1: Timing of the Game

3 The Benchmark: Symmetric Information

We first study the benchmark case with symmetric information. The firm observes each consumer's u only and can offer them personalized buy-now discounts $\tau(u)$. The discounts do not convey any information about the consumers' outside options. Each consumer can be treated as a separate market. Consider a consumer who values the firm's product at u . When facing a buy-now discount $\tau(u)$, she obtains a surplus $u - p + \tau(u)$ if she buys the firm's product immediately; while if she chooses to search for the outside option, her expected surplus is $\mathbb{E}[\max\{v, u - p\}] - s$.⁴

Therefore, the consumer will stop searching and buy immediately if and only if

$$u - p + \tau(u) \geq \mathbb{E}[\max\{v, u - p\}] - s \iff \mathbb{E}[\max\{v - (u - p), 0\}] \leq s + \tau(u).$$

Therefore, the consumer's search problem with a buy-now discount can be rephrased as a standard search problem with an amplified search cost: the left-hand side is the standard expected benefit from sampling the outside option in the absence of the buy-now discount, and the right-hand side is the amplified search cost due to the buy-now discount.

⁴This surplus must be positive under the assumptions of $\underline{v} \geq 0$ and $\mathbb{E}[v - \underline{v}] > s$. This ensures that in our model the consumer will never leave the market without purchasing anything.

Using integration by parts, we can rewrite the above condition as

$$B(u - p) \equiv \mathbb{E}[\max\{v - (u - p), 0\}] = \int_{u-p}^{\bar{v}} [1 - G(v)] dv \leq s + \tau(u), \quad (1)$$

where the function

$$B(x) \equiv \int_x^{\bar{v}} [1 - G(v)] dv,$$

represents the expected benefit of sampling the outside option in the standard search problem when the current available surplus is x . We define the reservation value r in the optimal stopping rule of the standard search problem as the solution to

$$B(r) = s. \quad (2)$$

This has a unique solution $r \in (\underline{v}, \bar{v})$ under the assumption $\mathbb{E}[v - \underline{v}] > s$.

To solve the firm's optimal pricing strategy, we need to consider two distinct cases: First, for a consumer with $u - p \geq r$ or equivalently $B(u - p) \leq s$, condition (1) holds and so she will purchase immediately even if $\tau(u) = 0$. Therefore, for such consumers, the firm's optimal pricing strategy is to offer no discounts.

Second, for a consumer with $u - p < r$, she will continue to search unless a sufficiently high discount is offered such that condition (1) holds. From condition (1), we can see that in this benchmark regime the discount

$$\tau^b(u) \equiv B(u - p) - s \quad (3)$$

makes the consumer just willing to stop searching and buy immediately. As expected, when the search cost increases, it gets easier to deter consumer search, so this discount decreases. Offering this discount to the consumer, the firm earns a profit

$$p - \tau^b(u). \quad (4)$$

On the other hand, if no discount is offered, the consumer will continue to search but eventually return to purchase at the regular price p if $v < u - p$, which occurs with probability $G(u - p)$. Thus, offering no discount yields a profit

$$p \cdot G(u - p). \quad (5)$$

Therefore, for a consumer with $u - p < \tau$, the firm prefers to offer the buy-now discount $\tau^b(u)$ to deter search if and only if

$$p - \tau^b(u) > p \cdot G(u - p) \iff p > \frac{B(u - p) - s}{1 - G(u - p)}. \quad (6)$$

Note that the right-hand side equals zero at $u - p = r$, so this condition must hold if $u - p$ is sufficiently close to r . We then have the following result (all the omitted proofs can be found in the Appendix):

Proposition 1. *In the benchmark case with symmetric information:*

- (1) *For a consumer with $u - p \geq r$, the firm offers no buy-now discount, and she purchases immediately without further search. The firm earns a profit $\Pi^b(u) = p$.*
- (2) *For a consumer with $u - p < r$, if condition (6) holds, the firm offers the buy-now discount $\tau^b(u)$ as defined in (3), and she purchases immediately without further search. The firm earns a profit $\Pi^b(u) = p - \tau^b(u)$.*
- (3) *For a consumer with $u - p < r$, if condition (6) does not hold, the firm offers no buy-now discount, and she continues to search. The firm earns a profit $\Pi^b(u) = p \cdot G(u - p)$.*

Moreover, if $1 - G(v)$ is log-concave, then condition (6) holds if and only if $u - p \in (k, r)$ where $k < r$ is the $u - p$ that uniquely solves (6).

Note that $1 - G(v)$ is log-concave for many often used distributions such as uniform, exponential, and (truncated) normal. Figure 2 illustrates the firm's optimal buy-now discount scheme and its optimal profit when $G(v)$ is a uniform distribution on $[0, 1]$. When $u - p > r$ (in which case there is no need to offer discounts) or $u - p < k$ (in which case it is too costly to use discounts to deter consumer search), the firm offers a zero discount. In between, a discount is offered and it decreases in $u - p$, meaning that a higher discount is needed to deter search when the consumer has a lower valuation for the firm's product.

4 The Main Case: Superior Information

We now turn to the regime of superior information where the firm knows consumers' outside options. The firm's buy-now discount $\tau(u, v)$ can now depend on v as well. We aim to investigate how this more refined price discrimination affects profit and consumer welfare.

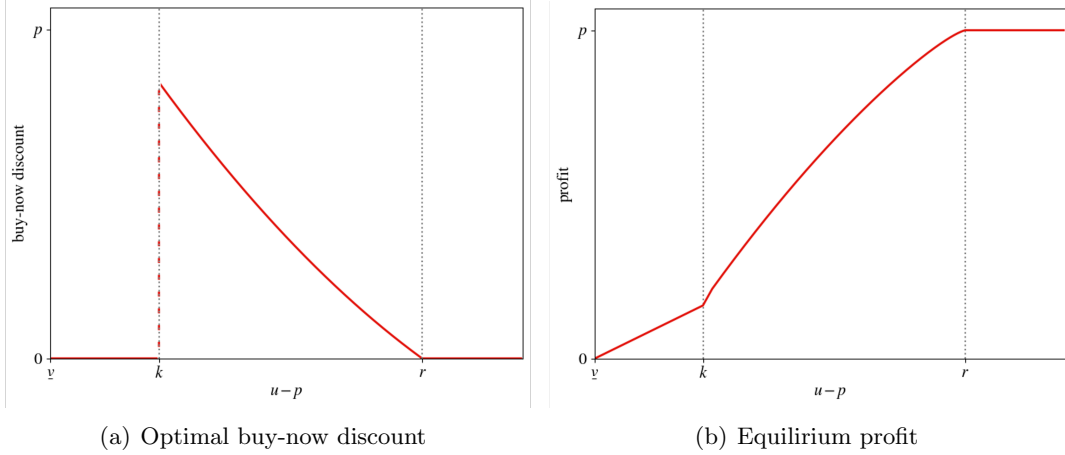


Figure 2: The benchmark with symmetric information
 $(v \sim U[0, 1], p = 0.3, s = 0.1)$

Naive consumers

It is helpful to start with the case when consumers are naive and do not attempt to infer information about their outside options from the firm's discount offers. Then, following the analysis in the benchmark case, the firm will offer $\tau(u, v) = 0$ to a consumer with $u - p \geq r$. For a consumer with $u - p < r$, the optimal discount depends on her v . If $v < u - p$, she will return to purchase even if she continues to search first; hence, the firm offers a zero discount $\tau(u, v) = 0$. If $v \geq u - p$, she will never return once she leaves the firm; therefore, the firm offers the discount $\tau(u, v) = \tau^b(u)$ defined in the benchmark case to deter search, whenever this discount does not exceed p .

This outcome, however, cannot be sustained in equilibrium when consumers are rational. If they understand that the firm employs the above pricing strategy, consumers will interpret $\tau = 0$ as a signal of $v < u - p$, and will therefore choose not to search. Consequently, the firm will have an incentive to offer a zero discount even to consumers with $v \geq u - p$. On the other hand, consumers will interpret the positive discount $\tau^b(u)$ as a signal of $v \geq u - p$, and will then find searching more attractive than in the benchmark case. As a result, the buy-now discount $\tau^b(u)$ will no longer be sufficient to deter search. This suggests that offering a buy-now discount can potentially *encourage* consumer search rather than prevent it.

Strategic consumers

Suppose from now on that consumers are strategic and fully rational. They may then interpret the buy-now discount as a potential signal about their outside option v .

We allow both the firm and the consumer to play mixed strategies. The firm's (potentially mixed) pricing strategy for a consumer of type (u, v) is denoted by a probability density function $\mu(\tau; u, v)$ for $\tau \in [0, p]$, which should satisfy $\int_0^p \mu(\tau; u, v) d\tau = 1$ for any (u, v) . For a given u , let $\Omega(\mu|u)$ denote the support of μ across all possible v 's, i.e., the set of all possible discount levels offered to a consumer of type u with a positive density under the pricing strategy μ . Whenever there is no confusion, for a given consumer we use the simpler notation $\Omega(\mu)$ by dropping the conditional variable u . Let $\tau_0 = 0$ denote the zero discount. As we will see below, $\tau_0 \in \Omega(\mu)$ in any equilibrium. Let $\sigma(\tau; u)$ denote the probability that a consumer of type u will continue to search for the outside option after receiving a discount τ .

When the firm uses pricing strategy μ and consumers follow search strategy σ , the firm's expected profit from a consumer of type (u, v) is:

$$\pi(\mu; \sigma, u, v) = \int_0^p [(p - \tau) (1 - \sigma(\tau; u)) + p \sigma(\tau; u) \mathbb{I}_{v < u - p}] \mu(\tau; u, v) d\tau, \quad (7)$$

where \mathbb{I} is an indicator function. If the consumer buys immediately after seeing a discount τ , the firm makes a profit $p - \tau$; if she chooses to search, she will return to buy at the regular price p when the outside option v turns out to be worse than $u - p$.

When a consumer of type u believes that the firm is using strategy μ and receives a discount $\tau \in \Omega(\mu)$, her decision of whether to search for the outside option depends on her posterior belief about v . Let

$$g(v; u, \tau) = \frac{\mu(\tau; u, v)g(v)}{\int_{\underline{v}}^{\bar{v}} \mu(\tau; u, v) dG(v)} \quad (8)$$

be the posterior density function after seeing discount τ whenever the denominator is strictly positive. Let $G(v; u, \tau)$ be the CDF of $g(v; u, \tau)$. Then, following a similar logic as in the benchmark analysis, the consumer will prefer to search if

$$B(u - p; \tau) \equiv \int_{u-p}^{\bar{v}} [1 - G(v; u, \tau)] dv > \tau + s \quad (9)$$

where $B(u - p; \tau)$ represents the expected benefit from searching for the outside option, if the current available surplus is $u - p$, under the updated belief after receiving discount τ . Conversely, the consumer will prefer not to search if $B(u - p; \tau) < \tau + s$. She may adopt a mixed search strategy if $B(u - p; \tau) = \tau + s$.

4.1 Equilibrium analysis

We use the solution concept of Perfect Bayesian Equilibrium (PBE) to solve the game. Let μ^* and σ^* denote the firm's equilibrium pricing strategy and the consumer's equilibrium search strategy, respectively. Since both the firm and consumers observe u , we can treat each consumer with a different u as a separate market. Depending on consumer valuation, we distinguish between two separate cases – starting with a simpler case, where consumers have sufficiently high valuations for the firm's product that no discount is necessary.

The case of $u - p \geq r$

For such a consumer with $u - p \geq r$, there is a pooling equilibrium in which the firm does not offer her any discount regardless of her v . If the consumer expects the firm to use this pricing strategy, her posterior belief of v remains as the prior after seeing no discount, so she will stop searching and buy immediately given $u - p \geq r$. The firm's profit from this consumer is thus p , the maximum possible profit. Therefore, the firm has no strict incentive to deviate and offer any discounts, regardless of the consumer's off-equilibrium beliefs. There can exist other equilibria, but none of them can generate a higher profit.

Lemma 1. *In the market for a consumer with $u - p \geq r$, the most profitable equilibrium is a pooling equilibrium in which the firm offers no discount (i.e., $\mu^*(\tau_0; u, v) = 1$ for any v) and the consumer does not search for the outside option (i.e., $\sigma^*(\tau_0; u) = 0$).*

The case of $u - p < r$

In the following, we analyze the more interesting case when a consumer has $u - p < r$. The simple observation is that if the firm observes that the consumer has an outside option $v < u - p$, it will offer her a zero discount in any equilibrium (i.e., $\mu^*(\tau_0; u, v) = 1$ for any $v < u - p$). Such a consumer will always return to buy the firm's product at the regular price even if she first chooses to search for the outside option. Therefore, offering no discount is a (weakly) dominant strategy.

This observation implies that in the market for a consumer with $u - p < r$, the only possible pooling equilibrium is that the firm offers no discount regardless of her v . The consumer then always chooses to search but will return to buy if $v < u - p$. The firm's profit will be $p \cdot G(u - p)$. Even if this pooling equilibrium exists, it is often dominated in terms of profit by some non-pooling equilibrium we will construct below.

We now consider non-pooling equilibrium. The following result reports several necessary properties of a non-pooling equilibrium:

Lemma 2. *In the market for a consumer with $u - p < r$, if a non-pooling equilibrium (μ^*, σ^*) exists where $\Omega(\mu^*)$ includes some $\tau \in (0, p)$, it must satisfy:*

- (1) $\mu^*(\tau_0; v, u) = 1$ for $v < u - p$;
- (2) $\mu^*(\tau_0; v, u) > 0$ for a positive measure of $v \geq u - p$;
- (3) $0 < \sigma^*(\tau_0; u) < 1$, and $B(u - p; \tau_0) = s$;
- (4) $\sigma^*(\tau; u) < 1$ for any positive $\tau \in \Omega(\mu^*)$, and $B(u - p; \tau) \leq \tau + s$;
- (5) $(p - \tau)(1 - \sigma^*(\tau; u))$ is a constant for any $\tau \in \Omega(\mu^*)$.

Here we offer a sketch of the proof and basic intuition for these results. Result (1) has been explained earlier, i.e., the firm offers a zero discount to a consumer with $v < u - p$. Result (2) implies that in a non-pooling equilibrium, in addition to the consumers with $v < u - p$, the firm must also sometimes offer a zero discount to certain consumers with $v \geq u - p$. Otherwise, a zero discount would be a perfect signal of $v < u - p$ and then consumers would always stop searching after receiving no discounts; as a result, the firm would never want to offer a positive discount.

Result (3) states that a consumer must adopt a mixed search strategy after receiving no discount. She cannot always stop searching as explained above, nor can she always continue to search because otherwise offering no discount yields a zero profit from her if $v > u - p$. Result (4) implies that a consumer should weakly prefer not to search after seeing a positive discount; otherwise, the firm would prefer a zero discount given our tie-breaking rule for the firm.

Note that $(p - \tau)(1 - \sigma^*(\tau; u))$ is the equilibrium profit from a consumer with $v \geq u - p$ when a discount τ is offered. Result (5) implies that this profit must be the same for any discounts on the equilibrium path. Otherwise, one of them would be offered all the time regardless of v , in which case either a pooling equilibrium would arise (if the best discount is zero), or the firm would never offer a zero discount for $v \geq u - p$ (if the best discount is positive), which violates result (2).

In any non-pooling equilibrium, the firm's expected profit from a consumer of type u is:

$$pG(u - p) + (p - \tau)(1 - \sigma^*(\tau; u))[1 - G(u - p)], \quad (10)$$

where τ is an equilibrium discount in $\Omega(\mu^*)$. When $v < u - p$, the firm offers no discount and the consumer always buys its product at the regular price (even after searching for the outside

option). This explains the first term. When $v \geq u - p$, no matter what equilibrium discount is offered, the firm should make the same profit according to result (5) in Lemma 2. This explains the second term. Whenever a non-pooling equilibrium exists, it generates a higher profit than a pooling equilibrium (where as explained before, the firm makes a profit $pG(u - p)$).

Even under the requirements in Lemma 2, our signaling game can admit many possible non-pooling equilibria, ranging from simple binary-discount schemes to more complex multi-level schemes (as illustrated in Appendix B). Substantial indeterminacy remains even within the class of binary schemes. For example, a binary scheme could entail offering a positive discount only when v exceeds a threshold, when v lies in an intermediate interval, or when v falls within several disconnected intervals. In some regions, the firm may even randomize between offering no discount and a positive one. Given the plethora of potential equilibria, we focus on the one that is the most profitable for the firm. This selection only serves to strengthen our main result that, relative to the benchmark case with symmetric information, the firm suffers from having superior information about consumers' outside options.

In the following, we will first construct the most profitable “binary-discount” equilibrium (if it exists) where the firm offers only two possible discounts, i.e., $\Omega(\mu^*) = \{\tau_0, \tau_1 \in (0, p)\}$. (Recall that result (1) in Lemma 2 implies that τ_0 must be in $\Omega(\mu^*)$.) We will then show that the constructed binary-discount equilibrium is also the most profitable among all possible non-pooling equilibria.

Binary-discount equilibrium We first examine a type of equilibrium in which there are only two levels of buy-now discounts, $\tau_0 = 0$ and $\tau_1 > 0$. In such a binary-discount equilibrium, the firm's pricing strategy is characterized by two variables: $\mu(\tau_0; u, v)$, the probability that the firm offers a zero discount to a consumer of type (u, v) , and τ_1 , the positive discount.

The following lemma reports an equilibrium where the firm offers a positive discount for intermediate $v \in [u - p, \hat{v}]$, where $\hat{v} \in (u - p, \bar{v})$, and a zero discount for other low or high v 's. This equilibrium, if exists, turns out to be the most profitable for the firm.

Let \hat{v} solve

$$\int_{\hat{v}}^{\bar{v}} (v - u + p - s) dG(v) = s \cdot G(u - p). \quad (11)$$

For any $u - p < r$, one can show that the above equation has a unique solution $\hat{v} \in (u - p, \bar{v})$.⁵

⁵The left-hand side of (11) is zero at $\hat{v} = \bar{v}$, smaller than the right-hand side; it equals $B(u - p) - s[1 - G(u - p)]$ at $\hat{v} = u - p$, greater than the right-hand side as $B(u - p) > s$ for $u - p < r$. Hence, (11) must have a solution $\hat{v} \in (u - p, \bar{v})$. To see the uniqueness, notice that the derivative of the left-hand side with respect to \hat{v} is $-[\hat{v} - u + p - s]g(\hat{v})$, which is positive at $\hat{v} = u - p$ and changes its sign at most once as \hat{v} increases. Therefore, the left-hand side of (11) must

Then, we define

$$\tau_1^* = \frac{B(u-p) - s}{G(\hat{v}) - G(u-p)}, \quad (12)$$

which must be positive given $B(u-p) > s$ for $u-p < r$.

Lemma 3 (Optimal binary-discount equilibrium). *In the market for a consumer with $u-p < r$, if τ_1^* defined in (12) is less than p , the following (μ^*, σ^*) is the most profitable binary-discount equilibrium for the firm where $\Omega(\mu^*) = \{\tau_0, \tau_1^*\}$:*

$$(1) \mu^*(\tau_0; u, v) = 1 \text{ if } v < u-p \text{ or } v > \hat{v} \text{ and } \mu^*(\tau_0; u, v) = 0 \text{ if } v \in [u-p, \hat{v}],$$

$$(2) \sigma^*(\tau_0; u) = \frac{\tau_1^*}{p} \text{ and } \sigma^*(\tau_1^*; u) = 0,$$

where \hat{v} is defined in (11). If $1 - G(v)$ is log-concave and s is sufficiently small, $\tau_1^* < p$ if and only if $u-p \in [\theta, r]$ where $\theta < r$ is the $u-p$ that uniquely solves $\tau_1^* = p$.

In this equilibrium, a positive discount is offered to a consumer with $u-p < r$ if and only if her v is intermediate. The consumer will stop searching after receiving the positive discount but will sometimes choose to search if no discount is offered.⁶

To sustain any non-pooling equilibrium, the firm must ensure that consumers who receive no discount remain uncertain about the value of their outside option. If consumers are certain that receiving no discount signals a low outside option, they would immediately buy without searching—prompting the firm to never offer positive discounts in the first place. To preserve this necessary uncertainty, the firm must occasionally withhold discounts even from consumers with relatively high outside values ($v > u-p$), allowing some of them to search. This approach inevitably entails potential losses: high- v consumers might discover attractive alternatives and choose not to return. To minimize this loss — specifically, the number of consumers who search and ultimately leave — the firm finds it optimal to withhold discounts only from those with the highest outside values ($v > \hat{v}$). By selectively withholding discounts only from these highest-value consumers, the firm can induce the desired consumer beliefs with minimal sacrifice. This occurs because even a small probability assigned to very high outside options can substantially raise the consumer's expectation about the benefit of searching, quickly achieving the necessary level of uncertainty. In effect, by

be single-peaked in the range of $\hat{v} \in [u-p, \bar{v}]$, implying the solution uniqueness. Notice also that the left-hand side must be decreasing at the solution \hat{v} , which implies $\hat{v} > u-p+s$.

⁶Note that $\tau_1^* \rightarrow 0$ as $u-p \rightarrow r$, and so $\tau_1^* < p$ must hold for $u-p$ is sufficiently close to r . However, compared to the benchmark regime with symmetric information, it is now harder to find a simple cutoff condition on $u-p$ for $\tau_1^* < p$. This is because \hat{v} decreases in $u-p$ and so it is harder for (12) to decrease in $u-p$ compared to the right-hand side of (6). We can derive a cutoff condition only when the search cost is small.

concentrating zero-discount offers on both ends of the v distribution (low and very high), the firm reduces the frequency and impact of costly search outcomes.

In the proof of Lemma 3, we first strengthen the result (4) in Lemma 2, which is critical for establishing the firm-optimal binary discount equilibrium. This enhanced result (Claim 1 in the proof) specifies that in the firm-optimal binary discount equilibrium, τ_1 must be such that the consumer is indifferent between buying immediately and continuing to search (i.e., $B(u - p; \tau_1) = \tau_1 + s$), but she chooses not to search (i.e., $\sigma(\tau_1; u) = 0$). Then, the profit equation in (10) and result (5) in Lemma 2 imply that the firm's profit from the consumer must be $p \cdot G(u - p) + (p - \tau_1)[1 - G(u - p)]$. We then choose the discount scheme to maximize this profit subject to the constraint $B(u - p; \tau_0) = s$ from result (3) in Lemma 2. Intuitively, to prevent consumer search, τ_1 needs to be higher when the posterior of v conditional on the discount τ_1 is higher. Therefore, to maximize profit, we should minimize the posterior of v conditional on τ_1 . Since the constraint $B(u - p; \tau_0) = s$ requires enough weight on $v > u - p$ where the firm offers no discounts, the best way is then to offer the discount when $v \geq u - p$ is below a certain threshold \hat{v} . (Recall that no discounts will be offered when $v < u - p$.) It is then straightforward to pin down \hat{v} from the constraint $B(u - p; \tau_0) = s$ and τ_1 from the indifference condition $B(u - p; \tau_1) = \tau_1 + s$. Finally, we show that the firm has no incentive to offer off-equilibrium discounts under an appropriate off-equilibrium belief.

Specifically, we adopt the following off-equilibrium belief:⁷ upon receiving an unexpected discount $0 < \tau < \tau_1^*$,⁸ a consumer of type u holds a posterior of v which has the following density function

$$g(v; u, \tau) = \begin{cases} \frac{g(v)}{1 - G(u - p)}, & \text{if } v \geq u - p, \\ 0, & \text{otherwise.} \end{cases} \quad (13)$$

That is, the off-equilibrium discount induces the consumer to believe that v follows a truncated distribution on $[u - p, \bar{v}]$ since any $v < u - p$ would have led to a zero discount. Given this off-equilibrium belief, we show in the proof that the firm has no incentive to deviate to a discount lower than τ_1^* .

⁷Many other off-equilibrium beliefs will also work for our argument. Note that our off-equilibrium belief used here survives both the Intuitive Criterion and the Divinity Criterion. For any $v \geq u - p$, a deviation to $0 < \tau < \tau_1^*$ would be profitable if the consumer would stop searching (which is an undominated consumer strategy). Therefore, no $v \geq u - p$ will be assigned a zero probability according to the Intuitive Criterion. Meanwhile, since both the equilibrium profit and the deviation profit are independent of v , if a deviation is profitable for v_1 under a certain consumer reaction, it should be also profitable for $v_2 \neq v_1$ under the same consumer reaction. Therefore, the Divinity Criterion has no selection power here.

⁸As shown in the proof, it is never profitable for the firm to offer a deviation discount above τ_1^* given the consumer is already willing to stop searching after seeing τ_1^* .

From equation (10), we can see that in the binary-discount equilibrium in Lemma 3 the firm makes a profit:

$$\Pi^*(u) = pG(u - p) + (p - \tau_1^*)[1 - G(u - p)] = p - \tau_1^*[1 - G(u - p)]. \quad (14)$$

Optimality of binary-discount equilibrium We now consider a more general case where the pricing strategy may include more than two discount levels. We show that the binary-discount equilibrium in Lemma 3 (whenever it exists) remains the most profitable among *all* possible non-pooling equilibria.

Lemma 4 (Optimality of binary-discount equilibrium). *In the market for a consumer with $u - p < r$, if there is a non-pooling equilibrium $(\tilde{\mu}, \tilde{\sigma})$ with at least three discount levels (i.e., if $|\Omega(\tilde{\mu})| \geq 3$), there must exist a binary-discount equilibrium which generates a higher profit.*

This lemma establishes one of our main results that even when multiple discount levels are permitted, a simple binary-discount strategy remains to be optimal for the firm. The basic idea of the proof is simple. Whenever there is an equilibrium with at least three discount levels, a more profitable binary-discount equilibrium can be structured. In this equilibrium, the firm offers a fixed discount across all cases where it previously offered varying positive discounts. We show that we can always set this new constant discount such that it is lower than the highest discount level in the original equilibrium, ensuring it induces consumers to just stop searching. This new discount should therefore generate a higher profit than the highest discount in the original equilibrium. According to result (5) in Lemma 2, any discount (including the highest one) in the original equilibrium should generate the same profit, and so our result follows.

Thus far, our equilibrium analysis has focused on the case where the constructed optimal binary-discount equilibrium exists, which requires $\tau_1^* < p$. To complete the analysis, we also need to deal with the case where $\tau_1^* \geq p$. We now turn to this remaining case.

Lemma 5. *In the market for a consumer with $u - p < r$, if τ_1^* defined in (12) exceeds p , the only equilibrium is a pooling equilibrium in which the firm offers no discounts (i.e., $\mu^*(\tau_0; u, v) = 1$ for any v) and the consumer continues to search for the outside option (i.e., $\sigma^*(\tau_0; u) = 1$).*

The lemma shows that if the firm is compelled to offer discount $\tau_1^* \geq p$ to deter consumer search, a pooling equilibrium where the firm uniformly offers a zero discount across all v is the unique equilibrium. We first show the existence of such a pooling equilibrium by using the same off-equilibrium belief in (13). We then show that there are no non-pooling equilibria in this case.

If a non-pooling equilibrium with at least three discount levels exists, according to Lemma 4, there must exist a binary-discount equilibrium. If a binary-discount equilibrium exists, we can also show that there must exist a (weakly) more profitable binary-discount equilibrium as characterized in Lemma 3 with $\tau_1^* < p$. This leads to a contradiction, confirming the non-existence of non-pooling equilibrium.

With Lemmas 1-5, we can conclude:

Proposition 2 (Equilibrium with superior information). *In the main case with superior information, the firm-optimal equilibrium involves a simple binary discount strategy. Specifically,*

- (1) *For a consumer with $u - p \geq r$, the firm offers no buy-now discounts regardless of her v , and she purchases immediately without further search. The firm earns a profit $\Pi^*(u) = p$.*
- (2) *For a consumer with $u - p < r$, if the required discount to deter search (i.e., τ_1^* defined in (12)) is less than p , in the most profitable equilibrium the firm offers binary discounts as characterized in Lemma 3. The firm earns a profit $\Pi^*(u)$ defined in (14).*
- (3) *For a consumer with $u - p < r$, if τ_1^* exceeds p , the firm offers no buy-now discounts regardless of her v , and the consumer continues to search. The firm earns a profit $\Pi^*(u) = p \cdot G(u - p)$.*

Proposition 2 highlights our main finding that despite the firm's ability to finely tailor prices, the optimal strategy under superior information boils down to a simple binary discount scheme. In equilibrium, only consumers with intermediate outside option values receive a discount. This coarse two-tier policy is profit-maximizing because it influences consumer behavior where it matters, while avoiding unnecessary discounts for consumers with low outside options (who would buy anyway) and unprofitable discounts for those with high outside options (who would require excessively large incentives to convert).

Figure 4(a) below depicts the equilibrium when G is a uniform distribution on $[0, 1]$. In this case, it can be shown that $\tau_1^* < p$ if and only if $u - p \in [\theta, r]$. The shaded area in the figure, bounded by \hat{v} defined in (11), represents the range of consumer types (u, v) to whom the firm offers a positive discount τ_1^* . Notably, the intensity of the shading indicates the magnitude of the discount offered, with darker shades corresponding to higher discounts, showing that the discount decreases in u but is independent of v . Furthermore, as u increases, \hat{v} concurrently decreases, narrowing the range over which discounts are applied, but the discount offered to consumers also decreases.

5 The Impact of Superior Information

When the firm possesses superior information, it can further tailor buy-now discounts based on consumers' outside options. However, this more refined price discrimination does not necessarily help deter consumer search or increase the firm's profit. Since the firm never offers a discount when $v < u - p$, offering a discount then signals a relatively high v , which may encourage consumer search and ultimately harm the firm. In this section, we examine the impacts of superior information on both the firm and consumers by comparing the two studied information regimes.

5.1 Search and profit

We first compare consumer search and profit by considering three different types of consumers in turn. First, for a consumer with $u - p \geq r$, the firm offers no discount in either regime and the consumer does not search, so the firm makes the same profit p .

Second, for a consumer with $u - p < r$ but $\tau_1^* \geq p$, in the superior information regime the firm offers no discount in the pooling equilibrium, prompting the consumer to search for the outside option. As a result, the firm makes a profit $pG(u - p)$. In the benchmark regime, if condition (6) holds, the firm offers a discount that prevents the consumer from searching, yielding a profit $p - \tau^b > pG(u - p)$. Otherwise, the firm offers no discount, the consumer continues to search, and the firm earns a profit $pG(u - p)$, the same as in the superior information regime.

Third, consider a consumer with $u - p < r$ and $\tau_1^* < p$. Notice first that

$$\tau_1^* = \frac{B(u - p) - s}{G(\hat{v}) - G(u - p)} > \frac{B(u - p) - s}{1 - G(u - p)}.$$

Therefore, whenever $\tau_1^* < p$, in which case the binary-discount equilibrium arises in the superior information regime, condition (6) must hold and so the firm offers a buy-now discount and the consumer buys immediately in the benchmark regime. Note that in the former case the consumer sometimes searches while she does not in the latter case, so the consumer searches more in the superior information regime. According to (14), the profit in the binary-discount equilibrium is

$$\begin{aligned} \Pi^*(u) &= p - \tau_1^*[1 - G(u - p)] = p - \frac{B(u - p) - s}{G(\hat{v}) - G(u - p)}[1 - G(u - p)] \\ &< \Pi^b(u) = p - \tau^b = p - [B(u - p) - s]. \end{aligned}$$

That is, when the binary-discount equilibrium arises in the market for a consumer with $u - p < r$,

the firm makes lower profit than in the benchmark case. Compared to the benchmark, the superior information regime has different impacts on profit, depending on the consumer's outside option as described in Figure 3: (i) When the consumer has $v < u - p$, the firm offers no discount and makes higher profit, as the consumer always buys at the regular price regardless of whether she searches or not. (ii) When the consumer has $v \in [u - p, \hat{v}]$, the firm offers a larger discount $\tau_1^* > \tau^b$ to deter her search, since the consumer makes a favorable inference on v upon receiving a discount—resulting in lower profit for the firm. (iii) When the consumer has $v > \hat{v}$, the firm also offers no discount, but outcomes vary: if the consumer does not search and then buys at the regular price, the firm benefits; if the consumer searches, the firm loses as she never returns.⁹ Taking together, the negative effects (deeper discounts in the intermediate range and lost sales at the high end) outweigh the positive effects (full-price purchases at the low end and sometimes at the high end), leading the firm to be worse off under superior information.




Figure 3: The impact of superior information on profit

The following result then follows:

Proposition 3. *Consumers search more and the firm makes a lower profit in the superior information regime than in the symmetric information regime.*

Proposition 3 highlights a counterintuitive outcome: having more consumer information and making more personalized offers actually make the firm worse off and induce consumers to search more. In the superior information regime, targeted discounts serve as signals that the consumer's outside option may be attractive. Anticipating this consumer inference, the firm must offer a larger discount ($\tau_1^* > \tau^b$) to deter search among intermediate v consumers, thereby reducing its profit margin on discounted sales. At the same time, the firm offers discounts less frequently—only when $v \in [u - p, \hat{v}]$ —which risks losing customers when no discount is given, as some consumers choose to search and do not return if $v > \hat{v}$. Indeed, consumers search more in the superior information

⁹In this last case, the loss from less sales dominates: the firm's profit is $p(1 - \sigma^*(\tau_0; u)) = p - \tau_1^*$, which is lower than $< p - \tau^b$ in the benchmark.

regime, precisely the opposite of the firm's intention. Thus, these two effects—deeper discounts where applied and more lost customers where discounts are withheld—jointly lead to lower profit under superior information, as illustrated by Figure 3.

The black and the red curve in Figure 4(b) illustrate this result by comparing profit across all possible $u - p$ when $G(v)$ is a uniform distribution. The green curve depicts the profit in the superior information regime when consumers are naive. As expected, without consumer strategic inference, the profit should be the higher when the firm can do more refined price discrimination. Our profit comparison suggests that it is profitable for a firm to acquire information on consumers' outside options only when there are enough naive consumers in the market.

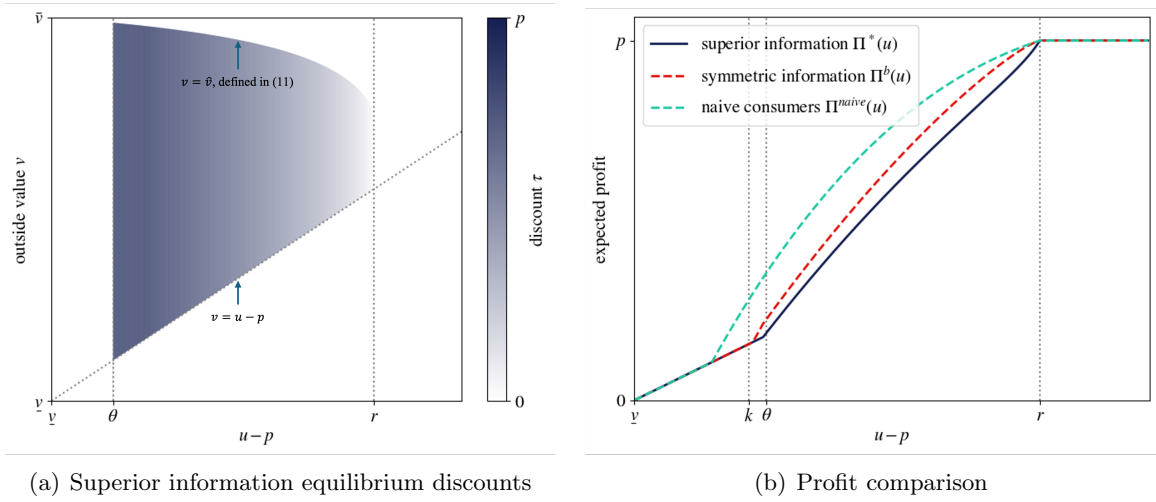


Figure 4: The main case with superior information
 $(v \sim U[0, 1], p = 0.3, s = 0.1)$

5.2 Consumer and total welfare

We now compare consumer and total welfare. The first type of consumers with $u - p > r$ obviously obtain the same surplus $u - p$ between the two regimes, as they receive no discounts and buy immediately in either case. The second type of consumers with $u - p < r$ but $\tau_1^* \geq p$ obtain the same surplus $u - p + B(u - p) - s$, as $\tau^b = B(u - p) - s$ in the benchmark is such that consumers are indifferent between buying immediately and continuing to search.¹⁰ The third type with $u - p < r$ and $\tau_1^* < p$ actually also obtain the same surplus between the two regimes. To understand that, note that their surplus is $u - p + B(u - p) - s$ in the benchmark case as they are indifferent buying

¹⁰Note that in either regime the expected surplus when a consumer chooses to search for the outside option is $\mathbb{E}[\max\{v, u - p\}] - s = u - p + \mathbb{E}[\max\{v - (u - p), 0\}] - s = u - p + B(u - p) - s$.

immediately and searching, while it is

$$[G(u - p) + 1 - G(\hat{v})](u - p) + [G(\hat{v}) - G(u - p)](u - p + \tau_1^*) = u - p + B(u - p) - s$$

in the superior information regime, where the equality used the expression for τ_1^* in (12). The consumers with $v < u - p$ or $v > \hat{v}$ receive no discounts and are indifferent between buying immediately and searching; those with $v \in [u - p, \hat{v}]$ receive a discount τ_1^* which is such that they are indifferent between buying immediately and searching. (Note that although the aggregate consumer surplus remains unchanged, the former group actually becomes worse off than in the benchmark while the latter group becomes better off given $\tau_1^* > \tau^b$.)

We then immediately have the following result:

Proposition 4. *Aggregate consumer surplus is the same in both the symmetric information and the superior information regimes. However, total welfare is lower in the superior information regime.*

To understand the intuition for the total welfare result, note first that in the socially optimal scenario a consumer should search if and only if $v > u + s$ given price is a pure transfer in our model. Consider now, for example, the third type of consumers with $u - p < r$ and $\tau_1^* < p$. They receive a discount and do not search in the benchmark. In the superior information regime, those with $v \in [u - p, \hat{v}]$ do not search either, so generate the same total welfare. For those with $v < u - p$, they should not search as $v < u + s$, but now they search with a certain probability, which leads to excessive search and an efficiency loss relative to the benchmark. For those with $v > \max\{\hat{v}, u + s\}$, they should search, and now they search more than in the benchmark and so yield an efficiency gain. Overall, the loss for unnecessary search from low- v consumers dominates and total welfare decreases.

5.3 Discussion: endogenous list prices

So far, we have assumed that the regular list price p is predetermined outside the model. This assumption is appropriate in many retail settings where salespeople do not control the regular price but have the discretion to offer discounts to customers to finalize transactions. The regular list price p is typically set by higher-level management, often at the national level. Conceptually it is not hard to endogenize the list price in our model: for each given $u - p$, we have characterized the equilibrium and the firm's profit; the optimal list price p in each regime should then maximize the profit integrated over all possible u . However, given the generality of our model, we do not have a

sharp characterization of the optimal p .¹¹

Nevertheless, the firm still suffers from having superior information even when the list price is endogenously determined. To see this, recall that for each given $u - p$ we have shown that the firm makes the same or strictly lower profit under the superior information regime. Consequently, for any fixed list price p , the firm's aggregate profit across all possible u must also be lower in the superior information regime, implying that the firm remains worse off even when it can optimally set its list price. On the other hand, allowing for endogenous list prices can influence our results concerning consumer and total welfare, but drawing general conclusions remains challenging. The impact depends on the details of the optimal list price in each regime.

6 Conclusion

Firms increasingly utilize detailed information about consumer preferences—both for their own products and competing alternatives—to tailor discount strategies and discourage consumer search. In a sequential search model, we have shown that when a firm possesses superior information about consumers' outside options, discounts can inadvertently signal attractive alternatives, thereby encouraging consumer search rather than deterring it. As a result of this signaling effect, the firm ultimately suffers from having more information and the ability to engage in more refined price discrimination, and total welfare declines as well.

Moreover, we have also demonstrate that, despite the firm's ability to tailor discounts across a continuum of consumer valuations, the firm-optimal equilibrium features a simple two-tier discount scheme, comprising a uniform positive discount when the consumer outside option is intermediate and no discount when the outside option is low or high. This challenges the usual view that the availability of more detailed consumer data leads to increasingly complex pricing strategies.

While this study provides valuable insights into personalized pricing in the presence of consumer search, it also opens several avenues for future research. A natural extension would be to

¹¹For example, in the benchmark regime with symmetric information, if $1 - G$ is log-concave, according to Proposition 1, the firm's aggregate profit is

$$p[1 - F(p + r)] + \int_{p+k}^{p+r} [p - \tau^b(u)]dF(u) + \int_{\underline{u}}^{p+k} pG(u - p)dF(u),$$

where k solves the equality of (6) and is also a function of p . Using the definition of $\tau^b(u)$, one can derive the first-order condition with respect to p :

$$1 - F(p + k) + \int_{\underline{u}}^{p+r} G(u - p)dF(u) - p \int_{\underline{u}}^{p+k} g(u - p)dF(u) = 0.$$

This determines the optimal list price in the benchmark if the aggregate profit function is well-behaved.

relax the assumption that the firm knows a consumer’s outside option accurately. In practice, firms often possess only a noisy signal of a consumer’s outside option. Analyzing such a setting with imperfect consumer information would test the robustness of our findings but also introduce significant analytical challenges. Would the simple two-tier structure remain optimal? And would the paradox of “more information leading to lower profit“ still hold? Exploring these questions could deepen our understanding of how imperfect information advantages shape the strategy of personalized discounts.

Another important direction for future research is to consider a competitive environment in which the consumer’s outside option is not fixed exogenously but instead determined by a rival firm engaged in its own strategic pricing. In this setting, the consumer’s search decision becomes part of a game between two sellers—the incumbent and a competitor—both potentially capable of offering targeted discounts. Understanding how the firms’ pricing strategies and informational advantages interact in equilibrium would provide valuable insights. We leave this rich and challenging extension for future work.

References

- Anderson, S. P. and R. Renault (1999). Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model. *The RAND Journal of Economics*, 719–735.
- Armstrong, M. and J. Zhou (2016). Search deterrence. *The Review of Economic Studies* 83(1), 26–57.
- Bagnoli, M. and T. Bergstrom (2005). Log-concave probability and its applications. *Economic Theory* 26(2), 445–469.
- Beaudry, P. (1994). Why an informed principal may leave rents to an agent. *International Economic Review* 35(4), 821–832.
- Bedard, N. C. (2017). The strategically ignorant principal. *Games and Economic Behavior* 102, 548–561.
- Choi, W. J., Q. Liu, and J. Shin (2024). Predictive analytics and ship-then-shop subscription. *Management Science* 70(2), 1012–1028.
- Dzyabura, D. and J. R. Hauser (2019). Recommending products when consumers learn their preference weights. *Marketing Science* 38(3), 417–441.
- Ellison, G. and A. Wolitzky (2012). A search cost model of obfuscation. *The RAND Journal of Economics* 43(3), 417–441.
- Groh, C.-C. (2024). Competitive price discrimination, imperfect information, and consumer search. *mimeo*.
- Lee, J.-Y., J. Shin, and J. Yu (2024). Communicating attribute importance under competition. *mimeo*.
- Li, X. and Z. Xu (2022). Superior knowledge, price discrimination, and customer inspection. *Marketing Science* 41(6), 1097–1117.
- Maskin, E. and J. Tirole (1992). The principal-agent relationship with an informed principal, ii: Common values. *Econometrica* 60(1), 1–42.
- Mauring, E. (2025). Search and price discrimination online. *Journal of Political Economy: Microeconomics* 3(1), 138–176.
- Myerson, R. B. (1983). Mechanism design by an informed principal. *Econometrica* 51(6), 1767–1797.
- Preuss, M. (2023). Search, learning, and tracking. *The RAND Journal of Economics* 54(1), 54–82.
- Rafieian, O. and S. Zuo (2024). Personalization, algorithm dependence, and learning. *mimeo*.
- Rogers, A. (2013). After you read the listings, your agent reads you. *New York Times*.

- Shin, J. and J. Yu (2021). Targeted advertising and consumer inference. *Marketing Science* 40(5), 900–922.
- Silvers, R. (2012). The value of information in a principal–agent model with moral hazard: The ex post contracting case. *Games and Economic Behavior* 74(1), 352–365.
- Wernerfelt, B. (1994). On the function of sales assistance. *Marketing Science* 13(1), 68–82.
- Wolinsky, A. (1986). True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics* 101(3), 493–511.
- Xu, Z. and A. Dukes (2019). Product line design under preference uncertainty using aggregate consumer data. *Marketing Science* 38(4), 669–689.
- Xu, Z. and A. Dukes (2022). Personalization from customer data aggregation using list price. *Management Science* 68(2), 960–980.

A Appendix: Omitted Proofs

A.1 Proof of Proposition 1

Results (1)-(3) have been explained in the main text. Here we prove the remaining cut-off result. When $1 - G(v)$ is log-concave (which must be the case, e.g., when the density function $g(v)$ is log-concave), it is known that $\int_x^{\bar{v}} [1 - G(v)] dv$ is also log-concave (see, e.g., Theorem 3 in Bagnoli and Bergstrom 2005). Then

$$\frac{-[1 - G(x)]}{\int_x^{\bar{v}} [1 - G(v)] dv} \text{ is decreasing in } x \iff g(x) \int_x^{\bar{v}} [1 - G(v)] dv < [1 - G(x)]^2.$$

Given $r < \bar{v}$, we then have

$$g(x) \int_x^r [1 - G(v)] dv < [1 - G(x)]^2 \iff \frac{-[1 - G(x)]}{\int_x^r [1 - G(v)] dv} \text{ is decreasing in } x.$$

Since $B(x) - s = \int_x^r [1 - G(v)] dv$, this implies that $\frac{B(x) - s}{1 - G(x)}$ is decreasing in x , and so we have the cutoff result. \square

A.2 Proof of Lemma 2

Result (1) has been explained in the main text. We will prove results (2)-(4) together. Notice first that it is impossible to have $\sigma^*(\tau_0; u) = 0$. If the consumer never chooses to search after seeing a zero discount, the firm would never offer a positive discount even when the consumer has $v \geq u - p$. The firm would only make a strictly less profit by doing so regardless of whether the consumer would search or not after seeing a positive discount. Therefore, for a non-pooling equilibrium to exist, we must have $\sigma^*(\tau_0; u) > 0$.

Given this result, it is easy to see result (2). Suppose in contrast that $\mu^*(\tau_0; v, u) = 0$ for all $v \geq u - p$, i.e., the firm never offers the consumer a zero discount when $v \geq u - p$. Then a zero discount becomes a perfect signal of $v < u - p$, and therefore the consumer would never search after seeing a zero discount. A contradiction with $\sigma^*(\tau_0; u) > 0$.

To see result (4), suppose in contrast that $\sigma^*(\tau; u) = 1$ for a positive $\tau \in \Omega(\mu^*)$, i.e., the consumer always chooses to search after seeing a positive discount on the equilibrium path. Note that a positive discount will be offered to the consumer only if she has $v \geq u - p$. The firm then must make a zero profit from this consumer since she will never come back. Then, offering such a positive discount is dominated by offering no discount according to our tie-breaking rule. To have

$\sigma^*(\tau; u) < 1$, consumers must weakly prefer not to search after receiving discount τ , i.e., we must have $B(u - p; \tau) \leq \tau + s$.

We can now show $\sigma^*(\tau_0; u) < 1$ and then result (3) follows. Again, suppose in contrast $\sigma^*(\tau_0; u) = 1$, i.e., the consumer always chooses to search after seeing a zero discount. Consider the case when she has a $v > u - p$ such that $\mu^*(\tau_0; v, u) > 0$. In this case, the firm makes a zero profit from her when it offers a zero discount. However, if it deviates to offer the positive discount $\tau < p$ which is in $\Omega(\mu^*)$, it would earn a profit $(p - \tau)(1 - \sigma^*(\tau; u)) > 0$ given $\sigma^*(\tau, u) < 1$ from result (4). This is a contradiction. Since $0 < \sigma^*(\tau_0; u) < 1$, the consumer must be indifferent between continuing to search and buying immediately after seeing a zero discount, which implies $B(u - p; \tau_0) = s$.

Finally, we prove result (5). For a consumer with $u - p < r$ and $v \geq u - p$, if the firm offers her a discount $\tau \in \Omega(\mu^*)$, it earns a profit $(p - \tau)(1 - \sigma^*(\tau; u))$ since she will not return once she continues to search. This profit is independent of v and so must be the same for all $\tau \in \Omega(\mu^*)$. Otherwise, some $\tilde{\tau}$ would be more profitable than other discounts and so would be chosen by the firm all the time. If $\tilde{\tau} = 0$, we would end up with having a pooling equilibrium; if $\tilde{\tau} > 0$, that would violate result (2). \square

A.3 Proof of Lemma 3

We first prove a result which strengthens result (4) in Lemma 2:

Claim 1. *In the market for a consumer with $u - p < r$, if a binary-discount equilibrium with $\Omega(\mu^*) = \{\tau_0, \tau_1\}$ is the best for the firm among all possible binary-discount equilibria, it must satisfy $B(u - p; \tau_1) = \tau_1 + s$ and $\sigma^*(\tau_1; u) = 0$.*

That is, in the firm-optimal binary-discount equilibrium, consumers must be indifferent between continuing to search and buying immediately after seeing the positive discount, and meanwhile they choose not to search.

Proof. Result (4) in Lemma 2 implies $\sigma^*(\tau_1; u) < 1$ in a binary-discount equilibrium. We first argue that $\sigma^*(\tau_1; u) = 0$ in the most profitable binary-discount equilibrium. Suppose in contrast that $0 < \sigma^*(\tau_1; u) < 1$. Then the consumer must be indifferent between continuing to search and buying immediately after seeing the positive discount, i.e., $B(u - p; \tau_1) = \tau_1 + s$. (Recall that result (3) in Lemma 2 is that the consumer must be also indifferent after seeing a zero discount and $0 < \sigma^*(\tau_0; u) < 1$.) We now construct a new binary-discount equilibrium where all the configurations remain unchanged except that we slightly reduce both $\sigma^*(\tau_0; u)$ and $\sigma^*(\tau_1; u)$ so

that we still have $p(1 - \sigma^*(\tau_0; u)) = (p - \tau_1)(1 - \sigma^*(\tau_1; u))$ which is required by result (5) in Lemma 2. Since this modification leads to a higher profit, therefore, off-equilibrium deviations, which are unprofitable under the original equilibrium, will be even less desirable. Therefore, the newly constructed equilibrium is sustained and is also more profitable. This is a contradiction.

We then show that $B(u - p; \tau_1) = \tau_1 + s$ in the most profitable binary-discount equilibrium. Suppose in contrast that $B(u - p; \tau_1) < \tau_1 + s$ (which implies $\sigma^*(\tau_1; u) = 0$). Then, there must exist another equilibrium where all the configurations remain unchanged except that when the firm offers a positive discount, it offers a slightly smaller one $\hat{\tau}_1 < \tau_1$ so that the above inequality still holds (note that the posterior associated with the discount is unchanged, so is $B(u - p; \tau_1)$), and at the same time we also set a slightly smaller $\hat{\sigma}^*(\tau_0; u)$ so that we still have $p(1 - \hat{\sigma}^*(\tau_0; u)) = p - \hat{\tau}_1$ which is required by result (5) in Lemma 2. It is clear that the firm makes a higher profit in this new equilibrium and the new equilibrium is also sustained. This is again a contradiction. \square

For notational simplicity, for a given consumer with $u - p < r$, we let $\mu_0(v) = \mu_0(v, u)$ in the remaining of the proof. We now show that $\mu_0(v)$ must be a step function in the firm-optimal binary discount equilibrium.

Claim 2. *In the firm-optimal binary discount equilibrium in the market for a consumer with $u - p < r$, $\mu_0(v)$ is a step function: $\mu_0(v) = 0$ for $v \in [u - p, \hat{v}]$ and $\mu_0(v) = 1$ otherwise, where \hat{v} solves the equation (11).*

Proof. Result (5) in Lemma 2 and $\sigma^*(\tau_1; u) = 0$ from Claim 1 jointly imply that in the optimal binary-discount equilibrium, the firm's profit from the consumer with $u - p < r$ must be $pG(u - p) + (p - \tau_1)[1 - G(u - p)]$. Therefore, we aim to

$$\min_{\mu_0(v)} \tau_1 = B(u - p; \tau_1) - s \quad \text{s.t.} \quad B(u - p; \tau_0) = s,$$

where the equality in the objective function is from Claim 1 and the constraint is from result (3) in Lemma 2.

Note that the law of iterated expectations implies that

$$\Pr(\tau_0)B(u - p; \tau_0) + [1 - \Pr(\tau_0)]B(u - p; \tau_1) = B(u - p),$$

where $\Pr(\tau_0) = G(u - p) + \int_{u-p}^{\bar{v}} \mu_0(v) dG(v)$ is the probability that the firm offers no discount to a

consumer of type u across all possible v 's. Then

$$\tau_1 = B(u - p; \tau_1) - s = \frac{B(u - p) - \Pr(\tau_0)B(u - p; \tau_0)}{1 - \Pr(\tau_0)} - s = \frac{B(u - p) - s}{1 - \Pr(\tau_0)}, \quad (15)$$

where the last equality used the constraint $B(u - p; \tau_0) = s$. Given $B(u - p) > s$ for $u - p < r$, minimizing τ_1 is thus equivalent to minimizing $\Pr(\tau_0)$ or $\int_{u-p}^{\bar{v}} \mu_0(v) dG(v)$. Therefore, we rewrite our problem as

$$\min_{\mu_0(v)} \int_{u-p}^{\bar{v}} \mu_0(v) dG(v) \quad \text{s.t.} \quad B(u - p; \tau_0) = s.$$

Since $\mu_0(v) = 1$ for $v < u - p$, to sustain the constraint we need to assign enough weight to $v \geq u - p$ for which $\mu_0(v) > 0$. Given our aim, the least “costly” way to do so is to put all the weight on the highest possible v . That is, the optimal $\mu_0(v)$ must take a cut-off form in the range of $v \in [u - p, \bar{v}]$: $\mu_0(v) = 0$ if $v < \hat{v}$ and $\mu_0(v) = 1$ otherwise.¹² \square

Given this step function result, it is straightforward to determine the remain equilibrium characteristics. First, given the step-function form of $\mu_0(v)$, the constraint can be written as

$$B(u - p; \tau_0) \equiv \frac{\int_{\hat{v}}^{\bar{v}} [v - (u - p)] dG(v)}{1 - G(\hat{v}) + G(u - p)} = s,$$

which is equivalent to (11). As we have proved in footnote 5, \hat{v} is unique in $(u - p, \bar{v})$. Second, given the step-function form of $\mu_0(v)$, we have $1 - \Pr(\tau_0) = G(\hat{v}) - G(u - p)$. Then we have (12) immediately from (15). Third, given $\sigma^*(\tau_1; u) = 0$, the firm's indifference condition in result (5) of Lemma 2 requires $p(1 - \sigma^*(\tau_0; u)) = p - \tau_1^*$, and so $\sigma^*(\tau_0; u) = \tau_1^*/p$. Lastly, to ensure $\sigma^*(\tau_0; u) < 1$, we also need $\tau_1^* < p$.

The final step in establishing the binary-discount equilibrium is to ensure that the firm has no incentive to deviate by offering off-equilibrium discounts. This is obviously true for a consumer with $v < u - p$, as the firm is already making the highest possible profit from her on the equilibrium path. Therefore, we only need to consider a consumer with $v \geq u - p$.

We first argue that the firm has no incentive to offer an off-equilibrium discount $\tau > \tau_1^*$. For a

¹²More formally, suppose in contrast that there exist $\tilde{v} \in (u - p, \bar{v})$ and two sets $V_1 \subset [u - p, \tilde{v}]$ and $V_2 \subset [\tilde{v}, \bar{v}]$ of positive measure, such that $\mu_0(v) > 0$ for $v \in V_1$ while $\mu_0(v) < 1$ for $v \in V_2$. Consider then a deviation with $\mu_0(v) - \epsilon$ for $v \in V_1$ and $\mu_0(v) + \alpha\epsilon$ for $v \in V_2$, where $\epsilon > 0$ is small and α satisfies $G(V_1) = \alpha G(V_2)$ where $G(V_i)$ is the probability measure of V_i . Such a deviation keeps our objective function unchanged, but it increases $B(u - p; \tau_0) = \mathbb{E}[\max\{u - p, v\} | \tau_0] - (u - p)$. (Notice that the deviation does not change the denominator in the explicit expression for $\mathbb{E}[\max\{u - p, v\} | \tau_0]$ but changes the numerator by $\epsilon[-\int_{V_1} v dG(v) + \alpha \int_{V_2} v dG(v)]$. This must be positive because $\alpha \int_{V_2} v dG(v) > \alpha \tilde{v} G(V_2) = \tilde{v} G(V_1) > \int_{V_1} v dG(v)$.) The existence of such a deviation immediately implies that we can always find another deviation which satisfies the constraint while at the same time reduces the objective function. This is a contradiction.

consumer with $v \in [u - p, \hat{v}]$, she is already willing to stop searching under the equilibrium discount τ_1^* , so offering a higher discount will only reduce the firm's profit (regardless of the consumer search behavior off the equilibrium path). For a consumer with $v > \hat{v}$, the firm makes a profit $p(1 - \sigma^*(\tau_0; u))$, which equals $p - \tau_1^*$ according to result (5) in Lemma 3, by offering a zero discount on the equilibrium path; if it offers a discount $\tau > \tau_1^*$, even if the consumer stops searching, its profit will be $p - \tau$, which is below the equilibrium profit.

Also, given our off-equilibrium belief in (13), the firm has no incentive to offer any off-equilibrium discount $\tau < \tau_1^*$ either. If the firm offers a discount $0 < \tau < \tau_1^*$, the consumer, regardless of her $v \geq u - p$, will choose to search under the above off-equilibrium belief. This is because the off-equilibrium belief is more optimistic than the equilibrium belief associated with τ_1^* , and so $B(u - p; \tau) > B(u - p; \tau_1^*) = \tau_1^* + s > \tau + s$, where the equality is an equilibrium condition. Thus, the firm will then make a zero profit from this consumer given $v \geq u - p$, thus eliminating any incentive to offer a discount different from τ_1^* .

Finally, we prove the cutoff result for $\tau_1^* < p$ when s is sufficiently small. Let $x = u - p$. From (11), we derive

$$-g(\hat{v})\hat{v}'(x) = \frac{1 - G(\hat{v}) + sg(x)}{\hat{v} - x - s} > 0.$$

(From footnote 5, we know $\hat{v} > x + s$, so the denominator is positive.) One can also check that τ_1^* decreases in x if and only if

$$[g(x) - g(\hat{v})\hat{v}'(x)] \int_x^{\tau} [1 - G(v)]dv < [1 - G(x)][G(\hat{v}) - G(x)].$$

When $s \rightarrow 0$, we have $\hat{v} \rightarrow \bar{v}$ and $-g(\hat{v})\hat{v}'(x) \rightarrow 0$, so the above inequality must hold under the log-concavity condition as we already see in the proof of Proposition 1. \square

A.4 Proof of Lemma 4

Suppose that there is an equilibrium $(\tilde{\mu}, \tilde{\sigma})$ with $|\Omega(\tilde{\mu})| \geq 3$. Let us construct a binary-discount equilibrium $(\hat{\mu}, \hat{\sigma})$ as follows: (i) When the firm offers no discount in the original equilibrium, it continues to do so in the new equilibrium. That is, for any v , we set

$$\hat{\mu}(\tau_0 = 0; u, v) = \tilde{\mu}(\tau_0 = 0; u, v).$$

(ii) When the firm offers any positive discount in the original equilibrium, it now offers a fixed discount $\hat{\tau} > 0$ instead. That is, for all v , we set

$$\hat{\mu}(\hat{\tau}; u, v) = \int_{\Omega(\tilde{\mu}) \setminus \{\tau_0\}} \tilde{\mu}(\tau; v, u) d\tau.$$

Such a construction obviously implies that

$$\mathbb{E}[v|\tau > \tau_0] = \mathbb{E}[v|\hat{\tau}].$$

More precisely, here the first expectation operator is based on the original equilibrium, and the second one is based on the newly constructed equilibrium. Therefore, there must exist some $\tilde{\tau} \in \Omega(\tilde{\mu}) \setminus \{\tau_0\}$ such that

$$\mathbb{E}[v|\tilde{\tau}] > \mathbb{E}[v|\hat{\tau}]. \quad (16)$$

According to result (4) in Lemma 2, $\tilde{\tau}$ in the original equilibrium must satisfy

$$B(u - p; \tilde{\tau}) \leq \tilde{\tau} + s \iff u - p + \tilde{\tau} \geq \mathbb{E}[v|\tilde{\tau}] - s,$$

i.e., the consumer should weakly prefer to stop searching after seeing $\tilde{\tau}$. Recall that in either equilibrium a positive discount is offered only if $v \geq u - p$, so $B(u - p; \tilde{\tau}) = \mathbb{E}[v - (u - p)|\tilde{\tau}]$. For the newly constructed equilibrium, following the idea in the proof of Claim 1, we can always set $\hat{\tau}$ such that

$$u - p + \hat{\tau} = \mathbb{E}[v|\hat{\tau}] - s.$$

Note that according to our construction $\mathbb{E}[v|\hat{\tau}]$ is fixed when we adjust $\hat{\tau}$. Then (16) implies that

$$\tilde{\tau} > \hat{\tau}. \quad (17)$$

Note that $\hat{\tau} < p$ given $\tilde{\tau} < p$. Moreover, in the newly constructed equilibrium, we can also set $\hat{\sigma}(\hat{\tau}; u) = 0$. By the same argument in the proof of Lemma 3, it is then easy to see that the newly constructed equilibrium is indeed an equilibrium under the off-equilibrium belief in (13).

From result (5) in Lemma 2, we know that the firm must make the same profit from the consumer when $v < u - p$. Hence, to compare profit between these two equilibria, we only need to consider the profit from the consumer when $v \geq u - p$. Per the same result in Lemma 2, conditional on $v \geq u - p$, the profit in the original equilibrium must be equal to $(p - \tilde{\tau})(1 - \tilde{\sigma}(\tilde{\tau}; u))$, while that

in the new equilibrium is $p - \hat{\tau}$. Then (17) immediately implies that the latter profit is higher. \square

A.5 Proof of Lemma 5

Let us first show that this pooling equilibrium exists. On the equilibrium path, given $u - p < r$ the consumer indeed prefers to search after receiving no discount. Suppose that the consumer holds the off-equilibrium belief specified in (13) upon seeing any positive discount $0 < \tau \leq p$. From the proof of Lemma 3, we already know that for a consumer with $u - p < r$ and $\tau_1^* \geq p$, she needs a discount greater than p to be willing to stop searching when believing $v \in [u - p, \hat{v}]$ where $\hat{v} \in (u - p, \bar{v})$ solves (11). Therefore, she must prefer to search upon seeing a discount $\tau \leq p$ given the same belief. Now the off-equilibrium belief is $v \in [u - p, \bar{v}]$, which is more optimistic than $v \in [u - p, \hat{v}]$, so no discount $\tau \leq p$ can stop her from searching. This implies that the firm has no incentive to deviate and offer any positive discount regardless of v .

We then show that there are no non-pooling equilibria. If a non-pooling equilibrium with $|\Omega(\tilde{\mu})| \geq 3$ exists, according to Lemma 4, there must exist a binary-discount equilibrium. If a binary-discount equilibrium exists, according to the proof of Lemma 3, there must exist a (weakly) more profitable binary-discount equilibrium as characterized in Lemma 3 with $\tau_1^* < p$. This is a contradiction. \square

Appendix B: Examples of potential equilibrium discount scheme

The firm has access to an extensive range of strategies for designing its discount schemes, each leading to numerous potential equilibria. We provide illustrative examples of such strategies below. Figure 5 visualizes four different equilibrium discount schemes $\mu^*(\tau; v, u)$ for a given u . Darker shades indicate a higher probability of the discount being offered.

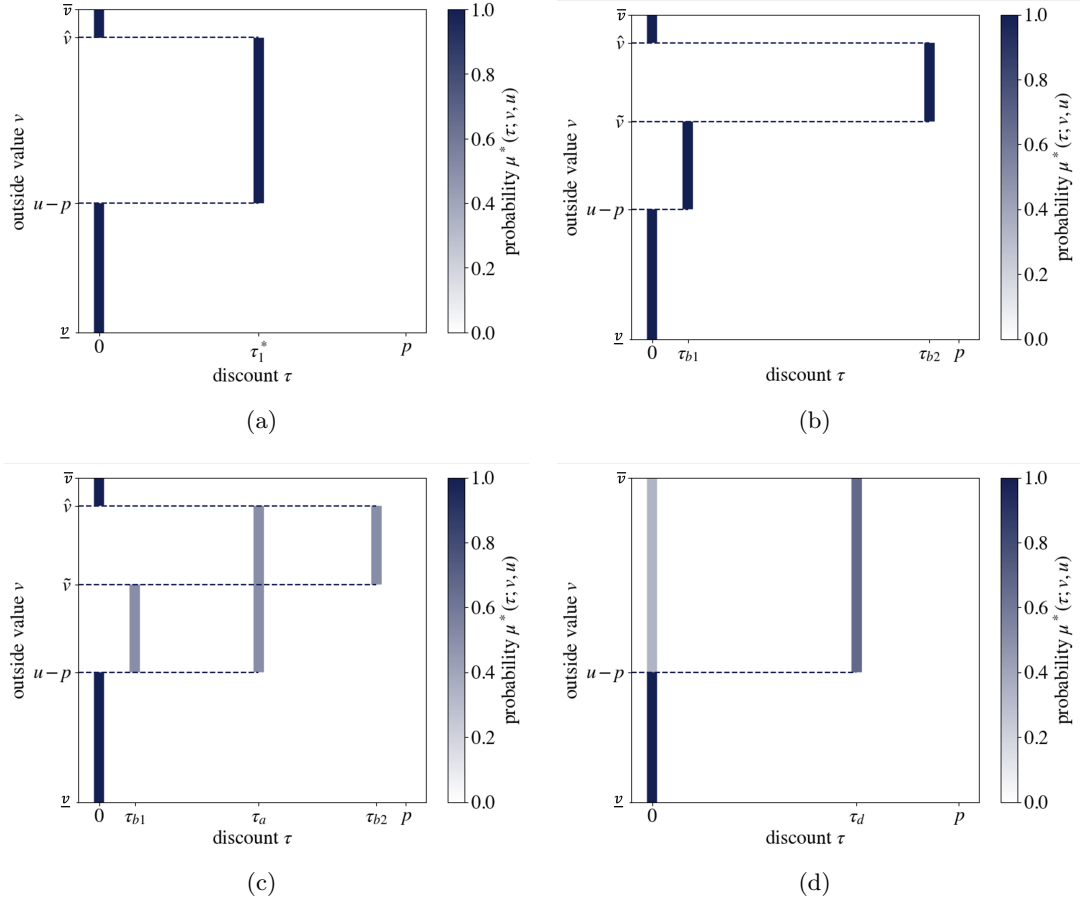


Figure 5: Examples of equilibrium discount scheme

In Figure 5(a), the firm employs a binary discount scheme, offering a fixed discount (τ_a) to deter search among consumers with moderately high outside values and no discount ($\tau_0 = 0$) to those with the highest outside values. This represents the optimal binary equilibrium analyzed in the main model. Equilibrium conditions ensure consumer indifference and firm indifference between offering and no discount. Figure 5(b) illustrates a multi-level discount strategy, assigning different discounts (τ_{b1} and τ_{b2}) to various segments within moderately high outside values. Such differentiation remains an equilibrium if consumer indifference conditions are met across discount levels. Further flexibility arises through linear combinations of equilibria. Figure 5(c) exemplifies

this by combining previously described binary and multi-level schemes, highlighting the substantial flexibility in equilibrium construction. Lastly, Figure 5(d) presents a semi-separating equilibrium, where the firm probabilistically withholds discounts from high-value consumers, thus maintaining consumer uncertainty and achieving equilibrium indifference conditions.

Beyond these finite discount examples, infinitely many equilibria exist, further illustrating the complexity and diversity of potential equilibrium discount strategies.