B Supplementary Appendix

B.1 Other Conditions Under Which The Greedy Algorithm is Optimal

Footnote 20 states that the greedy algorithm is optimal if F is convex, and the optimal price on an interval $[\underline{v}, \tilde{v}]$—denoted by $p(\tilde{v})$—has a slope bounded above by 1 and is weakly concave. We verify this statement here. As in the proof of Proposition 4, for any v in the support of F and $w \in [p(v), v]$, let

$$\overline{P}(v, w) \equiv (F(v) - F(w))w + (F(w) - F(p(v)))p(w).$$

This corresponds to the average price if the segment $[p(v), v]$ were to separate into an upper segment $[w, v]$ so that all types above w pay w, and all types in $[p(v), v]$ were to pay price $p(w)$. By direct calculation,

$$\frac{d^2\overline{P}}{dw^2} = 2f(w)(p'(w) - 1) + f'(w)(p(w) - w) + (F(w) - F(p(v)))p''(w).$$

Observe that this term is weakly negative if F is convex, and $p(\tilde{v})$ has a slope bounded above by 1 and is weakly concave; under these conditions $\overline{P}$ is concave in w. Because $\overline{P}(v, v) = \overline{P}(v, p(v))$, it follows that $\overline{P}(v, w)$ is minimized at $w = p(v)$. Therefore, as in Proposition 4, the greedy segmentation attains the lowest possible average price on each of its segments.

B.2 Multidimensional Types

Here we analyze the setting described in Section 3.5.3.

Proposition 11 (Proposition 1$'$)**.** *With simple evidence, across all equilibria, the consumer's interim payoff is bounded above by* $\max\{v(t) - p^*, 0\}$.

Proof. Consider an equilibrium. Let $\tilde{T}$ be the set of types that in equilibrium send the non-disclosure message, T. Thus, every type in $T \backslash \tilde{T}$ sends a message that fully reveals itself. Sequential rationality demands that the monopolist charges a price of $v(t)$ to every such type, leading to an interim payoff of 0. We prove below that the non-disclosure message must induce a price that is no less than p^*.

Suppose towards a contradiction that it leads to a price $\tilde{p}$ that is strictly less than p^*. In equilibrium, if $v(t) > \tilde{p}$, the consumer must be sending the non-disclosure message T (because sending the message $\{t\}$ leads to a payoff of 0). Therefore, in equilibrium,

$$\tilde{T} \supseteq \{t \in T : v(t) > \tilde{p}\} \supseteq \{t \in T : v(t) \geq p^*\}.$$

By charging a price of $\tilde{p}$, the firm's payoff is

$$\begin{aligned} \tilde{p}\mu(\{t \in \tilde{T} : v(t) \geq \tilde{p}\}) &\leq \tilde{p}\mu(\{t \in T : v(t) \geq \tilde{p}\}) \\ &< p^*\mu(\{t \in T : v(t) \geq p^*\}) \\ &= p^*\mu(\{t \in \tilde{T} : v(t) \geq p^*\}), \end{aligned}$$

where the weak inequality follows from $\tilde{T} \subseteq T$, the strict inequality follows from p^* being the (lowest) optimal price, and the equality follows from $\{t \in T : v(t) \geq p^*\} \subseteq \tilde{T}$. Therefore, the monopolist gains from profitably deviating from charging $\tilde{p}$ to a price of p^* when facing the non-disclosure measure, thereby rendering a contradiction. □

To describe the segmentation with rich evidence, consider a sequence of prices $\{p_s\}_{s=0,1,2,\ldots,S}$ where $S \leq \infty$, $p_0 = \overline{v}$, and for every s where $p_{s-1} > \underline{v}$, p_s is the (lowest) maximizer of $p_s(F(p_{s-1}) - F(p_s))$. If $p_{s'} = \underline{v}$ for some s', then we halt the algorithm and set $S = s'$; otherwise, $S = \infty$ and $p_\infty = \underline{v}$. We use these prices to construct sets of types, $(M_s)_{s=1,2,\ldots,S} \cup M_\infty$:

$$M_s \equiv \{t \in T : v(t) \leq p_{s-1}\}.$$
$$M_\infty \equiv \{t \in T : v(t) = \underline{v}\}$$

Because v is quasiconvex and T is convex, M_s is a convex set for every $s = 0, 1, 2, \ldots, S$, and therefore M_s is a feasible message.

Proposition 12 (Proposition 2′). *With rich evidence, there is a Pareto-improving equilibrium in which the consumer's disclosure strategy is*

$$M^*(t) = \begin{cases} M_s & \text{if } p_s < v(t) \leq p_{s-1}, \\ M_\infty & \text{if } t \in M_\infty. \end{cases}$$

When receiving an equilibrium disclosure of the form M_s, the seller charges a price of p_s and sells to all types that send that message.

Proof. We augment the description of the strategy-profile with the off-path belief system where when the seller receives a message $M \notin (\cup_{s=1,\ldots,S} M_s) \cup M_\infty$, she puts probability 1 on a type in M with the highest valuation (i.e. a type in $\arg\max_{t \in M} v(t)$), and charges a price equal to that valuation.

Observe that the seller has no incentive to deviate from this strategy-profile because for each (on- or off-path) message, the price that he is prescribed to charge in equilibrium is her optimal price given the beliefs that are induced by that message.

We consider whether the consumer has a strictly profitable deviation. Let us consider on-path messages first. Consider a consumer type t that is prescribed to send message M_s where $p_s < v(t) \leq p_{s-1}$. Sending any message of the form $M_{s'}$ where $s' < s$ results in a higher price and therefore is not a profitable deviation. All messages of the form $M_{s'}$ where $s' > s$ are infeasible because $t \notin M_{s'}$ for any $s' > s$. Finally, if the type t is such that she is prescribed to send message M_∞, her equilibrium payoff is 0, and sending any other message results in a weakly higher price. Thus, the consumer has no profitable deviation to any other on-path message. There is also no profitable deviation to any off-path message: because for any set M that contains t, $v(t) \leq \max_{t' \in M} v(t')$, any off-path message is guaranteed to result in a payoff of 0. □

Finally, the general notion of an evidence technology facilitating group pricing is the same as before, and the argument for Proposition 5 follows exactly as is.

B.3 Competition Between More than Two Firms

In this section, we show how our prior conclusions generalize to the case of multiple firms that produce differentiated products. Although the economic intuition for how disclosure amplifies competition remains the same, the arguments and notation are necessarily more involved to account for the higher dimensionality introduced by there being more firms.

Suppose that the set of firms is $N \equiv \{1, \dots, n\}$ where the number of firms, n, is at least 2. The consumer has a valuation for each of these products, encoded in its type, $t \equiv (t_1, \dots, t_n)$, where t_i is the consumer's value for the good produced by firm i. We assume that t is drawn according to distribution μ whose support is $T \equiv \left[\underline{t}, \bar{t}\right]^n$, where $0 < \underline{t} < \bar{t}$; for simplicity, we assume that μ has a strictly positive and continuous density f on its support. We say that firm i is type t's *favorite firm* if t_i is weakly higher than t_j for every firm j.

We first construct a partial pooling equilibrium using simple evidence. We then construct an equilibrium with rich evidence. Finally, we compare the equilibrium prices under the two technologies to the model without personalized pricing.

Simple Evidence: For each type t, the set of messages available to the consumer is $\mathcal{M}(t) = \{T, \{t\}\}$. Messages are private and $M_i(t)$ denotes the message sent to firm i. It is useful to define a demand function for firm i assuming that all other firms charge a price of 0. For every non-negative price p, let $Q^i(p) \equiv \mu\left(\{t : t_i - p \geq \max_{j \neq i} t_j\}\right)$ be the probability that the consumer purchases from firm i at a price of p when all other firms charge 0. Let p_1^i denote the lowest maximizer of $pQ^i(p)$. The following result constructs an equilibrium with simple evidence.

Proposition 13. *With simple evidence, there is an equilibrium in which the consumer's disclosure strategy is*

$$M_i^*(t) = \begin{cases} \{t\} & \textit{if } t_i - p_1^i < \max_{j \neq i} t_j, \\ T & \textit{otherwise.} \end{cases}$$

The prices charged by firm i are

$$p_i^*(M) = \begin{cases} \max\{0, t_i - \max_{j \neq i} t_j\} & \textit{if } M = \{t\}, \\ p_1^i & \textit{otherwise.} \end{cases}$$

In equilibrium, every consumer type purchases from its favorite firm.

This equilibrium generalizes the simple partitional form of Proposition 7 and Figure 2. The consumer sends the non-discriminatory message T only when she has strong preferences for the product of a particular firm. Following these messages, whenever firm i receives the message T, it infers that the consumer has a strong preference for its product and that the consumer has sent a fully revealing message to every other firm. Anticipating that every other firm is then charging a price of 0, firm i then charges its optimal local monopoly price, which is p_1^i. When the consumer has mild preferences for the products of each firm, she sends the fully revealing message to each. The firms then compete for her business using standard Bertrand prices (with differentiated products). We omit a proof of this result as it is identical to that of Proposition 7.

Rich Evidence: We now construct an equilibrium using the rich evidence structure of Section 3.5.3 where $\mathcal{M}(t)$, the set of messages that the consumer can send when her type is t, is the set of all closed and convex subsets of T that contain t. We partition T on the basis of each consumer type's favorite firm, and the closest competitor that would like to poach that type's business. For each consumer type t, let $\alpha(t)$ denote type t's favorite firm and $\beta(t)$ denote type t's favorite among the remaining firms, breaking ties in favor of the firm that has the lowest index in each case.[41] It is helpful to define a localized demand function: for every non-negative price p and every (Borel) set A, let $Q^i(p, A) \equiv \mu\left(\{t \in A : t_i - p \geq \max_{j \neq i} t_j\}\right)$ denote the probability that the consumer purchases from firm i at price p when the consumer's type is in A and every other firm is charging a price of 0.

For each firm i and competitor j, define a sequence of prices $\{p_s^{ij},\}_{s=0,1,2,\ldots}$ and sets of types $\{M_s^{ij}\}_{s=0,1,2,\ldots}$ such that $p_0^{ij} = \bar{t} - \underline{t}$, $M_s^{ij} \equiv \left\{t \in T : i = \alpha(t), j = \beta(t), t_i - t_j \leq p_{s-1}^{ij}\right\}$, and p_s^{ij} is the smallest maximizer of $pQ(p, M_s^{ij})$. The set M_s^{ij} denotes all consumer types for which i is the favorite, j is the second favorite, and each type is willing to pay as much as p_{s-1}^{ij} to obtain the product from firm i; this is a convex set. The price p_s^{ij} is the local monopoly price of firm i when it knows that the consumer's type is in M_s^{ij}. Denoting the set of types for whom firms i and j are two equally favorite firms by $M_\infty^{ij} \equiv \{t \in T : i = \alpha(t), j = \beta(t), t_i = t_j\}$ so that $p_\infty^{ij} = 0$, we now state and prove the following result.

Proposition 14. *With rich evidence, there is an equilibrium in which the consumer's disclosure strategy is*

$$M^*(t) = \begin{cases} M_s^{ij} & \textit{if } t \in M_s^{ij} \textit{ and } t_i - t_j > p_s^{ij}, \\ M_\infty^{ij} & \textit{if } t \in M_\infty^{ij}. \end{cases}$$

When receiving an equilibrium disclosure of the form M_s^{ij}, firm i charges a price of p_s^{ij}, and all other firms charge a price of 0. The consumer always purchases the good from her favorite firm.

Proof. We begin by proving that each M_s^{ij} is convex for each s. If t and t' are each in M_s^{ij}, then $t_i \geq t_j \geq \max_{k \in N \setminus \{i,j\}} t_k$, and $t'_i \geq t'_j \geq \max_{k \in N \setminus \{i,j\}} t'_k$. Therefore, for every $\rho \in (0, 1)$, the following are true:

$$\rho t_i + (1-\rho)t'_i \geq \rho t_j + (1-\rho)t'_j \geq \rho \max_{k \in N \setminus \{i,j\}} t_k + (1-\rho) \max_{k \in N \setminus \{i,j\}} t'_k \geq \max_{k \in N \setminus \{i,j\}} (\rho t_k + (1-\rho)t'_k)$$

$$\rho t_i + (1-\rho)t'_i - \rho t_j + (1-\rho)t'_j = \rho(t_i - t_j) + (1-\rho)(t_i - t_j) \leq \rho p_{s-1}^{ij} + (1-\rho)p_{s-1}^{ij} = p_{s-1}^{ij}.$$

Therefore, $\rho t + (1-\rho)t'$ is also an element of M_s^{ij}, and so M_s^{ij} is convex for every s. Clearly, M_∞^{ij} is also convex by the same argument.

[41] Formally, treating min as the lowest element of a set, $\alpha(t) = \min\{i \in \{1, \ldots, n\} : t_i \geq t_j \text{ for all } j\}$, and $\beta(t) = \min\{i \in \{1, \ldots, n\} \setminus \{\alpha(t)\} : t_i \geq t_j \text{ for all } j \neq \alpha(t)\}$. These tie-breaking rules are for completeness, but any tie-breaking rule that preserves $\alpha(t) \neq \beta(t)$ suffices.

The remainder of this argument follows the same logic as that of Proposition 8. If firm i receives an off-path message M, she holds degenerate beliefs $\delta_{\tau(i,M)}$ where type $\tau(i, M)$ is the type in M that has the highest value of $t_i - \max_{j \neq i} t_j$. Because M is a closed set, $\tau(i, M)$ is defined for every M. Given such beliefs, the firm charges a price $p_i(M) = \max\{t_i - \max_{j \neq i} t_j, 0\}$.

First, we prove that given the pricing strategies, no consumer type has an incentive to deviate. Consider a consumer type such that $t \in M_s^{ij}$ and $t_i - t_j > p_s^{ij}$. A message of M_s^{ij} to each firm induces all firms other than firm i to set a price of 0 and induces firm i to set a price of p_s^{ij}. No other message can lead to lower prices from any other firm. Moreover, the consumer cannot send any other equilibrium path message to firm i that leads to lower prices. Finally, every off-path message sent to firm i can only increase the price since for any off-path message $M \in \mathcal{M}(t)$, $p_i(M) \geq M_s^{ij}$.

Second, we argue that firms have no incentives to deviate in their pricing strategies. For any on-path message, the prices charged are (by construction) optimal. For any off-path message M, each firm assumes that the consumer sent the equilibrium path message to other firms. If Firm $i = \alpha(\tau(i, M))$ (i.e., Firm i is the favorite), then it assumes that other firms are charging a price of 0, in which case $p_i(M)$ is optimal. If Firm i is not the favorite, then it charges a price of 0, and anticipates any strictly positive price to be rejected. □

Benefits of Personalized Pricing: We show that the equilibria constructed in Propositions 13 and 14 improve consumer surplus relative to the benchmark of uniform pricing. For the benchmark model, to guarantee full market coverage, we assume that the consumer has to purchase the product from one firm. We say that a distribution is symmetric if whenever t' is a permutation of t, $f(t') = f(t)$.

Proposition 15. *If f is symmetric and log-concave, then every type has a strictly higher payoff in the equilibria constructed in Propositions 13 and 14 than in the benchmark setting with uniform pricing.*

Proof. Because f is log-concave, it follows directly from Caplin and Nalebuff (1991) that for each firm i, $Q^i(p)$ is log-concave. Since firms are symmetric, we drop the subscripts. Let p^* denote the price from the symmetric equilibrium of the game without personalized pricing. Observe that for each firm i,

$$p_1^i = -\frac{Q(p_1^i)}{q(p_1^i)} < -\frac{Q(0)}{q(0)} = p^*,$$

where the first equality follows from the first-order condition that p_1^i solves, the inequality follows from Q being strictly log-concave, and the second equality follows from the first-order condition that p^* solves. Now all consumers are necessarily better off in the equilibrium constructed with simple evidence. The argument is the same as in Proposition 9: all consumer types either pay a price of p_1^i or below. In the equilibrium constructed with rich evidence, it follows from symmetry that $p_1^{ij} = p_1^i$, and therefore, once again, all consumer types either pay a price of p_1^i or below. □