Online Appendix for "Personalized Discounts and Consumer Search"

A.1 Symmetric-Information Benchmark: Further Characterization

This section provides a sharper characterization of the benchmark policy under a mild shape restriction, together with a numerical illustration. Neither result is needed for Proposition 1; they are included only to further visualize the benchmark.

Lemma A1. *Suppose the survival function $1 - G(v)$ is log-concave on $[\underline{v}, \overline{v}]$. Then there exists a cutoff $k \in [\underline{v}, r)$ such that condition (5) holds if and only if*

$$u - p \in (k, r).$$

Equivalently, for $u - p < r$, the seller offers a positive buy-now discount in the symmetric-information benchmark if and only if $u - p \in (k, r)$, and offers no discount otherwise.

Proof. Let $x \equiv u - p < r$, and write $S(x) \equiv 1 - G(x)$. Since $B(r) = s$, condition (5) can be rewritten as

$$p > \frac{B(x) - s}{1 - G(x)} = \frac{\int_x^r S(v)\,dv}{S(x)}.$$

Define

$$A(x) \equiv \int_x^r S(v)\,dv.$$

Under log-concavity of S, a standard preservation result for log-concavity implies that A is also log-concave. Hence

$$\frac{d}{dx}\log A(x) = \frac{A'(x)}{A(x)} = -\frac{S(x)}{A(x)}$$

is decreasing in x, which implies that $S(x)/A(x)$ is increasing and therefore $A(x)/S(x)$ is decreasing. Thus the right-hand side of the condition above is decreasing in x.

Because $A(r) = 0$, the condition holds for all x sufficiently close to r. Since the right-hand side is monotone, the set of $x < r$ for which the condition holds must be an interval of the form (k, r) for some $k \in [\underline{v}, r)$. This proves the cutoff result. □

The log-concavity assumption is satisfied by many commonly used distributions, including the uniform, exponential, and truncated normal distributions. Figure 1 illustrates the benchmark policy when $G(v)$ is uniform on $[0, 1]$. When $u - p \geq r$, no discount is needed because the consumer does not search even at the regular price. When $u - p < k$, deterring search is too costly relative to the profit from letting the consumer search and possibly return. For intermediate values $u - p \in (k, r)$, the seller offers the search-deterring discount $\tau^b(u)$, which decreases in $u - p$.

Figure 1: Symmetric-Information Benchmark under Uniform $G(v)$.

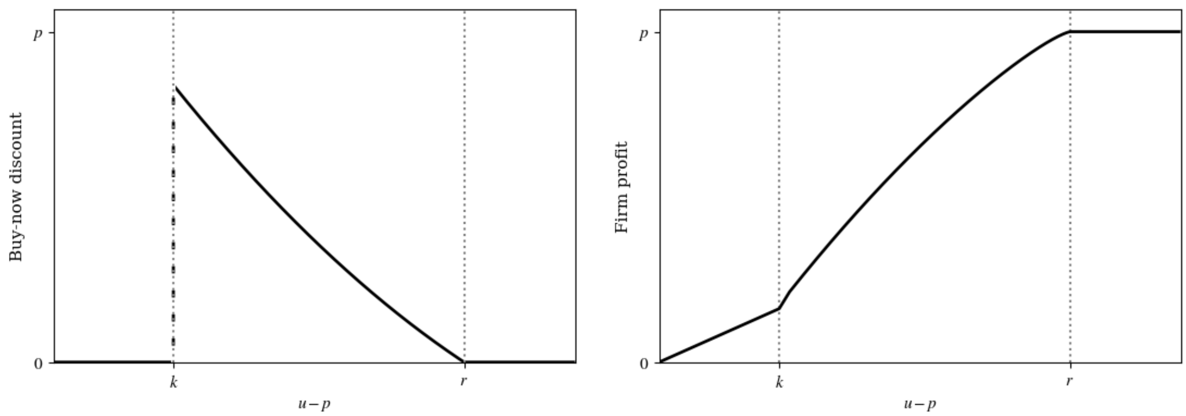

(a) Optimal buy-now discount (b) Equilibrium profit

A.2 Additional Equilibrium Results for Mixed Markets

This subsection collects the equilibrium constructions omitted from the main text for the case $\lambda \in (0,1)$. We first state the two inference-management equilibria that underlie Proposition 3, and then prove the regime-switch result.

For mixed markets, we say that an equilibrium is an *inference-management equilibrium* if the seller assigns τ_0 with positive probability to a positive measure of types with $v \geq u-p$, so that the zero discount remains partially non-revealing. It is a *naive-principle equilibrium* otherwise. In an inference-management equilibrium, the mixed-market analogues of Lemma 2 in the main text apply. Conditions (1) and (5) remain unchanged. The seller's on-path indifference condition becomes

$$\lambda(p-\tau)\big(1-\sigma_s(\tau;u)\big)+(1-\lambda)(p-\tau)\big(1-\sigma_n(\tau;u)\big)$$

constant across all on-path discounts $\tau \in \Omega(\mu)$. The analogue of condition (4) is weaker than in the case $\lambda = 1$: it requires only $0 \leq \sigma_s(\tau_0;u) < 1$ and $B(u-p;\tau_0) = s$, since the presence of naive consumers allows τ_0 to be sustained without interior mixing by strategic consumers.

Lemma A2 (Inference-management equilibrium for large λ)**.** *Suppose $u-p<r$ and $\tau_1^*(u)<p$. If*

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

where

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

then there exists a perfect Bayesian equilibrium with $\Omega(\mu^*) = \{\tau_0, \tau_1^*(u)\}$ and

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

(2) $\sigma_s^*(\tau_0;u) = \dfrac{\tau_1^*(u)-(1-\lambda)p}{\lambda p}$ and $\sigma_s^*(\tau_1^*(u);u)=0$;

(3) $\sigma_n^*(\tau_0;u)=1$ and $\sigma_n^*(\tau_1^*(u);u)=0$,

where $\tau_1^*(u)$ and $\hat{v}$ are defined in (10) *and* (9)*, respectively.*

Proof. Fix u with $u-p<r$ and suppose $\tau_1^*(u)<p$. Let $\hat{v}$ solve (9), and write $\tau_1^* \equiv \tau_1^*(u)$. Consider the candidate profile described in the lemma:

$$\Omega(\mu^*) = \{\tau_0, \tau_1^*\},$$

$$\mu^*(\tau_0 \mid u, v) = \mathbf{1}\{v < u-p \text{ or } v > \hat{v}\}, \qquad \mu^*(\tau_1^* \mid u, v) = 1 - \mu^*(\tau_0 \mid u, v),$$

$$\sigma_n^*(\tau_0; u) = 1, \qquad \sigma_n^*(\tau_1^*; u) = 0,$$

$$\sigma_s^*(\tau_1^*; u) = 0, \qquad \sigma_s^*(\tau_0; u) = \frac{\tau_1^* - (1-\lambda)p}{\lambda p}.$$

We verify that this profile constitutes a perfect Bayesian equilibrium when $\lambda \geq \max\{\lambda_1(u), \lambda_2(u)\}$.

Step 1. Beliefs and consumers' optimality. Take Bayes-consistent beliefs on the equilibrium path. After observing τ_1^*, a strategic consumer infers $v \in [u-p, \hat{v}]$; after observing τ_0, she infers $v \in [\underline{v}, u-p) \cup (\hat{v}, \bar{v}]$.

Naive consumers do not update from discounts. Since $u-p<r$ implies $B(u-p)>s$, a naive consumer strictly prefers to search after receiving τ_0, so $\sigma_n^*(\tau_0; u) = 1$. Because $\tau_1^*(u) > \tau^b(u)$, she strictly prefers to buy immediately after receiving τ_1^*, so $\sigma_n^*(\tau_1^*; u) = 0$.

For strategic consumers, (9) implies $B(u-p; \tau_0) = s$, so the consumer is indifferent after receiving τ_0. Likewise, (10) implies $B(u-p; \tau_1^*) = s + \tau_1^*$, so she is also indifferent after receiving τ_1^*. Hence the proposed search probabilities are sequentially rational.

Step 2. Firm optimality on the equilibrium path. If $v < u-p$, offering τ_0 weakly dominates any positive discount, since the consumer returns to buy at price p even if she searches. Hence $\mu^*(\tau_0 \mid u, v) = 1$ is optimal.

Fix $v \geq u-p$. If the seller offers τ_1^*, both naive and strategic consumers buy immediately, yielding profit $p - \tau_1^*$. If it offers τ_0, naive consumers search and never return, while strategic consumers buy immediately with probability $1 - \sigma_s^*(\tau_0; u)$. Thus the profit from offering τ_0 equals

$$\lambda\big(1 - \sigma_s^*(\tau_0; u)\big)p.$$

The choice

$$\sigma_s^*(\tau_0; u) = \frac{\tau_1^* - (1-\lambda)p}{\lambda p}$$

imposes the firm-indifference condition

$$\lambda\big(1 - \sigma_s^*(\tau_0; u)\big)p = p - \tau_1^*,$$

so τ_0 and τ_1^* are both best replies for all $v \geq u-p$. Therefore, the cutoff structure of μ^* is sequentially rational.

Step 3. Feasibility and the condition $\lambda \geq \lambda_1(u)$. We require $\sigma_s^*(\tau_0; u) \in [0,1]$. Since

$$\sigma_s^*(\tau_0; u) = 1 - \frac{p - \tau_1^*}{\lambda p} \leq 1$$

always holds, feasibility requires $\sigma_s^*(\tau_0; u) \geq 0$, which is equivalent to

$$\lambda \geq 1 - \frac{\tau_1^*}{p} \equiv \lambda_1(u).$$

Step 4. Belief-robust deviation and the condition $\lambda \geq \lambda_2(u)$. Because naive consumers do not infer from discounts, the seller can secure positive deviation profit from them regardless of strategic consumers' off-path beliefs. Consider a deviation in which the seller offers the benchmark discount $\tau^b(u)$ to a consumer with $v \geq u - p$. Even if all strategic consumers search and never return, naive consumers buy immediately, yielding deviation profit at least

$$(1-\lambda)\big(p - \tau^b(u)\big).$$

In the candidate equilibrium, the seller earns $p - \tau_1^*$ from any $v \geq u - p$. Thus, to rule out this deviation independently of beliefs, it is necessary that

$$p - \tau_1^* \geq (1-\lambda)\big(p - \tau^b(u)\big),$$

which is equivalent to $\lambda \geq \lambda_2(u)$.

Step 5. Remaining off-path deviations and D1. Consider any off-path discount $\tau \notin \{\tau_0, \tau_1^*\}$. Construct off-path beliefs for strategic consumers following the proof of Proposition 2. Under those beliefs, no profitable off-path deviation exists, and the constructed beliefs are D1-consistent.

Combining Steps 1–5, when $\lambda \geq \max\{\lambda_1(u), \lambda_2(u)\}$, the candidate strategies and beliefs constitute a perfect Bayesian equilibrium with the properties stated in the lemma. □

When $\lambda_2(u) < \lambda < \lambda_1(u)$, the equilibrium in Lemma A2 may fail because the seller-indifference condition would require $\sigma_s^*(\tau_0; u) < 0$. This does not rule out inference management. Because naive consumers still search after receiving τ_0, the seller can instead sustain a binary-discount equilibrium with $\sigma_s^*(\tau_0; u) = 0$ and positive strategic search after the high discount.

Lemma A3 (Inference-management equilibrium in the intermediate region)**.** *Suppose $u - p < r$ and $\tau_1^*(u) < p$. If*

$$\lambda_3(u) < \lambda < \lambda_1(u), \qquad \lambda_3(u) \equiv \frac{p - \tau^b(u)}{2p - \tau^b(u)},$$

then there exists a perfect Bayesian equilibrium with $\Omega(\mu^*) = \{\tau_0, \tau_1^*(u)\}$ and

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

(2) $\sigma_s^*(\tau_0; u) = 0$ and $\sigma_s^*(\tau_1^*(u); u) = \dfrac{p - \tau_1^*(u) - \lambda p}{\lambda p - \lambda \tau_1^*(u)}$;

(3) $\sigma_n^*(\tau_0; u) = 1$ and $\sigma_n^*(\tau_1^*(u); u) = 0$,

where $\tau_1^*(u)$ and $\hat{v}$ are defined in (10) *and* (9)*, respectively.*

Proof. Fix u with $u - p < r$ and suppose $\tau_1^*(u) < p$. Let $\hat{v}$ solve (9), and write $\tau_1^* \equiv \tau_1^*(u)$. Consider the candidate profile

$$\Omega(\mu^*) = \{\tau_0, \tau_1^*\},$$

$$\mu^*(\tau_0 \mid u, v) = \mathbf{1}\{v < u - p \text{ or } v > \hat{v}\}, \qquad \mu^*(\tau_1^* \mid u, v) = 1 - \mu^*(\tau_0 \mid u, v),$$

$$\sigma_n^*(\tau_0;u)=1,\qquad \sigma_n^*(\tau_1^*;u)=0,$$
$$\sigma_s^*(\tau_0;u)=0,\qquad \sigma_s^*(\tau_1^*;u)=\frac{p-\tau_1^*-\lambda p}{\lambda(p-\tau_1^*)}.$$

We verify that this profile constitutes a perfect Bayesian equilibrium for $\lambda_3(u)<\lambda<\lambda_1(u)$.

Step 1. Beliefs and consumers' optimality. Take Bayes-consistent beliefs on the equilibrium path. After observing τ_1^*, a strategic consumer infers $v\in[u-p,\hat{v}]$; after observing τ_0, she infers $v\in[\underline{v},u-p)\cup(\hat{v},\bar{v}]$.

Naive consumers do not update from discounts. Since $u-p<r$ implies $B(u-p)>s$, a naive consumer strictly prefers to search after τ_0, so $\sigma_n^*(\tau_0;u)=1$. Because $\tau_1^*(u)>\tau^b(u)$, she strictly prefers to buy immediately after τ_1^*, so $\sigma_n^*(\tau_1^*;u)=0$.

For strategic consumers, (9) implies $B(u-p;\tau_0)=s$, so the consumer is indifferent after τ_0; we select $\sigma_s^*(\tau_0;u)=0$. Likewise, (10) implies $B(u-p;\tau_1^*)=s+\tau_1^*$, so the consumer is indifferent after τ_1^* and any $\sigma_s^*(\tau_1^*;u)\in[0,1]$ is sequentially rational.

Step 2. Firm optimality on the equilibrium path. If $v<u-p$, offering τ_0 weakly dominates any positive discount, since the consumer returns to buy at price p even if she searches.

Fix $v\geq u-p$. If the seller offers τ_0, naive consumers search and never return, while strategic consumers buy immediately, so profit equals λp. If it offers τ_1^*, naive consumers buy immediately and yield $(1-\lambda)(p-\tau_1^*)$. Strategic consumers buy immediately with probability $1-\sigma_s^*(\tau_1^*;u)$ and otherwise search and never return, so their contribution is $\lambda(1-\sigma_s^*(\tau_1^*;u))(p-\tau_1^*)$. Hence the profit from τ_1^* equals
$$(p-\tau_1^*)\big(1-\lambda\,\sigma_s^*(\tau_1^*;u)\big).$$
Imposing firm indifference between τ_0 and τ_1^* for $v\geq u-p$ yields
$$\lambda p=(p-\tau_1^*)\big(1-\lambda\,\sigma_s^*(\tau_1^*;u)\big),$$
which implies
$$\sigma_s^*(\tau_1^*;u)=\frac{p-\tau_1^*-\lambda p}{\lambda(p-\tau_1^*)}.$$
Therefore τ_0 and τ_1^* are both best replies for all $v\geq u-p$, and the cutoff form of μ^* is sequentially rational.

Step 3. Feasibility. The condition $\sigma_s^*(\tau_1^*;u)>0$ is equivalent to
$$\lambda<1-\frac{\tau_1^*}{p}\equiv\lambda_1(u).$$
To ensure $\sigma_s^*(\tau_1^*;u)\leq 1$, it is enough that
$$\lambda\geq\frac{p-\tau_1^*}{2p-\tau_1^*}.$$
Since $\tau_1^*(u)>\tau^b(u)$ and the function $x\mapsto\frac{p-x}{2p-x}$ is decreasing, we have
$$\frac{p-\tau_1^*(u)}{2p-\tau_1^*(u)}<\frac{p-\tau^b(u)}{2p-\tau^b(u)}=\lambda_3(u).$$

Hence $\lambda > \lambda_3(u)$ implies $\sigma_s^*(\tau_1^*; u) \leq 1$.

Step 4. Belief-robust deviation to $\tau^b(u)$. Because naive consumers do not infer from discounts, the seller can secure deviation profit from them regardless of strategic consumers' off-path beliefs. Consider a deviation to the benchmark discount $\tau^b(u)$. Even if all strategic consumers search and never return, naive consumers buy immediately under $\tau^b(u)$, so the deviation profit is at least

$$(1-\lambda)\big(p - \tau^b(u)\big).$$

In the candidate equilibrium, the seller earns λp from any $v \geq u - p$. Therefore, ruling out this deviation independently of beliefs requires

$$\lambda p \geq (1-\lambda)\big(p - \tau^b(u)\big) \quad \iff \quad \lambda \geq \frac{p - \tau^b(u)}{2p - \tau^b(u)} \equiv \lambda_3(u).$$

Step 5. Remaining off-path deviations and D1. Consider any off-path discount $\tau \notin \{\tau_0, \tau_1^*\}$. Under the candidate equilibrium, the seller's equilibrium payoff equals λp for all $v \geq u - p$ and equals p for all $v < u - p$. If the seller deviates to τ, the difference in deviation gains between types $v \geq u - p$ and $v < u - p$ satisfies

$$\Delta_1(\tau, \sigma) - \Delta_2(\tau, \sigma) = (1-\lambda)p > 0.$$

Hence, for any off-path τ, types $v \geq u - p$ strictly dominate types $v < u - p$ in deviation gains. The D1 criterion therefore requires that any admissible posterior assign probability only to $v \geq u - p$.

Fix such a posterior and choose it so that $B(u-p;\tau) > \tau + s$, implying that the strategic consumer strictly prefers to search. Then deviating to τ yields zero profit from all $v \geq u - p$, which is strictly below the equilibrium profit λp. Deviations are never profitable for $v < u - p$ because offering τ_0 already yields profit p. Therefore, no profitable off-path deviation exists.

Combining Steps 1–5, the proposed strategies and beliefs form a perfect Bayesian equilibrium whenever

$$\lambda_3(u) < \lambda < \lambda_1(u).$$

□

A.2.1 Proof of Proposition 3 in the main text

Proof. Fix u with $u - p < r$ and $\tau_1^*(u) < p$. For all equilibria considered here, the seller offers τ_0 whenever $v < u - p$, and therefore earns profit p from these types. It is therefore sufficient to compare the seller's profit conditional on $v \geq u - p$.

Step 1. Inference-management equilibria when $\lambda \geq \lambda_1(u)$. Fix any non-pooling perfect Bayesian equilibrium in which the seller engages in inference management and uses at least one positive on-path discount. Let $\Omega(\mu)$ be the on-path support and let $\bar{\tau} \equiv \max \Omega(\mu)$ denote the largest on-path discount. For any type $v \geq u - p$, the seller's realized profit from any positive on-path discount $\tau \in \Omega(\mu) \setminus \{\tau_0\}$ is at most $p - \tau$, since immediate purchase yields $p - \tau$ and any search response weakly lowers profit. Hence the seller's conditional profit from such types is bounded above by $p - \bar{\tau}$.

By on-path indifference, the seller must earn the same conditional profit from every discount in $\Omega(\mu)$ used against types $v \geq u - p$. Therefore the conditional equilibrium profit is bounded above by $p - \bar{\tau}$. Maximizing profit within this class is therefore equivalent to minimizing the largest on-path positive discount subject to keeping τ_0 non-revealing. The same logic as in the proof of Proposition 2 then implies that the smallest feasible largest discount is $\tau_1^*(u)$, attained by the binary equilibrium in Lemma A2. Hence, when $\lambda \geq \lambda_1(u)$, the maximal attainable conditional profit from types $v \geq u - p$ among all inference-management equilibria is

$$p - \tau_1^*(u).$$

Step 2. Inference-management equilibria when $\lambda \in (\lambda_3(u), \lambda_1(u))$. Fix $\lambda \in (\lambda_3(u), \lambda_1(u))$ and consider any inference-management equilibrium in which τ_0 is used on path for a positive measure of types with $v \geq u - p$. By on-path indifference, the seller must earn the same conditional profit from all discounts used against those types. In particular, the equilibrium profit from $v \geq u - p$ must equal the profit from offering τ_0.

Under τ_0, naive consumers search and never return when $v \geq u - p$, so the seller earns zero from naive consumers. Strategic consumers buy immediately with probability at most one, yielding at most λp. Hence any inference-management equilibrium in this region yields conditional profit at most λp from $v \geq u - p$. The equilibrium constructed in Lemma A3 attains this bound, since it sets $\sigma_s^*(\tau_0; u) = 0$ and therefore yields exactly λp from types $v \geq u - p$.

Step 3. Comparison with the naive-principle equilibrium. In a naive-principle equilibrium, the seller offers $\tau^b(u)$ to types with $v \geq u - p$ and earns conditional profit

$$(1 - \lambda)\big(p - \tau^b(u)\big)$$

from those types, since naive consumers buy immediately while strategic consumers search and never return.

If $\lambda_2(u) < \lambda_1(u)$, then $\hat{\lambda}(u) = \lambda_3(u)$. By Step 2, any inference-management equilibrium yields conditional profit at most λp, while the naive-principle equilibrium yields conditional profit $(1 - \lambda)(p - \tau^b(u))$. By the definition of $\lambda_3(u)$, for all $\lambda < \lambda_3(u)$ we have

$$\lambda p < (1 - \lambda)\big(p - \tau^b(u)\big),$$

so inference management is strictly dominated.

If instead $\lambda_2(u) \geq \lambda_1(u)$, then $\hat{\lambda}(u) = \lambda_2(u)$. By Step 1, the best attainable inference-management conditional profit is $p - \tau_1^*(u)$, while the naive-principle equilibrium yields $(1 - \lambda)(p - \tau^b(u))$. By the definition of $\lambda_2(u)$, for all $\lambda < \lambda_2(u)$ we have

$$p - \tau_1^*(u) < (1 - \lambda)\big(p - \tau^b(u)\big),$$

so inference management is again strictly dominated.

Conversely, when $\lambda \geq \hat{\lambda}(u)$, Steps 1–2 together with Lemmas A2 and A3 show that a corresponding inference-management equilibrium exists and achieves the maximal feasible inference-management profit in the relevant region, which weakly dominates the naive-principle equilibrium.

Combining the above establishes the regime switch stated in Proposition 3. □

A.3 Endogenous List Price

In the main analysis, we take the list price as exogenously given. This assumption is common in models of search and pricing, as prices are often set in advance and cannot be adjusted flexibly at the individual level. In many applications, list prices are determined at the national or regional level and change infrequently, whereas discounts can be personalized or adjusted in real time.

Nevertheless, one may ask how asymmetric information and consumer inference affect the firm's *pricing decision itself*. In this appendix, we briefly discuss how the optimal list price differs across informational regimes.

Conceptually, endogenizing the list price is straightforward. For any given consumer valuation u and list price p, we have characterized the equilibrium and the firm's profit under symmetric information, denoted by $\Pi^b(u;p)$, and under superior information, denoted by $\Pi^*(u;p,\lambda)$ for $\lambda \in \{0,1\}$. The firm's expected profit under list price p is obtained by integrating over the distribution of u. Accordingly, the optimal list prices solve

$$p^b \in \arg\max_p \ \mathbb{E}_u\big[\Pi^b(u;p)\big],$$

and

$$p^*(\lambda) \in \arg\max_p \ \mathbb{E}_u\big[\Pi^*(u;p,\lambda)\big].$$

Because our model allows for a general joint distribution $H(u,v)$ and arbitrary search costs s, we do not obtain a sharp analytical characterization of the optimal list prices. Moreover, the effect of superior information and consumer inference on optimal pricing depends on the underlying distribution of valuations and model parameters. In particular, the ordering of p^b, $p^*(0)$, and $p^*(1)$ is not generally monotone in λ.

Figure 2 illustrates this point using two examples with $v \sim U[0,\bar{v}]$ and $u \sim U[0,\bar{u}]$. In panel (a), we have $p^*(0) < p^b < p^*(1)$, whereas in panel (b) the ordering reverses to $p^*(0) > p^b > p^*(1)$. These examples show that allowing the firm to optimally adjust its list price can either attenuate or amplify the pricing response to superior information, depending on the environment.

Figure 2: Optimal List Prices under Endogenous Pricing.

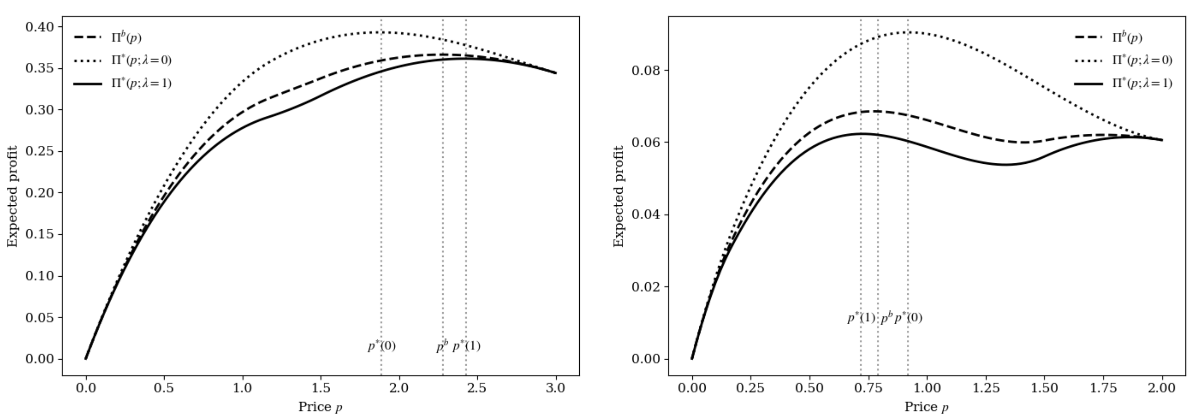

(a) $s=1,\ \bar{v}=5,\ \bar{u}=3$ (b) $s=1,\ \bar{v}=5,\ \bar{u}=2$

Despite this ambiguity in pricing responses, the firm remains worse off from having superior information when consumer inference is strong, even when the list price is endogenously chosen.

Recall that for any given realization of (u, p), the firm earns the same or strictly lower profit under the superior-information regime with strong inference ($\lambda = 1$) than under symmetric information, and the same or strictly higher profit under weak inference ($\lambda = 0$). Consequently, for any fixed list price p, integrating over the distribution of u preserves this ordering in expected profits. Optimizing over p therefore does not overturn it.

Proposition A1. *Let $\Pi^b(u; p)$ denote the firm's profit conditional on u at list price p under symmetric information, and let $\Pi^*(u; p, \lambda)$ denote the corresponding profit under superior information with inference parameter $\lambda \in \{0, 1\}$. Let*

$$\Pi^b(p) \equiv \mathbb{E}_u[\Pi^b(u; p)] \quad \text{and} \quad \Pi^*(p; \lambda) \equiv \mathbb{E}_u[\Pi^*(u; p, \lambda)].$$

Then the optimal expected profits satisfy

$$\Pi^*\big(p^*(0); 0\big) \;\geq\; \Pi^b\big(p^b\big) \;\geq\; \Pi^*\big(p^*(1); 1\big).$$

Proof. From the equilibrium analysis in the main text, for any realization (u, p) we have

$$\Pi^*(u; p, 0) \;\geq\; \Pi^b(u; p) \;\geq\; \Pi^*(u; p, 1),$$

with at least one inequality strict on a positive-measure set of (u, p) whenever a non-pooling equilibrium exists. Taking expectations over u preserves this ordering for any fixed p, implying

$$\Pi^*(p; 0) \;\geq\; \Pi^b(p) \;\geq\; \Pi^*(p; 1) \quad \text{for all } p.$$

Optimizing both sides over p yields

$$\Pi^*\big(p^*(0); 0\big) = \max_p \Pi^*(p; 0) \;\geq\; \max_p \Pi^b(p) = \Pi^b\big(p^b\big) \;\geq\; \max_p \Pi^*(p; 1) = \Pi^*\big(p^*(1); 1\big),$$

which establishes the result. □

Therefore, allowing the firm to optimally choose its list price does not eliminate the profitability disadvantage induced by consumer inference. The loss arises from the same mechanism emphasized in the main analysis: inference management constrains the firm's ability to exploit its informational advantage, forcing it to either deepen discounts or withhold them from certain consumers.

Finally, while endogenous pricing may affect consumer surplus and total welfare, drawing general conclusions is difficult. The welfare implications depend on how the optimal list price adjusts in each regime, which in turn depends on the distribution of valuations and the magnitude of search costs. A full welfare analysis would therefore require additional structure and is beyond the scope of this appendix.

A.4 Duopoly with Endogenous Outside Option

This appendix extends the baseline model to a duopoly environment in which the consumer visits Firm 1 first and can then search and visit Firm 2. The goal is to verify that explicitly modeling product-market competition on the outside option side does not alter the main insights in the text: (i) consumer inference compresses equilibrium pricing toward coarse menus; and (ii) when inference is sufficiently strong, superior information can reduce Firm 1's equilibrium profit relative to the symmetric-information benchmark.

Although this duopoly extension explicitly models Firm 1 and Firm 2 and solves for equilibrium pricing and consumer behavior under competition, the reason the main insights remain unchanged can be understood at a higher level. From the consumer's perspective, competition primarily reshapes the distribution of continuation payoffs following search by censoring and shifting outside options, while preserving the monotonic relationship between search incentives and the underlying type. As a result, the consumer's inference problem and the firm's incentive to manage beliefs through coarse pricing take a form that is qualitatively similar to the baseline model with a modified outside–option distribution. This analogy helps explain why explicitly modeling competition does not overturn the inference–management mechanism emphasized in the main text, even though equilibrium prices and profits differ.

A.4.1 Environment and Timing

Fix a realization of u, the consumer's valuation for Firm 1's product. The consumer's valuation for Firm 2's product is $v \in [\underline{v}, \bar{v}]$. The pair $(u, v) \sim H(u, v)$. Since u is commonly observed upon the first visit, we condition on u throughout and write the (conditional) prior of v as G with density g.

Firm 1 posts a list price p_u and can offer a buy-now discount $\tau_u \in [0, p_u]$ when the consumer arrives. If the consumer searches, she incurs search cost $s > 0$ and visits Firm 2. Upon visiting Firm 2, both u and v are revealed to the consumer and Firm 2. Firm 2 posts list price p_v and can offer a buy-now discount $\tau_v \in [0, p_v]$. As in the main text, buy-now discounts are committed and expire upon search/visiting the rival. Thus, if the consumer returns to Firm 1 after visiting Firm 2, she pays p_u.

A fraction $\lambda \in [0, 1]$ of consumers are strategic and update beliefs about v from τ_u using Bayes' rule; the remaining fraction $1-\lambda$ are naive and do not infer. In this appendix we focus on $\lambda \in \{0, 1\}$.

A.4.2 Subgame at Firm 2 (Given Search)

Suppose the consumer searches and arrives at Firm 2. Since u and v are revealed to Firm 2, Firm 2's discounting problem is static. If the consumer buys from Firm 2 at discount τ_v, she pays $p_v - \tau_v$ and obtains surplus $v - p_v + \tau_v$. If she returns to Firm 1, she pays p_u and obtains $u - p_u$.

Firm 2 can acquire the consumer at the lowest discount that weakly beats Firm 1:

$$v - p_v + \tau_v \ \geq \ u - p_u \quad \Longleftrightarrow \quad \tau_v \ \geq \ (u - v) - (p_u - p_v).$$

Therefore Firm 2's optimal discount upon arrival is

$$\tau_v^b(u, v) = \min\Big\{p_v,\ \max\{0,\ (u - v) - (p_u - p_v)\}\Big\}. \tag{A1}$$

Under (A1), the consumer's continuation surplus from searching equals

$$\max\{u - p_u,\ v - p_v\}. \tag{A2}$$

Intuitively, if $v - p_v \geq u - p_u$, Firm 2 need not discount and the consumer gets $v - p_v$; if $v - p_v < u - p_u$, Firm 2 discounts just enough to make the consumer indifferent, and the consumer gets $u - p_u$.

$$\mathbb{E}[\max\{u - p_u,\ v - p_v\}] = u - p_u + \int_{u - p_u + p_v}^{\bar{v}} [1 - G(v)]\, dv.$$

Define the (duopoly) expected benefit of searching when the current outside-option cutoff is $x + p_v$:

$$B_2(x) \ \equiv\ \int_{x+p_v}^{\bar{v}} [1 - G(v)]\, dv. \tag{A3}$$

This is identical to the baseline $B(\cdot)$ in the main text, except that the relevant cutoff is shifted by p_v (because the rival's list price enters the consumer's net outside option).

Define the reservation threshold r by

$$B_2(r) = s. \tag{A4}$$

We focus on the economically relevant region

$$u - p_u < r \quad \Longleftrightarrow \quad B_2(u - p_u) > s,$$

where absent sufficient discounts the consumer strictly prefers to search.

A.4.3 Symmetric-Information Benchmark

In the symmetric-information benchmark, Firm 1 does not observe v and shares the same prior G with the consumer. Firm 1 chooses τ_u as a function of u only.

Consumer search decision. If the consumer does not search, she buys immediately and obtains $u - p_u + \tau_u$. If she searches, she obtains $\mathbb{E}[\max\{u - p_u,\ v - p_v\}] - s = u - p_u + B_2(x) - s$. Therefore she searches iff

$$B_2(u - p_u) > s + \tau_u. \tag{A5}$$

Hence, when $u - p_u < r$, Firm 1 can deter search at minimal cost by setting

$$\tau_u^b(u) \ \equiv\ B_2(u - p_u) - s. \tag{A6}$$

Firm 1 profit. If Firm 1 deters search via τ_u^b, its profit is $\Pi_u^{b,\text{deter}}(u) = p_u - \tau_u^b(u)$. If Firm 1 offers no discount, the consumer searches. She returns to Firm 1 only if $u - p_u \geq v$. Thus Firm 1's profit is $G(u - p_u)\, p_u$. Therefore,

$$\Pi_u^b(u) = \max\Big\{G(u - p_u)\, p_u,\ \ p_u - \tau_u^b(u)\Big\}. \tag{A7}$$

Firm 2 profit. If Firm 1 deters search, Firm 2 is never visited and $\Pi_v^b(u) = 0$.

If Firm 1 does not deter search (i.e., $\tau_u = 0$ and the consumer searches), Firm 2 is visited for sure. Conditional on (u,v), Firm 2 earns

$$p_v - \tau_v^b(u,v) = \begin{cases} 0, & v \le u - p_u, \\ v - x, & u - p_u < v < u - p_u + p_v, \\ p_v, & v \ge u - p_u + p_v, \end{cases}$$

where the middle region uses $\tau_v^b = u - p_u + p_v - v$ from (A1). Thus, benchmark Firm 2 profit (conditional on search) is

$$\Pi_v^b(u) = \int_{u-p_u}^{u-p_u+p_v} (v - u + p_u)\, dG(v) + \int_{u-p_u+p_v}^{\bar{v}} p_v\, dG(v). \tag{A8}$$

(And $\Pi_v^b(u) = 0$ when Firm 1 deters search.)

A.4.4 Superior Information

We now assume Firm 1 observes v upon the first visit and can condition τ_u on (u,v). We restrict attention to $u - p_u < r$ and to $\lambda \in \{0,1\}$.

Case $\lambda = 0$. When consumers are naive, τ_u does not affect beliefs. The consumer searches if and only if (A5) holds. Hence the same logic as in the main text applies, with B_2 replacing B.

Lemma A4 (Naive-consumer optimal pricing in duopoly)**.** *When $\lambda = 0$, the profit-maximizing pricing rule for Firm 1 is*

$$\tau_u^*(u,v) = \begin{cases} 0, & v < u - p_u, \\ \tau_u^b(u), & v \ge u - p_u, \end{cases}$$

where τ_u^b is given in (A6)*, provided $\tau_u^b(u) < p_u$.*

The proof is identical to Lemma 1 in the main text: for $v < x$ the consumer will return to Firm 1 even after search, so discounting is weakly dominated; for $v \ge x$, the cheapest way to deter naive search is the indifference discount $\tau_u^b(u)$.

Case $\lambda = 1$. Strategic consumers update beliefs about v from τ_u. Let $\mu(\tau_u \mid u, v)$ be Firm 1's (possibly mixed) discount strategy, and let $\Omega(\mu \mid u)$ denote its on-path support for a given u. For $\tau_u \in \Omega(\mu \mid u)$, Bayes' rule implies the posterior density

$$g(v \mid u, \tau_u) = \frac{\mu(\tau_u \mid u, v)\, g(v)}{\int \mu(\tau_u \mid u, v')\, dG(v')}. \tag{A9}$$

Given posterior $G(\cdot \mid u, \tau_u)$, the expected benefit of search (relative to the current cutoff $u - p_u + p_v$) is

$$B_2(u - p_u; \tau_u) \;\equiv\; \int_{u-p_u+p_v}^{\bar{v}} \big[1 - G(v \mid u, \tau_u)\big]\, dv. \tag{A10}$$

The consumer searches after observing τ_u if and only if

$$B_2(u - p_u; \tau_u) > s + \tau_u, \tag{A11}$$

and she is indifferent when equality holds.

As in the main text, a "naive-style" separating policy (offering $\tau_u = 0$ only when $v < u - p_u$ and $\tau_u = \tau_u^b$ otherwise) cannot be sustained when consumers are strategic, because $\tau_u = 0$ would reveal $v < u - p_u$ and eliminate search, giving Firm 1 incentives to deviate.

The same inference-management logic implies three robust properties in any non-pooling equilibrium in which some $\tau_u > 0$ is used on path: (i) Firm 1 never discounts when $v < u - p_u$; (ii) Firm 1 must sometimes offer $\tau_u = 0$ even when $v \geq u - p_u$; and (iii) Firm 1 must be indifferent across on-path discounts for $v \geq u - p_u$.

Lemma A5 (Necessary conditions in any non-pooling equilibrium, duopoly)**.** *Suppose $\lambda = 1$ and $u - p_u < r$. Consider any non-pooling PBE in which $\Omega(\mu \mid u)$ contains some $\tau_u > 0$. Then:*

1. *$\mu(0 \mid u, v) = 1$ for all $v < u - p_u$.*
2. *$\mu(0 \mid u, v) > 0$ for a positive-measure set of $v \geq u - p_u$.*
3. *$(p_u - \tau_u)\big(1 - \sigma(\tau_u; u)\big)$ is constant across all $\tau_u \in \Omega(\mu \mid u)$, where $\sigma(\tau_u; u)$ is the strategic consumer's equilibrium search probability after τ_u.*
4. *$0 < \sigma(0; u) < 1$ and $B_2(u - p_u; 0) = s$.*
5. *For any $\tau_u \in \Omega(\mu \mid u) \setminus \{0\}$, $\sigma(\tau_u; u) < 1$ and $B_2(u - p_u; \tau_u) \leq s + \tau_u$.*

The proof is a verbatim adaptation of Lemma 2 in the main text, replacing B by B_2.

We now characterize the natural analogue of the firm-optimal binary-discount equilibrium (Proposition 2 in the main text) with $\Omega(\mu \mid u) = \{\tau_0, \tau_1\}$ for some $\tau_1 > 0$, with a cutoff $\hat{v} \in (x, \bar{v})$ such that

$$\text{Firm 1 offers } \tau_1 \text{ iff } v \in [u - p_u, \hat{v}], \quad \text{and offers } \tau_0 \text{ iff } v < u - p_u \text{ or } v > \hat{v}.$$

As in the main text, concentrating τ_0 on the highest v types minimizes the mass of profitable ($v \geq u - p_u$) consumers from whom Firm 1 must withhold discounts in order to keep τ_0 non-revealing. Indifference after τ_0 requires

$$B_2(u - p_u; 0) = s, \tag{A12}$$

which defines $\hat{v}$.

Under the cutoff rule, the consumer's posterior after τ_1 is supported on $v \in [x, \hat{v}]$. In the firm-optimal binary equilibrium we choose τ_1 so that the search-deterrence constraint binds:

$$B_2(u - p_u; \tau_1) = s + \tau_1, \qquad \sigma(\tau_1; u) = 0, \tag{A13}$$

and then use firm indifference between τ_0 and τ_1 to pin down $\sigma(\tau_0; u)$. We obtain

$$\tau_1^*(u) = \frac{B_2(u - p_u) - s}{G(\hat{v}) - G(u - p_u)}. \tag{A14}$$

Finally, firm indifference across on-path discounts implies

$$p_u\big(1 - \sigma(\tau_0; u)\big) = p_u - \tau_1^*(u) \quad \Longleftrightarrow \quad \sigma(\tau_0; u) = \frac{\tau_1^*(u)}{p_u}.$$

Proposition A2. *Fix u and suppose $u - p_u < r$. If $\tau_1^*(u)$ defined in* (A14) *satisfies $\tau_1^*(u) < p_u$, then there exists a binary-discount PBE with $\Omega(\mu^* \mid u) = \{0, \tau_1^*(u)\}$ such that:*

1. *Firm 1's pricing rule is a cutoff rule:*

$$\mu^*(0 \mid u, v) = 1 \quad \textit{iff} \quad v < u - p_u \textit{ or } v > \hat{v}, \qquad \mu^*(\tau_1^*(u) \mid u, v) = 1 \quad \textit{iff} \quad v \in [u - p_u, \hat{v}],$$

   where $\hat{v}$ solves the indifference equation (A12)*.*

2. *The strategic consumer's search strategy satisfies*

$$\sigma^*(\tau_1^*(u); u) = 0, \qquad \sigma^*(0; u) = \frac{\tau_1^*(u)}{p_u}.$$

Moreover, among all non-pooling equilibria (including equilibria with richer discount supports), this binary equilibrium maximizes Firm 1's profit whenever it exists.

The proof follows the same two steps as Proposition 2 in the main text: (i) any equilibrium with $|\Omega| > 2$ can be weakly improved upon by collapsing all positive discounts into a single level that binds the search-deterrence constraint; (ii) among binary equilibria, allocating the no-discount signal to the highest v types minimizes the probability of τ_0 subject to $B_2(u - p_u; 0) = s$, thereby minimizing the required τ_1 in (A14).

As in Lemma 3 in the main text, if $\tau_1^*(u) \geq p_u$, then no non-pooling inference-management equilibrium with a positive discount can exist because the minimal positive discount required to keep τ_0 non-revealing would be infeasible. In that case, the remaining equilibrium is pooling at $\tau_0 = 0$ with consumer search, yielding Firm 1 profit $G(u - p_u)p_u$ (and Firm 2 profit equal to the benchmark expression (A8), multiplied by the probability of search, which is one under pooling at τ_0 when $u - p_u < r$).

Next, we compare Firm 1's profits $\Pi_u^b(u)$ and $\Pi_u^*(u; \lambda)$ for $\lambda \in \{0, 1\}$, as well as Firm 2's expected profits $\Pi_v^b(u)$ and $\Pi_v^*(u; \lambda)$.

Firm 1 profit. In the symmetric-information benchmark (when Firm 1 deters search), Firm 1's profit is

$$\Pi_u^b(u) = p_u - \tau_u^b(u) = p_u - \big(B_2(u - p_u) - s\big).$$

In the superior-information case, following Lemma A4 and Proposition A2, Firm 1's profit is

$$\Pi_u^*(u; 0) = G(u - p_u)\, p_u + \big(1 - G(u - p_u)\big)\big(p_u - \tau_u^b(u)\big),$$

and

$$\Pi_u^*(u; 1) = G(u - p_u)\, p_u + \big(1 - G(u - p_u)\big)\big(p_u - \tau_1^*(u)\big).$$

Therefore,

$$\Pi_u^*(u; 0) \;\geq\; \Pi_u^b(u) \;\geq\; \Pi_u^*(u; 1). \tag{A15}$$

Relative to the symmetric-information benchmark, inference management forces Firm 1 to use a strictly larger positive discount whenever a non-pooling equilibrium exists. As in the main text, the resulting profit loss under strong inference is driven by two forces: deeper discounts for intermediate v and the need to withhold discounts from some high-v consumers in order to keep τ_0 non-revealing.

Firm 2 profit. When $\lambda = 0$, Firm 1 deters search whenever $v \geq u - p_u$; only consumers with $v < u - p_u$ search, and Firm 2 earns zero from such consumers. Hence, $\Pi_v^*(u; 0) = 0$.

When $\lambda = 1$, the consumer searches only after observing τ_0, and does so with probability $\sigma^*(0; u) = \tau_1^*(u)/p_u$. Conditional on visiting Firm 2, Firm 2 earns zero when $v < u - p_u$, earns

$v-(u-p_u)$ when $v \in (\hat{v},\, u-p_u+p_v)$, and earns p_v when $v \geq u-p_u+p_v$. Therefore, Firm 2's equilibrium profit is

$$\Pi_v^*(u;1) = \frac{\tau_1^*(u)}{p_u}\left[\int_{\hat{v}}^{u-p_u+p_v} \big(v-(u-p_u)\big)\, dG(v) + \int_{u-p_u+p_v}^{\bar{v}} p_v \, dG(v)\right]. \tag{A16}$$

This expression highlights what is new relative to the monopoly case: inference management by Firm 1 endogenously generates a flow of consumers to Firm 2 through τ_0-induced search.

Finally, when $G(u-p_u)\, p_u > p_u - \tau_u^b(u)$, Firm 1 optimally does not deter search in the benchmark, in which case Firm 2's profit is given by (A8) and exceeds $\Pi_v^*(u;1)$. Otherwise, the benchmark deters search and Firm 2 earns zero, which is below $\Pi_v^*(u;1)$. Hence, Firm 1's superior information may raise its competitor's profit.

A.5 No Commitment to Buy–Now Discounts

In the main text, we assume that a buy–now discount expires once the consumer searches and therefore cannot be accessed upon return. In this appendix, we relax this assumption and argue that doing so does not eliminate the equilibrium characterized in the main text. Rather, allowing the firm to revise its pricing policy after search weakly enlarges the equilibrium set while preserving the existence of the main–text equilibrium.

Formally, suppose that in addition to the initial discount $\tau^{(1)}$, the firm may offer a second–round discount $\tau^{(2)}$ if the consumer searches and returns. The timing and information structure are otherwise identical to the main text, and no additional tie–breaking assumptions are imposed.

To see why the main–text equilibrium remains valid, consider the continuation subgame in which the consumer has searched and decides whether to return to the firm. If $v < u - p$, then upon return the firm strictly prefers not to offer any discount, and the consumer strictly prefers to return and purchase at the regular price. If $v \geq u - p$, then the firm can at most offer a second–round discount that makes the consumer indifferent between purchasing and taking the outside option. Hence, conditional on $v \geq u - p$, the consumer's continuation payoff from returning is weakly equal to her outside option. In this continuation subgame, one equilibrium is that the consumer does not return whenever she is indifferent. This continuation behavior is consistent with the equilibrium described in the main text, in which the consumer never returns after search. Embedding this continuation equilibrium into the full game yields exactly the same equilibrium path and payoffs as in the commitment model. Therefore, the equilibrium characterized in the main text remains a Perfect Bayesian Equilibrium of the no–commitment environment.

Nonetheless, relaxing commitment does enlarge the set of equilibria. In particular, when the consumer is indifferent upon return, she may return with positive probability. In such equilibria, the firm's continuation value changes, as it can acquire some consumers with $v \in [u - p, u]$ who search and then return by matching their outside option. However, acquiring these consumers does not eliminate the need for inference management in the initial stage. To deter search by consumers with higher outside options, or to manage the beliefs induced by different first–round discounts, the firm must still design its pricing policy following the logic developed in the main text. Thus, while continuation play may differ, the core inference–management mechanism remains operative.

This interpretation is consistent with Armstrong and Zhou (2016), who study environments without exogenous commitment. They show that returning after search reveals unfavorable information about the consumer's outside option, giving the seller an incentive to worsen continuation terms. Anticipating this, consumers may optimally choose not to return after search in equilibrium, so that front–loaded pricing arises endogenously as an equilibrium outcome rather than as a contractual restriction.

Taken together, these observations imply that the commitment assumption in the main text should be interpreted as a reduced–form representation of equilibrium behavior rather than a literal institutional constraint. Whether or not commitment is imposed does not alter the main economic insights of the paper: consumer inference constrains the firm's use of superior information and shapes optimal pricing through the need to manage search incentives and beliefs.

A.6 Informed and Uninformed Consumers

This appendix extends the baseline model by allowing some consumers to know their outside option v prior to search. A fraction $\beta \in [0, 1]$ of consumers are *uninformed*—they do not know v and must incur the search cost s to learn it, exactly as in the main model. The remaining fraction $1 - \beta$ are *informed*—they already know v but must still incur s to access the outside option. All consumers are strategic ($\lambda = 1$). The firm does not observe whether a given consumer is informed or uninformed; it knows only the population share β.

Since an informed consumer already knows v, the discount τ affects her behavior only through its direct price effect. Specifically, an informed consumer with known v buys immediately if $u - p + \tau \geq v - s$, i.e., she searches if and only if $v > u - p + s + \tau$. Hence informed consumers with $v \leq u - p + s$ never search regardless of the discount. By contrast, uninformed consumers face the same inference problem as in the main text: they update beliefs about v from the observed discount and decide whether to search based on their posterior.

When $\beta = 1$, the model reduces to the main case studied in the text, in which all consumers are uninformed and the firm's superior information about v can backfire through the signaling channel. We therefore begin by analyzing the opposite extreme, $\beta = 0$, in which all consumers are informed. Comparing the two extremes yields a unified insight: the firm is harmed whenever its information about the outside option diverges from consumers' information—whether the firm knows more or less than they do. We then briefly discuss the intermediate case $\beta \in (0, 1)$, where the firm-optimal equilibrium can be characterized using the same inference-management approach as in the main text, though the simultaneous presence of informed and uninformed consumers introduces an additional trade-off that slightly complicates the optimal discount structure.

A.6.1 The Case $\beta = 0$: All Consumers Informed

When all consumers know v, no signaling problem arises: the discount carries no additional information for consumers. The economic question is instead whether the firm benefits from also knowing v. We compare two information regimes: one in which the firm observes v and can personalize the discount, and one in which the firm does not observe v and must offer a uniform discount.[1]

The firm does not observe v. The firm must offer a uniform discount $\tau(u)$. An informed consumer with known v searches if and only if $v > u - p + s + \tau$. The firm's expected profit is

$$\Pi^{-}(u; \beta = 0) = \max_{\tau \geq 0} \; G(u - p + s + \tau)\,(p - \tau).$$

The optimal discount $\tau^{-}(u)$ satisfies the first-order condition $g(u - p + s + \tau)(p - \tau) = G(u - p + s + \tau)$, provided an interior solution exists. Note that even at $\tau = 0$ the firm earns $G(u - p + s)\,p > 0$, because consumers with $v \leq u - p + s$ buy at full price without any inducement. The trade-off is that any positive discount deters additional consumers from searching but simultaneously reduces the price paid by all consumers with $v \leq u - p + s + \tau$, including those who would have bought at full price.

[1] Relative to informed consumers, the regime in which the firm also observes v is one of symmetric information—both sides know v—while the regime in which only consumers observe v places the firm at an informational disadvantage. We avoid using the terms "symmetric information" and "superior information" from the main text to prevent confusion, as their meaning would be reversed relative to the case $\beta = 1$.

The firm observes v. Because consumers already know v, the firm's pricing problem is one of standard personalized pricing with no signaling considerations.

Lemma A6. *Suppose $\beta = 0$ and $u - p < r$. When the firm observes v, the optimal discount is*

$$\tau^+(u,v) = \begin{cases} 0, & v \le u-p+s, \\ v-(u-p)-s, & v \in (u-p+s,\ u+s], \\ 0, & v > u+s, \end{cases}$$

yielding profit p for $v \le u-p+s$, profit $u+s-v$ for $v \in (u-p+s, u+s]$, and profit 0 for $v > u+s$. The firm's expected profit is therefore

$$\Pi^+(u;\beta=0) = G(u-p+s)\,p + \int_{u-p+s}^{u+s} (u+s-v)\,dG(v).$$

Proof. If $v \le u-p+s$, the consumer does not search at any $\tau \ge 0$, so $\tau = 0$ is optimal. If $v \in (u-p+s, u+s]$, the minimum discount that deters search is $v-(u-p)-s$, yielding $p-[v-(u-p)-s] = u+s-v \ge 0$. Not offering a discount yields zero (the consumer searches and takes the outside option). If $v > u+s$, deterrence requires $\tau > p$, which is infeasible. □

Comparison. Knowing v strictly benefits the firm when all consumers are informed, see Figure 3. The uniform discount under the regime in which the firm does not observe v creates two inefficiencies: it discounts unnecessarily to consumers with low v who would have bought at full price, and it applies the same discount to all high-v consumers rather than calibrating to each consumer's exact willingness to search. Observing v eliminates both inefficiencies.

Figure 3: Profit Comparison When All Consumers Are Informed ($\beta = 0$).

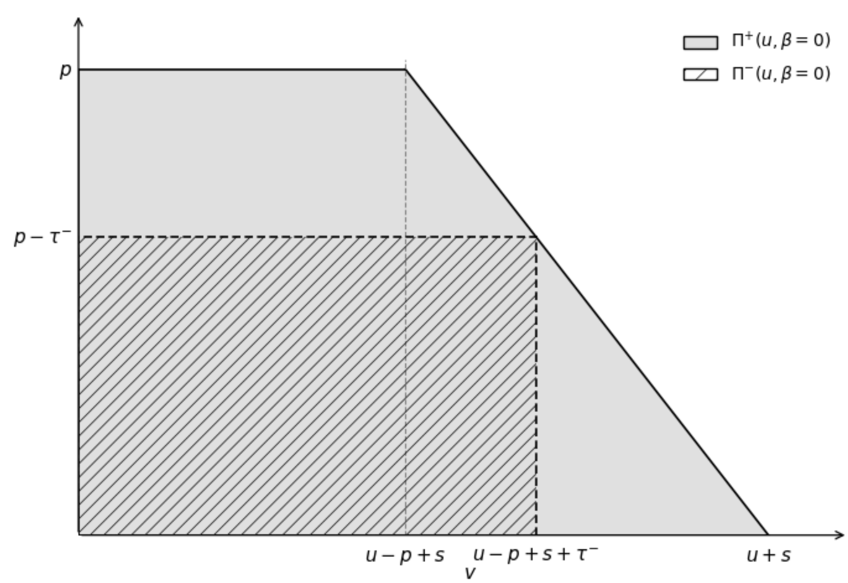

Proposition A3. *Suppose $\beta = 0$. Then $\Pi^+(u;\beta=0) > \Pi^-(u;\beta=0)$.*

Proof. Fix u and write $x \equiv u - p$. For any uniform discount $\tau \geq 0$ offered when the firm does not observe v, the realized profit from a consumer with outside option v is

$$\pi^{-}(v;\tau) = (p-\tau)\,\mathbf{1}\{v \leq x+s+\tau\}.$$

By Lemma A6, when the firm observes v, the realized profit is

$$\pi^{+}(v) = \begin{cases} p, & v \leq x+s, \\ u+s-v, & v \in (x+s,\, u+s], \\ 0, & v > u+s. \end{cases}$$

We compare these two profit schedules pointwise. For any fixed $\tau \geq 0$:

1. If $v \leq x+s$, then
$$\pi^{+}(v) = p \geq p - \tau = \pi^{-}(v;\tau).$$
2. If $x+s < v \leq x+s+\tau$, then
$$\pi^{+}(v) = u+s-v = p-(v-x-s) \geq p-\tau = \pi^{-}(v;\tau),$$
because $v-x-s \leq \tau$ on this region.
3. If $x+s+\tau < v \leq u+s$, then
$$\pi^{+}(v) = u+s-v > 0 = \pi^{-}(v;\tau).$$
4. If $v > u+s$, then
$$\pi^{+}(v) = 0 = \pi^{-}(v;\tau).$$

Thus, for every $\tau \geq 0$ and every v,

$$\pi^{+}(v) \geq \pi^{-}(v;\tau).$$

Taking expectations with respect to G gives

$$\Pi^{+}(u;\beta=0) \geq G(x+s+\tau)(p-\tau) \qquad \text{for every } \tau \geq 0.$$

Maximizing the right-hand side over τ yields

$$\Pi^{+}(u;\beta=0) \geq \Pi^{-}(u;\beta=0).$$

To see that the inequality is strict, note that under the maintained support and density assumptions there is positive probability on the interval $(x+s,\, u+s]$. On that interval, the third case above implies

$$\pi^{+}(v) > \pi^{-}(v;\tau) \qquad \text{for every } \tau \geq 0$$

whenever $v > x+s+\tau$, and the first case is strict whenever $\tau > 0$. Hence at least one of the pointwise inequalities is strict on a positive-measure set, so

$$\Pi^{+}(u;\beta=0) > \Pi^{-}(u;\beta=0).$$

□

Thus, when consumers are already informed about their outside options, the firm unambiguously benefits from possessing the same information.

A.6.2 Comparing the Two Extremes

The two polar cases yield contrasting conclusions about the value of the firm's information.

When $\beta = 0$ (all consumers informed), the firm benefits from knowing v. Lacking this information places the firm at an informational disadvantage relative to consumers: it cannot tailor discounts to individual outside options and must rely on a uniform discount that wastes surplus on low-v consumers and fails to retain some high-v consumers. Information parity—the firm knowing what consumers know—eliminates these inefficiencies.

When $\beta = 1$ (all consumers uninformed), as shown in the main text, the firm is harmed by knowing v when consumer inference is strong. Superior information triggers a signaling problem: discounts reveal the firm's private information about v, prompting consumers to search more aggressively. The firm would be better off if it, like consumers, did not know v, so that discounts carried no informational content.

Combining the two cases reveals a unifying principle: *the firm's profit is maximized when its information about the outside option is aligned with consumers' information.* When consumers are informed, the firm benefits from also being informed; when consumers are uninformed, the firm benefits from also being uninformed. Any informational asymmetry—whether the firm knows more or less than consumers—creates frictions that reduce profit. In one direction, the friction is a standard adverse-selection problem (the firm cannot price-discriminate without knowing v); in the other, it is a signaling problem (the firm's pricing reveals too much about v). While the mechanisms differ, both reflect the same underlying force: misaligned information distorts the pricing problem and constrains the firm's ability to extract surplus.

A.6.3 Intermediate values of $\beta \in (0, 1)$: profit reversal and persistence of coarse pricing

We now consider the case in which both informed and uninformed consumers are present, that is, $\beta \in (0, 1)$. This environment interpolates between the two polar cases discussed above. When $\beta = 0$, all consumers are informed, there is no signaling problem, and the seller strictly benefits from observing v. When $\beta = 1$, the extension reduces to the main case with only strategic uninformed consumers, so under condition (5) the seller's superior information is strictly less profitable than the benchmark without observing v. Thus, as β rises, the value of observing v must eventually reverse.

At the same time, the same inference-management force remains present for intermediate β. The complication is that informed consumers respond to discounts through their direct price effect, while uninformed consumers respond through inference. This additional layer breaks the clean v-independence that underlies Proposition 2 and makes a full characterization of the optimal discount assignment substantially more involved. The next proposition shows that two conclusions nevertheless go through under natural regularity assumptions: the profit effect changes sign at some interior threshold, and for sufficiently large β the coarse binary pricing structure from the main model re-emerges.

Proposition A4 (Profit reversal and persistence of coarse pricing)**.** *Fix $u - p < r$ and suppose condition (5) holds. Let $\Pi^+(u; \beta)$ denote the seller's optimal profit in the informed–uninformed extension when it observes v, and let $\Pi^-(u; \beta)$ denote the seller's optimal profit when it does not observe v.*

Assume that (1) the value functions $\Pi^+(u;\beta)$ and $\Pi^-(u;\beta)$ are continuous in $\beta \in [0,1]$; and (2) $\beta = 1$ is a strict isolated optimum: there exists $\eta > 0$ such that its profit exceeds the profit from every feasible non-binary equilibrium by at least η, and the profit of each candidate equilibrium varies continuously in β in a neighborhood of $\beta = 1$. Then:

(i) there exists at least one threshold $\bar{\beta}(u) \in (0,1)$ such that

$$\Pi^+(u;\bar{\beta}(u)) = \Pi^-(u;\bar{\beta}(u)).$$

If, in addition, $\Delta(u;\beta) \equiv \Pi^+(u;\beta) - \Pi^-(u;\beta)$ is strictly decreasing in β, $\bar{\beta}(u)$ is unique and

$$\Pi^+(u;\beta) > \Pi^-(u;\beta) \quad \textit{for } \beta < \bar{\beta}(u), \qquad \Pi^+(u;\beta) < \Pi^-(u;\beta) \quad \textit{for } \beta > \bar{\beta}(u).$$

(ii) there exists $\underline{\beta}(u) < 1$ such that for all $\beta \in [\underline{\beta}(u), 1]$, the firm-optimal equilibrium remains binary. In particular, for all such β, the support of the firm-optimal pricing rule is $\{\tau_0, \tau_1(u;\beta)\}$, with the positive discount assigned to an intermediate region of outside-option values and the zero discount assigned otherwise. Moreover,

$$\tau_1(u;\beta) \to \tau_1^*(u) \quad \textit{and} \quad [v_L(u;\beta),\, v_H(u;\beta)] \to [u-p, \hat{v}(u)] \qquad \textit{as } \beta \to 1.$$

Proof. When $\beta = 0$, all consumers are informed. By Proposition A3, the seller strictly benefits from observing v, so

$$\Pi^+(u;0) - \Pi^-(u;0) > 0.$$

When $\beta = 1$, all consumers are uninformed and the extension reduces to the main case with only strategic consumers. Under condition (5), the main-text profit comparison implies that superior information is strictly less profitable than the benchmark, so

$$\Pi^+(u;1) - \Pi^-(u;1) < 0.$$

Define

$$\Delta(u;\beta) \equiv \Pi^+(u;\beta) - \Pi^-(u;\beta).$$

By assumption (A1), $\Delta(u;\beta)$ is continuous on $[0,1]$. Since $\Delta(u;0) > 0$ and $\Delta(u;1) < 0$, the Intermediate Value Theorem implies that there exists at least one $\bar{\beta}(u) \in (0,1)$ such that

$$\Delta(u;\bar{\beta}(u)) = 0.$$

If $\Delta(u;\beta)$ is strictly decreasing in β, then this zero is unique, and the sign comparisons follow immediately.

For part (ii), let $V_B(\beta)$ denote the profit from the continuous continuation of the binary equilibrium at $\beta = 1$, and let $V_N(\beta)$ denote the supremum profit over all feasible non-binary equilibria. By assumption (A2), at $\beta = 1$ we have

$$V_B(1) \geq V_N(1) + \eta.$$

By continuity of both objects in β near $\beta = 1$, there exists $\underline{\beta}(u) < 1$ such that for all $\beta \in [\underline{\beta}(u), 1]$,

$$|V_B(\beta) - V_B(1)| < \eta/3 \qquad \text{and} \qquad |V_N(\beta) - V_N(1)| < \eta/3.$$

Hence, for all such β,

$$V_B(\beta) - V_N(\beta) \geq \big(V_B(1) - \eta/3\big) - \big(V_N(1) + \eta/3\big) \geq \eta/3 > 0.$$

Therefore the binary equilibrium remains strictly optimal for all sufficiently large β. The convergence claims follow from continuity of the equilibrium continuation. □

Proposition A4 shows that the two main qualitative results of the paper survive in this richer environment. First, the profit effect of observing v reverses as the seller's superior information relative to consumers changes: when few consumers are uninformed, observing v helps; when most consumers are uninformed, the signaling problem dominates and observing v can backfire. Second, although the exact assignment of v-types to discount levels becomes more involved for intermediate β, the coarse binary pricing logic from the main model re-emerges once the share of uninformed consumers is sufficiently large.

The reason the binary characterization no longer follows mechanically for all $\beta \in (0, 1)$ is that informed consumers' search behavior depends directly on realized v. In particular, for $v \in [u - p, u - p + s]$, informed consumers never search regardless of the discount, so offering them a positive discount is a pure transfer. Yet from the perspective of managing uninformed consumers' beliefs, the seller may still wish to place some of these consumers in the positive-discount region. This is the extra trade-off that makes the full intermediate-β characterization more complicated than in the main model.