Anticipatory Shipping and Price Dynamics

Konstantinos Serfes* and Veikko Thiele†

VERY INCOMPLETE AND PRELIMINARY – PLEASE DO NOT CIRCULATE

May 2019

Abstract

We analyze optimal price dynamics when retailers ship products even before customers order these products (anticipatory selling). More specifically, we develop a model where a retailer can ship a product sequentially to potential customers until a customer accepts the product (or the product is returned to the warehouse). We show that it can be optimal for the retailer to initially offer products at high prices, and then to reduce prices over time to increase the probability of acceptance. Moreover, with strategic customers we find that offering price discounts upon product rejection is more likely to be optimal for products that appeal to many potential customers, in areas with low transportation costs (e.g. urban areas). We then analyze the effects of customer-specific and product-specific learning through sales on optimal price dynamics.

Keywords: Price dynamics, price discounts, strategic customers, artificial intelligence.

JEL classification: D4.

*School of Economics, LeBow College of Business, Drexel University, E-mail: ks346@drexel.edu.

†Queen's University, Smith School of Business, Goodes Hall, 143 Union Street, Kingston, Ontario, Canada K7L 3N6, e-mail: thiele@queensu.ca.

1 Introduction

Online shopping continues to replace traditional brick and mortar shopping. While some retailers use online shopping as an additional channel for generating sales (e.g. Home Depot, Walmart, and Bass Pro Shop), some retailers only sell their products online (e.g. Amazon (non-food items), iTunes, and Groupon). Nonetheless all retailers still use what we call a *reactive* selling strategy: customers first select the products (and pay for them) before they can take the products home or have them shipped to their address.

The recent rise of artificial intelligence (AI) enables retailers to better predict the preferences of individual consumers, and therefore the products they may want to buy in the near future. This allows retailers to use a proactive selling strategy, called *anticipatory shipping*: shipping some products to customers that they may want to purchase in the near future, i.e., when the products are not yet ordered. For example, suppose a family is using a specific brand of a laundry detergent, and ordered one bottle at the beginning of every month from an online retailer for the past two years. The retailer can now predict (with almost certainty) that the family needs a bottle of the laundry detergent at the beginning of the next month. And instead of waiting for the actual order, the online retailer could simply proactively ship a bottle to the family's doorstep. The only thing left to the family is to decide whether to keep the bottle of detergent (the amount is then charged on the family's credit card), or to reject the product, in which case it is picked up by the retailer (and perhaps send to the next potential customer).

Interestingly, Amazon filed a patent for such anticipatory shipping already in 2012. According to Amazon's patent #US 8,615,473 B2 (see Figure 7 in this patent, granted on Dec. 24, 2013), the procedure includes the following four main steps:[1]

(i) Ship package to geographic area without specifying delivery address.

(ii) Identify potential customer nearby, and send package to this customer.

(iii) If package is not accepted, compare cost of return vs. redirecting package to another potential customer nearby.

(iv) Offer package to the next customer, potentially at a discounted price.

We are convinced that *anticipatory shipping* will be the next major disruption in the retail industry (after online shopping), for the following two reasons: First, artificial intelligence

[1]The patent can be viewed and downloaded from https://patents.google.com/patent/US8615473B2/en.

becomes exponentially more powerful, therefore enabling retailers to predict preferences of individual consumers much more accurately. Second, with the latest technological developments, such as self-driving cars and autonomous drones, we expect shipping costs to drop significantly over the next few years. In fact, Amazon Air did its first successful delivery of a package with a fully autonomous drone on Dec. 7, 2016 in Cambridgeshire, England.

In this paper we analyze optimal price dynamics when a retailer uses anticipatory shipping to sell its products. Specifically we ask whether and when the retailer engages in price discrimination based on the sequence of potential customers (e.g. offering the same product to the first customers at a higher price, and then reducing the price for the last potential customers). Moreover, we identify conditions so that customers, who initially reject a product, get the same product reoffered at a discounted price, so the retailer can avoid additional shipping costs. We also ask how customer-specific and product-specific learning (e.g. through improved artificial intelligence) affect pricing strategies under anticipatory selling.

To answer these research question we develop a dynamic model with a retailer and several potential customers for a product in a given geographic region. Customers have private information about their own reservation prices for the product. The retailer then uses anticipatory shipping to sell a product: the product is shipped to the first customer (without being ordered), and if this customer rejects the product, the product is potentially being send to the next customer, and so on. Moreover, we allow the retailer's pricing strategy to be two-dimensional. First, the retailer can choose the price at which the product is initially being offered to a given customer (and this price can change when shipping the product to the next customers). Second, the retailer can decide whether to reoffer a selected customer the same product at a discount price, should this customer (and obviously all previous customers) reject the product – we call the *discount policy*. Alternatively, the retailer can choose to never offer such discounts upon product rejection, and instead, to always ship the product to the next customer, or back to the warehouse. We call this the *no-discount policy*. Finally, an important feature of our model ist that customers are strategic (or sophisticated): when deciding whether to accept or reject the product, they take into account the chance of getting the product reoffered at a discounted price (if the retailer adopted the discount policy).

We first consider the benchmark with non-strategic (myopic) customers. We then show that it is always optimal for the retailer to initially offer the product at a high price, and then offering a price discount to a selected customer, should he (and all previous customers) reject the product. Put differently, adopting the discount policy is always optimal when customer are non-strategic, as this avoids additional shipping costs when a customer refuses to accept the product. However, we show that this a no longer true with strategic customers. In fact, we find that adopting the discount policy remains only optimal for sufficiently low shipping costs, and for a sufficiently large customer base for a product; otherwise it is optimal to never offer any ex post price discount upon product rejection.

The intuition is as follows. The discount and the no-discount policies each have a distinct disadvantage: When adopting the discount policy in the presence of strategic customers, the retailer needs to reduce the initial price to prevent strategic product rejections. And the no-discount policy leads to higher expected shipping costs. With a larger number of potential customers, each customer knows that getting the price discount after rejecting the product, becomes less likely (discount policy). This in turn allows the retailer to offer the product initially at a higher price. Additionally we show that including all potential customers in the shipping sequence, and therefore realizing a higher retail price, is only optimal for the retailer when the shipping cost is sufficiently low. In other words, our model predicts that lower retail prices and ex post price discounts are more likely to be observed for products that appeal to many potential customers (e.g. laundry detergent), in areas with low transportation costs (e.g. urban areas).

This remainder of this paper is structured as follows. In the next section we discuss the contribution of this paper to the literature. Section 3 introduces our main model. In Section 4 we consider the benchmark with non-strategic customers, and identify the optimal pricing strategy for a retailer when using anticipatory shipping. We then analyze our main model with strategic customers in Section 5. In Section 6 we examine the effects of customer-specific and product-specific learning on the optimal dynamic pricing strategy. Section 7 discusses the main empirical predictions from our model. It is followed by a brief conclusion that summarizes the main results and explores avenues for future theory work. All proofs are in the Appendix.

2 Related Literature

To be written...

3 The Base Model

Consider a firm, called *Firm A*, trying to proactively selling a product in a given geographic region. Specifically, Firm A considers shipping the product to potential customers before they even order the product (*anticipatory shipping*). The cost of the product is normalized the zero.[2]

[2]Allowing for a positive cost adds more complexity to the model, but does not affect the nature of our results.

There are m potential customers in the geographic region that would consider buying the product. A potential customer has either a high reservation price r_H for the product, or a low reservation price r_L, with $r_H > r_L > 0$. Firm A cannot observe reservation prices, but knows that a given customer has a high reservation price r_L with probability $\lambda \in (0, 1)$.

A customer, who received the product, can either pay the price P and keep the product, or reject the product. Firm A has two options to deal with a product rejection: (i) reoffering the product to the same customer at a discounted price, or (ii) shipping the product to the next customer. Firm A's dynamic pricing strategy therefore entails two dimensions: First, choosing the price P when initially offering the product to a customer – and this price may change over time. Second, deciding whether to offer some selected customers a price discount after product rejection (discount policy), or to commit to never reoffering the product at a discounted price (no-discount policy). We assume that such commitment is possible.

The cost of shipping the product to the next customer is $c > 0$. The cost of shipping the product from and to the warehouse is $(1+\theta)c$, with $\theta \geq 1$. We assume that a product, eventually shipped back to the warehouse, has no resale value (or the cost of refurbishing the product would exceed its value). For parsimony we assume zero discounting.

As a benchmark we first consider the case with non-strategic (myopic) customers. A customer then always keeps the product as long as the price P is weakly lower than his reservation price, i.e., when $P \leq r_i$, $i = L, H$. For our main model we allow for strategic (sophisticated) customers. These customers then also take into account that they may get a price discount when initially rejecting the product. In other words, these customers may strategically reject the product, even if the originally offered price P is (weakly) below their reservation prices.

4 Benchmark – Non-Strategic Customers

We first consider the benchmark with non-strategic customers. This allows us to identify the fundamental drivers for Firm A's optimal dynamic pricing strategy and discount policy.

Suppose Firm A considers offering some customer $j \in [1, m]$ the price discount ΔP_j after the product is rejected (discount policy). This would avoid shipping the product to the next customer (at cost c), or back to the warehouse (at cost $(1 + \theta)c$). It is then always optimal for Firm A to initially offer the product at the price $P = r_H$, which a customer with the high reservation price r_H would accept. And only after rejection potentially lowering the price to $P = r_L$, so that even a customer with the low reservation price r_L would accept the product. The optimal price discount for some customer $j \in [1, m]$ is therefore given by $\Delta P_j = r_H - r_L$.

Firm A's cost of shipping the product to the first customer is $(1+\theta)c$. With probability λ, customer 1 has the high reservation price r_H, and therefore buys the product at the price $P = r_H$. With probability $1-\lambda$, customer 1 has the low reservation price r_L and rejects the product. Firm A can then either ship the product to the next customer, customer 2, or reoffer the product to customer 1 at the discounted price $P = r_L$ (which customer 1 would accept). In case of shipping to the next customer, the new customer 2 accepts the product with $P = r_H$ with probability λ, and rejects with probability $1-\lambda$, and so on. Note that it is always optimal for Firm A to offer the price discount $\Delta P = r_H - r_L$ at least to the last customer m (in case all previous $m-1$ customers rejected the product with $P = r_H$), as shipping the product back to the warehouse imposes the cost $(1+\theta)c$, and leads to a product without resale value. Thus, when offering the price discount ΔP_j to customer j, $j \in [1, m]$, Firm A's expected profit is given by

$$\pi^d(\Delta P_j) = -(1+\theta)c + \lambda r_H + \sum_{i=1}^{j-1}(1-\lambda)^i\left[-c + \lambda r_H\right] + (1-\lambda)^j r_L.$$

We derive the following lemma in the Appendix:

Lemma 1 *Consider the discount policy with non-strategic customers. It is optimal to offer only the last customer m the price discount $\Delta P_m = r_H - r_L$ (i.e., $j^* = m$) after rejection if $c < \widehat{c}^{\,d} \equiv \lambda\left[r_H - r_L\right]$. Otherwise it is optimal to offer customer 1 the price discount $\Delta P_1 = r_H - r_L$ (i.e., $j^* = 1$).*

The lemma essentially says that shipping the product to the next customer, and not offering a discount after rejection, is optimal as long as the expected price premium $\lambda\left[r_H - r_L\right]$ exceeds the incremental shipping cost c. The discount is then possibly only offered to the last customer m in order to avoid the cost of shipping the product back to the warehouse.

Alternatively, Firm A can commit to never offer a price discount to a customer after rejecting the product (no-discount policy). In this case the product is always either shipped to the next customer (at cost c), or back to the warehouse (at cost $(1+\theta)c$). When ruling out possible discounts, it is not necessarily optimal for Firm A to offer the product a the high price $P = r_H$. In fact, it can be optimal to offer a customer the low price $P = r_L$ to eliminate the risk of product rejection, and therefore to avoid additional shipping costs.

One possible pricing strategy is to offer the first customer the product at the low price $P_1 = r_L$, who would then always accept. Firm A's expected profit is then given by

$$\pi^n(P_1 = r_L) = -(1+\theta)c + r_L.$$

We call this the *low-price strategy*.

Similarly, Firm A can initially offer the product at the high price $P = r_H$, and then lower the price to $P_j = r_L$ for some customer j, $2 \leq j \leq m$ (should all previous customer have rejected the high price). We call this the *price-reduction strategy*, which leads to the following expected profit:

$$\pi^n(P_j = r_L) = -(1+\theta)c + \lambda r_H + \sum_{i=1}^{j-2}(1-\lambda)^i \left[-c + \lambda r_H\right] + (1-\lambda)^{j-1}\left[-c + r_L\right].$$

Another possible strategy for Firm A is to offer the product to customers only at the high price $P = r_H$ – we call this the *high-price strategy*. This strategy clearly entails the risk that all customers have a low valuation and reject the product, so that Firm A may incur the cost $(1+\theta)c$ of shipping the product back to the warehouse. Firm A's expected profit is then

$$\pi^n(P = r_H) = -(1+\theta)c + \lambda r_H + \sum_{i=1}^{m-1}(1-\lambda)^i \left[-c + \lambda r_H\right] - (1-\lambda)^m(1+\theta)c.$$

The next lemma compares the different pricing strategies when Firm A commits to never reoffer the product to the same customer at a discounted price after rejection (no-discount policy).

Lemma 2 *There exists two threshold shipping costs, $\widehat{c}^{\,n}$ and $\widehat{\widehat{c}}^{\,n}$, with*

$$\widehat{c}^{\,n} = \frac{\lambda r_H - r_L}{(1-\lambda)(1+\theta)} \qquad \widehat{\widehat{c}}^{\,n} = \frac{\lambda}{1-\lambda}\left[r_H - r_L\right]$$

and $\widehat{c}^{\,n} < \widehat{\widehat{c}}^{\,n}$, such that the optimal pricing strategy under the no-discount policy is as follows:

- *For $c < \widehat{c}^{\,n}$ offer all customers the high price $P = r_H$ (high-price strategy).*
- *For $\widehat{c}^{\,n} \leq c < \widehat{\widehat{c}}^{\,n}$ offer last potential customer m the low price $P_m = r_L$ (price-reduction strategy).*

- *For $c \geq \widehat{\widehat{c}}^{n}$ offer the first customer the low price $P_1 = r_L$ (low-price strategy).*

If the shipping cost is sufficiently low ($c < \widehat{c}^{n}$) – which includes the potential cost of shipping the product back to the warehouse – it is optimal for Firm A to pursue the high-price strategy. The chance of getting the high price $P = r_H$ from a customer (which occurs with probability λ) then outweighs the risk of incurring additional shipping costs when the product is rejected, even for the last customer m. For intermediate shipping costs ($\widehat{c}^{n} \leq c < \widehat{\widehat{c}}^{n}$) it is still optimal to offer each customer the product at the high price $P = r_H$, except for the last potential customer m. Reducing the price to $P_m = r_L$ for customer m eliminates the risk of having the product shipped back to the warehouse. Finally, if the cost of shipping is too high ($c \geq \widehat{\widehat{c}}^{n}$), it is optimal for Firm A to avoid additional shipping costs entirely by offering the low price $P_1 = r_L$ already to the first customer (who would then always accept).

We can now identify the optimal strategy for Firm A.

Proposition 1 *With non-strategic customers, it optimal to use the discount policy (as specified in Lemma 1) for all $c \geq 0$.*

It is optimal for Firm A to initially offer each customer the product for the high price $P = r_H$, and then offer the price discount $\Delta P = r_H - r_L$ to select customers should they choose to reject the product. Using the discount policy essentially allows Firm A to test the high price for each customer at no cost, as the product can always be reoffered at the discounted price (which would avoid additional shipping costs). And with non-strategic customers Firm A knows that every customer with a high valuation would accept and pay the price $P = r_H$. This also implies that Firm A does not only learn the types of its customer (i.e., the customers who received the product before it is eventually purchased), but also extract the entire surplus when selling the product to a customer. However, as we will show in the next section, the discount policy is not necessarily optimal when customers are strategic.

5 Pricing with Strategic Customers

We now turn to our main model and allow customers to be strategic (or sophisticated). With a discount policy, customers with a high reservation price may then initially reject a product in hope of getting it reoffered at a discounted price. Naturally the decision to strategically reject a product depends on the probability that a given customer is selected by Firm A to receive the price discount upon rejection.

Suppose Firm A initially offers the product at the price $P(j)$, and plans to offer customer j, $1 \leq j \leq m$, the price discount $\Delta P(j) = P(j) - r_L$, should he, and all previous $j-1$ customers, reject the project. A given customer does not know whether he is indeed customer j, and therefore selected by Firm A for a possible discount. However, if a customer receives the product then he knows that he must belong to the selected customer group $[1, 2, ..., j]$. The unconditional probability that he is customer j, is then $1/j$. Moreover, the probability that all previous $j-1$ customers had a valuation (r_L), and therefore rejected the product, is given by $(1-\lambda)^{j-1}$. Thus, the conditional probability that a given customer gets the price discount $\Delta P(j)$ upon rejection, is $\frac{1}{j}(1-\lambda)^{j-1}$.

A customer with a high reservation price accepts the product and pays the price $P(j) \in [r_L, r_H]$ if

$$r_H - P(j) \geq \frac{1}{j}(1-\lambda)^{j-1}\left[r_H - r_L\right].$$

Solving for $P(j)$ we get the maximum price that a high-valuation customer is willing to pay without first strategically rejecting the product:

$$P(j) = r_H - \frac{1}{j}(1-\lambda)^{j-1}\left[r_H - r_L\right].$$

In the Appendix we show that $P(j)$ is increasing and concave in j. This is very intuitive, as including a larger number of potential customers (higher j) reduces the probability that a given customer is indeed customer j, and therefore the likelihood that he receives the price discount upon rejection. This also implies that the surplus that a high-valuation customer can extract, is decreasing in the number of potential customers j. For the extreme case with $j \to \infty$ we have $\lim_{\to\infty} P(j) = r_H$, so that even high-valuation customers do not get any surplus. In contrast, when $j = 1$, and therefore $P(1) = r_L$, a high-valuation customer gets the maximum possible surplus $r_H - r_L$.

Firm A's expected profit, when offering the product at the price $P(j)$, and potentially offering the price discount $\Delta P(j)$ to customer j upon product rejection, is given by

$$\pi^d(\Delta P(j)) = -(1+\theta)c + \lambda P(j) + \sum_{i=1}^{j-1}(1-\lambda)^i\left[-c + \lambda P(j)\right] + (1-\lambda)^j r_L.$$

The next proposition characterizes the optimal discount policy with strategic customers.

Proposition 2 *Consider the discount policy with strategic customers. It is then optimal to offer the discount* $\Delta P(j^*)$ *to customer* j^* *upon rejection, with* $j^* = \text{argmax}_{k\in[1,m]}\{\widehat{c}^{\,d}(k) \mid \widehat{c}^{\,d}(k) \geq c\}$*, where*

$$\widehat{c}^{\,d}(k) = \sum_{i=1}^{k-1}(1-\lambda)^{i-k}\lambda\left(P(k)-P(k-1)\right)+\lambda\left(P(k)-r_L\right), \quad k \in [1,m].$$

The threshold cost $\widehat{c}^{\,d}(k)$ *is decreasing in* k*.*

By choosing j Firm A essentially decides how many customers it wants to potentially send the product to, as the last customer j will always accept (at least after being offered the price discount $\Delta P(j^*)$). And including more potential customers (by choosing a higher j) has two positive effects on Firm A's expected profit: First, it allows Firm A to initially offer its product at a higher price without the risk of customers strategically rejecting the product in order to get a price discount (as $P(j)$ is increasing in j). Second, including more potential customers increases the chance of finally selling the product to a high-valuation customer at the price $P(j)$. However, including more potential customers also increases the total expected shipping cost.

According to Proposition 2 there exists a threshold shipping cost $\widehat{c}^{\,d}(k)$ so that it is only optimal for Firm A to send the product to customer k, after the previous customer $k-1$ rejected, as long as $c \leq \widehat{c}^{\,d}(k)$. In fact, it turns out that it is optimal for Firm A to include the largest possible number of customers so that the actual shipping cost c is still (weakly) smaller than the threshold $\widehat{c}^{\,d}(k)$. However, the marginal benefit of including another customer gets relatively smaller (as $P(j)$ is concave increasing in j), which requires a lower shipping cost c. Formally this implies that the threshold shipping cost $\widehat{c}^{\,d}(k)$ is decreasing in k.

In other words, Proposition 2 implies that including all m potential customers ($j^* = m$), and only offering the last customer m the price discount $\Delta P(m)$, is optimal for Firma A, as long as its shipping cost c is sufficiently low ($c \leq \widehat{c}^{\,d}(m)$). Firm A then pursues a high-price strategy, with $P(m)$, $P(m) < r_H$, as the highest attainable price when using a discount policy with strategic customers. For some intermediate levels of c ($\widehat{c}^{\,d}(2) \leq c < \widehat{c}^{\,d}(m)$) it is optimal for Firm A to restrict the number of potential customers in order to curb its total expected shipping costs. We then have an interior solution for j ($1 < j^* < m$) – contrary to the case with non-strategic customers where it is optimal to offer the discount either to customer 1 or to customer m; see Lemma 1. This also forces Firm A to reduce the original price for its product down to $P(j^*)$, with $P(j^*) < P(m)$. Finally, if the shipping cost is sufficiently high ($c > \widehat{c}^{\,d}(2)$), it is

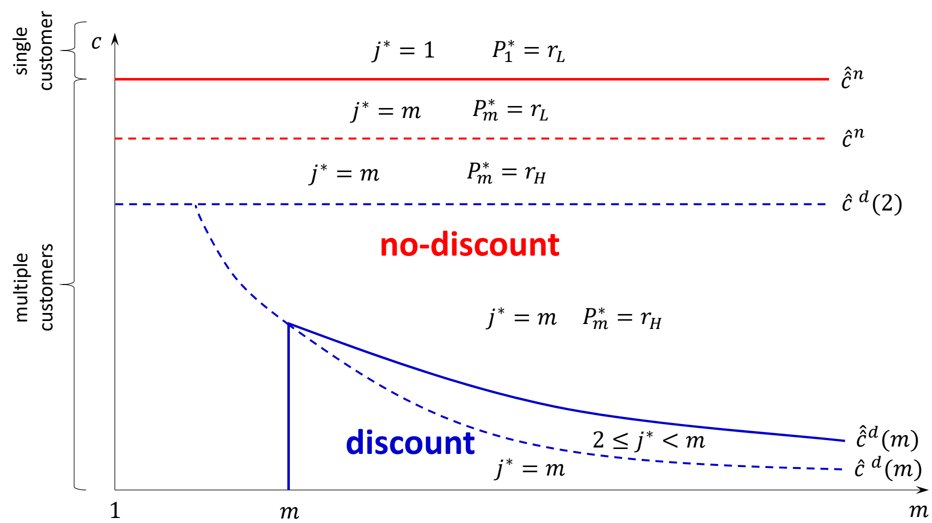

Figure 1: Optimal Pricing with Strategic Customers

optimal to only ship the product to one customer ($j^* = 1$). The product is then offered at the lowest price $P(1) = r_L$, which the only customer will always accept.

As far as the no-discount policy is concerned, it does not matter for the optimal pricing strategy whether customers are strategic or non-strategic (as considered in Section 4). A customer with a high reservation price would always accept and pay the price $P = r_H$, because he would never get the product reoffered at a lower price after rejection. The optimal pricing strategy, as summarized by Lemma 2 therefore also applies to the case with strategic customers.

The next proposition identifies the optimal strategy for Firm A when customers are strategic.

Proposition 3 *There exists a shipping cost threshold $\widehat{\widehat{c}}^{\,d}(m)$ so that a discount policy is only optimal for $c < \widehat{\widehat{c}}^{\,d}(m)$ and $m \geq \underline{m} > 1$. Otherwise using a no-discount policy is optimal.*

Figure 1 illustrates the key insights from Proposition 3, as well as from Lemma 2 and Proposition 2. We put the number of potential customers m on the horizontal axis, and the shipping cost c on the vertical axis. For convenience we treat m as a continuous parameter in Figure 1, even though by assumption m is discrete.

We know from Proposition 1 that using a discount policy is always optimal when customers are nonstrategic (or myopic). However, this is no longer true with strategic customers. In fact, a discount policy can then only be optimal as long as there is a sufficiently large number of potential customers ($m \geq \underline{m} > 1$), and the actual shipping cost is sufficiently low ($c < \widehat{\widehat{c}}^{\,d}(m)$).

At first glance one would expect that a discount policy is always optimal when shipping costs are *high*, as offering discounts eliminates the need for Firm A to incur additional (high) shipping costs when the product is rejected. The intuition, why we obtain the opposite result, is as follows. When using a no-discount policy, we know from Lemma 2 that it is optimal for Firm A to pursue a high-price strategy and to include all m potential customers, as long as $c < \widehat{c}^{\,n}$. Moreover, when adopting the discount policy, it is optimal to also include all m customers when $c \leq \widehat{c}^{\,d}(m)$; see Proposition 2 (reflected by the blue dashed line in Figure 1). However, each policy has its own drawback: The no-discount policy imposes additional shipping costs, while the discount policy requires a reduced initial price to prevent strategic product rejections.

We know that a larger potential customer base m allows Firm A to increase the price $P(m)$, which in turn makes the discount policy more profitable. And when the customer base m exceed the critical level $\underline{m}$, the cost of using the discount policy (lower price) becomes relatively less severe compared to the cost of using the no-discount policy (extra shipping costs). Overall this implies that for $m \geq \underline{m}$ and $c \leq \widehat{c}^{\,d}(m)$ the discount policy with $j^* = m$ is optimal.

For $\widehat{c}^{\,d}(m) \leq c < \widehat{\widehat{c}}^{\,d}(m)$ it is optimal for Firm A to include many, but not all, potential customers, i.e., $2 \leq j^* < m$; see Proposition 2. This also implies a lower initial price $P(j^*)$ for the product, with $P(j^*) < P(m)$, $j^* < m$. Nonetheless Firm A's lost profit because of the reduced price $P(j^*)$ under the discount policy, is still lower than the extra shipping cost c (or $(1+\theta)c$) associated with the no-discount policy. Adopting the discount policy with $2 \leq j^* < m$ in this region therefore remains optimal. In contrast, for $c < \widehat{\widehat{c}}^{\,d}(m)$ the optimal number of customers j^*, and therefore the attainable price $P(j^*)$, is too low under the no discount policy. Committing to not offer any ex post price discounts is then the optimal strategy for Firm A.

The main result so far is that with strategic customers, offering an ex post price discount can only be optimal if there is a sufficiently large customer base, and shipping costs are sufficiently low. However, Figure 1 provides some additional interesting insights. Sending a product sequentially to multiple customers is only optimal as long as the shipping cost is not too high ($c < \widehat{\widehat{c}}^{\,n}$). In this case Firm A either reduces the price over time to avoid having to send the product back to the warehouse ($\widehat{c}^{\,n} \leq c < \widehat{\widehat{c}}^{\,n}$), or maintains a high price ($P = r_H$) for all customers ($\widehat{\widehat{c}}^{\,d} < c < \widehat{c}^{\,n}$). Moreover, unless Firm A adopts the discount policy, it is always optimal to include all m potential customers as long as the shipping cost is not too high ($c < \widehat{\widehat{c}}^{\,n}$).

6 Learning

6.1 Customer-specific Learning

To be written...

6.2 Product-specific Learning

To be written...

7 Empirical Predictions

- Interpretation of key parameters:
  - Potential customers m:
    - niche product (low m) vs. mass product (high m)
    - durability of product
  - Shipping cost c:
    - urban areas (low c) vs. rural areas (high c)
    - technology improvements (drone delivery etc.)
  - Distribution of reservation prices (λ): low average income areas (low λ) vs. high average income areas (high λ)
- Key variables of interest:
  - Prices (dynamics; low/high)
  - Price discounts (early/late)
- Price discounts more likely for mass products (high m) and in urban areas (low c)
- No price discounts for niche products (low m) and in rural areas (high c)
- More efficient delivery methods (e.g. autonomous drones) make price discounts more likely

8 Conclusion

To be written...

Appendix

Proof of Lemma 1.

It is optimal to offer the discount to customer j, instead of to customer $j-1$, if $\pi^d(\Delta P_j) > \pi^d(\Delta P_{j-1})$, which is equivalent to

$$\begin{aligned} \sum_{i=1}^{j-1}(1-\lambda)^i[-c+\lambda r_H] + (1-\lambda)^j r_L &> \sum_{i=1}^{j-2}(1-\lambda)^i[-c+\lambda r_H] + (1-\lambda)^{j-1} r_L \\ \Leftrightarrow \quad (1-\lambda)^{j-1}[-c+\lambda r_H] + (1-\lambda)^j r_L &> (1-\lambda)^{j-1} r_L \\ \Leftrightarrow \quad c &< \lambda[r_H - r_L]. \end{aligned}$$

Note that this condition is independent of j. Thus, if $c < \lambda[r_H - r_L]$ we have $j^* = m$; otherwise $j^* = 1$. □

Proof of Lemma 2.

Consider first the *price-reduction strategy* with the expected profit $\pi^n(P_j = r_L)$. It is optimal for Firm A to offer customer j the low price $P = r_L$, instead of customer $j-1$, if $\pi^n(P_j = r_L) > \pi^n(P_{j-1} = r_L)$. This condition can be written as

$$\begin{aligned} \sum_{i=1}^{j-2}(1-\lambda)^i[-c+\lambda r_H] + (1-\lambda)^{j-1}[-c+r_L] &> \sum_{i=1}^{j-3}(1-\lambda)^i[-c+\lambda r_H] \\ &\quad +(1-\lambda)^{j-2}[-c+r_L] \\ \Leftrightarrow \quad (1-\lambda)^{j-2}[-c+\lambda r_H] + (1-\lambda)^{j-1}[-c+r_L] &> (1-\lambda)^{j-2}[-c+r_L] \\ \Leftrightarrow \quad c &< \widehat{\widehat{c}}^{\,n} \equiv \frac{\lambda}{1-\lambda}[r_H - r_L]. \end{aligned}$$

Notice that this condition does not depend on j. Consequently, for $c < \widehat{\widehat{c}}^{\,n}$ we have $j^* = m$, and for $c \geq \widehat{\widehat{c}}^{\,n}$ we have $j^* = 2$.

Next, the *low-price strategy* is weakly more profitable than the *price-reduction strategy* with $j^* = 2$ if $\pi^n(P_1 = r_L) \geq \pi^n(P_2 = r_L)$, i.e., if

$$r_L \geq \lambda r_H + (1-\lambda)[-c+r_L] \quad \Leftrightarrow \quad c \geq \widehat{\widehat{c}}^{\,n} = \frac{\lambda}{1-\lambda}[r_H - r_L].$$

Overall this implies that $j^* = m$ for $c < \widehat{\widehat{c}}^{\,n}$, and $j^* = 1$ otherwise.

Moreover, the *high-price strategy* is more profitable than the *price-reduction strategy* with $j^* = m$ if $\pi^n(P = r_H) > \pi^n(P_j = r_L)$, which is equivalent to

$$\begin{aligned}
\sum_{i=1}^{m-1}(1-\lambda)^i\left[-c+\lambda r_H\right] - (1-\lambda)^m(1+\theta)c &> \sum_{i=1}^{m-2}(1-\lambda)^i\left[-c+\lambda r_H\right] \\
&\quad +(1-\lambda)^{m-1}\left[-c+r_L\right] \\
\Leftrightarrow \quad (1-\lambda)^{m-1}\left[-c+\lambda r_H\right] - (1-\lambda)^m(1+\theta)c &> (1-\lambda)^{m-1}\left[-c+r_L\right] \\
\Leftrightarrow \quad c &< \widehat{c}^{\,n} \equiv \frac{\lambda r_H - r_L}{(1-\lambda)(1+\theta)}.
\end{aligned}$$

Finally we have $\widehat{c}^{\,n} < \widehat{\widehat{c}}^{\,n}$ because

$$\frac{\lambda r_H - r_L}{(1+\theta)} < \lambda\left[r_H - r_L\right] \quad \Leftrightarrow \quad 0 < (1-\lambda)\, r_L + \theta\lambda\left[r_H - r_L\right].$$

□

Proof of Proposition 1.

First we can immediately see that $\widehat{c}^{\,d} < \widehat{\widehat{c}}^{\,n}$. For $c > \widehat{\widehat{c}}^{\,n}$ the relevant expected profits are

$$\begin{aligned}
\pi^d(\Delta P_1) &= -(1+\theta)c + \lambda r_H + (1-\lambda)r_L \\
\pi^n(P_1 = r_L) &= -(1+\theta)c + r_L.
\end{aligned}$$

It is easy to show that $\pi^d(\Delta P_1) > \pi^n(P_1 = r_L)$, which implies that the discount policy is optimal for $c > \widehat{\widehat{c}}^{\,n}$.

For $c < \widehat{c}^{\,d}$ Firm A's expected profit, when adopting the discount policy, is given by

$$\pi^d(\Delta P_m) = -(1+\theta)c + \lambda r_H + \sum_{i=1}^{m-1}(1-\lambda)^i\left[-c+\lambda r_H\right] + (1-\lambda)^m r_L.$$

We have $\pi^d(\Delta P_m) > \pi^n(P = r_H)$ because

$$\begin{aligned} \sum_{i=1}^{m-1}(1-\lambda)^i\left[-c+\lambda r_H\right]+(1-\lambda)^m r_L &> \sum_{i=1}^{m-1}(1-\lambda)^i\left[-c+\lambda r_H\right]-(1-\lambda)^m(1+\theta)c \\ \Leftrightarrow \quad r_L &> -(1+\theta)c. \end{aligned}$$

Furthermore, we have $\pi^d(\Delta P_m) > \pi^n(P_m = r_L)$, with

$$\pi^n(P_m = r_L) = -(1+\theta)c + \lambda r_H + \sum_{i=1}^{m-2}(1-\lambda)^i\left[-c+\lambda r_H\right] + (1-\lambda)^{m-1}\left[-c+r_L\right],$$

because

$$\begin{aligned} \sum_{i=1}^{m-1}(1-\lambda)^i\left[-c+\lambda r_H\right]+(1-\lambda)^m r_L &> \sum_{i=1}^{m-2}(1-\lambda)^i\left[-c+\lambda r_H\right]+(1-\lambda)^{m-1}\left[-c+r_L\right] \\ \Leftrightarrow \quad (1-\lambda)^{m-1}\left[-c+\lambda r_H\right]+(1-\lambda)^m r_L &> (1-\lambda)^{m-1}\left[-c+r_L\right] \\ \Leftrightarrow \quad \lambda\left(r_H-r_L\right) &> 0. \end{aligned}$$

The discount policy is therefore optimal for $c < \widehat{c}^{\,d}$.

Finally suppose that $c \in [\widehat{c}^{\,d}, \widehat{\widehat{c}}^{\,n})$. Recall that $\pi^d(\Delta P_m) > \pi^n(P = r_H)$, $\pi^n(P_m = r_L)$ for $c < \widehat{c}^{\,d}$. Moreover, for $c \geq \widehat{c}^{\,d}$ we have $\pi^d(\Delta P_1) > \pi^d(\Delta P_m)$. Thus, for $c \in [\widehat{c}^{\,d}, \widehat{\widehat{c}}^{\,n})$ we have $\pi^d(\Delta P_1) > \pi^n(P_m = r_L)$, $\pi^n(P = r_H)$, i.e., the discount policy is optimal. Consequently, the discount policy is optimal for all $c > 0$. $\square$

Change of $P(j)$ in j.

Let $\widetilde{j}$ be a continuous parameter, with $\widetilde{j} \in [1, m]$, so that $P(\widetilde{j}) = r_H - \frac{1}{\widetilde{j}}(1-\lambda)^{\widetilde{j}-1}\left[r_H - r_L\right]$. Taking the first derivative of $P(\widetilde{j})$ w.r.t. $\widetilde{j}$ we get

$$\frac{dP(\widetilde{j})}{d\widetilde{j}} = \frac{1}{\widetilde{j}}(1-\lambda)^{\widetilde{j}}\left(\frac{1}{\widetilde{j}} - \ln(1-\lambda)\right)(1-\lambda)^{-1}\left[r_H - r_L\right].$$

Note that $\lambda \in (0,1)$, so that $\ln(1-\lambda) < 0$. Thus, $dP(\widetilde{j})/d\widetilde{j} > 0$. Moreover,

$$\begin{aligned}\frac{d^2P(\widetilde{j})}{d\widetilde{j}^2} &= -\left[\frac{1}{\widetilde{\widetilde{j}}}\left(\frac{1}{\widetilde{\widetilde{j}}} - \underbrace{\ln(1-\lambda)}_{<0}\right)\left(\frac{1}{\widetilde{\widetilde{j}}} - \underbrace{\ln(1-\lambda)}_{<0}\right)(1-\lambda)^{\widetilde{j}} + \frac{1}{\widetilde{\widetilde{j}}^3}(1-\lambda)^{\widetilde{j}}\right](1-\lambda)^{-1}[r_H - r_L] \\ &< 0\end{aligned}$$

Proof of Proposition 2.

To be written...

Proof of Proposition 3.

To be written...

References