Appendix A – Proofs

Theorem 1

Proof. The existence of the threshold follow because $U_t(v)$ is increasing convex and with slope no greater than one, which is easily shown by induction on t. (The fact that $U_t(v)$ is increasing with slope no greater than one ensures that each threshold v_t is uniquely defined.)

We now prove that $v_t \leq \bar{v}$ together with $\beta U_{t+1}^{new} \leq \bar{v}$ by induction. This is true for $t = T$, because from (7), $\beta U_{T+1}^{new} = \beta v_{out} \leq \bar{v}$ and $v_T = \beta v_{out} \leq \bar{v}$. Assume the result is true for $t+1 \leq T$ and let us show that it holds for t. First,

$$U_{t+1}^{new} = \max\{E \max\{V, \beta U_{t+2}(V)\} - k, \beta U_{t+2}^{new}\}$$

If the maximum is achieved at $\beta U_{t+2}^{new} \leq \bar{v}$, then $\beta U_{t+1}^{new} \leq \bar{v}$. Otherwise,

$$\begin{aligned} E \max\{V, \beta U_{t+2}(V)\} - k \leq\ & E \max\{V, v_{t+1}\} - k \\ =\ & v_{t+1} + E(V - v_{t+1})^+ - E(V - \bar{v})^+ \\ \leq\ & v_{t+1} + \bar{v} - v_{t+1} \\ =\ & \bar{v}. \end{aligned}$$

This proves that $\beta U_{t+1}^{new} \leq \bar{v}$ always. Moreover,

$$\begin{aligned} \beta U_{t+1}(\bar{v}) =\ & \alpha_{t+1} \beta U_{t+1}^{new} + (1 - \alpha_{t+1}) \beta U_{t+1}^{old}(\bar{v}) \\ \leq\ & \alpha_{t+1} \bar{v} + (1 - \alpha_{t+1}) \beta \max\{\bar{v} - k, \beta U_{t+2}(\bar{v})\} && \text{from the induction property} \\ \leq\ & \alpha_{t+1} \bar{v} + (1 - \alpha_{t+1}) \beta \bar{v} && \text{because } \bar{v} \geq v_{t+1} \text{ implies } \bar{v} \geq \beta U_{t+2}(\bar{v}) \\ \leq\ & \alpha_{t+1} \bar{v} + (1 - \alpha_{t+1}) \bar{v} \leq \bar{v}. \end{aligned}$$

Hence $\beta U_{t+1}(\bar{v}) \leq \bar{v}$, which itself implies that $v_t \leq \bar{v}$.

Finally, when $\beta = 1$, we prove that $v_t \geq v_{t+1}$ by induction. We show this by proving that, in each period, $v_t \geq v_{t+1}$ and $U_{t+1}^{new} \geq v_t$. This is true for $t = T$ since $v_T = v_{T+1} = v_{out}$. Assume it is true for t, and let us show it for $t - 1$. First, we have that $U_t^{old}(v_t) = \max\{v_t - k, v_t\} = v_t$ and $U_t^{new} \geq U_{t+1}^{new} \geq v_t$. This implies that $U_t(v_t) \geq v_t$ and hence, since $U_t(\cdot)$ is non-decreasing with slope no greater than one, $v_{t-1} \geq v_t$. Second, we have that $U_t(v_{t-1}) = \alpha_t U_t^{old}(v_{t-1}) + (1 - \alpha_t) U_t^{new} = v_{t-1}$. Since $v_{t-1} \geq v_t$, $U_t^{old}(v_{t-1}) = \max\{v_{t-1} - k, U_{t+1}(v_{t-1})\} \leq v_{t-1}$. Hence, $U_t^{new} \geq v_{t-1}$. ∎

Theorem 2

Proof. We begin by showing that there exists an optimal policy under which store visits only occur in periods in which there is a positive probability that the product will change. To see this, take t such that $\alpha_t = 0$, i.e., the product between periods $t-1$ and t does not change. Then, by its definition, $U_t(v) = \alpha_t U_t^{new} + (1-\alpha_t)U_t^{old}(v) = U_t^{old}(v)$. This implies that $U_t^{old}(v) = \max\{v-k, \beta U_{t+1}(v)\} = U_t(v)$ and $U_t^{new} = \max\{E\max\{V, \beta U_{t+1}(V)\} - k, \beta U_{t+1}^{new}\}$.

Consider two possible scenarios for period $t-1$.

Suppose first that the customer knew the product's valuation in period $t-1$. Then, in period $t-1$, the customer decided whether or not to visit the store by solving $U_{t-1}^{old}(v) = \max\{v - k, \beta U_t(v)\} = \max\{v-k, \beta \max\{v-k, \beta U_{t+1}(v)\}\} = \max\{v-k, \beta^2 U_{t+1}(v)\}$. Clearly, the maximum will never be achieved by visiting and buying in period t, which would generate a profit of $\beta(v-k)$.

Suppose now that the customer did not know the product's valuation in period $t-1$. In that case, the customer would have visited the store in period $t-1$ if and only if $\beta U_t^{new} \leq EJ_{t-1}(V) - k = E\max\{V, \beta U_t(V)\} - k = E\max\{V, \beta(V-k), \beta^2 U_{t+1}(V)\} - k = E\max\{V, \beta^2 U_{t+1}(V)\} - k$. Again, the maximum cannot be reached by visiting and buying in period t, yielding a profit of $\beta(V-k)$. Finally, the remaining case involves a situation in which the customer did not visit the store in period $t-1$, but then visited the store in period t. We claim that this is again not possible. In this scenario, it must be that $\beta U_t^{new} > E\max\{V, \beta^2 U_{t+1}(V)\} - k$; and because the product was not changed in period t, the state in that period is *new*, so $U_t^{new} = E\max\{V, \beta U_{t+1}(V)\} - k$. This means that

$$E\max\{\beta V, \beta^2 U_{t+1}(V)\} - \beta k > E\max\{V, \beta^2 U_{t+1}(V)\} - k$$

which yields

$$\begin{aligned} k &> \frac{E\max\{V, \beta^2 U_{t+1}(V)\} - E\max\{\beta V, \beta^2 U_{t+1}(V)\}}{1-\beta} \\ &\geq E\left(V - v_t\right)^+. \end{aligned}$$

The last inequality follows by considering each possible sample path. If $V \leq \beta^2 U_{t+1}(V)$, then $V \leq v_t$ so $E\left(V-v_t\right)^+ = 0$ and the inequality holds. If $\beta^2 U_{t+1}(V) \leq V \leq \beta U_{t+1}(V)$, then again $V \leq v_t$ so $E\left(V-v_t\right)^+ = 0$. Finally, when $V \geq \beta U_{t+1}(V)$, the left-hand side of the inequality is equal to $V \geq E\left(V-v_t\right)^+$. As a result, $k = E\left(V-\bar{v}\right)^+ > E\left(V-v_t\right)^+$, implying that $v_t \geq \bar{v}$. This contradicts the result in Theorem 1.

We then have that, if $\alpha_t = 1$, then $U_t(v) = U_t^{new}$; if $\alpha_t = 0$, then we can run the recursion with $U_t(v) = \beta U_{t+1}(v)$ because no visit takes place in that period. Let $t+\tau$ be the next period in which the product changes. Then, $U_t(v) = U_t^{new} = \max\{E\max\{V, \beta^d U_{t+\tau}^{new}\} - k, \beta^\tau U_{t+\tau}^{new}\}$. Hence, letting $\tilde{U}_s^{new} = \beta^s U_s^{new}$, we have

$$\tilde{U}_t^{new} = \tilde{U}_{t+\tau}^{new} + \beta^t \left[E\left(V - \frac{\tilde{U}_{t+\tau}^{new}}{\beta^t}\right)^+ - k\right]^+. \tag{13}$$

This recursion implies that $\tilde{U}_t^{new} = \beta^t U_t^{new}$ (for the periods where the product has changed) is non-increasing over time. This monotonicity, together with assumption (7), allows us to completely characterize the optimal visit pattern. We next show that under a deterministic product rotation policy (i.e., when $\alpha_t \in \{0,1\}$ for all t), it is optimal to visit the store in period t if and only if the product has changed, i.e., $\alpha_t = 1$.

We already know that if the customer visits the store in period t then it cannot be that $\alpha_t = 0$, so it must be that $\alpha_t = 1$ in view of the deterministic product rotation policy. Next, we prove that the customer is actually better off visiting the store when the product changes.

Equation (13) implies that $\tilde{U}_t^{new}$ is non-increasing over time (for the periods in which new products are introduced). Let $t_1 < \cdots < t_N$ be the N periods in which a new product is introduced. Let $\beta_n = \beta^{t_n}$, $w_{N+1} = \beta^{T+1} v_{out}$ and $w_n = w_{n+1} + \beta_n \Big[E(V - w_{n+1}/\beta_n)^+ - k\Big]^+$. It first follows that w_n is equal to $\tilde{U}_{t_n}^{new}$.

We show that $w_{n+1}/\beta_n < \bar{v}$ for all $n = 1, \cdots, N$ by backward induction. Indeed, this is true for $n = N$ because $w_{N+1}/\beta_N = \beta^{T+1-t_N} v_{out} \le \beta v_{out} < \bar{v}$, where the last inequality follows from (7). Assume that it is valid for n, i.e., $w_{n+1}/\beta_n < \bar{v}$. If $w_n = w_{n+1}$, then the inequality holds because $\beta_{n-1} = \beta^{t_{n-1}} > \beta^{t_n} = \beta_n$ as $t_{n-1} < t_n$, so $w_n/\beta_{n-1} = w_{n+1}/\beta_{n-1} < w_{n+1}/\beta_n < \bar{v}$ by induction. Otherwise, $w_n = w_{n+1} + \beta_n \left[E(V - w_{n+1}/\beta_n)^+ - k\right]$ so $w_n/\beta_{n-1} = \Big[E \max(V, w_{n+1}/\beta_n) - k\Big]\beta_n/\beta_{n-1} < \Big[E\max\{V, \bar{v}\} - k\Big]\beta_n/\beta_{n-1} = \bar{v}\beta_n/\beta_{n-1} < \bar{v}$. Thus, $w_{n+1}/\beta_n < \bar{v}$ for all n. This implies that $E(V - w_{n+1}/\beta_n)^+ > E(V - \bar{v})^+ = k$ for all n. That is, the customer will visit the store as many times as the firm introduces a new product. This proves the result.

Note that if the condition in (7) does not hold, then $w_{N+1}/\beta_N \ge \bar{v}$, implying that $E(V - w_{N+1}/\beta_N)^+ \le E(V - \bar{v})^+ = k$. This, in turn, implies that $w_N = w_{N+1}$ and, repeating the argument, we have that $w_n = w_{N+1}$ for all n. In that case, the customer never visits the store.

To conclude, we have that $w_{N+1} = \beta^{T+1} v_{out}$ and $w_n = w_{n+1} + \beta_n \Big[E(V - w_{n+1}/\beta_n)^+ - k\Big]$. Recall also that w_n is equal to $\tilde{U}_t^{new}$ when there have been n product changes between the beginning of the selling season until period t_n (included). Hence, v_{t_n} is the unique solution to $v = \beta U_{t_n+1}(v) = \beta^{t_{n+1}-t_n} U_{t_{n+1}}^{new} = \tilde{U}_{t_{n+1}}^{new}/\beta_n = w_{n+1}/\beta_n$.

We finally show that the thresholds v_{t_i} are decreasing as long as β is sufficiently large. This guarantees that the retailer's profit is concave. It follows from the recursion in (11) that $w_1 > w_2 > \cdots > w_n > w_{n+1}$. Therefore, there exists β_0 such that $\beta^{T-1} w_i > w_{i+1}$ for all $\beta > \beta_0$. Then, for $\beta > \beta_0$,

$$v_{t_i} = \frac{w_{i+1}}{\beta^{t_i}} < \beta^{T-1} \frac{w_i}{\beta^{t_i}} \frac{\beta^{t_{i-1}}}{\beta^{t_{i-1}}} = \beta^{T-1} \frac{v_{t_{i-1}}}{\beta^{t_i - t_{i-1}}} = \beta^{(T-1)-(t_i - t_{i-1})} v_{t_{i-1}} < v_{t_{i-1}},$$

and the result follows. ∎

Theorem 3

Proof. Because $\beta = 1$, we have that $z_{i+1}^n \stackrel{def}{=} v_{t_i}$, for $i = 1, \ldots, n$, satisfies $z_1^n \geq \cdots \geq z_n^n \geq z_{n+1}^n = v_{out}$. That is, z_i^n is a customer's threshold associated to the i^{th} product rotation when there are a total of n product changes throughout the season (so, for example, $z_2^{n+1} = z_1^n$). Also, $v_{t_i}^{myopic} = v_{out}$ for all t_i. To compare the optimal values n^{myopic} and n^{strat}, we first note that

$$\begin{aligned} n^{myopic} \leq n^{strat} \quad & \text{if} \;\; (1 - v_{out}) v_{out}^{n-1} \leq \bar{F}(z_1^n) \textstyle\prod_{i=2}^n F(z_i^n) \text{ at } n = n^{strat} \\ n^{myopic} \geq n^{strat} \quad & \text{if} \;\; \bar{F}(z_1^n) \textstyle\prod_{i=2}^n F(z_i^n) \leq (1 - v_{out}) v_{out}^{n-1} \text{ at } n = n^{myopic}. \end{aligned}$$

We define

$$P_n \stackrel{def}{=} \frac{\bar{F}(z_1^n) \prod_{i=2}^n F(z_i^n)}{v_{out}^{n-1}}.$$

Then,
$$P_{n+1} = \frac{\bar{F}(z_1^{n+1}) \prod_{i=2}^{n+1} F(z_i^{n+1})}{v_{out}^n} = P_n \times \frac{F(z_1^n) \bar{F}(z_1^{n+1})}{v_{out} \; \bar{F}(z_1^n)}.$$

First, note that when $n = 1$, $P_1 = \bar{F}(z_1^1) = 1 - v_{out}$. Consider now the factor

$$\frac{F(z_1^n) \bar{F}(z_1^{n+1})}{v_{out} \; \bar{F}(z_1^n)}. \tag{14}$$

Because $z_i^n = E \max\{V, z_{i+1}^n\} - k$ is increasing in n for fixed i and it is also bounded above by $\bar{v}$, we have in particular that $z_1^n, z_1^{n+1} \to \bar{v} > v_{out}$ as $T = n \to \infty$, where recall that $\bar{v}$ solves the equation $\bar{v} = E \max\{V, \bar{v}\} - k$. Therefore,

$$\frac{F(z_1^n) \bar{F}(z_1^{n+1})}{v_{out} \; \bar{F}(z_1^n)} \to \frac{F(\bar{v})}{F(v_{out})} > 1 \;\; \text{as} \;\; T = n \to \infty.$$

It follows that there exists $\hat{n}$ such that the factor in (14) is greater than 1 for all $n \geq \hat{n}$. This implies that P_n is increasing for $n \geq \hat{n}$. Moreover,

$$\frac{P_{n+1}}{P_n} = \frac{F(z_1^n) \bar{F}(z_1^{n+1})}{v_{out} \; \bar{F}(z_1^n)} \to \frac{F(\bar{v})}{F(v_{out})} > 1 \;\; \text{as} \;\; T = n \to \infty.$$

Then, we must have that P_n is unbounded and therefore $P_n \to \infty$ as $T = n \to \infty$. Thus, there exist $1 \leq \hat{n}_1 \stackrel{def}{=} \max\{n : P_m \leq 1 - F(v_{out}) \text{ for all } 0 \leq m \leq n\}$ and $\hat{n}_2 \stackrel{def}{=} \min\{n : P_m > 1 - F(v_{out}) \text{ for all } m \geq n\}$. The result follows from the definitions of P_n, $\hat{n}_1$ and $\hat{n}_2$. ∎

Theorem 4

Proof. The existence of a threshold v_t is a direct extension of Theorem 1, where the only change is that valuations are no longer stationary. The rest of the proof is similar to that of Theorem

2. First, it is clear that, if there is no change in the assortment ($d_{t_i} = 0$ and $\alpha_{t_i} = 0$), then the customer should never visit the store. As a result, we only need to define the thresholds for $t = t_i$. Moreover, $v_{t_i} = \beta^{t_{i+1}-t_i} U^{new}_{t_{i+1}}$. However, since valuations are not stationary, it is possible that even when $d_{t_i} > 0$ and $\alpha_{t_i} = 1$ it may still better not to visit the store. Indeed, when a new assortment with depth $d_{t_i} > 0$ is introduced, the customer has two choices: to visit the store, thereby obtaining an expected utility equal to $E \max\{V_{d_{t_i}}, v_{t_i}\} - k$; or to not visit the store, getting an expected utility $\beta^{t_{i+1}-t_i} U^{new}_{t_{i+1}} = v_{t_i}$. Hence,

$$v_{t_{i-1}} = \beta^{t_i - t_{i-1}} U^{new}_{t_i} = \beta^{t_i - t_{i-1}} \max\left\{E \max\{V_{d_{t_i}}, v_{t_i}\} - k, v_{t_i}\right\} = \beta_{t_i - t_{i-1}} \left\{v_{t_i} + \left[E\left(V_{d_{t_i}} - v_{t_i}\right)^+ - k\right]^+\right\}.$$

A customer's visit takes place if and only if $E\left(V_{d_{t_i}} - v_{t_i}\right)^+ - k \geq 0$. Moreover, the customer makes a purchase if and only if $V_{d_{t_i}} \geq v_{t_i}$. ■

Proposition 1

Proof. For $T = 2$, the retailer profit is given by

$$\pi_R = r - rF(v_1)^{d_1} F(v_{out})^{d_2} - c(d_1) - c(d_2),$$

with the recursion $v_1 = v_{out} + E\left(V_{d_2} - v_{out}\right)^+ - k$.

We prove the result assuming that d_1, d_2 are continuous. The argument is similar for integer variables. In that case, we would show that if $d_1 > d_2$, then the retailer would be better off using $d_1 - 1, d_2 + 1$.

The first-order conditions with respect to d_1 and d_2 are given by

$$\begin{aligned} c'(d_1) &= -\ln(F(v_1) r F(v_1)^{d_1} F(v_{out})^{d_2}, \\ c'(d_2) &= -\Big[\ln(F(v_{out})) + d_1 \frac{f(v_1)}{F(v_1)} \frac{\partial v_1}{\partial d_2}\Big] r F(v_1)^{d_1} F(v_{out})^{d_2}. \end{aligned}$$

Log-concavity of $F(v_1(d_2))$ with respect to d_2 yields

$$d_1 \frac{f(v_1)}{F(v_1)} \frac{\partial v_1}{\partial d_2} \leq \ln\left(\frac{F(v_1)}{F(v_{out})}\right).$$

This directly implies $c'(d_1) \leq c'(d_2)$. Because $c(\cdot)$ is convex, it follows that, at optimality, we must have that $d_1^* \leq d_2^*$.

We next check the log-concavity condition for the case of uniform valuations. We provide this condition for general $T \geq 2$.

The log-concavity condition on $F(v_t(\mathbf{d}_{t+1}))$ is always satisfied for the case of uniform valuations in $[0,1]$ with $\beta = 1$ and $k = 0$. In this setting, the threshold recursion is given by $v_T = v_{out}$ and

$$v_t = 1 - \frac{1 - v_{t+1}^{1+d_{t+1}}}{1 + d_{t+1}}$$

for $t = 1, ..., T-1$. As a result, we have

$$\frac{\partial v_t}{\partial d_{t+1}} = \frac{1 - v_{t+1}^{1+d_{t+1}}}{(1 + d_{t+1})^2} + \frac{v_{t+1}^{1+d_{t+1}} \ln v_{t+1}}{1 + d_{t+1}}.$$

The log-concavity condition of $F(v_t(\mathbf{d}_{t+1}))$ thus becomes

$$\ln\left(\frac{F(v_t)}{F(v_{t+1})}\right) \geq d_{t+1}^* \frac{f(v_t)}{F(v_t)} \frac{\partial v_t}{\partial d_{t+1}}$$

or equivalently

$$\ln\left(\frac{1 - \frac{1 - v_{t+1}^{1+d_{t+1}}}{1+d_{t+1}}}{v_{t+1}}\right) \geq d_{t+1} \frac{1}{1 - \frac{1 - v_{t+1}^{1+d_{t+1}}}{1+d_{t+1}}} \left(\frac{1 - v_{t+1}^{1+d_{t+1}}}{(1 + d_{t+1})^2} + \frac{v_{t+1}^{1+d_{t+1}} \ln v_{t+1}}{1 + d_{t+1}}\right)$$

which only depends on v_{t+1} and d_{t+1}. For ease of exposition, let's write this condition as

$$\ln\left(\frac{1 - \frac{1 - z^{1+d}}{1+d}}{z}\right) \geq d \frac{1}{1 - \frac{1 - z^{1+d}}{1+d}} \left(\frac{1 - z^{1+d}}{(1 + d)^2} + \frac{z^{1+d} \ln z}{1 + d}\right),$$

which is equivalent to

$$\ln\left(\frac{d + z^{1+d}}{z(1 + d)}\right) - \frac{d}{1 + d} \frac{1 - z^{1+d} + (1 + d) z^{1+d} \ln z}{d + z^{1+d}} \geq 0.$$

The derivative of the left-hand side is

$$-\frac{d^2 (1 - z^{1+d} + (1 + d) z^{1+d} \ln z)}{z(d + z^{1+d})},$$

which one can verify that is negative for $z \in (0, 1)$. That is, the left-hand size of the inequality is decreasing and its limit as $z \to 1$ is zero. As a result, the condition is verified for any values of v_{t+1} and d_{t+1}. ■

Appendix B – Randomized Assortment Changes

In this section, we examine the case of general randomized product rotation policy. In the presence of myopic customers, we show that it is optimal for the retailer to restrict attention to a deterministic product rotation policy, i.e., one with $\alpha_t = 0$ or 1 for all $t = 2, ..., T$.

Myopic Customers

We first revisit the sequence of events from the point of view of a single myopic customer. Prior to period t, there is uncertainty regarding the action of the retailer: with probability α_t a new product is introduced and with probability $1 - \alpha_t$ the same product from the last period is kept on offer. At the beginning of period t, the customer learns the realization of A_t – i.e., whether the product was refreshed or not. (In practice, this may be facilitated by emails, catalogs or other promotional materials indicating that a new product line is in stores.) At that point, the myopic customer decides whether or not to visit the store (incurring a cost k) and, if a visit takes place, whether or not to purchase the product. Because the customer is myopic, her decision is the same as what is described in §4.2 – i.e., the myopic customer only visits when a new product is introduced and purchases the product if her realized valuation is no smaller than the discounted value of the outside option.

Following a similar analysis as in §4.3, the probability that a purchase takes place throughout the selling season is

$$S(\alpha) = \sum_{l=1}^{T} p_l(\alpha) = 1 - \prod_{s=1}^{T} \left(1 - \alpha_s \bar{F}(v_s)\right)$$

(we stress the dependence on α in this expression, but we will omit it in the subsequent analysis unless noted). The retailer maximizes $rS(\alpha) - c\sum_{l=2}^{T} \alpha_l$. We next show that the optimal rotation policy is deterministic. To that end, we first calculate

$$\frac{\partial S}{\partial \alpha_t} = \bar{F}(v_t) \prod_{\substack{s=1 \\ s \neq t}}^{T} \left(1 - \alpha_s \bar{F}(v_s)\right).$$

Let α^* be the optimal solution to the retailer's maximization problem. Suppose that there exist a $\alpha_n^*, \alpha_m^* \in (0, 1)$. Let $A = \alpha_n^* + \alpha_m^*$. Let us now consider the optimization problem in which we fix the policy for all periods except for n and m to the values α_i^* (for $i \neq n, m$). Then, α_n^* and α_m^* maximize the remaining optimization problem:

$$\max_{\alpha_n, \alpha_m} \quad rS(\alpha_n, \alpha_m, \{\alpha_i^*, i \neq n, m\}) - c \sum_{i=2, i \neq n, m}^{T} \alpha_i^* - c\alpha_n - c\alpha_m,$$

subject to $\alpha_n + \alpha_m = A$. Note that the cost remains equal to that in the original solution. We can restate this optimization problem as:

$$\max_{\alpha_n, \alpha_m} \quad rS(\alpha_n, A - \alpha_n, \{\alpha_i^*, i \neq n, m\}) - c \sum_{i=2}^{T} \alpha_i^*.$$

The derivative of the sales function with respect to α_n is given by

$$\bar{F}(v_n) \prod_{\substack{s=1 \\ s \neq n,m}}^{T} \left(1 - \alpha_s^* \bar{F}(v_s)\right) \left(1 - (A - \alpha_n)\bar{F}(v_m)\right) -$$

$$\bar{F}(v_m) \prod_{\substack{s=1 \\ s \neq n,m}}^{T} \left(1 - \alpha_s^* \bar{F}(v_s)\right) \left(1 - \alpha_n \bar{F}(v_n)\right).$$

Let $K = \prod_{\substack{s=1 \\ s \neq n,m}}^{T} \left(1 - \alpha_s^* \bar{F}(v_s)\right)$, a constant with respect to α_n. We now compute the second-order derivative of the sales function with respect to α_n, which is equal to $2\bar{F}(v_n)\bar{F}(v_m)K > 0$. Therefore, this restricted sales function is convex in α_n, indicating that the optimal α_n is either $\max\{-1+A, 0\}$ or $\min\{A, 1\}$. If $\alpha_n = 0$ or 1, that is a contradiction. If $\alpha_n = -1 + A$ or A, then $\alpha_m = A - \alpha_n = 1$ or 0, respectively – again a contradiction. Suppose now that there exists a single $\alpha_n^* \in (0, 1)$, while all other periods' policies are deterministic. Then, if $\mathcal{S} = \{i : \alpha_i^* = 1\}$, we have

$$rS(\alpha) - c \sum_{l=2}^{T} \alpha_l = r \left[1 - \prod_{i \in \mathcal{S}} F(v_i) \left(1 - \alpha_n^* \bar{F}(v_n)\right)\right] - c|\mathcal{S}| - c\alpha_n^*.$$

This function is linear in α_n^*, so we must have that $\alpha_n^* = 0$ or 1 – a contradiction. Thus, when customers are myopic, the optimal product rotation policy is deterministic.

Strategic Customers

The same sequence of events applies when customers are strategic and the product rotation policy is random. In particular, at the beginning of each period t, the customer learns the realization of A_{t-1} – i.e., whether the product was refreshed or not – while the future product rotation policy is still unknown to the customer. In the case of strategic customers, the thresholds v_t depend on the firm's product rotation policy $\alpha_1, \ldots, \alpha_{T-1}$. The analysis quickly becomes intractable. Instead, we performed a numerical study covering over 600 scenarios and determined the retailer's optimal refreshment policy in each scenario via a grid search of all possible policies. We tested different instances by combining the following parameter values: $T = 5, \beta = 1, c \in [0, 0.2], k \in [0, 0.5], v_{out} \in [0, 0.9]$ and $V \sim \mathcal{U}[0, 1]$. In all scenarios, the optimal solution was deterministic: $\alpha_t^* = 0, 1$. This suggests that randomized policies hurt the retailer, even when customers are strategic. Indeed, adding uncertainty into the customer's buying process seems to reduce the sales probability. The following example illustrates this effect.

Consider the customer's problem with $T = 3$, $\beta = 1$ and $k = 0$ (so that visits always take place). In period $t = 1$, the customer will always see a new product. Let us compare a policy where the retailer only introduces a product in period $t = 2$, i.e., $\alpha_2 = 1, \alpha_3 = 0$, to one where a product is introduced in periods $t = 2$ and/or $t = 3$ with 50% probability, respectively. That is, $\alpha_2 = \alpha_3 = 0.5$. The retailer incurs the same product rotation cost c under the two policies. In the first scenario, the customer may purchase the product in periods 1 or 2. Let p_t be the probability that the customer buys the product in period t and p_{last} the probability that it buys in $t = 3$ when a product is introduced then. In the second scenario, one has to consider the possible product rotation realizations from the perspective of the customer. There are four such possible realizations, each occurring with a probability 0.25:

- No new product is introduced (we call this scenario NN). Then, $p_2^{NN} = p_3^{NN} = 0$.
- A new product is released only in $t = 3$. Then, $p_2^{NY} = 0$ and $p_3^{NY} = p_{last}$.
- A new product is released only in $t = 2$. Then, $p_3^{YN} = 0$ and $p_2^{YN} < p_{last}$ because the customer's purchase threshold in period $t = 2$ is higher relative to a setting in which there is certainty about the realization of A_2 as a refreshed version of the product may still appear in period $t = 3$.
- Two new products are released, in $t = 2$ and $t = 3$. Then, $p_3^{YY} = p_{last}$ and $p_2^{YY} < p_{last}$.

Overall, the expected sales probability in periods $t = 2, 3$ is $(p_3^{NY} + p_2^{YN} + p_2^{YY} + (1 - p_2^{YY})p_3^{YY})/4 \leq p_{last}$. Hence, a deterministic product rotation policy yields higher expected profit for the retailer compared to a randomized policy with $\alpha_i = 0.5$. A similar reasoning applies to different randomized policies and longer time horizons.

Appendix C – Extensions

Customer Heterogeneity in Arrivals and Departures

Our model can be extended to incorporate different sales windows for each customer class. Specifically, we assume here that a customer of type j "discovers" the need or interest in buying a product at a time $A_j \geq 1$ (arrival time) and is willing to explore different product offers until time period D_j (departure time). The time window in which customers remain in the system can be regarded as the customers' patience, as in Besbes and Lobel (2015). Customer types may also differ in terms of their search costs k_j, distribution of valuations F_j, unit revenue rates r_j, and their value for the outside option $v_{out,j}$. Each customer type is defined by a frequency within the entire population given by θ_j. This description of customer heterogeneity with respect to arrival and departure times is broad, and it includes, for example, the case of random arrivals and random stays in the system. If arrivals occur uniformly over the selling season $[1, T]$ (e.g., arrivals follow a stationary Poisson process), and customers remain in the system a random number of periods that follows a geometric distribution with rate μ, then the proportion of customers with $A_j = a$ and $D_j = d$, $1 \leq a \leq d \leq T$, is $\theta_j = \frac{1}{T} \frac{(1-\mu)\mu^{d-a}}{1-\mu^{T+1-a}}$.

Now the profit associated with customer type j only depends on n_j, the number of new products offered during the customer's stay in the system. We define $n_j = 1 + \sum_{t=A_j+1}^{D_j} \alpha_t$ and write the retailer's revenue from type j as $R_j^i(n_j) = r_j \times SALES_{n_j}^i$, for $i = myopic, strat$. We next present the retailer's optimization problem. The retailer maximizes the expected revenue from sales minus the cost of renewing the assortment:

$$\begin{aligned} \max_{\alpha_2,\ldots,\alpha_T \in \{0,1\}} \quad & \pi_R(n) \stackrel{def}{=} \sum_{j=1}^{J} \theta_j R_j(n_j) - c(n-1) \\ \text{s.t.} \quad & n = 1 + \sum_{t=2}^{T} \alpha_t \\ & n_j = 1 + \sum_{t=A_j+1}^{D_j} \alpha_t \text{ for } j = 1, \ldots, J. \end{aligned} \tag{15}$$

This is an integer maximization program with a concave objective function and linear constraints. To solve this integer program, consider a continuous extension of $R_j(n_j)$, defined as follows: for

$x_j \in \mathbb{R}$, let $\hat{R}_j(x_j) = R_j(\lfloor x_j \rfloor) + (x_j - \lfloor x_j \rfloor)\Big(R_j(\lfloor x_j \rfloor + 1) - R_j(\lfloor x_j \rfloor)\Big)$. Consider the linear relaxation:

$$
\begin{aligned}
\max_{\alpha_2,\ldots,\alpha_T \in [0,1]} \quad & \sum_{j=1}^{J} \theta_j \hat{R}_j(x_j) - c(x-1) \\
\text{s.t.} \quad & x = 1 + \sum_{t=2}^{T} \alpha_t \\
& x_j = 1 + \sum_{t=A_j+1}^{D_j} \alpha_t \text{ for } j = 1, \ldots, J.
\end{aligned}
\tag{16}
$$

As we show next, such linear relaxation yields an integer solution, making it quite easy to solve.

Proposition 2 *The extremes of the polytope* $P = \{\alpha_t | 0 \leq \alpha_t \leq 1, n_j \leq 1 + \sum_{t=A_j+1}^{D_j} \alpha_t \leq n_j + 1, j = 1, \cdots, J\}$ *are integral. As a result, the optimal solution to the linear relaxation in* (16) *is integer.*

Proof. The argument is as follows. In each region where x_j remains in the interval $[n_j, n_j + 1]$, the objective is linear, and as a result, within each of those regions, there is an extreme optimum. Moreover, As a result, we know that the optimal product timing decision can be found by solving (16). the extremes of the polytope $P = \{\alpha_t | 0 \leq \alpha_t \leq 1, n_j \leq 1 + \sum_{t=A_j+1}^{D_j} \alpha_t \leq n_j + 1\}$ turn out to be integral because the matrix $M \in \mathbb{R}^{(T-1)\times J}$ with entries $M_{ij} = 1_{A_j+1 \leq i \leq D_j}$ is a totally unimodular matrix. Indeed, consider S a square sub-matrix of M, and let us show that the determinant is $1, 0$ or -1. The rows of S are made of zeros, then ones, and then zeros again. Since permutating the rows of S only changes the sign of the determinant, we can assume without loss of generality that the rows are sorted so that $i < j$ implies that the lowest column in which a one is present is lower in i than in j, and, when it is equal, the highest column in which a one is present is lower in i than in j. One such matrix could for example be:

$$
\begin{pmatrix}
1 & 1 & 0 & 0 \\
1 & 1 & 1 & 0 \\
0 & 1 & 0 & 0 \\
0 & 1 & 1 & 1
\end{pmatrix}.
$$

We claim that the determinant of such matrix is always $1, 0$ or -1. We can prove it by induction. It is clearly true when the matrix has size 1×1. If it is true for sizes $k \leq n-1$, then let us prove that it is also true for matrices of size $n \times n$. Indeed, if $S_{11} = 0$, then it means that the first column is made of zeros and the determinant must be zero. Otherwise, $S_{11} = 1$, in which case we can subtract the first row from all the rows such that $S_{i1} = 1$, without changing the determinant. In this case the determinant of S is equal to the determinant of a matrix that only has a one in the first column ($S_{11} = 1$), which means that it is equal to the determinant of the $(n-1) \times (n-1)$ matrix made of rows and columns from 2 to n. An example of this operation is the following:

$$\begin{vmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 \end{vmatrix} = \begin{vmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 \end{vmatrix} = \begin{vmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 1 & 1 & 1 \end{vmatrix}$$

Since this smaller matrix has the same properties as the larger one (made of zeros, then ones, and then zeros again), we can use the induction hypothesis to know that its determinant is $1, 0$ or -1. This completes the proof. ■

The result in Proposition 2 allows us to efficiently calculate the retailer's optimal product rotation policy following a simple greedy algorithm. In each step, we move from a given vector $(\alpha_2, \ldots, \alpha_T)$ to the neighboring vector (where each coordinate is increased or decreased by one unit) that results in the highest profit. This is akin to using the simplex method on the hypercube $0 \leq \alpha_t \leq 1$.

The expected number of customers in the system affects the retailer's optimal product rotation decision. Moreover, the customers' arrival pattern affects the optimal timing of product rotation. Consider a scenario with $\beta = 1$, and $A_j = 1$ for all customer types (while their departure times are random). Suppose there is a solution $\{\alpha^*\}$ with $\alpha^*_\tau = 0$ and $\alpha^*_{\tau+l} = 1$ for some $l \geq 1$. Consider an alternate solution $\{\hat{\alpha}^*\}$ with $\hat{\alpha}^*_\tau = 1$ and $\hat{\alpha}^*_{\tau+l} = 0$ and $\hat{\alpha}^*_i = \alpha^*_i$ otherwise. The total number of product changes is the same as under $\{\alpha^*\}$, so the associated cost is the same. However, the number of product changes for all customer types with $D_j < \tau + l$ is one unit higher than under $\{\alpha^*\}$, while the number of product changes for the other customer types does not change. Thus, the revenue associated with the customer types with $D_j < \tau + l$ is higher under $\{\hat{\alpha}^*\}$ than under $\{\alpha^*\}$. Therefore, in such setting, it is best to concentrate all product introductions at the beginning of the season. Conversely, consider a setting in which customers arrive randomly throughout the selling season, but they all remain in the system until the end of the selling horizon (i.e., they have unlimited patience). In such setting, one can similarly prove that it is optimal to concentrate all product introductions towards the end of the selling season. In general, longer stays in the system (i.e., more patient customers) always lead to higher retailer profit.

Figure 8 plots the optimal timing of product changes in two settings. In the example shown in the graphs on the left, all customers arrive at the beginning of the selling season, but they remain in the system for a duration that is geometrically distributed with an expected length of stay (or window of patience) equal to $1/\mu = 4$ periods, either for strategic or myopic customers. The graphs on the right of Figure 8 correspond to a setting in which customers arrive uniformly over the selling season, following a stationary Poisson process, and their stay in the system is geometrically distributed with mean $1/\mu = 4$. In this case, it is best to spread product introductions throughout the season.

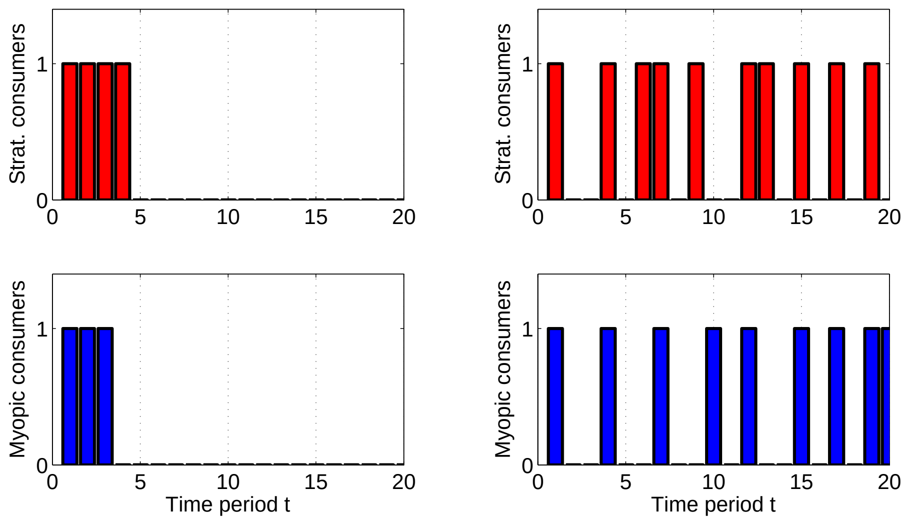

Figure 8: Optimal time of product rotation, with an expected length of stay $1/\mu = 4$. The top figures correspond to strategic customers while the bottom ones to myopic customers. In the graphs on the left, all customers arrive at $t = 1$, while on the right, customers arrive uniformly over $[1, T]$. Other parameters: $T = 20, \beta = 1, v_{out} = 0.1, k = 0.1, c = 0.002, V$ uniform in $[0, 1]$.

Increasing Revenue

So far, we have focused on stationary margins, i.e., $r_t = r$. This implies that revenue is proportional to the overall purchase probability of the customer.

In some circumstances, r_t may vary over time, even though product prices may be the same from the customer perspective. Indeed, the retailer may capture additional revenue in every visit that the customers make to the store. For instance, they may buy other items while shopping for the main product; or there may be a value directly generated from traffic, such as advertising revenue in online retailing. We focus here on the case where the retailer receives a revenue r when the customer purchases the main item, and also collects $\hat{r} \geq 0$ in every visit. We can define the expected number of visits for strategic customers, as follows:

$$VISITS_n^{strategic} = \sum_{i=1}^{n} i \left(\prod_{j=1}^{i-1} F(v_{t_j}) - \prod_{j=1}^{i} F(v_{t_j}) \right) + n \prod_{j=1}^{n} F(v_{t_j}),$$

and for myopic customers:

$$VISITS_n^{myopic} = \sum_{i=1}^{n} i \left(\prod_{j=1}^{i-1} F\Big(\beta^{T-t_j+1} v_{out}\Big) - \prod_{j=1}^{i} F\Big(\beta^{T-t_j+1} v_{out}\Big) \right) + n \prod_{j=1}^{n} F(\beta^{T-t_j+1} v_{out}).$$

Note that the last term in these sums corresponds to the fraction of customers that do not purchase the product, but nevertheless visits the store n times.

The revenue of the retailer is therefore $rSALES_n + \hat{r}VISITS_n$. A simple transformation shows that the revenue in period t_i is given by $r_{t_i} = r + \hat{r}i$. When all the products are introduced in the earlier periods, we thus have that $r_t = r + \hat{r}t$ increases over time. We focus on the case of $\beta = 1$ for tractability. In §4.4, we have shown that $SALES_n$ is concave in n. We show here that $VISITS_n$ is also concave.

Theorem 5 *Assume $\beta = 1$. $VISITS_n^{myopic}$ is concave in n. Moreover, when $\frac{\bar{F}(x+E(V-x)^+-k)}{\bar{F}(x)}$ is non-decreasing in x (e.g., when the distribution of V is uniform, exponential, or normal), then $VISITS_n^{strat}$ is also concave in n.*

Proof. Let us first consider the myopic case. Given n products being introduced in $t = 1, \ldots, n$, the purchase probability at t is equal to $(1-\gamma)\gamma^{t-1}$, where $\gamma = F(v_{out})$. As a result,

$$\begin{aligned} VISITS_n^{myop} = & \sum_{t=1}^{n} t(1-\gamma)\gamma^{t-1} + n\gamma^n = \sum_{t=0}^{n-1}(t+1)\gamma^t - \sum_{t=0}^{n} t\gamma^t + n\gamma^n \\ = & \sum_{t=1}^{n} \gamma^{t-1} = \frac{1-\gamma^n}{1-\gamma} \end{aligned}$$

This is clearly concave increasing in n.

For strategic customers, we have

$$\begin{aligned} VISITS_n^{strat} = & \sum_{t=1}^{n} t \left(\prod_{j=1}^{t-1} F(w_j^n) - \prod_{j=1}^{t} F(w_j^n) \right) + n \prod_{j=1}^{n} F(w_j^n) \\ = & \sum_{t=0}^{n-1} (t+1) \left(\prod_{j=1}^{t} F(w_j^n) \right) - \sum_{t=0}^{n} t \left(\prod_{j=1}^{t} F(w_j^n) \right) + n \prod_{j=1}^{n} F(w_j^n) \\ = & \sum_{t=0}^{n-1} \left(\prod_{j=1}^{t} F(w_j^n) \right) \end{aligned}$$

When adding another product, we have $w_{j+1}^{n+1} = w_j^n$, for all $j = 1, \ldots, n$ (keep in mind that we start the recursion with $w_{n+1}^n = v_{out}$). Thus, by adding the $n+1$-th product, there is one visit and, if the product was not purchased in the first visit, the additional number of visits is equal to $VISITS_n^{strat}$:

$$VISITS_{n+1}^{strat} = 1 + F(w_1^{n+1})VISITS_n^{strat}$$

and thus $VISITS_{n+1}^{strat} - VISITS_n^{strat} = 1 - \bar{F}(w_1^{n+1})VISITS_n^{strat}$. We claim that this expression is positive and non-increasing. We show this by induction that $\bar{F}(w_1^{n+1})VISITS_n^{strat} \leq 1$ and decreasing in n. To establish the result, we use that $w_1^{n+1} = w_1^n + E(V - w_1^n)^+ - k$, from Theorem 1. We initiate the recursion for $n = 1$: $VISITS_1^{strat} = 1$ while $\bar{F}(w_1^2) < 1$, so $\bar{F}(w_1^2)VISITS_1^{strat} \leq 1$.

Assume the induction property is true for $n - 1 \geq 1$. For n,

$$\begin{aligned}
\bar{F}(w_1^{n+1})VISITS_n^{strat} = & \ \bar{F}(w_1^{n+1})(1 + F(w_1^n)VISITS_{n-1}^{strat}) \\
= & \ \bar{F}(w_1^{n+1})\left(1 + \frac{F(w_1^n)}{\bar{F}(w_1^n)}\bar{F}(w_1^n)VISITS_{n-1}^{strat}\right) \\
\leq & \ \bar{F}(w_1^{n+1})\left(1 + \frac{F(w_1^n)}{\bar{F}(w_1^n)}\right) \text{ using the induction property} \\
\leq & \ \frac{\bar{F}(w_1^{n+1})}{\bar{F}(w_1^n)} \\
\leq & \ 1 \text{ because } w_1^{n+1} \geq w_1^n.
\end{aligned}$$

We now need to show that $\bar{F}(w_1^{n+1})VISITS_n^{strat} \geq \bar{F}(w_1^n)VISITS_{n-1}^{strat}$ or equivalently

$$\frac{\bar{F}(w_1^{n+1})}{\bar{F}(w_1^n)} \geq \frac{VISITS_{n-1}^{strat}}{VISITS_n^{strat}} = \frac{1 + \frac{F(w_1^{n-1})}{\bar{F}(w_1^{n-1})}\bar{F}(w_1^{n-1})VISITS_{n-2}^{strat}}{1 + \frac{F(w_1^n)}{\bar{F}(w_1^n)}\bar{F}(w_1^n)VISITS_{n-1}^{strat}}$$

Using the induction property, the right-hand side is no greater than

$$\frac{1 + \frac{F(w_1^{n-1})}{\bar{F}(w_1^{n-1})}\bar{F}(w_1^n)VISITS_{n-1}^{strat}}{1 + \frac{F(w_1^n)}{\bar{F}(w_1^n)}\bar{F}(w_1^n)VISITS_{n-1}^{strat}} \leq \frac{1 + \frac{F(w_1^{n-1})}{\bar{F}(w_1^{n-1})}}{1 + \frac{F(w_1^n)}{\bar{F}(w_1^n)}} = \frac{\bar{F}(w_1^n)}{\bar{F}(w_1^{n-1})}.$$

Thus, it is sufficient to show that $\frac{\bar{F}(w_1^n)}{\bar{F}(w_1^{n-1})} \leq \frac{\bar{F}(w_1^{n+1})}{\bar{F}(w_1^n)}$, which is true if $\frac{\bar{F}(x+E(V-x)^+-k)}{\bar{F}(x)}$ is increasing in $x \leq \bar{v}$, because $\bar{v} \geq w_1^n \geq w_1^{n-1}$. This completes the induction. Hence, $VISITS_n^{strat}$ is concave. ∎

Thus, finding the optimal number of product refreshments is easy, since we are maximizing a concave function $rSALES_n + \hat{r}VISITS_n$. Figure 9 displays the optimal solution for a setting with $T = 20$. When traffic revenue $\hat{r}$ is high (vertical axis), the optimal solution is to set $n^{strat} = 10$, at which point the additional traffic revenue does not recover the cost c. Interestingly, the optimal solution is more sensitive to $\hat{r}$ than to r, which suggests that the impact of n on the number of visits may be more significant than on expected sales. However, in any practical setting, $\hat{r}$ is an order of magnitude lower than r and therefore the impact of n on the probability of purchase prevails.

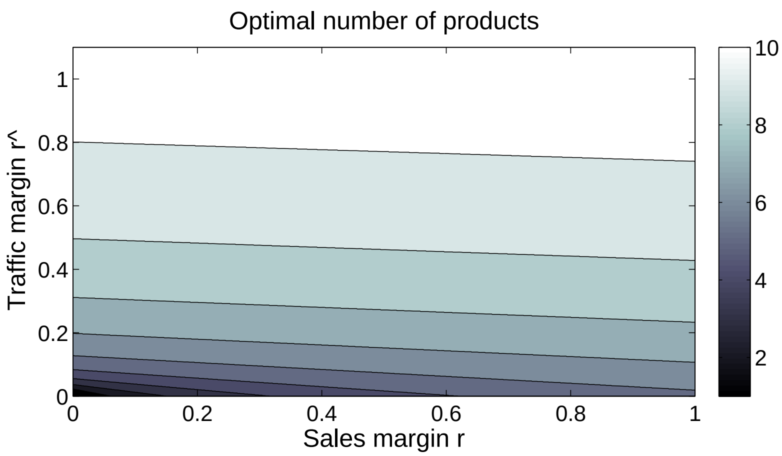

Figure 9: Optimal number of refreshed products for strategic customers, as a function of sales margin r and traffic margin $\hat{r}$. $T = 20, \beta = 1, v_{out} = 0.3, k = 0.1, c = 0.01, V$ uniform in $[0, 1]$.

Another observation from Figure 9 is that the regions where n^{strat} is constant seem to have linear boundaries. This is indeed the case for myopic customers where $rSALES_n^{myopic} + \hat{r}VISITS_n^{myopic} = [r + \hat{r}/(1 - F(v_{out}))][1 - F(v_{out})^n]$. Similarly, for strategic customers, we know that for large n, $SALES_n^{strat}$ and $VISITS_n^{strat}$ are approximately geometric, with parameter $F(\bar{v})$. As a result, $rSALES_n^{strat} + \hat{r}VISITS_n^{strat}$ is approximately equal to $[r + \hat{r}/(1 - F(\bar{v}))][1 - F(\bar{v})^n]$. This provides a rationale for the linear shape of the boundaries when customers are strategic.

Partial Assortment Changes

In this section, we explore the possibility of making partial assortment updates. Through a relatively simple setting consisting of two products, we illustrate the complexity of the analysis in the case of partial assortment updates and derive a few insights. Specifically, we find that:

- A strategic customer's purchasing decision in a given period is based on a threshold that depends on the number of products that will be updated in the following period. As a result, the recursion necessary to compute the thresholds depends on the entire rotation policy of the retailer, so it is computationally intractable. That is, there is no simple recursion for the thresholds as that shown in (11).

- The visit decision also depends on the retailer's rotation policy and it is therefore complex, in contrast to the result in Theorem 2 for the case of a single product. Nevertheless, we find that the customer does not visit the store if no product has been updated from the previous period. In addition, the customer always visits the store if all products are new, regardless of the retailer's future rotation policy. The complexity of the customer's visit decision problem arises in scenarios in which only a subset of the products is new. Indeed, we find that a customer may visit the store even if only a subset of the products has been updated – this appears to be the case if the search cost is not too high or if the observed valuation of the products that will remain on display is sufficiently high.

Setting. We consider a setting with two products. The retailer can rotate both products, just one, or none, from period to period. In particular, we consider a scenario with two products and focus on the last two periods of the selling season, i.e., periods $T-1$ and T. Moreover, to facilitate the derivation of results, we assume that $\beta = 1$ and that a customer's valuation for either product follows a uniform distribution in the interval $[0, 1]$. As in the case of a single product, we further assume that $E[V - v_{out}]^+ > k$, which reduces to $(1 - v_{out})^2 > 2k$ in this setting.

Purchase decision at T. We begin by considering the last period T and consider four possible states of the system: both products are new (which we will denote as nn), only the first product is new (which we will denote as no), only the second product is new (which we will denote as on), and neither product is new (which we will denote as oo). The profit-to-go function in period T depends on the state of the system, as follows:

$$\begin{aligned}
U_T^{nn} &= \max\{E_{V_1,V_2} J_T(V_1, V_2) - k, v_{out}\} \\
U_T^{no}(v_2) &= \max\{E_{V_1} J_T(V_1, v_2) - k, v_{out}\} \\
U_T^{on}(v_1) &= \max\{E_{V_2} J_T(v_1, V_2) - k, v_{out}\} \\
U_T^{oo}(v_1, v_2) &= \max\{\max\{v_1, v_2\} - k, v_{out}\}
\end{aligned}$$

where

$$J_T(v_1, v_2) = \max\{v_1, v_2, v_{out}\}.$$

If a visit to the store takes place, then the customer purchases the product if the realized valuations $\max\{v_1, v_2\} \geq v_{out}$, i.e., v_{out} is the purchase threshold in the last period.

Visit decision at T. We now examine the visit decision in period T under the state nn. Note that

$$E_{V_1,V_2} J_T(V_1, V_2) = \frac{2 + v_{out}^3}{3}.$$

Therefore, a visit takes place if $\frac{2+v_{out}^3}{3} - k \geq v_{out}$, which is true because $(1-v_{out})^2 > 2k$. Consider now the state *no*. In this case,

$$E_{V_1} J_T(V_1, v_2) = \frac{1 + (\max\{v_2, v_{out}\})^2}{2}.$$

One can verify that $\frac{1+(\max\{v_2,v_{out}\})^2}{2} - k \geq v_{out}$ because $(1-v_{out})^2 > 2k$, which implies that the customer should visit the store in state *no*. A similar analysis follows for state *on*. We conclude that, as long as at least one of the products is new in period T, then it is optimal for the customer to visit the store and a purchase takes place if the maximum realized valuation of the two products on display is no smaller than v_{out}. As for the case of a single product, if neither product is new, then the customer will not visit the store because a potential purchase would have occurred upon the earlier visit to the store when the items were first observed.

We next proceed to period $T-1$. We first define

$$U_T(v_1, v_2) \quad = \quad \alpha_T^1 \alpha_T^2 U_T^{nn} + \alpha_T^1 (1-\alpha_T^2) U_T^{no}(v_2) + (1-\alpha_T^1)\alpha_T^2 U_T^{on}(v_1) + (1-\alpha_T^1)(1-\alpha_T^2) U_T^{oo}(v_1, v_2) \tag{17}$$

and $J_T(v_1, v_2) = \max\{v_1, v_2, U_T(v_1, v_2)\}$, where $\alpha_T^j = 1$ if product j is renewed in period T and $\alpha_T^j = 0$ otherwise.

Purchase decision at $T-1$. We now consider a customer's purchasing decision, given that she has decided to visit the store in period $T-1$. (Note again that, if neither product is renewed in period $T-1$, then it is not optimal for the customer to visit the store in that period. We therefore ignore the scenario *oo*.) If a customer visits the store in period $T-1$, then we need to compare the maximum realized valuation $\max\{v_1, v_2\}$ to the expected utility derived from waiting until period T, which is given by $U_T(v_1, v_2)$. This comparison depends on the rotation policy of the retailer. We therefore consider each case separately.

(i) $\alpha_T^1 = 1; \alpha_T^2 = 1$. In this case, $U_T(v_1, v_2) = \frac{2+v_{out}^3}{3} - k$ so the customer buys if

$$\max\{v_1, v_2\} \geq \frac{2 + v_{out}^3}{3} - k.$$

(ii) $\alpha_T^1 = 1; \alpha_T^2 = 0$. In this case, $U_T(v_1, v_2) = \frac{1+(\max\{v_2,v_{out}\})^2}{2} - k$ so the customer buys if

$$\max\{v_1, v_2\} \geq \frac{1 + (\max\{v_2, v_{out}\})^2}{2} - k.$$

In other words, letting $v_2^{1,0}$ be the unique solution to $v_2 = \frac{1+(\max\{v_2,v_{out}\})^2}{2} - k$, a purchase takes place if $v_2 \geq v_2^{1,0}$ or $v_1 \geq \frac{1+(\max\{v_2,v_{out}\})^2}{2} - k$.

(iii) $\alpha_T^1 = 0; \alpha_T^2 = 1$. In this case, $U_T(v_1, v_2) = \frac{1+(\max\{v_1, v_{out}\})^2}{2} - k$ so the customer buys if

$$\max\{v_1, v_2\} \geq \frac{1 + (\max\{v_1, v_{out}\})^2}{2} - k.$$

(ii) $\alpha_T^1 = 0; \alpha_T^2 = 0$. No product is updated, and the customer buys if $\max\{v_1, v_2\} \geq v_{out}$.

The right-hand side values in the above four inequalities determine the purchase thresholds in each scenario.

Visit decision at $T - 1$. We next examine the decision to visit the store in period $T - 1$ and start with state nn. In this case, the customer decides between visiting the store and discovering two new products or waiting until period T in which again the two products will be new. That is, the profit-to-go in that case is $U_{T-1}^{nn} = \max\{E_{V_1,V_2} J_{T-1}(V_1, V_2) - k, U_T^{nn}\}$, where $U_T^{nn} = \frac{2+v_{out}^3}{3} - k$. We next compute $E_{V_1,V_2} J_{T-1}(V_1, V_2)$, which depends on the retailer's rotation policy for period T. Suppose first that $\alpha_T^1 = 1, \alpha_T^2 = 1$. In that case, letting

$$E_{V_1,V_2} J_{T-1}(V_1, V_2) = \frac{1}{3}\left(\frac{2 + v_{out}^3}{3} - k\right)^3 + \frac{2}{3}.$$

As a result, the customer visits the store in this case if and only if

$$\frac{1}{3}\left(\frac{2 + v_{out}^3}{3} - k\right)^3 + \frac{2}{3} \geq \frac{2 + v_{out}^3}{3} \iff \frac{2 + v_{out}^3}{3} - k \geq v_{out},$$

which holds because $(1 - v_{out})^2 > 2k$. That is, if $\alpha_T^1 = 1, \alpha_T^2 = 1$, then the customer visits the store in period $T - 1$. Suppose now that $\alpha_T^1 = 1; \alpha_T^2 = 0$. In that case,

$$\begin{aligned} E_{V_1,V_2} J_{T-1}(V_1, V_2) &= E_{V_1,V_2} \max\left\{V_1, V_2, \frac{1 + (\max\{V_2, v_{out}\})^2}{2} - k\right\} \\ &\geq E_{V_1,V_2} \max\left\{V_1, V_2, \frac{1 + v_{out}^2}{2} - k\right\} \geq E_{V_1,V_2} \max\{V_1, V_2, v_{out}\} \end{aligned}$$

Therefore, a visit always takes place. A similar analysis applies to the case $\alpha_T^1 = 0; \alpha_T^2 = 1$. That is, if a single product is renewed in period T, then the customer visits the store in period $T - 1$. Finally, if $\alpha_T^1 = 0, \alpha_T^2 = 0$, then one can easily verify that the customer is indifferent between visiting the store in periods $T - 1$ or T. In sum, if the state is nn, a customer visits the store in period $T - 1$ as long as at least one product will be renewed in period T.

We next consider state no. In this case, the customer's decision depends on the outcome that maximizes the profit-to-go in period $T - 1$ given by

$$U_{T-1}^{no}(v_2) = \max\left\{E_{V_1} J_T(V_1, v_2) - k, \alpha_T^2 U_T^{nn} + (1 - \alpha_T^2) U_T^{no}(v_2)\right\}.$$

Consider first the retailer's rotation policy given by $\alpha_T^1 = 1, \alpha_T^2 = 1$. In that case,

$$E_{V_1} J_T(V_1, v_2) = E_{V_1} \max\left\{V_1, v_2, \frac{2+v_{out}^3}{3} - k\right\} = \frac{1 + \left(\max\left\{v_2, \frac{2+v_{out}^3}{3} - k\right\}\right)^2}{2}.$$

Therefore, the customer visits the store if and only if

$$\frac{1 + \left(\max\left\{v_2, \frac{2+v_{out}^3}{3} - k\right\}\right)^2}{2} \geq \frac{2+v_{out}^3}{3}.$$

This holds if (i) $k \leq \frac{2+v_{out}^3}{3} - \sqrt{\frac{1+2v_{out}^3}{3}}$, i.e, if the cost of visit k is low enough, regardless of the value of v_2, or (ii) if $v_2 \geq \sqrt{\frac{1+2v_{out}^3}{3}}$, i.e., if the valuation of the incumbent product v_2 is high enough. Consider now the retailer's rotation policy given by $\alpha_T^1 = 1, \alpha_T^2 = 0$. In that case,

$$E_{V_1} J_T(V_1, v_2) = E_{V_1} \max\left\{V_1, v_2, \frac{1 + \max\{v_2, v_{out}\}^2}{2} - k\right\} = \frac{1 + \left(\max\left\{v_2, \frac{1+\max\{v_2, v_{out}\}^2}{2} - k\right\}\right)^2}{2}.$$

Therefore, the customer visits the store if and only if

$$\frac{1 + \left(\max\left\{v_2, \frac{1+\max\{v_2, v_{out}\}^2}{2} - k\right\}\right)^2}{2} \geq \frac{1 + \max\{v_2, v_{out}\}^2}{2},$$

which always holds. Finally, consider the retailer's rotation policy given by $\alpha_T^1 = 0, \alpha_T^2 = 1$. In that case,

$$E_{V_1} J_T(V_1, v_2) = E_{V_1} \max\left\{V_1, v_2, \frac{1 + (\max\{V_1, v_{out}\})^2}{2} - k\right\}.$$

Therefore, the customer visits the store if v_2 is high enough or k is low enough. Finally, it is easy to verify that the customer is indifferent between visiting the store in periods $T-1$ or T when $\alpha_T^1 = 0, \alpha_T^2 = 0$.

A similar analysis applies when the state is *on*.

The retailer's problem. We finally explore the retailer's optimal rotation policy in period T. Suppose that the initial state in period $T-1$ is nn, in which case the customer always visits the store. Further, suppose that the assortment rotation cost is additive, i.e., the cost is $c(\alpha_T^1 + \alpha_T^2)$. If both products are renewed in period T, then the retailer's profit is

$$\pi_R = \left(1 - \left(\frac{2+v_{out}^3}{3} - k\right)^2\right) + \left(\frac{2+v_{out}^3}{3} - k\right)^2 (1 - v_{out}^2) - 2c,$$

while if only product j is renewed in period T, the retailer's profit is

$$\pi_R = \int_0^1 \left[\left(1 - \left(\frac{1 + (\max\{v_j, v_{out}\})^2}{2} - k \right)^2 \right) + \left(\frac{1 + (\max\{v_j, v_{out}\})^2}{2} - k \right)^2 (1 - v_{out}^2) - c \right] dv_j.$$

If no product is renewed in period T, then the retailer's profit is $\pi_R = 1 - v_{out}^2$.

We conclude that, depending on the value of c, it may be optimal to rotate one product, two, or none. A similar analysis applies if the state is either *no* or *on*, although in those cases, the rotation cost and other parameters would also determine whether or not a visit takes place in period $T - 1$.