

# Personalized Discounts and Consumer Search

Zikun Liu\*

Jiwoong Shin<sup>†</sup>

Jidong Zhou<sup>‡</sup>

April 23, 2026

## Abstract

The growing availability of consumer data and predictive analytics enables firms to predict consumers' outside options more accurately than consumers themselves. This paper studies how a firm can use 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 using consumer-specific information, the firm-optimal equilibrium features a simple two-tier discount scheme: the seller either offers no discount or a single positive discount, with discounts targeted only to consumers with intermediate outside options. Thus, signaling compresses an otherwise rich pricing problem into a coarse pricing scheme. We further show that the profit effect of superior information is conditional on consumer inference. Superior information raises profit when consumer inference is limited, but can reduce profit when consumers are sufficiently strategic because discounts induce additional search. In the regime in which superior information reduces profit, consumer surplus remains unchanged while total welfare declines because equilibrium search becomes excessive. Our results rationalize why firms with rich data often use simple pricing formats and why superior information does not necessarily increase profitability.

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

---

\*School of Management, Yale University. Email: [zikun.liu@yale.edu](mailto:zikun.liu@yale.edu). <sup>†</sup>School of Management, Yale University. Email: [jiwoong.shin@yale.edu](mailto:jiwoong.shin@yale.edu). <sup>‡</sup>School of Management, Yale University. Email: [jidong.zhou@yale.edu](mailto:jidong.zhou@yale.edu). This is a chapter of the first author's PhD dissertation at Yale University. We thank seminar participants at City University of Hong Kong, Chinese University of Hong Kong, University of Minnesota, and Yale University, as well as participants at the Marketing Science Conference.

# 1 Introduction

Consumers often face imperfect information about product qualities, fit, and prices, and therefore rely on search to make better purchase decisions. Firms, in turn, increasingly use consumer data and predictive analytics to anticipate how consumers evaluate their products and what they are likely to find elsewhere. A natural use of such information is to offer a targeted buy-now discount that makes immediate purchase more attractive and thereby discourages further search. However, such tactics can backfire. Once consumers recognize that firms may know more than they do, a substantial discount may arouse consumer suspicion about the firm’s motives and prompt even more extensive exploration of alternatives. A discount is then not just an inducement to stop searching; it can also be a signal that attractive outside options are likely to exist.

Consider a consumer deliberating between the BMW 3-series, Mercedes C-class, and other models. Because she is still learning which model best fits her preferences, she visits several dealerships. At the BMW dealership, her interest in the 3-series becomes clear. In response, the salesperson offers a \$5,000 discount if she purchases immediately. 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.

A similar logic can arise in other settings in which firms use simple but targeted concessions to accelerate purchase after observing a signal that the consumer may continue searching or leave. In real estate, for example, an agent with deep market knowledge and detailed property information, often unavailable to average buyers, may use this informational advantage strategically. If the agent learns that a family places exceptional value on school quality, 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. Likewise, an online merchant such as Shopify may follow up on some abandoned checkouts with a coupon code or reminder after observing hesitation at checkout<sup>1</sup>; and some insurance companies such as State Farm keep a standard retention offer in reserve for some cancellation encounter and offer it to limit potential customer defection.<sup>2</sup> The institutional details differ, but the pattern is the same: after observing a consumer-specific signal, the firm intervenes with a relatively simple concession. In each case, the offer is intended

---

<sup>1</sup>Shopify lets merchants manually send abandoned-checkout emails and also choose which visitor category receives automated recovery emails. See Shopify Help Center, “Recovering Abandoned Checkouts” <https://help.shopify.com/en/manual/promoting-marketing/create-marketing/abandoned-checkouts>

<sup>2</sup>LexisNexis describes data-driven outreach to strengthen retention when carriers observe life-event or shopping signals, and State Farm offers simple discounts, including loyalty discounts in California. See LexisNexis Risk Solutions, “Life Events Happen: When an Agent Should Reach Out,” <https://risk.lexisnexis.com/insights-resources/article/life-happens-when-an-agent-should-reach-out>; and State Farm, “Car Insurance Discounts California,” <https://www.statefarm.com/insurance/auto/discounts/california>.

to reduce further search or defection, yet its very presence may suggest that attractive alternatives exist. Notably, even in data-rich environments, the intervention often takes a simple form, such as a limited-time concession, a coupon code, or a standard retention discount, rather than a fully elaborate individualized price schedule.

These examples illustrate the central tension of the paper. Buy-now discounts are meant to discourage search, but they can instead stimulate search when consumers interpret them as informative. This possibility is economically important because personalized offers are now widely understood to reflect firms' predictive capabilities. Through direct interactions, firms learn about consumers' tastes and priorities (Wernerfelt, 1994; Rogers, 2013), and those informational advantages are increasingly sharpened by predictive algorithms and consumer-level data (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 advantage. Consumers may therefore interpret why a particular offer is made, not just how attractive it is. This signaling channel fundamentally alters the role of personalization: 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 studies personalized discounts in the presence of consumer inference. We consider a monopolist seller facing a consumer who is uncertain about the value of her outside option and must incur a search cost to learn it. The seller observes the consumer's valuation for its own product and, under superior information, also observes the consumer's outside option. Before the consumer decides whether to search, the seller can offer a buy-now discount. A strategic consumer understands that the offer may depend on the seller's private information and updates beliefs accordingly. We characterize the perfect Bayesian equilibrium of this game and compare it with a benchmark in which the seller lacks superior information about the outside option and therefore cannot condition discounts on it.

We first find that, perhaps surprisingly, the firm's profit-maximizing equilibrium 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.

We then show that the profit effect of superior information is conditional on consumer inference. To make this explicit, we allow the market to contain both strategic and naive consumers. When inference is weak, superior information improves the seller's ability to target concessions and can raise profit. When inference is strong, however, the same information can become difficult to exploit. Discounts invite search by revealing that attractive alternatives are likely, and the seller may be forced either to concede too much or to lose consumers who search and do not return. Superior information can therefore reduce profit relative to a benchmark in which the seller lacks that informational advantage. This result also helps reconcile the model with the firms' incentives to invest in consumer data: information is valuable when consumers do not strongly infer from targeted offers, but its value erodes as consumers become more inference-aware.

The model also yields important welfare implications. In the regime in which superior information reduces profit, it induces additional consumer search but leaves aggregate consumer surplus unchanged. Gains from better matching are offset by losses from excessive search. Total welfare nevertheless falls because the extra search induced by signaling is socially inefficient. The analysis also suggests a clear empirical implication. As firms are perceived to have stronger predictive capabilities, and as consumers become more attentive to the informational content of offers, equilibrium pricing should become coarser and more transparent rather than more finely individualized. This also points to the settings in which the mechanism is most likely to matter, namely markets where firms interact closely with consumers, observe rich signals during the selling process, and can deploy simple buy-now concessions without rewriting the entire price schedule

We also examine several extensions that broaden the scope of the analysis. We show that the core pricing logic survives when the seller chooses the list price endogenously, when exogenous commitment is relaxed, and when the outside option is generated by a rival seller in a duopoly environment. These variants change the environment in economically meaningful ways, but they do not overturn the central mechanism linking targeted offers, consumer inference, and coarse equilibrium pricing.

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 consumer search models with product differentiation (Wolinsky, 1986; Anderson and Renault, 1999) and to the literature on search deterrence. Firms may discourage search by increasing consumers’ information-acquisition costs (Ellison and Wolitzky, 2012) or by using exploding offers and buy-now discounts (Armstrong and Zhou, 2016).<sup>3</sup> Armstrong and Zhou (2016) show how a buy-now discount can deter search by making immediate purchase more attractive. We keep that search-deterrence motive but add an inference channel: when the seller has superior information about the consumer’s likely outside option, the discount itself becomes informative. A buy-now offer can then encourage search rather than deter it.<sup>4</sup>

We also relate to a broader literature on seller information, targeting, and consumer inference. The key mechanism in our model is the signaling effect of targeted offers: a buy-now discount may reveal what the firm believes about the consumer’s likely outside option. Wernerfelt (1994) studies sales assistance as a way for a seller to use superior knowledge about a consumer’s situation, and Shin and Yu (2021) shows that consumers may draw inferences from the mere fact that they are targeted. Recent work on prediction and recommendation similarly studies how firms or platforms use data to predict product fit or likely purchase across available options (Dzyabura and Hauser, 2019; Choi et al., 2024; Rafeian and Zuo, 2024; Ning et al., 2025). Much of this work focuses on settings in which the firm’s informational advantage concerns consumers’ valuations for its own products or options within its selling environment. Relatedly, Xu and Dukes (2022) studies personalization from aggregated customer data and the role of list prices in addressing consumer suspicion, while Lee et al. (2024) studies how firms communicate attribute importance under competition. Our paper instead studies buy-now pricing when the seller’s private information concerns the consumer’s outside option or likely search outcome, and when consumers infer from the offer itself.

The closest paper to ours is Li and Xu (2022), which also studies personalized pricing with seller superior information and consumer inference. In their model, the firm privately knows the consumer’s valuation for its own product, the consumer can infer that valuation from the offered price, and the consumer may inspect the product at a cost. Both papers therefore feature signaling through prices and show that superior information need not straightforwardly benefit the seller. The key difference lies in what the consumer learns. In their model, inspection resolves uncertainty about the seller’s own product and can make purchase more likely; in our model, search reveals

---

<sup>3</sup>A broader recent literature studies tracking, segmentation, and price discrimination in search markets; see, for example, Preuss (2023); Groh (2024); Mairing (2025).

<sup>4</sup>Armstrong and Zhou (2016) note this possibility, but do not analyze the equilibrium pricing implications when discounts depend on the consumer’s outside option.

the attractiveness of an outside option and therefore draws demand away from the focal seller.<sup>5</sup> This difference helps explain why, in our setting, superior information always hurts the firm once consumers are strategic, whereas in their model the firm can sometimes benefit. Also, unlike their model, which studies binary consumer valuations and therefore binary discount levels, we allow a continuum of valuations and flexible discount conditioning. This richer environment lets us show that a coarse two-tier discount scheme can emerge endogenously as the firm-optimal policy, rather than being driven mechanically by a binary type space.

Finally, our paper is related to the informed principal literature (Myerson, 1983; Maskin and Tirole, 1992). In our setting, the seller privately knows payoff-relevant information and commits to a pricing mechanism from which the consumer may infer that information. Our setting differs because, after observing the offer, the consumer can take a costly action to learn about a competing alternative. We also compare the seller’s outcome under superior information with a benchmark without that information, allowing us to assess the value of being informed. In this respect, our paper is closest to the branch of the informed principal literature that asks whether a principal would want to become informed in the first place (Beaudry, 1994; Silvers, 2012; Bedard, 2017). They show, in settings quite different from ours, that a principal may prefer to remain uninformed. Our paper brings that question into a consumer-search setting with endogenous consumer inference.

## 2 The Model

There is a monopoly seller (e.g., a local car dealership) offering a single product.<sup>6</sup> A consumer has already reached this focal seller and must decide whether to buy immediately or engage in additional search. If she buys from the focal seller, she obtains gross surplus  $u$ . If she continues to search, she may discover an outside option that yields gross surplus  $v$ . Across consumers, the pair  $(u, v)$  is drawn i.i.d. from a joint distribution  $H(u, v)$  with full support on  $[\underline{u}, \bar{u}] \times [\underline{v}, \bar{v}]$ . Correlation between  $u$  and  $v$  is fully allowed.

When the consumer arrives at the focal seller, the realization of  $u$  is observed by both the consumer and the seller. By contrast, the realization of  $v$  is initially unknown to the consumer and can be learned only by incurring a search cost  $s > 0$ .<sup>7</sup> Search should therefore be interpreted as

---

<sup>5</sup>This learning structure also separates our paper from deliberation and recommendation models such as Guo and Zhang (2012), Dzyabura and Hauser (2019), and Xu and Dukes (2019), where consumers learn their fit with the seller’s offerings and firms respond through product design or recommendation.

<sup>6</sup>Appendix ?? extends the model to a duopoly environment in which the outside option is endogenously generated by a competing seller. Explicitly modeling competition does not alter the main inference-management mechanism or the qualitative comparisons in the text.

<sup>7</sup>We take the search cost  $s$  to be common and known for tractability. If search costs were heterogeneous or privately

*additional* search after the consumer has already reached the focal seller, such as visiting another dealership, checking additional listings, or otherwise learning whether an alternative is a better match. Consumers cannot access the outside option without first searching, and after searching they have free recall and can return to the focal seller without an additional search cost.

Throughout the paper, we solve the seller’s pricing problem conditional on the realized value of  $u$ . This is without loss of generality. Once  $u$  is observed, each realized  $u$  defines a local pricing problem, and all equilibrium objects depend only on the induced conditional distribution of  $v$  given  $u$ . We therefore write that conditional distribution as  $G(v | u)$  and, when no confusion arises, suppress the dependence on  $u$  and write  $G(v)$ . All expectations over  $v$  should be understood as taken with respect to this conditional distribution. This is simply a notational convenience; the analysis does not rely on independence between  $u$  and  $v$ .

To rule out degenerate cases, we assume  $0 \leq \underline{v} < \bar{u} - p$  and  $\mathbb{E}[v - \underline{v}] > s$ , where the expectation is taken with respect to the conditional distribution of  $v$  given the realized  $u$ . The first condition ensures the seller’s product is not always ex post dominated at the regular price, so a consumer who searches may still prefer to return to the seller for some realizations. The second ensures that search has enough upside to be worth considering for some consumers. This guarantees a nontrivial search problem.

The seller’s pricing policy has two layers. First, there is a regular list price  $p$ , which is posted *ex ante*, before the seller observes the consumer-specific private information that will later matter for the personalized offer.<sup>8</sup> This captures settings in which the headline price is set in advance or adjusted infrequently, while personalized concessions are decided later. Second, after the consumer arrives, the seller may offer a buy-now discount  $\tau \in [0, p]$ . If the consumer buys immediately, she pays  $p - \tau$ . If she searches, the initial discount is no longer available, and a returning consumer pays the regular list price  $p$ .<sup>9</sup> This timing distinction is important. In our model, the post-visit discount can serve as the signaling instrument because it is chosen after the relevant consumer-specific information is realized, whereas the ex ante list price cannot.

---

known to consumers, the seller would face a distribution of search propensities rather than a single reservation cutoff. This would affect the exact cutoff structure and comparative statics, but not the central mechanism: buy-now discounts would still both deter search and reveal information, so the firm would still face the same basic trade-off between search deterrence and information revelation.

<sup>8</sup>Appendix ?? allows the seller to choose the list price endogenously and shows that doing so changes levels and cutoff regions but does not overturn the core inference-based logic or the resulting two-tier discount structure.

<sup>9</sup>Appendix ?? shows that the main equilibrium logic survives when exogenous commitment is relaxed. The expiring discount should therefore be interpreted as a reduced-form representation of front-loaded pricing incentives rather than a literal institutional assumption that must hold in every market. See also Armstrong and Zhou (2016) for related discussions of front-loaded pricing without exogenous commitment.

## Two information regimes

During the selling process, firms often learn about consumers' needs and preferences through direct interaction and sales assistance (Wernerfelt, 1994; Rogers, 2013). We therefore assume that, when the consumer visits the seller, the seller observes the consumer's valuation  $u$  for its own product.

We consider two information regimes depending on whether the seller also has superior information about the consumer's outside option. In the benchmark regime, the seller observes  $u$  but not  $v$ . Seller and consumer therefore face the same uncertainty about the outside option prior to search, so the seller's buy-now discount may depend on  $u$  but not on  $v$ . We refer to this as the *symmetric-information* regime. In the main regime of interest, the seller observes both  $u$  and  $v$  when the consumer arrives.<sup>10</sup> We refer to this as the *superior-information* regime. Under superior information, the seller can condition the buy-now discount on both  $u$  and  $v$ . A consumer who recognizes this informational advantage may then rationally infer information about  $v$  from the discount itself and adjust her search strategy accordingly.

Consumers differ in how they process that information. A fraction  $\lambda \in [0, 1]$  of consumers are *strategic*. These consumers correctly anticipate the seller's discounting incentives and update their beliefs about  $v$  from the observed discount using Bayes' rule. The remaining fraction  $1 - \lambda$  are *naive*. Naive consumers do not infer information about  $v$  from the discount and instead evaluate the search decision based solely on the unconditional distribution  $G(v|u)$ . The seller does not observe a consumer's type and sets discounts based on the market composition, but  $\lambda$  is common knowledge to both the seller and consumers.<sup>11</sup>

Both the seller and consumers are risk neutral. Consumers choose whether to search to maximize expected surplus given their beliefs, and the seller chooses discounts to maximize expected profit given rational expectations about consumer search behavior. If offering no discount yields the same expected profit as offering a positive discount, we assume the seller prefers not to discount. Without loss of generality, the seller's unit production cost is normalized to zero.

---

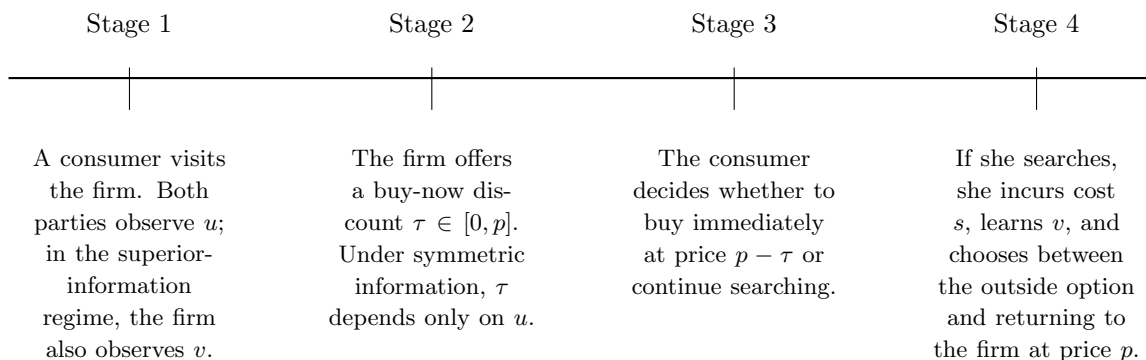
<sup>10</sup>This assumption should be interpreted as a tractable reduced form of predictive superiority about the consumer's likely outside option or search outcome, rather than a literal claim that firms directly observe competitors' realized offers in every market.

<sup>11</sup>If the firm could observe a consumer's type, it would optimally condition discounts on the type and solve two independent pricing problems. In that case, equilibrium profits would be a weighted average of the profits obtained in the two benchmark environments with fully strategic and fully naive consumers, respectively. Allowing for such observability therefore does not affect the mechanisms or insights emphasized in the paper.

## Timing of the game

As summarized in Figure 1, the game unfolds as follows. The list price  $p$  is set ex ante and taken as given. At Stage 1, a consumer visits the firm. Both the consumer and the seller observe the consumer's valuation  $u$ . In the superior information regime, the firm also observes the consumer's outside-option value  $v$ . At Stage 2, the seller offers a personalized buy-now discount  $\tau(u, v) \in [0, p]$ . In the symmetric information regime, the discount cannot depend on  $v$  and is therefore a function of  $u$  alone. At Stage 3, after observing the discount, the consumer decides whether to buy immediately at the discounted price  $p - \tau$  or to continue searching for the outside option. If she chooses to search, then at Stage 4 she incurs the search cost, learns her valuation  $v$ , 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 begin with the symmetric-information benchmark, in which the seller observes each consumer's  $u$  but not  $v$ . Because the buy-now discount  $\tau$  cannot depend on  $v$ , it conveys no information about the outside option. Strategic and naive consumers therefore behave identically in this benchmark, so all outcomes are independent of  $\lambda$ .

Each consumer can be treated as a separate market. Fix a realized  $u$ , then  $u - p$  is the consumer's surplus from buying immediately at the list price. If the seller offers a buy-now discount  $\tau(u)$ , immediate purchase yields surplus  $u - p + \tau(u)$ . If instead the consumer searches, she incurs cost  $s$ , learns  $v$ , and then chooses between the outside option and returning to the seller at price  $p$ . The expected benefit from search, given the current surplus upon returning and paying the full list

price, is

$$B(u - p) \equiv \mathbb{E}[\max\{v - (u - p), 0\}] = \int_x^{\bar{v}} [1 - G(v)] dv. \quad (1)$$

The consumer buys immediately if and only if

$$B(x) \leq s + \tau. \quad (2)$$

Therefore, the consumer's search problem with a buy-now discount can be rephrased as a standard search problem with an amplified search cost:  $\tau$  only raises the effective cost of search from  $s$  to  $s + \tau$ . Let  $r$  denote the reservation value in the standard search problem that solves

$$B(r) = s. \quad (3)$$

Because  $B(x)$  is continuous and decreasing, with  $B(\underline{v}) = \mathbb{E}[v - \underline{v}]$  and  $B(\bar{v}) = 0$ , the condition  $\mathbb{E}[v - \underline{v}] > s$  ensures that the equation  $B(r) = s$  has a unique solution  $r \in (\underline{v}, \bar{v})$ .

For consumers with  $u - p < r$ , the smallest discount that just deters search is

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

Offering this discount yields profit  $p - \tau^b(u)$ . Offering no discount yields profit  $p \cdot G(u - p)$ , since the consumer searches and returns to the seller only if  $v < u - p$ . The seller therefore deters search if and only if<sup>12</sup>

$$p - \tau^b(u) > p \cdot G(u - p) \iff p > \frac{B(u - p) - s}{1 - G(u - p)}. \quad (5)$$

We summarize the benchmark policy as follows.

**Proposition 1.** *In the symmetric-information benchmark, let  $r$  solve  $B(r) = s$ .*

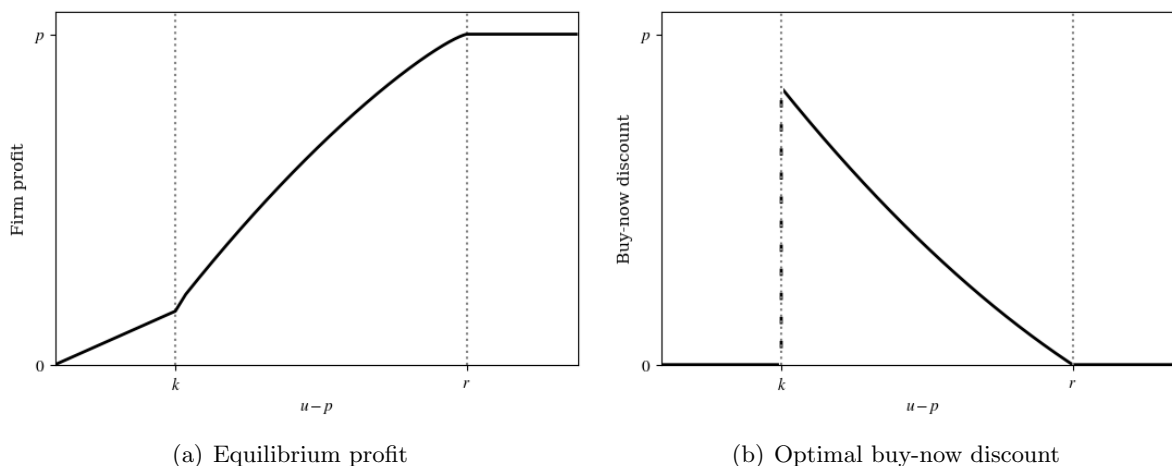
- (1) *If  $u - p \geq r$ , the seller offers no discount, the consumer buys immediately, and the seller earns  $\Pi^b(u) = p$ .*
- (2) *If  $u - p < r$ , the minimal discount that deters search is  $\tau^b(u) = B(u - p) - s$ . The seller offers this discount if condition (5) holds; otherwise it offers no discount and lets the consumer search. Thus,  $\Pi^b(u) = \max\{p - \tau^b(u), p \cdot G(u - p)\}$*

Figure 2 provides a numerical illustration for the case in which  $G(v)$  is uniform on  $[0, 1]$ . In this example, the seller offers no discount when  $u - p$  is sufficiently high, because search is already

---

<sup>12</sup>Since we assume the seller prefers not to discount when it is indifferent, the condition for deterring search should be strict, not weak.

Figure 2: Symmetric-Information Benchmark.



unattractive, and also when  $u - p$  is sufficiently low, because deterring search is too costly. For intermediate values of  $u - p$ , the seller offers a positive buy-now discount, and that discount decreases in  $u - p$ .<sup>13</sup>

This benchmark serves as the reference point for the rest of the paper. Absent superior information, a buy-now discount is purely a search-deterrence instrument. The benchmark therefore provides the reservation cutoff  $r$ , the search-deterrence discount  $\tau^b(u)$ , and the baseline profit comparison  $\Pi^b(u)$  that we use in the next section.

## 4 The Main Case: Superior Information

We now turn to the regime of superior information, in which the seller observes both the consumer's valuation  $u$  for its own product and her valuation  $v$  for the outside option. Unlike the symmetric-information benchmark, a buy-now discount may now convey information about the attractiveness of the outside option and therefore affect behavior through beliefs as well as prices. A discount that is intended to deter search may thus have the opposite effect if consumers interpret it as evidence that the outside option is strong. The seller's problem is therefore no longer simply how much to discount, but how to use discounts without revealing too much.

To see the underlying force of this problem, suppose first that consumers are naive and do not

<sup>13</sup>In the uniform example, the discounting region is bounded below by the cutoff  $k$  shown in the figure. Online Appendix ?? shows that, under a mild regularity condition of log-concavity of  $1 - G(v)$ , the region in which the seller offers discounts can be represented by a single cutoff  $k$ , so that discounts are offered if and only if  $u - p \in (k, r)$ . Without this condition, the optimal policy need not admit such a simple cutoff representation, but Proposition 1 remains unchanged. The log-concavity condition is used only to obtain this simple cutoff structure, as illustrated for the uniform case in Figure 2.

infer from discounts. Then the seller would naturally like to use its information in the most direct way: offer no discount when  $v < u - p$ , because search cannot draw demand away, and offer the benchmark discount  $\tau^b(u)$  when  $v \geq u - p$ , because that is the smallest discount that deters search. We show below that this policy is indeed optimal when consumers are naive, but it breaks down once consumers are strategic. A zero discount then signals that the outside option is weak, while a positive discount signals that it is relatively attractive. As a result, the seller must trade off the direct search-deterrence effect of a discount against the inference it induces.

#### 4.1 Strategies, beliefs, and payoffs

We allow both the firm and consumers to use mixed strategies. A firm's pricing strategy specifies, for each consumer type  $(u, v)$ , a probability distribution over buy-now discounts. Formally, let  $\mu(\tau | u, v)$  denote the density of the discount  $\tau \in [0, p]$  offered to a consumer of type  $(u, v)$ , where  $\int_0^p \mu(\tau | u, v) d\tau = 1$  for all  $(u, v)$ . For a given  $u$ , let

$$\Omega(\mu | u) \equiv \{\tau \in [0, p] : \exists v \text{ such that } \mu(\tau | u, v) > 0\}$$

denote the set of discounts that the seller uses with positive probability for at least one realization of  $v$ . Whenever there is no ambiguity, we suppress the dependence on  $u$  and write  $\Omega(\mu)$ . Also, we let  $\tau_0 \equiv 0$  denote the zero discount. As we show below,  $\tau_0 \in \Omega(\mu)$  in any equilibrium.

On the consumer side, consumers differ in sophistication. A fraction  $\lambda$  of consumers are strategic and update their beliefs about  $v$  from the observed discount, while the remaining fraction  $1 - \lambda$  are naive and do not. The seller does not observe a consumer's sophistication type. Let  $\sigma_n(\tau; u)$  and  $\sigma_s(\tau; u)$  denote the probabilities that a naive or strategic consumer with valuation  $u$ , respectively, continues to search for the outside option after observing discount  $\tau$ .

**Naive consumers.** A naive consumer does not infer information about  $v$  from the discount. She therefore evaluates search exactly as in the benchmark. Thus, her search probability satisfies

$$\sigma_n(\tau; u) = \begin{cases} 1, & B(u - p) > s + \tau, \\ 0, & B(u - p) < s + \tau, \end{cases}$$

with  $\sigma_n(\tau; u) \in [0, 1]$  when  $B(u - p) = s + \tau$ .

**Strategic consumers.** A strategic consumer, by contrast, understands that the discount may depend on the seller's private information about  $v$ . Given the seller's pricing strategy  $\mu$ , after observing a discount  $\tau$  in the support of the seller's strategy, she updates her belief about  $v$  by Bayes' rule:

$$g(v | u, \tau) = \frac{\mu(\tau | u, v) g(v)}{\int_{\underline{v}}^{\bar{v}} \mu(\tau | u, \tilde{v}) dG(\tilde{v})}, \quad (6)$$

where  $g$  and  $G$  denote the prior conditional density and distribution of  $v$  given  $u$ . Let  $G(v | u, \tau)$  denote the corresponding posterior CDF. As in the benchmark, the discount raises the effective cost of search from  $s$  to  $s + \tau$ , but the expected benefit from search is now computed under the posterior induced by the observed discount:

$$B(u - p; \tau) \equiv \int_{u-p}^{\bar{v}} [1 - G(v | u, \tau)] dv. \quad (7)$$

Her search probability therefore satisfies

$$\sigma_s(\tau; u) = \begin{cases} 1, & B(u - p; \tau) > s + \tau, \\ 0, & B(u - p; \tau) < s + \tau, \end{cases}$$

with  $\sigma_s(\tau; u) \in [0, 1]$  when  $B(u - p; \tau) = s + \tau$ .

**Firm payoff.** Aggregating over these two consumer groups, the seller's expected profit from a consumer of type  $(u, v)$  is

$$\pi(u, v; \mu, \sigma_n, \sigma_s) = \int_0^p \left[ (1 - \lambda) \pi_n(\tau; u, v) + \lambda \pi_s(\tau; u, v) \right] \mu(\tau | u, v) d\tau, \quad (8)$$

where, for  $j \in \{n, s\}$ ,

$$\pi_j(\tau; u, v) = (p - \tau)(1 - \sigma_j(\tau; u)) + p \sigma_j(\tau; u) \mathbb{I}\{v < u - p\}.$$

This expression is straightforward. If the consumer buys immediately after observing discount  $\tau$ , the seller earns  $p - \tau$ . If she searches, the seller earns  $p$  only if she later returns, which occurs exactly when  $v < u - p$ . Since the seller does not observe whether the consumer is naive or strategic, expected profit averages over the two groups using the population shares  $1 - \lambda$  and  $\lambda$ . We now turn to the equilibrium analysis.

## 4.2 Equilibrium Analysis

We focus on the economically relevant case  $u - p < r$ , where search and discounting are nontrivial. When  $u - p \geq r$ , the consumer would not search even in the absence of a discount, so the seller can pool on  $\tau_0$  and earn profit  $p$ .<sup>14</sup> Let  $\mu^*$  denote the seller's equilibrium pricing strategy, and let  $\sigma_n^*$  and  $\sigma_s^*$  denote the equilibrium search strategies of naive and strategic consumers. Together with consumers' beliefs, these objects constitute a Perfect Bayesian Equilibrium (PBE).<sup>15</sup> A simple observation will be useful throughout. If the seller observes  $v < u - p$ , then offering a positive discount is weakly dominated. Even if the consumer searches, she would return to buy from the seller at the regular price  $p$ . Hence  $\tau_0 = 0$  is optimal for such consumers in any equilibrium.

We begin with the case where all consumers are naive, which reveals the separating policy the seller would like to use when discounts do not trigger inference. We then show why this policy cannot survive once consumers are strategic, and characterize the firm-optimal equilibrium that emerges under inference. Finally, we extend the analysis to markets in which naive and strategic consumers coexist.

### 4.2.1 Naive consumers: $\lambda = 0$

Suppose first that all consumers are naive, so  $\lambda = 0$ . Because discounts do not affect beliefs, the seller can use its information about  $v$  in the most direct way. When  $v < u - p$ , no discount is needed. When  $v \geq u - p$ , the seller would like to deter search at the lowest possible cost, which means offering the benchmark discount  $\tau^b(u)$  whenever it is feasible. The next lemma formalizes this separating policy.

**Lemma 1** (Naive consumers). *Suppose  $\lambda = 0$  and  $u - p < r$ . If  $\tau^b(u) < p$ , a profit-maximizing pricing strategy is deterministic ( $\mu^*(\tau^*(u, v) \mid u, v) = 1$ ) and given by*

$$\tau^*(u, v) = \begin{cases} 0, & v < u - p, \\ \tau^b(u), & v \geq u - p, \end{cases}$$

where  $\tau^b(u)$  is defined in (4). If  $\tau^b(u) \geq p$ , no feasible discount deters search, so the seller offers  $\tau_0 = 0$  to all types.

This lemma identifies the policy the seller would like to use in the absence of inference: fully

---

<sup>14</sup>As in the benchmark, there exists a pooling equilibrium in which the seller offers no discount regardless of  $v$  and all consumers purchase immediately, yielding profit  $p$ .

<sup>15</sup>A formal definition of PBE is provided in Appendix B.

separate consumers by their outside-option values and offer discounts only when search can actually draw demand away.

#### 4.2.2 Strategic consumers: $\lambda = 1$

Suppose now that all consumers are strategic, so  $\lambda = 1$ . The separating policy in Lemma 1 cannot be sustained. If the seller were to offer  $\tau_0$  only when  $v < u - p$ , then observing no discount would reveal that the outside option is weak, so the consumer would strictly prefer not to search. Anticipating this response, the seller would then have an incentive to offer  $\tau_0$  even to some consumers with  $v \geq u - p$ . Conversely, if a positive discount were offered only when  $v \geq u - p$ , then observing that discount would signal that the outside option is relatively attractive, making search more appealing. As a result, the buy-now discount  $\tau^b$  is no longer sufficient to deter search. Therefore, once discounts become informative, the naive separating policy unravels.

To mitigate consumer inference, the firm must deliberately blur the informational content of discounts. In particular, any non-pooling equilibrium must therefore leave some uncertainty after the zero discount and must ensure that positive discounts do not trigger even more search. The next lemma summarizes the necessary properties that any non-pooling equilibrium must satisfy:

**Lemma 2** (Necessary properties of non-pooling equilibria). *Suppose  $\lambda = 1$  and  $u - p < r$ . If a non-pooling equilibrium  $(\mu^*, \sigma_s^*)$  exists such that  $\Omega(\mu^*)$  contains some  $\tau > 0$ , then it must satisfy:*

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

The lemma has a simple interpretation. Part (1) says that discounting is never useful when  $v < u - p$ . Part (2) says that the zero discount must remain non-revealing, so some consumers with relatively attractive outside options must also receive  $\tau_0$ . In this sense, the seller deliberately blurs the informational content of discounts. Part (3) is the seller's on-path indifference condition: for any discount  $\tau$  that is used on path, the seller must be indifferent across all such discounts in the equilibrium support  $\Omega(\mu^*)$ ; otherwise, it would profitably shift probability mass toward the most profitable one. Part (4) says that the zero-discount signal must leave the consumer exactly

indifferent between searching and not searching, and thus, a consumer must adopt a mixed search strategy after receiving no discount, so search occurs with positive but not unit probability after  $\tau_0$ . Part (5) says that any positive discount used on path must weakly deter search, and must leave some chance of immediate purchase. If a positive discount instead made search strictly optimal, it would never generate an immediate sale and so could not be worth using in equilibrium.

Even under the requirements in Lemma 2, our signaling game admits many possible non-pooling equilibria, ranging from simple binary-discount schemes to more complex multi-level schemes (as illustrated in Appendix C).<sup>16</sup> We therefore focus on the seller's most profitable equilibrium. This selection only serves to strengthen our main result that, when consumer inference is strong, the firm is worse off having superior information about consumers' outside options relative to the symmetric-information benchmark.

### The firm-optimal equilibrium when $\lambda = 1$

Among all equilibria consistent with Lemma 2, the seller finds it optimal to use only two discounts: no discount and a single positive discount. The logic is straightforward. To keep  $\tau_0$  non-revealing, the seller must assign it to some consumers with  $v \geq u - p$ , thereby allowing some potentially costly search. The seller therefore wants to minimize the mass of such consumers. Because very high outside-option values contribute the most to the expected gain from search, assigning  $\tau_0$  to the highest  $v$  types is the least costly way to preserve the uncertainty needed after observing no discount.<sup>17</sup> This logic implies a cutoff  $\hat{v} \in (u - p, \bar{v})$  such that the seller offers the positive discount to consumers with intermediate outside-option values  $v \in [u - p, \hat{v}]$  and offers no discount to consumers with  $v < u - p$  or  $v > \hat{v}$ . Also, conditions (2) and (4) in Lemma 2 determine how much uncertainty must remain after  $\tau_0$ . Condition (2) requires that some consumers with  $v \geq u - p$  also receive  $\tau_0$ , so that the zero discount does not fully reveal a weak outside option. Condition (4) then pins down how much uncertainty must remain after  $\tau_0$  by requiring the consumer to be indifferent between searching and buying immediately. This yields the cutoff equation

$$\int_{\hat{v}}^{\bar{v}} (v - u + p - s) dG(v) = s G(u - p). \quad (9)$$

---

<sup>16</sup>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. See Appendix C for more detailed examples and discussion.

<sup>17</sup>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 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.

For any  $u - p < r$ , this equation has a unique solution  $\hat{v} \in (u - p, \bar{v})$ .<sup>18</sup>

Given  $\hat{v}$ , condition (5) in Lemma 2 also implies that any positive discount used on path must weakly deter search. In the most profitable equilibrium, the seller therefore chooses the smallest positive discount that satisfies this requirement:

$$\tau_1^*(u) = \frac{B(u - p) - s}{G(\hat{v}) - G(u - p)}, \quad (10)$$

which is strictly positive since  $u - p < r$  implies  $B(u - p) > s$ . When  $\tau_1^*(u) < p$ , this construction yields the seller's most profitable equilibrium.

**Proposition 2** (Firm-optimal equilibrium). *Suppose  $\lambda = 1$  and  $u - p < r$ . If  $\tau_1^*(u) < p$ , the following strategy profile  $(\mu^*, \sigma_s^*)$  is the most profitable equilibrium for the seller, with  $\Omega(\mu^*) = \{\tau_0, \tau_1^*(u)\}$ :*

(1)  $\mu^*(\tau_0 | u, v) = 1$  if  $v < u - p$  or  $v > \hat{v}$ , and  $\mu^*(\tau_1^*(u) | u, v) = 1$  if  $v \in [u - p, \hat{v}]$ ;

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

(3) Moreover, the optimal discount is strictly larger than the benchmark discount:  $\tau_1^*(u) > \tau^b(u)$ ,

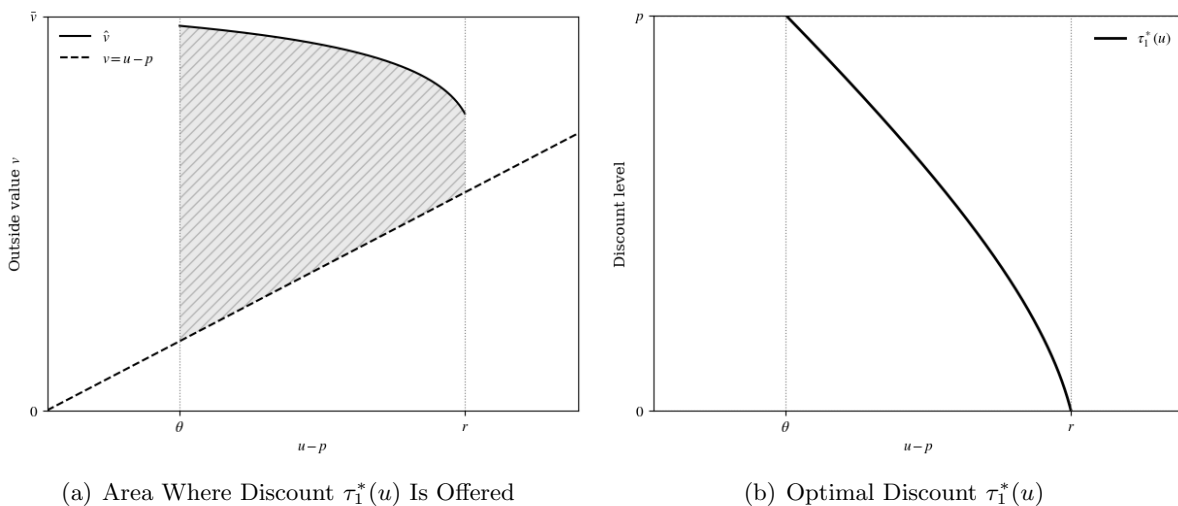
where  $\hat{v}$  is defined in (9).

Proposition 2 delivers the central equilibrium result of the paper. Even though the seller can condition discounts on a continuum of outside-option values, the most profitable equilibrium is remarkably coarse: the seller either offers no discount or a single positive discount. Consumer inference therefore makes a simple binary-discount scheme optimal rather than a richer menu of offers. Appendix B shows why. First, any equilibrium with more than two discount levels can be improved upon by collapsing all positive discounts into a single level. Second, within the class of binary-discount equilibria, the seller maximizes profit by assigning the zero discount only to the highest outside-option types. The same proof also constructs off-path beliefs under which the equilibrium survives the D1 refinement. More broadly, the proposition highlights the central tension in pricing under superior information: although the seller's information would in principle permit fine-grained price discrimination, strategic consumer inference limits the profitable use of such complexity. As a result, the seller optimally adopts a coarse pricing structure rather than a richer menu, because finer discount variation would reveal too much and induce more search.

---

<sup>18</sup>The proof is given in Appendix B.

Figure 3: Firm-optimal binary-discount equilibrium when  $\lambda = 1$ .



*Note.* Panel (a) shows the region  $v \in [u - p, \hat{v}(u)]$  in which the seller offers the positive discount  $\tau_1^*(u)$ ; outside that region the seller offers  $\tau_0 = 0$ . Panel (b) plots the positive discount level  $\tau_1^*(u)$  as a function of  $u - p$ , conditional on a positive offer. In this numerical illustration,  $\theta$  denotes the value of  $u - p$  that solves  $\tau_1^*(u) = p$ , so the binary equilibrium is plotted only for  $u - p \geq \theta$ .

Figure 3 illustrates the resulting discount scheme. It separates the two ingredients of Proposition 2. Panel (a) first shows the binary assignment rule: for each  $u - p$ , the seller offers the positive discount only to consumers with intermediate outside-option values  $v \in [u - p, \hat{v}(u)]$  and offers no discount otherwise. As shown, the range of positive discount decreases with  $u - p$ . Moreover, panel (b) plots the corresponding positive discount level  $\tau_1^*(u)$  for the numerical specification used in the figure. In this illustration, the positive discount decreases with  $u - p$ . In particular,  $\tau_1^*$  converges to 0 as  $(u - p) \rightarrow r$ , and so  $\tau_1^* < p$  must hold when  $(u - p)$  is sufficiently close to  $r$ .

Finally, the last part of the proposition shows that consumer inference forces the firm to offer a larger discount to prevent consumer search. It follows immediately from (10), since  $0 < G(\hat{v}) - G(u - p) < 1$  the positive discount in Proposition 2 is strictly larger than the benchmark discount:

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

Thus, consumer inference forces the seller to use a strictly larger positive discount than in the symmetric-information benchmark whenever a non-pooling equilibrium exists.

Thus far, our equilibrium analysis has focused on the case where the constructed optimal binary-discount equilibrium exists, which requires  $\tau_1^*(u) < p$ . When the positive discount required to

sustain inference management exceeds the list price, no non-pooling equilibrium can be sustained.

**Lemma 3.** *Suppose  $\lambda = 1$  and  $u - p < r$ . If  $\tau_1^*(u) > p$ , the only equilibrium is a pooling equilibrium in which the seller offers no discount to any consumer, that is,  $\mu^*(\tau_0 | u, v) = 1$  for all  $v$ , and the strategic consumer searches after observing  $\tau_0$ , that is,  $\sigma_s^*(\tau_0; u) = 1$ .*

Taken together, these results show how strategic inference compresses the seller’s pricing problem. Relative to the naive case, the seller can no longer exploit its information advantage in a fully separating way. To preserve enough uncertainty after the zero discount, it must deliberately forgo profits from some consumers with high outside options, accepting that some of them may search and not return. Yet the response to this constraint is not a richer menu of discounts. On the contrary, the seller optimally resorts to a simple binary scheme that balances the direct search-deterrence effect of discounts against the need to keep them from revealing too much. This reflects a fundamental trade-off: while more refined pricing could in principle better screen consumers, it would also intensify consumer inference and induce additional search. As a result, optimal pricing under strong consumer inference is characterized by simplicity rather than sophistication.

#### 4.2.3 Mixed market with both strategic and naive consumers: $\lambda \in (0, 1)$

We now extend the analysis to an environment in which naive and strategic consumers coexist. The parameter  $\lambda$  measures the prevalence of strategic consumers and therefore the strength of consumer inference in the market. As  $\lambda$  rises, a larger fraction of consumers interpret discounts as informative signals about their outside options, which increasingly constrains the seller’s ability to exploit its superior information.

For clarity, we distinguish between two equilibrium types. In an *inference-management equilibrium*, the seller keeps the zero discount partially non-revealing by assigning  $\tau_0$  to some consumers with  $v \geq u - p$ , thereby deliberately blurring the informational content of discounts to mitigate consumer inference.<sup>19</sup> In a *naive-principle equilibrium*, the seller does not do so and instead designs its pricing policy primarily to extract profit from naive consumers. The omitted equilibrium constructions and proofs are provided in Online Appendix ??.

The seller’s optimal response is governed by three threshold shares of strategic consumers:

$$\lambda_1(u) \equiv 1 - \frac{\tau_1^*(u)}{p}, \quad \lambda_2(u) \equiv 1 - \frac{p - \tau_1^*(u)}{p - \tau^b(u)}, \quad \lambda_3(u) \equiv \frac{p - \tau^b(u)}{2p - \tau^b(u)}.$$

---

<sup>19</sup>The mixed-market inference-management equilibria also must satisfy analogues of Lemma 2, which are straightforward. Conditions (1) and (5) remain unchanged. The seller’s on-path indifference condition now averages profits over strategic and naive consumers, and the indifference condition at  $\tau_0$  is weakened because naive consumers continue to search after receiving no discount.

These thresholds play different roles. The first two determine when the original binary inference-management equilibrium analogous to Proposition 2 can be sustained:  $\lambda_1(u)$  is a feasibility threshold, while  $\lambda_2(u)$  rules out profitable deviations to the benchmark discount  $\tau^b(u)$ . When  $\lambda_2(u) < \lambda < \lambda_1(u)$ , the original binary inference-management equilibrium may fail because it would require an infeasible strategic search probability after  $\tau_0$ . This does not rule out inference management. Instead, a modified binary equilibrium can still be sustained in which strategic consumers optimally choose not to search after  $\tau_0$  and search with positive probability after the positive discount  $\tau_1^*(u)$ ;  $\lambda_3(u)$  then determines when that modified equilibrium remains more profitable than reverting to naive-principle pricing.<sup>20</sup>

**Proposition 3** (Firm-optimal equilibrium with  $\lambda$ ). *Fix  $u$  and suppose  $u - p < r$  and  $\tau_1^*(u) < p$ .*

*Let*

$$\hat{\lambda}(u) \equiv \begin{cases} \lambda_3(u), & \text{if } \lambda_2(u) < \lambda_1(u), \\ \lambda_2(u), & \text{if } \lambda_2(u) \geq \lambda_1(u). \end{cases}$$

*Then the firm-optimal perfect Bayesian equilibrium exhibits a regime switch in  $\lambda$ :*

- (i) *If  $\lambda \geq \hat{\lambda}(u)$ , the firm's optimal equilibrium is an inference-management equilibrium with on-path support  $\Omega(\mu^*) = \{\tau_0, \tau_1^*(u)\}$ . In this equilibrium, the seller offers  $\tau_1^*(u)$  to consumers with  $v \in [u - p, \hat{v}]$  and offers  $\tau_0$  otherwise.*
- (ii) *If  $\lambda < \hat{\lambda}(u)$ , the firm's optimal equilibrium is a naive-principle equilibrium: the firm offers  $\tau^b(u)$  to consumers with  $v \geq u - p$  and offers  $\tau_0$  to consumers with  $v < u - p$ ; strategic consumers search upon receiving  $\tau^b(u)$  and buy immediately upon receiving  $\tau_0$ , while naive consumers buy immediately upon receiving  $\tau^b(u)$  and search upon receiving  $\tau_0$ .*

Proposition 3 shows that  $\lambda$  governs a regime switch in the seller's optimal pricing logic. When  $\lambda$  is small, the seller behaves largely as if the market were naive and uses discounts mainly for direct targeting. As  $\lambda$  rises, belief management becomes increasingly important; once  $\lambda$  exceeds  $\hat{\lambda}(u)$ , the seller optimally moves to an inference-management equilibrium. When  $\lambda$  is sufficiently large, the firm-optimal equilibrium coincides with the binary-discount equilibrium in Proposition 2. Overall, stronger consumer inference increasingly constrains the seller's use of personalized discounts, forcing

---

<sup>20</sup>More specifically, when  $\lambda < \lambda_1(u)$ , the strategic consumer's equilibrium search probability in the original binary inference-management equilibrium becomes infeasible. When  $\lambda < \lambda_2(u)$ , the seller profitably deviates to the benchmark discount  $\tau^b(u)$  and thereby extracts profit only from naive consumers. Comparing the profit from the modified inference-management equilibrium with that of the naive-principle equilibrium then yields  $\lambda_3(u)$ . See Online Appendix B for the formal constructions and proofs.

it to sacrifice surplus extraction from some high- $v$  consumers in order to keep discounts from revealing too much and thereby deter search.

## 5 The Impact of Superior Information

When the seller has superior information, it can condition buy-now discounts on consumers' outside options. Yet the value of this flexibility depends on how strongly consumers infer from the offers they receive. In this section, we compare profit and welfare across the symmetric- and superior-information regimes, and show how outcomes within the superior-information regime vary with the share of strategic consumers  $\lambda$ .

### 5.1 Firm Profit

We first compare the implications of superior information for firm profit across information regimes. For any consumer with  $u - p \geq r$ , the outcome is identical across regimes: the seller offers the zero discount  $\tau_0$ , the consumer does not search, and the seller earns profit  $p$ . We therefore focus on the economically relevant case  $u - p < r$ .

#### Benchmark profit

In the symmetric-information benchmark, as shown in Section 3, the seller either offers the search-detering discount  $\tau^b(u)$  or lets the consumer search. The corresponding equilibrium profit is

$$\Pi^b(u) = \max\{p - \tau^b(u), pG(u - p)\}. \quad (11)$$

When  $\tau_1^*(u) < p$ , condition (5) must hold, so in the region where inference management under superior information is feasible, the benchmark simplifies to

$$\Pi^b(u) = p - \tau^b(u). \quad (12)$$

#### Superior information and the role of $\lambda$

Under superior information, equilibrium pricing depends on the share  $\lambda$  of strategic consumers. When  $\lambda$  is small, the seller optimally adopts a naive-principle equilibrium and uses discounts mainly to target naive consumers. When  $\lambda$  is sufficiently large, the seller instead manages consumer inference and adopts one of the binary-discount equilibria characterized in Section 4.2.

Conditional on  $\tau_1^*(u) < p$ , the seller's equilibrium profit under superior information is

$$\Pi^*(u; \lambda) = \begin{cases} G(u-p)p + (1-G(u-p))(p - \tau_1^*(u)), & \lambda \geq \max\{\lambda_1(u), \lambda_2(u)\}, \\ G(u-p)p + (1-G(u-p))\lambda p, & \lambda_2(u) < \lambda_1(u) \text{ and } \lambda \in [\lambda_3(u), \lambda_1(u)), \\ G(u-p)p + (1-G(u-p))(1-\lambda)(p - \tau^b(u)), & \lambda < \hat{\lambda}(u). \end{cases} \quad (13)$$

The first line corresponds to the original binary inference-management equilibrium, in which strategic consumers do not search after the positive discount. The second line arises only when  $\lambda_2(u) < \lambda_1(u)$ , so that the original binary equilibrium becomes infeasible for intermediate values of  $\lambda$ , but a modified binary inference-management equilibrium can still be sustained. The third line corresponds to the naive-principle equilibrium, in which the seller effectively targets naive consumers and disregards strategic consumers' inference.<sup>21</sup>

Note that the three branches in (13) are exhaustive. When  $\lambda_2(u) < \lambda_1(u)$ , the firm-optimal equilibrium passes through three regions as  $\lambda$  rises: a naive-principle equilibrium for  $\lambda < \lambda_3(u)$ , a modified binary inference-management equilibrium for  $\lambda \in [\lambda_3(u), \lambda_1(u))$ , and the original binary inference-management equilibrium for  $\lambda \geq \lambda_1(u)$ . When  $\lambda_2(u) \geq \lambda_1(u)$ , the intermediate region disappears and the seller switches directly from naive-principle pricing to the original binary inference-management equilibrium at  $\lambda = \lambda_2(u)$ .

### When superior information helps and when it backfires

Equation (13) makes the role of  $\lambda$  transparent. When  $\lambda = 0$ , all consumers are naive, and superior information strictly raises profit because the seller can condition discounts on  $v$  and avoid unnecessary concessions to consumers with weak outside options:

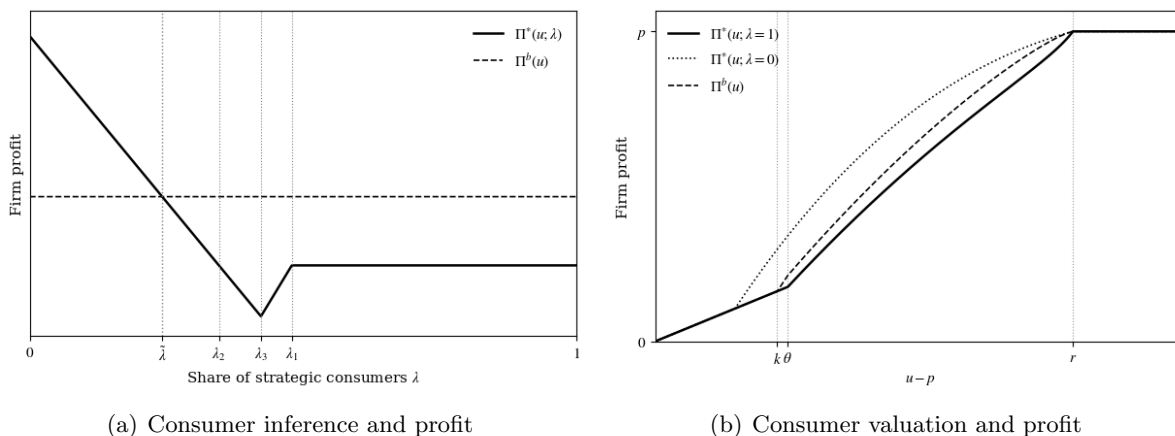
$$\Pi^*(u; 0) = G(u-p)p + (1-G(u-p))(p - \tau^b(u)) > \Pi^b(u).$$

As  $\lambda$  rises within the naive-principle region, profit falls because a larger fraction of consumers interpret the discount strategically and continue to search rather than buy immediately. Once  $\lambda$  is large enough, inference management becomes necessary. In that regime, the seller must preserve enough uncertainty after  $\tau_0$  and, as shown in Proposition 2, it must also use a strictly larger positive discount than in the symmetric-information benchmark. Profit can therefore be lower than

---

<sup>21</sup>If  $\tau_1^*(u) \geq p$ , the inference-management equilibria characterized in Section 4.2 are infeasible. In that case, when  $\lambda$  is sufficiently large the seller pools on  $\tau_0$  and earns profit  $G(u-p)p$ ; for smaller  $\lambda$ , it may still adopt a naive-principle equilibrium that uses  $\tau^b(u)$  to target naive consumers.

Figure 4: Profit consequences of superior information.



under symmetric information, even though the seller has more information. The next proposition summarizes the comparison.

**Proposition 4** (Profit impact of superior information). *Suppose  $u-p < r$ . There exists a threshold  $\tilde{\lambda}(u) \in (0, 1)$  such that*


$$\Pi^*(u; \lambda) > \Pi^b(u) \quad \text{if } \lambda < \tilde{\lambda}(u), \quad \Pi^*(u; \lambda) \leq \Pi^b(u) \quad \text{if } \lambda \geq \tilde{\lambda}(u).$$

Furthermore,  $\Pi^*(u; \lambda) = \Pi^b(u)$  for all  $\lambda \geq \tilde{\lambda}(u)$  when condition (5) does not hold. In contrast, when condition (5) holds,  $\Pi^*(u; \lambda) < \Pi^b(u)$  for all  $\lambda > \tilde{\lambda}(u)$ .

Figure 4 highlights complementary dimensions of Proposition 4. Panel (a) shows the regime switch in  $\lambda$ : superior information improves targeting and raises profit when inference is weak ( $\lambda < \tilde{\lambda}(u)$ ), but can backfire once consumers become sufficiently strategic ( $\lambda \geq \tilde{\lambda}(u)$ ). When the intermediate region exists, profit may rise with  $\lambda$  within that region, that is, for  $\lambda \in (\lambda_3(u), \lambda_1(u))$ , because inference management becomes more effective as a larger share of consumers is strategic. This increase, however, should not be interpreted as a return to the benchmark: once  $\lambda$  exceeds  $\tilde{\lambda}(u)$ , profit under superior information remains below the symmetric-information benchmark even if it is locally increasing in  $\lambda$ . Once  $\lambda$  is high enough, the firm must either offer larger discounts or allow more consumers to search and not return, both of which reduce profit relative to the symmetric-information benchmark.

Panel (b) shows the same comparison across consumer valuations. The cutoff  $k$  marks the lower boundary of the benchmark discounting region, while  $r$  marks the reservation cutoff above which search is already unattractive. Under weak inference, superior information improves targeting and

Figure 5: The Impact of Superior Information on Profit



raises profit over the nontrivial region where discounts matter. Under strong inference, by contrast, superior information lowers profit over that same interior region because consumer inference limits the seller’s ability to exploit its informational advantage. Taken together, the figure shows that acquiring superior information is profitable only when there are sufficiently many naive consumers in the market.

The mechanism behind this profit reversal is shown in Figure 5. Once the seller engages in inference management, superior information affects profit differently across realizations of the outside option. For consumers with  $v < u - p$ , the seller benefits from preserving full-price sales. For consumers with  $v \in [u - p, \hat{v}]$ , however, it must offer a strictly larger discount,  $\tau_1^*(u) > \tau^b(u)$ , to deter search. For consumers with  $v > \hat{v}$ , it withholds discounts to keep  $\tau_0$  non-revealing, accepting that some of these consumers will search and not return. The losses from deeper discounts and forgone sales can therefore outweigh the gains from full-price purchases.

This finding is also consistent with the widespread adoption of superior information and personalized pricing by firms, which also suggests a stage-zero interpretation. When few firms engage in such practices, consumers are less attentive to inference, and superior information enhances the seller’s ability to deter search and extract surplus. As such pricing strategies become more prevalent, however, consumers grow increasingly aware of the underlying tactics and respond more strategically. Heightened consumer inference then constrains pricing behavior, causing the profitability of superior information to diminish and eventually reverse. In this sense, superior information is most valuable when consumer inference is limited, and can become a liability when inference is strong.

## 5.2 Consumer and Total Welfare

We now turn to welfare implications. As in the profit analysis, we focus on the nontrivial region  $u - p < r$ , since for  $u - p \geq r$  all regimes coincide. Consumer welfare is the consumer’s expected utility net of prices and search costs, and total welfare is the sum of consumer welfare and firm profit.

## Consumer welfare

The welfare comparison is closely tied to search behavior. In the symmetric-information benchmark, the consumer's equilibrium utility is the value of entering the standard search problem with current surplus  $u - p$ . Under superior information, consumer welfare depends on whether the seller engages in inference management. In equilibrium, naive and strategic consumers respond differently to pricing signals, resulting in potentially heterogeneous welfare effects. The following proposition characterizes consumer welfare across information regimes and consumer types.

**Proposition 5** (Consumer welfare). *Suppose  $u - p < r$ . Consumer welfare satisfies:*

- (i) *under symmetric information,  $CS^b(u) = B(u - p) + u - p - s$ ;*
- (ii) *under superior information with inference management ( $\lambda \geq \hat{\lambda}(u)$ ), both naive and strategic consumers obtain  $CS_n^*(u; \lambda) = CS_s^*(u; \lambda) = B(u - p) + u - p - s$ ;*
- (iii) *under superior information without inference management ( $\lambda < \hat{\lambda}(u)$ ), naive consumers obtain  $CS_n^*(u; \lambda) = B(u - p) + u - p - s - G(u - p)B(u - p)$ , whereas strategic consumers obtain  $CS_s^*(u; \lambda) = B(u - p) + u - p - s + sG(u - p)$ .*

Proposition 5 highlights two welfare forces. First, when inference management is operative, the seller chooses discounts so that the signal-induced change in beliefs is exactly offset by the price incentive, leaving consumer welfare at the symmetric-information benchmark for both naive and strategic consumers. Second, when inference management is not operative, consumer welfare diverges by type. Naive consumers are worse off because they search inefficiently after receiving  $\tau_0$ , which generates an expected welfare loss of  $G(u - p)B(u - p)$ . Strategic consumers, by contrast, correctly infer that  $v < u - p$  after observing  $\tau_0$  and avoid unnecessary search, which generates an expected welfare gain of  $sG(u - p)$ .

## Total welfare

We next consider total welfare. Because prices are transfers in this model, the welfare comparison turns on how superior information affects search behavior. When  $\lambda < \hat{\lambda}(u)$ , the comparison is generally ambiguous. Superior information raises firm profit by allowing the seller to extract more surplus from naive consumers, but it also distorts search in opposite directions: naive consumers may over-search (a welfare loss), whereas strategic consumers avoid inefficient search (a welfare gain). As  $\lambda$  varies below  $\hat{\lambda}(u)$ , the relative strength of these forces changes, so the sign of  $W^*(u; \lambda) - W^b(u)$  is not pinned down in general.

When  $\lambda \geq \hat{\lambda}(u)$ , the seller engages in inference management and consumer welfare coincides with the symmetric-information benchmark. In that region, total-welfare differences are entirely driven by firm profit:

$$W^*(u; \lambda) - W^b(u) = \Pi^*(u; \lambda) - \Pi^b(u), \quad \text{for } \lambda \geq \hat{\lambda}(u). \quad (14)$$

**Proposition 6** (Total welfare under inference management). *Suppose  $\lambda \geq \hat{\lambda}(u)$  so that the seller engages in inference management. Aggregate consumer welfare is the same under symmetric and superior information. Moreover, total welfare under superior information is weakly lower than under symmetric information, and is strictly lower whenever condition (5) holds.*

The intuition is straightforward. Since prices are transfers, a planner would want the consumer to search if and only if  $v > u + s$ . Under inference management, the seller preserves uncertainty after  $\tau_0$  by inducing some consumers to search even when their outside options are low. In particular, consumers with  $v < u - p$  may search in equilibrium even though such search is socially wasteful, because the outside option is too weak to justify the search cost. This creates an efficiency loss relative to the symmetric-information benchmark. By contrast, higher- $v$  consumers may search more often, which can improve matching, but those gains are not large enough to offset the excessive search induced among low- $v$  consumers. In particular, the seller optimally assigns  $\tau_0$  only to a small tail of the highest- $v$  types in order to keep this pooling as cheap as possible. As a result, the beneficial matching gains are limited, while the induced wasteful search by low- $v$  consumers remains. As a result, once inference management becomes necessary, total welfare falls.

## 6 Extensions

The main model identifies the pricing consequences of consumer inference in a tractable search environment. We extend our main model into several different dimensions that place this mechanism in richer settings and clarify the scope of the results. Some preserve the main economics almost unchanged, while others add new implications without overturning the paper's central findings. Across these extensions, the same basic force remains: once consumers infer from targeted offers, the profitable use of superior information becomes constrained and pricing is pushed toward coarse menus. In this section, we briefly summarize these extensions, the details of which are relegated to the Online Appendix.

1. **Endogenous list price.** In the main model, the list price is chosen ex ante, while the buy-

now discount is chosen only after the consumer arrives and therefore serves as the signaling instrument. In Online Appendix A.3, we extend the model to allow the seller to choose the list price endogenously. This extension shows that the level of the optimal list price need not move monotonically with the information regime: depending on the distribution of valuations and the underlying parameter values, the superior-information seller may optimally choose either a higher or a lower list price than in the symmetric-information benchmark. The intuition, however, remains the same as in the main model. Because the list price is posted before the seller observes the consumer-specific information relevant for the personalized offer, endogenizing it changes price levels and cutoff regions but does not alter the informational role of the buy-now discount. As a result, the core comparative-profit insight survives: after optimizing over the list price, superior information remains more valuable when inference is weak than when it is strong.

2. **Duopoly with an endogenous outside option.** The main model treats the outside option in reduced-form terms. In Online Appendix A.4, we extend the model to a parsimonious duopoly environment in which the outside option is generated by a rival seller. Competition changes continuation values and price levels, but it does not overturn the main mechanism. The key reason is that the seller still needs to manage what consumers infer from receiving  $\tau_0$  versus a positive discount, so the incentive to compress pricing toward coarse menus survives even when the alternative is strategic rather than exogenous. This extension also delivers a new implication that is absent from the monopoly model: inference management by Firm 1 can send consumers to Firm 2 through  $\tau_0$ -induced search, so Firm 1's superior information may even raise its competitor's profit.
3. **No exogenous commitment.** The main model assumes that a buy-now discount expires once the consumer searches. In Online Appendix A.5, we extend the model to allow the seller to revise terms if the consumer returns after search. This extension does not eliminate the main equilibrium. The reason is that returning after search itself reveals a weak outside option, so continuation pricing need not be attractive to the consumer. In that sense, front-loaded pricing can arise endogenously even without literal commitment, in line with the logic emphasized by Armstrong and Zhou (2016). Relaxing exogenous commitment enlarges the equilibrium set, because some searching consumers may now return with positive probability, but the central inference-management mechanism remains valid: the seller still faces a trade-off between deterring search and avoiding excessive information revelation through discounts.

4. **Informed and uninformed consumers.** Finally, we consider a different form of consumer heterogeneity in Online Appendix A.6, motivated by the possibility that some consumers may already have learned their outside option before visiting the focal seller, for example because they previously visited another seller or acquired reliable information about relevant alternatives. A fraction  $\beta$  of consumers are uninformed and face the same inference problem as in the main model with strategic consumers, while the remaining fraction  $1 - \beta$  are already informed about  $v$  and therefore respond to discounts only through their direct price effect. This extension is useful because it varies the degree of the seller’s superior information relative to consumers. When  $\beta = 0$ , all consumers are already informed, there is no signaling problem, and the seller strictly benefits from observing  $v$ : knowing  $v$  eliminates the inefficiencies created by a uniform discount when the seller is informationally disadvantaged relative to consumers. When  $\beta = 1$ , the extension reverts to the main case with only strategic uninformed consumers, in which the signaling channel reappears and superior information can backfire. For intermediate  $\beta$ , the same inference-management logic remains present, and we show, under regularity assumptions, that the value of observing  $v$  reverses at an interior threshold and that, when  $\beta$  is sufficiently large, the coarse binary pricing structure from the main model re-emerges.

Taken together, these extensions show that the paper’s central conclusions do not hinge on exogenous list pricing, strong commitment, or a reduced-form monopoly interpretation of the outside option. They also clarify where richer environments add new implications without overturning the main result: once consumers infer from targeted offers, the profitable use of superior information becomes constrained, and pricing is pushed toward coarse menus.

## 7 Conclusion

This paper studies personalized buy-now discounts when firms possess superior information and consumers understand that targeted offers may reveal what firms know. Our main result is that, even though the seller can in principle condition discounts on a continuum of consumer-specific information, the firm-optimal equilibrium is remarkably coarse: for a given valuation, the seller either offers no discount or a single positive discount, with positive discounts targeted only to consumers with intermediate outside options. Consumer inference therefore compresses an otherwise rich pricing problem into a simple two-tier structure.

The paper also shows that the profit effect of superior information is conditional on inference

intensity. When consumer inference is weak, superior information improves targeting and can raise profit. When inference is strong, however, the same information becomes harder to exploit because discounts themselves trigger additional search. In that regime, the seller must either offer deeper discounts or allow more consumers to search and not return. As a result, superior information can backfire even though it expands the seller's informational advantage. The welfare analysis further shows that, once inference management becomes necessary, consumer surplus remains at the benchmark while total welfare falls because equilibrium search becomes excessive.

These results have both managerial and empirical implications. Managerially, they suggest that firms with strong predictive capabilities should not automatically expect more data to justify more complex price discrimination. When consumers are inference-aware, simpler and more transparent pricing formats may be more effective than finely tailored menus. Empirically, the model predicts that as firms are perceived to have stronger consumer information, and as consumers become more attentive to the informational content of targeted offers, equilibrium pricing should become coarser rather than more finely individualized.

The model focuses on the interaction between superior information, targeted discounts, and additional search after the consumer reaches the focal seller. Several extensions show that the core mechanism is robust to endogenous list pricing, relaxed commitment, competition, and heterogeneous consumer information. A natural next step would be to relax the assumption that the seller observes a consumer's outside option accurately. In practice, firms often possess only a noisy signal of a consumer's outside option. Noisier signals should attenuate the mechanism in two ways: they reduce the seller's ability to target discounts precisely and weaken the informational content of those discounts for consumers. Extending the model in that direction would test the robustness of our findings, but it would also introduce substantial analytical challenges because posterior beliefs and incentive constraints would have to be characterized jointly over noisy signals rather than realized outside options. It would also be useful to study institutional constraints on how firms can commit to or revise targeted offers. Finally, empirical work on when firms adopt coarse discount formats, and how those formats change as consumer inference becomes more salient, remains a promising direction for future research.

## 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.
- 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*.
- Guo, L. and J. Zhang (2012). Consumer deliberation and product line design. *Marketing Science* 31(6), 995–1007.
- 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.
- Ning, Z. E., J. Shin, and J. Yu (2025). Targeted advertising as an implicit recommendation and personal data opt-out. *Marketing Science* 44(2), 390–410.
- Preuss, M. (2023). Search, learning, and tracking. *The RAND Journal of Economics* 54(1), 54–82.
- Rafieian, O. and S. Zuo (2024). Personalization, algorithmic 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.

# Appendix

## A Definitions

**Definition 1** (Perfect Bayesian Equilibrium). *Fix  $u$ . A Perfect Bayesian Equilibrium (PBE) consists of a firm pricing strategy  $\mu^*(\tau | u, v)$  over  $\tau \in [0, p]$ , a search strategy  $\sigma_n^*(\tau; u)$  for naive consumers, a search strategy  $\sigma_s^*(\tau; u)$  for strategic consumers, and a belief system  $G(\cdot | u, \tau)$  for strategic consumers such that the following conditions are satisfied:*

- **(Sequential rationality of consumers)** *For every  $\tau$ ,  $\sigma_n^*(\tau; u)$  maximizes the naive consumer's expected utility, and  $\sigma_s^*(\tau; u)$  maximizes the strategic consumer's expected utility given belief  $G(\cdot | u, \tau)$ .*
- **(Sequential rationality of the firm)** *For every  $(u, v)$ , the pricing strategy  $\mu^*(\cdot | u, v)$  assigns positive probability only to discounts that maximize the firm's expected profit, taking as given the consumers' equilibrium search strategies  $\sigma_n^*$  and  $\sigma_s^*$ .*
- **(Belief consistency)** *The belief  $G(\cdot | u, \tau)$  is derived from Bayes' rule using the prior  $G$  and the equilibrium pricing strategy  $\mu^*$ .*

**Definition 2** (PBE satisfying D1). *Fix  $u$ . Let  $\pi^*(u, v)$  be the firm's equilibrium expected profit from type  $v$ . For any off-path discount  $\tau' \notin \Omega(\mu^*)$ , let  $G(\cdot | u, \tau')$  be the strategic consumer's belief and  $\sigma_s(\tau'; u)$  be the corresponding search strategy.*

*Define the deviation gain of type  $v$  from offering  $\tau'$  as*

$$\Delta(u, v; \tau') \equiv \pi(u, v; \sigma_s) - \pi^*(u, v).$$

*A PBE satisfies the D1 refinement if for every  $\tau' \notin \Omega(\mu^*)$ ,*

$$\text{supp } G(\cdot | u, \tau') \subseteq \arg \max_v \Delta(u, v; \tau').$$

*Under D1, after any off-path discount, beliefs may assign positive probability only to types with the strongest incentive to deviate, measured by deviation gains relative to equilibrium payoffs.*

## B Omitted Proofs

### B.1 Proof of Proposition 1

Results (1)-(2) have been explained in the main text.  $\square$

### B.2 Proof of Lemma 1

Suppose  $\lambda = 0$ . If  $v < u - p$ , the consumer strictly prefers the firm's product at price  $p$  even after searching, so any positive discount weakly reduces profit. If  $v \geq u - p$ , offering  $\tau^b(u)$  deters search at minimal cost, while any smaller discount induces search and any larger discount lowers profit. Since consumers are naive, no inference constraints apply.  $\square$

### B.3 Proof of Lemma 2

Fix  $\lambda = 1$  and  $u - p < r$ , and consider any non-pooling PBE  $(\mu^*, \sigma_s^*)$  such that  $\Omega(\mu^*)$  contains some  $\tau > 0$ .

**(1)  $\tau_0$  dominance for  $v < u - p$ .** If  $v < u - p$ , then even after searching the consumer prefers purchasing the firm's product at the regular price  $p$  to taking the outside option. Offering any  $\tau > 0$  strictly lowers the firm's revenue whenever a sale occurs and does not increase the probability of sale. Hence  $\mu^*(\tau_0 | u, v) = 1$  for all  $v < u - p$ .

**(5) Search deterrence for  $\tau > 0$ .** Take any  $\tau \in \Omega(\mu^*) \setminus \{\tau_0\}$ . If  $\sigma_s^*(\tau; u) = 1$ , then after observing  $\tau$  the consumer surely searches and (since  $\tau$  is offered only when  $v \geq u - p$ ) never returns, so the firm's profit from offering  $\tau$  is zero. Given our tie-breaking rule for the firm,  $\tau$  is then weakly dominated by offering  $\tau_0$ , contradicting that  $\tau$  is used on path. Therefore  $\sigma_s^*(\tau; u) < 1$ . Moreover, consumer optimality implies that not searching must be weakly optimal after observing  $\tau$ , which yields

$$B(u - p; \tau) \leq \tau + s.$$

**(4) Consumer indifference at  $\tau_0$ .** First,  $\sigma_s^*(\tau_0; u) > 0$ . If instead  $\sigma_s^*(\tau_0; u) = 0$ , then the firm can guarantee an immediate sale at price  $p$  by offering  $\tau_0$  to any type with  $v \geq u - p$ , and thus would never optimally offer any  $\tau > 0$ , contradicting non-pooling. Second,  $\sigma_s^*(\tau_0; u) < 1$ . If  $\sigma_s^*(\tau_0; u) = 1$ , then for any  $v \geq u - p$  with  $\mu^*(\tau_0 | u, v) > 0$  the firm earns zero profit when offering  $\tau_0$ , while deviating to any  $\tau > 0 \in \Omega(\mu^*)$  yields strictly positive profit  $(p - \tau)(1 - \sigma_s^*(\tau; u)) > 0$  by the

previous paragraph, a contradiction. Hence  $0 < \sigma_s^*(\tau_0; u) < 1$ . Consumer mixing at  $\tau_0$  then implies indifference between searching and not searching, so  $B(u - p; \tau_0) = s$ .

**(2)  $\tau_0$  non-revelation.** If  $\mu^*(\tau_0 \mid u, v) = 0$  for all  $v \geq u - p$ , then observing  $\tau_0$  would reveal  $v < u - p$ . In that case a strategic consumer would strictly prefer not to search after  $\tau_0$ , implying  $\sigma_s^*(\tau_0; u) = 0$ , contradicting  $0 < \sigma_s^*(\tau_0; u)$  established above. Therefore  $\mu^*(\tau_0 \mid u, v) > 0$  for a positive measure of  $v \geq u - p$ .

**(3) Firm indifference across on-path discounts.** For any  $v \geq u - p$ , if the firm offers  $\tau \in \Omega(\mu^*)$ , its expected profit is

$$(p - \tau)(1 - \sigma_s^*(\tau; u)),$$

since the consumer does not return after searching. This expression is independent of  $v$ . In a non-pooling equilibrium, all discounts in  $\Omega(\mu^*)$  that are chosen with positive probability must yield the same profit; otherwise the firm would profitably shift probability mass toward a more profitable on-path discount. Hence  $(p - \tau)(1 - \sigma_s^*(\tau; u))$  is constant for all  $\tau \in \Omega(\mu^*)$ .  $\square$

## B.4 Proof of Proposition 2

Fix  $\lambda = 1$  and  $u - p < r$ . Let  $\hat{v}$  solve (9) and let  $\tau_1^*$  be defined by (10). Assume  $\tau_1^* < p$ .

**Preliminary: existence and uniqueness of  $\hat{v}$ .** Define

$$\Phi(z) \equiv \int_z^{\bar{v}} (v - u + p - s) dG(v) - sG(u - p).$$

Since  $u - p < r$ , we have  $B(u - p) > s$ , and therefore

$$\Phi(u - p) = B(u - p) - s > 0.$$

Also,

$$\Phi(\bar{v}) = -sG(u - p) \leq 0.$$

Moreover,

$$\Phi'(z) = -(z - u + p - s)g(z),$$

so  $\Phi$  is single-peaked: it increases for  $z < u - p + s$  and decreases for  $z > u - p + s$ . Hence  $\Phi$  crosses zero at most once on  $(u - p, \bar{v}]$ . Since it is positive at  $u - p$  and weakly negative at  $\bar{v}$ ,

equation (9) has a unique solution  $\hat{v} \in (u - p, \bar{v}]$ . In the non-pooling equilibrium characterized below, this solution is interior, so  $\hat{v} \in (u - p, \bar{v})$ .

**Step 1. Existence: the proposed profile is a PBE.** We specify beliefs on path as follows. After observing  $\tau_1^*$ , the consumer assigns probability one to  $v \in [u - p, \hat{v}]$  (since  $\mu^*(\tau_1^* | u, v) = 1$  on this region and zero elsewhere). After observing  $\tau_0$ , the consumer assigns probability one to  $v \in (-\infty, u - p) \cup (\hat{v}, \bar{v}]$  (since  $\mu^*(\tau_0 | u, v) = 1$  on these regions).

*Consumer optimality.* By construction, (9) is exactly the indifference condition at  $\tau_0$ , i.e.  $B(u - p; \tau_0) = s$ . Hence the strategic consumer is willing to mix after  $\tau_0$ . Moreover, (10) is chosen so that the search-deterrence constraint at the positive discount binds, i.e.  $B(u - p; \tau_1^*) = \tau_1^* + s$ . Under the tie-breaking rule, the consumer does not search after  $\tau_1^*$ , so  $\sigma_s^*(\tau_1^*; u) = 0$  is optimal.

*Firm optimality on path.* For  $v < u - p$ , offering  $\tau_0$  weakly dominates any  $\tau > 0$  because the consumer returns to buy at price  $p$  even if she searches; thus  $\mu^*(\tau_0 | u, v) = 1$  is optimal. For any  $v \geq u - p$ , if the firm offers  $\tau_1^*$  it earns  $p - \tau_1^*$  (since  $\sigma_s^*(\tau_1^*; u) = 0$ ). If it offers  $\tau_0$ , it earns  $p(1 - \sigma_s^*(\tau_0; u)) = p - \tau_1^*$  because  $\sigma_s^*(\tau_0; u) = \tau_1^*/p$ . Hence the firm is indifferent between  $\tau_0$  and  $\tau_1^*$  for  $v \geq u - p$ , and the proposed  $\mu^*$  is a best response (with  $\tau_0$  selected on  $v > \hat{v}$  by tie-breaking). This verifies sequential rationality on path.

**Step 2. Off-path deviations and D1.** Fix  $u$  and consider any off-path discount  $\tau \notin \{\tau_0, \tau_1^*\}$ . Let  $G(\cdot | u, \tau)$  denote the strategic consumer's posterior belief and let  $\sigma \in [0, 1]$  be her optimal search probability under this belief.

For a consumer with  $v \geq u - p$ , the firm's equilibrium profit equals  $p - \tau_1^*$ , while deviating to  $\tau$  yields profit  $(1 - \sigma)(p - \tau)$ . The deviation gain is therefore

$$\Delta_1(\tau, \sigma) = (1 - \sigma)(p - \tau) - (p - \tau_1^*).$$

For  $v < u - p$ , the equilibrium profit equals  $p$ , and deviating to  $\tau$  yields  $(1 - \sigma)(p - \tau) + \sigma p$ , with deviation gain

$$\Delta_2(\tau, \sigma) = (1 - \sigma)(p - \tau) + \sigma p - p.$$

A direct comparison shows that  $\Delta_2(\tau, \sigma) > \Delta_1(\tau, \sigma)$  if and only if  $\sigma p > \tau_1^*$ . Under the D1 criterion, if  $\sigma p > \tau_1^*$  were to hold, beliefs must assign probability one to  $v < u - p$ . Under such beliefs, however, the consumer strictly prefers not to search, implying  $\sigma = 0$  and hence  $\sigma p \leq \tau_1^*$ , a contradiction. Therefore, D1-consistent beliefs must induce  $\sigma p \leq \tau_1^*$ .

We now construct off-path beliefs consistent with D1. Consider first  $\tau > \tau_1^*$ . There exists a posterior belief supported on  $v \geq u - p$  such that  $B(u - p; \tau) < \tau + s$ . One convenient choice is to let the off-path posterior coincide with the equilibrium posterior induced by  $\tau_1^*$ , i.e.,  $G(\cdot | u, \tau) = G(\cdot | u, \tau_1^*)$ . Since  $B(u - p; \tau_1^*) = \tau_1^* + s < \tau + s$ , this belief induces  $\sigma = 0$ . Conditional on this response,

$$\Delta_2(\tau, 0) < \Delta_1(\tau, 0).$$

Thus types  $v \geq u - p$  have strictly higher deviation gains than types  $v < u - p$ , so the constructed posterior is D1-consistent. Under this belief, deviating to  $\tau$  is strictly unprofitable.

If instead  $\tau < \tau_1^*$ , we choose a posterior satisfying  $B(u - p; \tau) = \tau + s$ , which renders the strategic consumer indifferent between searching and buying immediately; such a belief always exists. Selecting the mixed response with  $\sigma = \tau_1^*/p$  yields  $\sigma p = \tau_1^*$ , and hence

$$\Delta_1(\tau, \sigma) = \Delta_2(\tau, \sigma) = -\frac{(p - \tau_1^*)\tau}{p} \leq 0,$$

with strict inequality for  $\tau < \tau_1^*$ . Since all types attain the same deviation gain, the posterior is D1-consistent.

In all cases, there exists a D1-consistent posterior belief under which deviating to any off-path discount  $\tau$  yields weakly lower profit than the equilibrium payoff. Hence, the equilibrium survives the D1 refinement.

**Step 3. Optimality among binary-discount equilibria.** Consider any non-pooling binary-discount PBE with support  $\{\tau_0, \tau_1\}$  where  $\tau_1 > 0$ . By Lemma 2, such an equilibrium must satisfy: (i)  $B(u - p; \tau_0) = s$  and  $0 < \sigma_s(\tau_0; u) < 1$ ; (ii)  $\sigma_s(\tau_1; u) < 1$  and  $B(u - p; \tau_1) \leq \tau_1 + s$ ; and (iii) firm indifference across on-path discounts.

In a firm-optimal binary-discount equilibrium, we may, without loss of generality, impose

$$\sigma_s(\tau_1; u) = 0 \quad \text{and} \quad B(u - p; \tau_1) = \tau_1 + s.$$

If  $\sigma_s(\tau_1; u) > 0$ , then reducing  $\sigma_s(\tau_1; u)$  (and adjusting  $\sigma_s(\tau_0; u)$  to maintain firm indifference) strictly increases profit. If instead  $B(u - p; \tau_1) < \tau_1 + s$  with  $\sigma_s(\tau_1; u) = 0$ , then the firm can reduce  $\tau_1$  slightly without inducing search, again strictly increasing profit. Hence both equalities must bind at the optimum.

Under these conditions, the firm's expected profit in the market  $u - p < r$  equals

$$\Pi = p \cdot G(u - p) + (1 - G(u - p))(p - \tau_1),$$

so maximizing profit is equivalent to minimizing  $\tau_1$ .

Let  $\Pr(\tau_0)$  denote the ex ante probability that the firm offers  $\tau_0$  to a consumer of type  $u$ , and let  $\mu_0(v) \equiv \mu(\tau_0 | u, v)$  for  $v \geq u - p$ . Then

$$\Pr(\tau_0) = G(u - p) + \int_{u-p}^{\bar{v}} \mu_0(v) dG(v).$$

By the law of iterated expectations,

$$\Pr(\tau_0) B(u - p; \tau_0) + (1 - \Pr(\tau_0)) B(u - p; \tau_1) = B(u - p).$$

Using  $B(u - p; \tau_0) = s$  and  $B(u - p; \tau_1) = \tau_1 + s$ , we obtain

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

Since  $B(u - p) > s$  for  $u - p < r$ , minimizing  $\tau_1$  is equivalent to minimizing  $\Pr(\tau_0)$ , that is, minimizing  $\int_{u-p}^{\bar{v}} \mu_0(v) dG(v)$  subject to  $B(u - p; \tau_0) = s$ .

Because  $\mu(\tau_0 | u, v) = 1$  for all  $v < u - p$ , satisfying  $B(u - p; \tau_0) = s$  requires assigning  $\tau_0$  to a positive measure of types with  $v \geq u - p$ . Among all such assignments, the way to minimize  $\Pr(\tau_0)$  is to allocate  $\tau_0$  only to the highest  $v$  types. Consequently, the minimizing  $\mu_0(\cdot)$  takes a cutoff form:

$$\mu_0(v) = 0 \quad \text{for } v \in [u - p, \hat{v}], \quad \mu_0(v) = 1 \quad \text{for } v > \hat{v},$$

where  $\hat{v}$  is uniquely determined by the indifference condition  $B(u - p; \tau_0) = s$ , which is exactly (9). Substituting  $\Pr(\tau_0) = G(u - p) + 1 - G(\hat{v})$  into (15) yields (10). Therefore, among all binary-discount equilibria, the profile in Proposition 2 attains the smallest feasible  $\tau_1$  and hence the highest profit.

**Step 4. Binary discounts dominate equilibria with three or more discounts.** Let  $(\tilde{\mu}, \tilde{\sigma}_s)$  be any non-pooling equilibrium with  $|\Omega(\tilde{\mu})| \geq 3$ . By Lemma 2, all on-path discounts must deliver the same profit from  $v \geq u - p$ . Let  $\tilde{\tau}$  be a positive discount in  $\Omega(\tilde{\mu})$  that minimizes the firm's revenue loss among positive discounts (equivalently, that maximizes the firm's profit conditional on using a positive discount). Construct a binary-discount profile that (i) keeps the same probability

of  $\tau_0$  as under  $(\tilde{\mu}, \tilde{\sigma}_s)$  and (ii) replaces all positive discounts with a single discount  $\hat{\tau}$  chosen so that  $B(u - p; \hat{\tau}) = \hat{\tau} + s$  and  $\sigma_s(\hat{\tau}; u) = 0$ . Since  $\hat{\tau}$  is chosen to be just sufficient to deter search given the induced posterior, it weakly increases the firm's profit relative to using any positive discount that deters search only weakly (i.e., with slack  $B(u - p; \tau) < \tau + s$ ), and it strictly increases profit whenever the original equilibrium uses more than one positive discount level. The resulting binary profile can be supported by the same off-path belief construction as in Step 2 and therefore constitutes a feasible equilibrium with weakly higher profit.

Combining Steps 3 and 4, the binary-discount equilibrium characterized in the proposition is the most profitable equilibrium for the firm whenever  $\tau_1^* < p$ .

Finally, part (3) follows directly from (10). Because  $\hat{v} \in (u - p, \bar{v})$ , we have  $0 < G(\hat{v}) - G(u - p) < 1$ , which implies  $\tau_1^*(u) > B(u - p) - s = \tau^b(u)$ .  $\square$

## B.5 Proof of Lemma 3

Fix  $\lambda = 1$  and  $u - p < r$ , and suppose  $\tau_1^* > p$ , where  $\tau_1^*$  is defined in (10).

**Step 1: No non-pooling equilibrium can exist.** Suppose, toward a contradiction, that there exists a non-pooling PBE  $(\mu, \sigma_s)$  with some  $\tau > 0$  in  $\Omega(\mu)$ . By Lemma 2, we must have

$$B(u - p; \tau_0) = s \quad \text{and} \quad \mu(\tau_0 \mid u, v) > 0 \quad \text{for a positive measure of } v \geq u - p,$$

and firm indifference implies that the firm earns the same profit from any on-path discount offered to types  $v \geq u - p$ . In particular, we can w.l.o.g. focus on an on-path positive discount  $\bar{\tau} \in \Omega(\mu) \setminus \{\tau_0\}$  such that  $\sigma_s(\bar{\tau}; u) = 0$  and  $B(u - p; \bar{\tau}) = \bar{\tau} + s$  (otherwise, either  $\sigma_s(\bar{\tau}; u) > 0$  or  $B(u - p; \bar{\tau}) < \bar{\tau} + s$  would allow lowering  $\bar{\tau}$  while preserving incentive compatibility, increasing profit).

Let  $\Pr(\tau_0)$  denote the equilibrium probability (over  $v$ ) that the firm offers  $\tau_0$  to a consumer of type  $u$ . By the law of iterated expectations,

$$\Pr(\tau_0) B(u - p; \tau_0) + (1 - \Pr(\tau_0)) B(u - p; \bar{\tau}) = B(u - p).$$

Substituting  $B(u - p; \tau_0) = s$  and  $B(u - p; \bar{\tau}) = \bar{\tau} + s$  yields

$$\bar{\tau} = \frac{B(u - p) - s}{1 - \Pr(\tau_0)}. \tag{16}$$

Given  $B(u - p) > s$  for  $u - p < r$ , (16) implies that  $\bar{\tau}$  is decreasing in  $1 - \Pr(\tau_0)$ , i.e., it is minimized

by maximizing  $1 - \Pr(\tau_0)$ .

By construction,  $\tau_1^*$  in (10) is the minimal feasible positive discount that can satisfy  $B(u - p; \tau_0) = s$  while keeping the on-path positive discount just sufficient to deter search. Hence any non-pooling with a positive discount must satisfy  $\bar{\tau} \geq \tau_1^*$ . Since  $\tau_1^* > p$ , this implies  $\bar{\tau} > p$ , contradicting the constraint that discounts lie in  $[0, p]$ . Thus, no non-pooling equilibrium exists.

**Step 2: The remaining equilibrium is pooling with  $\tau_0$  and  $\sigma_s(\tau_0; u) = 1$ .** It follows that any equilibrium must be pooling with  $\mu^*(\tau_0 | u, v) = 1$  for all  $v$ . Given  $u - p < r$ , we have  $B(u - p) > s$ , so under pooling at  $\tau_0$  the strategic consumer strictly prefers to search, implying  $\sigma_s^*(\tau_0; u) = 1$ .

Finally, to support pooling, assign off-path beliefs so that any  $\tau > 0$  is attributed to types with  $v \geq u - p$ , making search optimal. Any such deviation yields zero profit and is strictly worse than the pooling payoff  $G(u - p)p$ . Hence, the pooling profile is a PBE.  $\square$

## B.6 Proof of Proposition 4

Let  $x \equiv u - p$ ,  $G_x \equiv G(u - p)$ ,  $\tau^b \equiv \tau^b(u)$ ,  $\tau_1^* \equiv \tau_1^*(u)$ , and define  $\Delta(\lambda) \equiv \Pi^*(u; \lambda) - \Pi^b(u)$ .

**Case 1: condition (5) holds.** Then  $\Pi^b(u) = p - \tau^b$ .

For  $\lambda < \hat{\lambda}(u)$ , (13) implies

$$\Delta(\lambda) = G_x \tau^b - (1 - G_x)(p - \tau^b)\lambda.$$

Hence  $\Delta(\lambda)$  is continuous and strictly decreasing on the naive-principle region, with

$$\Delta(0) = G_x \tau^b > 0.$$

If  $\lambda \geq \max\{\lambda_1(u), \lambda_2(u)\}$ , then

$$\Delta(\lambda) = G_x p + (1 - G_x)(p - \tau_1^*) - (p - \tau^b) = \tau^b - (1 - G_x)\tau_1^*.$$

Using (10),

$$\tau_1^* = \frac{\tau^b}{G(\hat{v}) - G_x},$$

and since  $G(\hat{v}) - G_x < 1 - G_x$ , we obtain

$$\Delta(\lambda) < 0.$$

If  $\lambda_2(u) < \lambda_1(u)$  and  $\lambda \in [\lambda_3(u), \lambda_1(u))$ , then

$$\Delta(\lambda) = G_x p + (1 - G_x) \lambda p - (p - \tau^b),$$

which is increasing in  $\lambda$ . Moreover,

$$\Delta(\lambda_1(u)) = G_x p + (1 - G_x)(p - \tau_1^*) - (p - \tau^b) < 0,$$

because  $\lambda_1(u)p = p - \tau_1^*$ . Hence  $\Delta(\lambda) < 0$  throughout this intermediate region as well.

Therefore, regardless of whether the intermediate region exists,  $\Delta(\lambda)$  is positive at  $\lambda = 0$  and negative once  $\lambda \geq \hat{\lambda}(u)$ . Since it is continuous and strictly decreasing on the naive-principle region, there exists a unique threshold  $\tilde{\lambda}(u) \in (0, \hat{\lambda}(u))$  such that

$$\Pi^*(u; \lambda) > \Pi^b(u) \quad \text{for } \lambda < \tilde{\lambda}(u), \quad \Pi^*(u; \lambda) < \Pi^b(u) \quad \text{for } \lambda > \tilde{\lambda}(u).$$

**Case 2: condition (5) does not hold.** Then  $\Pi^b(u) = G_x p$ .

If  $\tau^b(u) < p$ , the naive-principle equilibrium yields

$$\Pi^*(u; \lambda) = G_x p + (1 - G_x)(1 - \lambda)(p - \tau^b),$$

which is strictly above  $\Pi^b(u)$  for sufficiently small  $\lambda$ . For sufficiently large  $\lambda$ , inference management is infeasible or unprofitable, so the seller pools on  $\tau_0$  and earns

$$\Pi^*(u; \lambda) = G_x p = \Pi^b(u).$$

Hence there exists a threshold  $\tilde{\lambda}(u)$  such that

$$\Pi^*(u; \lambda) = \Pi^b(u) \quad \text{for all } \lambda \geq \tilde{\lambda}(u).$$

If  $\tau^b(u) \geq p$ , no feasible search-deterring discount exists, so both regimes pool on  $\tau_0$  and earn  $G_x p$ . This proves Proposition 4.  $\square$

## B.7 Proof of Proposition 5

**(i) Symmetric information.** In the symmetric-information benchmark, the consumer obtains the same payoff whether the seller deters search with  $\tau^b(u) = B(x) - s$  (where  $x \equiv u - p$ ) or lets

her search. In either case, her payoff is  $x + B(x) - s$ .

Hence

$$CS^b(u) = B(u - p) + u - p - s.$$

**(ii) Superior information with inference management.** Suppose  $\lambda \geq \hat{\lambda}(u)$ . In any inference-management equilibrium, the seller offers  $\tau_1^*$  if and only if  $v \in [x, \hat{v}]$ , and the equilibrium conditions imply

$$B(x; \tau_0) = s \quad \text{and} \quad B(x; \tau_1^*) = s + \tau_1^*.$$

Thus a strategic consumer is indifferent between searching and not searching after either on-path discount, so her ex ante payoff is

$$CS_s^*(u; \lambda) = x + (G(\hat{v}) - G(x))\tau_1^*.$$

Using (10),

$$(G(\hat{v}) - G(x))\tau_1^* = B(x) - s,$$

and therefore

$$CS_s^*(u; \lambda) = x + B(x) - s.$$

A naive consumer searches after  $\tau_0$  and buys immediately after  $\tau_1^*$ . Her ex ante payoff is

$$CS_n^*(u; \lambda) = G(x)(x - s) + (G(\hat{v}) - G(x))(x + \tau_1^*) + \int_{\hat{v}}^{\bar{v}} (v - s) dG(v).$$

Rearranging,

$$CS_n^*(u; \lambda) = x + (G(\hat{v}) - G(x))\tau_1^* + \int_{\hat{v}}^{\bar{v}} (v - x) dG(v) - s(1 - G(\hat{v}) + G(x)).$$

By (9),

$$\int_{\hat{v}}^{\bar{v}} (v - x) dG(v) = s(1 - G(\hat{v}) + G(x)),$$

so the last two terms cancel. Hence

$$CS_n^*(u; \lambda) = x + (G(\hat{v}) - G(x))\tau_1^* = x + B(x) - s.$$

Therefore,

$$CS_n^*(u; \lambda) = CS_s^*(u; \lambda) = B(u - p) + u - p - s.$$

**(iii) Superior information without inference management.** Suppose  $\lambda < \hat{\lambda}(u)$ , so the seller uses the naive-principle equilibrium.

A naive consumer searches after  $\tau_0$  and buys immediately after  $\tau^b(u)$ . Hence

$$CS_n^*(u; \lambda) = G(x)(x - s) + (1 - G(x))(x + \tau^b(u)) = x + B(x) - s - G(x)B(x).$$

A strategic consumer does the opposite: she buys immediately after  $\tau_0$  and searches after  $\tau^b(u)$ . Relative to the benchmark, she therefore saves the search cost exactly when  $v < x$ , which occurs with probability  $G(x)$ . Thus

$$CS_s^*(u; \lambda) = CS^b(u) + sG(x) = x + B(x) - s + sG(x). \quad \square$$

## B.8 Proof of Proposition 6

Suppose  $\lambda \geq \hat{\lambda}(u)$ . By Proposition 5, aggregate consumer welfare is the same under symmetric and superior information in this region. Hence, by (14),

$$W^*(u; \lambda) - W^b(u) = \Pi^*(u; \lambda) - \Pi^b(u).$$

Proposition 4 therefore implies that total welfare under superior information is weakly lower than under symmetric information. Moreover, when condition (5) holds, the proof of Proposition 4 shows that  $\tilde{\lambda}(u) < \hat{\lambda}(u)$ , so every  $\lambda \geq \hat{\lambda}(u)$  lies in the region where  $\Pi^*(u; \lambda) < \Pi^b(u)$ . Therefore,

$$W^*(u; \lambda) < W^b(u).$$

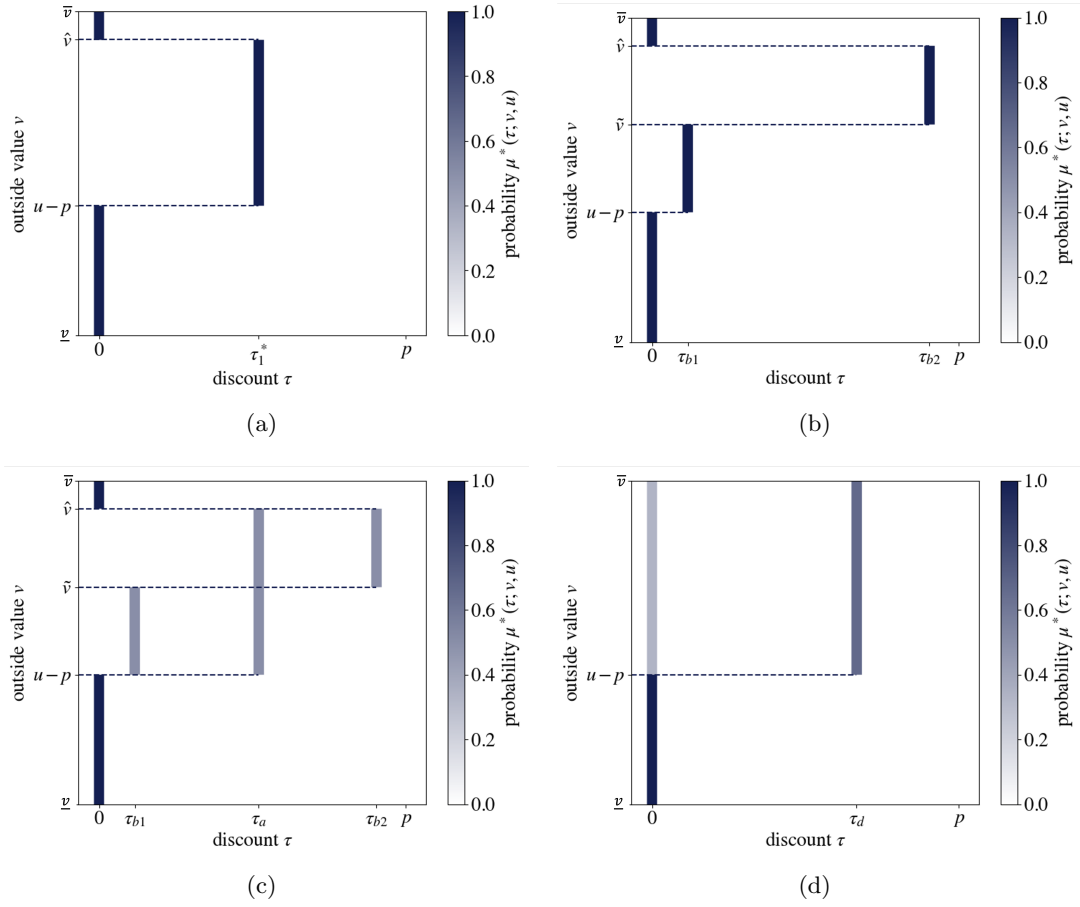
This proves Proposition 6. □

## C Examples of potential equilibrium discount schemes

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 6 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.

In Figure 6(a), the firm employs a binary discount scheme, offering a fixed discount ( $\tau_a$ ) to deter search among consumers with moderately high outside values and no discount ( $\tau_0 = 0$ ) to

Figure 6: Examples of equilibrium discount scheme



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 6(b) illustrates a multi-level discount strategy, assigning different discounts ( $\tau_{b1}$  and  $\tau_{b2}$ ) 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 6(c) exemplifies this by combining previously described binary and multi-level schemes, highlighting the substantial flexibility in equilibrium construction. Lastly, Figure 6(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.