

Strategic Implications of Alternative Learning Approaches in the Personalization Process

Scott Fay

Department of Marketing
University of Florida
P.O. Box 117155
Gainesville, FL 32611-7155
Phone: 352-273-3268
Fax: 352-846-0457
e-mail: scott.fay@cba.ufl.edu

Deb Mitra

Department of Marketing
University of Florida
P.O. Box 117155
Gainesville, FL 32611-7155
Phone: 352-273-3284
Fax: 352-846-0457
email: deb.mitra@cba.ufl.edu

Qiong Wang

Smeal College of Business
Pennsylvania State
University
University Park, PA 16802
Phone: 814-867-0226
Fax: 814-865-3015
email: qiongwang@psu.edu

April 2008

Strategic Implications of Alternative Learning Approaches in the Personalization Process

Abstract

A critical step in a product's personalization process is the learning stage during which a firm discovers its customers' preferences. We categorize two alternative learning approaches: (1) soliciting preference information directly from the customer at the time of purchase (*S-Learning*) and (2) inferring preference information indirectly based on past observations of the customer (*O-Learning*). We explore how a firm's optimal learning approach depends on (i) the relative adoption cost of *S-Learning* and *O-Learning*, (ii) the size of its customer base, and (iii) the learning approach(es) adopted by its rival. In particular, we show how different learning approaches adopted by a firm impact price competition in the marketplace and its rival's choice between these alternative approaches. A key insight is that *O-Learning* provides a credible signal for relaxing price competition, while *S-Learning* does not. Furthermore, *S-Learning* creates a disincentive for competitors to also invest in the *S-Learning* approach. We survey business customers on the learning approaches adopted by their supplier firms and find significant evidence supporting our theory.

Keywords: personalization, learning, competitive strategy, customer relationship management, customization

INTRODUCTION

It is generally accepted that personalizing products and/or services enable firms to increase their profits and customer loyalty (Peppers and Rogers 1997, Brady et al. 2000, Ansari and Mela 2003). However, personalization is only possible once a firm knows what a customer's "ideal" product is. Thus, a critical step in the personalization process is the learning method with which a firm discovers customers' preferences (Murthi and Sarkar 2003). Prior literature has explored various ways a firm might learn about its customer's preferences (e.g., Randall, Terwiesch, and Ulrich 2005, 2007; Toubia et al. 2003; Srikumar and Bhasker 2004). However, the strategic implications of different learning approaches remain unclear. In particular, does the way of learning adopted by a firm influence a rival's optimal way of learning or the degree of price competition in the marketplace? In this paper, we categorize the different ways of learning into two broad approaches and examine how the two learning approaches impact inter-firm competition.

The following examples help illustrate what we view as two fundamentally different approaches to learning customer preferences. At Nikeid.com, a customer designs an athletic shoe to his/her specifications, selecting each element of the shoe from the material of the sole to the color of the shoelace (Randall, Terwiesch, and Ulrich 2005). At Pandora.com, based on the user's previous listening pattern, personalized recommendations are made as to which new releases he would most enjoy (Moser 2006). In both examples, the seller helps a customer identify his/her most preferred product. And, both examples require substantial investment on the part of the seller to learn about customer's preferences – Nikeid.com designs an interface that solicits information from customers in an efficient and effective manner; Pandora.com creates a database to track listening behavior, hires a team to classify new music as it is released, and develops an algorithm to ensure accurate recommendations. We label Nikeid.com to be using *S-Learning*, where the firm

relies on “solicited” information, i.e., products are personalized on the basis of information collected at the time of purchase. Solicitation can be done directly – as at Dell.com where the customer designs his own laptop configuration – or indirectly, as Eleuria does by surveying a customer’s preferences for fragrances and then offering a perfume that best satisfies her reported tastes (Randall et al. 2005). On the other-hand, we label Pandora.com to be using *O-Learning*, where the firm solely relies on previous “observations” about the customers’ preferences. This information may be gathered through previous interactions with that particular customer and/or through other ‘similar’ customers. For instance, a firm that maintains a database consisting of personal purchase histories and click-stream data can use this knowledge to identify a potential user’s most preferred product offering. Of course, over time, the firms will update its knowledge of the user’s preferences based on the user’s purchases. Other examples based on *O-Learning* include Amazon.com offering personalized book recommendations and the Ritz-Carlton anticipating a return guest’s preferred snack (Court 2005).ⁱ

Our principal thesis is that the two learning approaches *S-Learning* and *O-Learning* impact firms within a market in three different ways. First, a firm cannot possibly use *O-Learning* to personalize products for new customers (i.e., those with no previous interaction with the firm). In contrast, a firm could use *S-Learning* to benefit both existing and new customers. Second, and related to the first, *O-Learning*, but not *S-Learning*, may relax competition since an *O-Learning* firm focuses on retaining its own customers rather than acquiring customers of competing firms. This means *O-Learning* by a firm bestows a positive externality upon its rival since the relaxed competition increases the profit of both. Third, *S-Learning*, but not *O-Learning*, imposes costs on the customers in obtaining personalized products. In view of these differences, our objectives are to analyze:

1. The firms' incentives to invest in these learning approaches
2. The strategic impact of a firm choosing a specific learning approach

In section 2, we describe previous literature on alternative learning approaches in the personalization process and provide anecdotal evidence suggesting the strategic consequences of this choice on firms. In section 3, we introduce our model under monopoly and competitive conditions. The monopoly model highlights how each learning approach affects the customer's decision-making process. The competitive model highlights the pricing equilibrium and the reaction decision, i.e., whether firm *B* should invest in *O-Learning* or *S-Learning* (or neither) given firm *A*'s previous decision. In section 4, we use a survey to examine firms' choice of learning approach and test whether these choices are consistent with our model. In section 5, we explore how a firm can use its investment decisions in order to influence a rival's decision and the degree of price rivalry in the market. Section 6 concludes with a summary of the results and suggestions for future research.

LEARNING APPROACHES IN THE PERSONALIZATION PROCESS

Drawing extensively upon the framework introduced by Zipkin (2001), Figure 1 identifies three elements of the personalization process and several representative papers from each of these literatures. Our paper addresses the first element, Learning – a mechanism for discovering customers' preferences, which is discussed in greater detail below and in Figure 2. The latter two elements, product process personalization and marketing mix personalization, comprise what is traditionally associated with the term customization – developing a production technology to efficiently produce products to match these preferences and tailoring the distribution of products to customers, e.g., via customized prices or communications.

Figures 1 and 2 here please

But, notice that to produce and distribute the “right” products, a firm must somehow learn

about customers' preferences. The learning process is defined as "an artful means of leading customers through the process of identifying exactly what they want ... [and] reduc[ing] the costs associated with customers' laborious searching" (Zipkin 2001). The limited published research on learning in the personalization context (Murthi and Sarkar 2003) focuses upon identifying and improving upon various techniques for learning customers' preferences. Figure 2 lists a few representative papers from this literature. A more extensive review can be found in Murthi and Sarkar (2003). In particular, the *S-Learning* approach can be implemented by gathering information at the time of purchase and then using various methodologies to infer preferences based on the customer's responses, e.g., asking customers to respond to prototypes or to participate in a series of preference tests and then employing conjoint analysis. Other papers have examined some potential requirements for this approach to be feasible, e.g., consumers must see enough benefit to warrant their participation (Dallaert and Stremersch 2005) and a firm must have the technology to implement this approach (Gilmore and Pine 1997).

On the other hand, *O-Learning* enables a firm to discover customers' preferences via passive observation. Many analytical tools can be employed in order to infer a customer's preferences, e.g., a variety of recommendation systems based on customer profiles using case-based reasoning, collaborative filtering, dynamic taxonomy hierarchy, or fuzzy logic. But, such methods can only be employed if customers trust the firm (Resnick and Varian 1997) and the firm has the technological sophistication to implement this approach (Gilmore and Pine 1997).

Choosing a Learning Approach. Although many papers study the various tools that can be used to implement learning approaches based on either *S-Learning* or *O-Learning* (see Figure 2 for additional citations), there has been very little research that considers WHICH approach should be used by a firm. This is an issue of high practical importance because a firm may need to make a substantial financial investment and/or acquire additional technical expertise in order to implement

these various technologies. Thus, firms cannot adopt all possible technologies but must instead focus upon the methodology that is most appropriate given their particular market situation. Simonson (2005) weighs these various learning approaches in the context of customer-level variables (e.g., stability of customers' preferences, responsiveness). Furthermore, customers' willingness and ability to participate in the firm's learning process will also affect which approach is optimal for a firm (Gilmore and Pine 1997). In contrast, we focus on competitive factors. For example, we are interested in how a firm's optimal learning approach is affected by the approach adopted by a rival.ⁱⁱ To illustrate that such strategic considerations can play an important role in the evolution of markets, we now provide information on the history of personalization in two different product categories. Past research has fruitfully used such a qualitative approach in motivating research on strategic implications (Bohlmann, Golder, and Mitra 2002; Zucker, Darby, and Armstrong 2002).

Case 1: Personal Computers. In 1984, the precursor of the present day Dell Corporation, PC's Unlimited started selling the first personalized personal computers, e.g., empowering customers to include features and components of one's own choice. At first, orders for personalized computers were placed via an 800 number. Later, in 1996, Dell launched its Web site which allowed customers to design their own PC's. To facilitate its personalization process, Dell devoted considerable resources to training its customer service staffs on how to answer questions, resolve complaints, take orders, and help clients select the best options for their computing needs. Since Dell solicited information from its customers about their preferences at the time of purchase, Dell used an *S-Learning* approach. This personalization approach helped Dell consistently improve its market share until it became the market leader of the PC industry that it is today.

Dell's success attracted the intense scrutiny of its competitors. In 1989, Gateway began to imitate Dell's customer-build *S-Learning* strategy using telephone and, later, the Internet. Then, in

1998, both HP and Compaq unveiled their own efforts to allow their customers to design and purchase PCs on the Internet. By 2000, all Dell's competitors had started using *S-Learning* based personalization. A savage price war took place in the year of 2001 that slashed profit margins for all PC makers. Gateway started losing money (EPS: -3.14) and market share (- 3%) leading to a change in its strategy from selling personalized products directly to selling standardized products using indirect channels. Thus, direct sales of personalized computers plummeted from a high of 90% in 1995 to less than 5% by 2005. Similarly, Compaq and HP decided to merge in 2002 and turn their focus back to selling more standardized products using traditional resellers. In 2005, IBM exited the business, selling its PC operations to Lenovo.

Case 2: Internet Search. During 1994, Jerry Yang and David Filo, started *Yahoo!*. Initially, Yahoo! took a look of a directory that enabled users to navigate the Internet by self-selecting the different options under indexes and sub-indexes. A family of specialized indexes for geographic audiences (e.g., Yahoo! Japan, UK & Ireland, Chicago, and Boston), shared-interest audiences (e.g., Yahoo! Internet Life magazine and website, Yahoo! Finance, and Yahoo! News), and demographic audiences (e.g., Yahoo!ligans!, a Web guide for kids; and Beatrice's Web Guide for women) were developed. Therefore, for Yahoo!, the personalization of a search is achieved using *S-Learning*, i.e., the user self-selects among the available indexes and navigates through these indexes according to his or her own personal preferences.

Google was launched in 1997 by Sergey Brin and Larry Page, and used a different type of learning approach to personalization to produce its search results. At first, Google did not have any established customers, and thus could not introduce *O-Learning* based personalization. Instead, it slowly but steadily built its capability to implement *O-Learning*. In particular, it began by trying to infer potential users' preferences by analyzing the structure of web-pages (i.e., the number and importance of other pages that link to the page) via its 'PageRank' algorithm. A

breakthrough for Google occurred when AOL/ Netscape and Yahoo! decided to appoint Google to administer the direct search option on their web portals.ⁱⁱⁱ These deals helped Google gain users and more importantly, crucial information on the preferences of these users. As its information improved, Google improvised on its search results by incorporating additional information based on users' preferences. Yahoo! improved its profitability (EPS 0.06 in 2000 to 0.58 in 2004) but kept losing market share to Google (from a high of 55% in 2000 to 26% in 2004). In 2004, amidst the growth of Google, Yahoo! dropped Google and launched its own search engine. Google responded by taking the final step of *O-Learning* and launching Google Personalized Search. Google Personalized Search uses personal information from “*My Search History*” data to refine the search, making it more individual-specific and unique. Since then, Yahoo! has started focusing on *O-Learning*-based Internet search rather than its original *S-Learning*-based Internet directory, which signifies an imitation strategy (of responding to *O-Learning* by *O-Learning*).

The success of Dell and Google has been the topic of much analysis. While we do not dispute alternative interpretations, in our view, these cases illustrate the general importance of considering the strategic implications of the choice of a specific learning approach by a firm on its competitors. For instance, these cases seem to illustrate (i) the failure of Gateway's *S-Learning* as a response to Dell's *S-Learning* approach, and (ii) the success of Google's *O-Learning* as a response to Yahoo's *S-Learning* approach. Nevertheless, while using these cases to motivate our research objective, we make no assertions on their representative nature and also acknowledge that factors other than the learning approach may also have played an important role. In order to formalize the comparison between these two learning approaches and to develop a more complete understanding, in the sections that follow, we present an analytical model and empirically examine the postulates that come out of this model.

MODEL

Approach

The fundamental aspect of the learning stage during the personalization process is that it provides a mechanism for the firm to help a customer locate her most preferred product (Zipkin 2001). For example, with Dell Computers, the key advantage of the learning stage is that it allows Dell to provide a customized notebook with minimal effort, i.e., without requiring the customer to visit numerous retailers or just “settling” for the most suitable configuration available at a particular local retailer.

Following this reasoning, and consistent with the idea of personalized recommendations (e.g., Srikumar and Bhasker 2004, Randall et al. 2007), we introduce a model based on search. Each firm offers an extensive product line. Customers differ in which product from this line that they prefer. In the absence of any effort by the firm to learn about individual’s preferences, each customer engages in costly search. At each point in time, the customer faces a decision of whether to stop searching and choose the best alternative she has encountered thus far or to continue searching for a better alternative. Learning (by firms) is a way to help customers avoid search costs. If the firm employs *S-Learning*, customers can avoid sequential search, but will incur additional effort (e.g., filling out surveys and/or acquiring the expertise needed to design their own products). In contrast, a firm that uses *O-Learning* eliminates all transaction costs since the firm already knows what customers want without any additional effort on the part of customers. However, the firm is only able to offer personalized recommendations for customers for whom it has historical data.

We start by building a monopoly model. Shugan (2002) argues that the central value of a model stems from its simplification which enables researchers to produce testable implications

and thus advance knowledge. Accordingly, our monopoly model allows us to convey our intuition clearly, make predictions, and serve as a benchmark that can be compared to the competitive model. Such a comparison allows us to identify actions taken for competitive reasons, i.e., the strategic implications of learning approaches, which are the focus of this paper.

Monopoly

Suppose a single firm offers J horizontally-differentiated products, each of which is sold at a price p . This is consistent with the observation from Syam et al. (2005) that many customizers do not price discriminate.^{iv} The unit cost of production is constant across these products and, for analytical convenience, is assumed to equal zero. We restrict our attention to the situation where J approaches infinity. The number of customers is normalized to one. Customer i values good j at V_{ij} and will consume at most one item. For all i , V_{ij} is drawn independently from the uniform distribution on the interval $[0, 1]$. Furthermore, customers cannot costlessly observe their valuations. Instead, valuations are revealed only as a result of search effort. The required effort depends on the type of learning approach used by the firm (if any) and is outlined below. In the next section, we extend the analysis to a duopoly setting. We also extend the model to show that our key results are robust to a setting where prices are customized and a firm can adopt both learning approaches simultaneously.

No Learning. If the firm does not invest in either learning approach, customers must screen the product offerings on their own. Suppose that each customer incurs a cost c for revealing the valuation of any particular product. A customer who purchases after sampling n products will earn a net customer surplus of:^v

$$(1) \quad CS_n = \text{Max}[V_1, V_2, \dots, V_n] - n \cdot c - p$$

For notational convenience, we have dropped the customer subscripts.

Standard analysis (e.g., Lippman and McCall 1976, McCall 1965) reveals that a customer maximizes her expected surplus by following an optimal stopping rule, i.e., continuing to search until she receives a value of at least \bar{V} . The optimal threshold is found by equating the cost of searching one more time (c) to the expected increase in value created by the additional search. Thus, the optimal stopping rule is:

$$(2) \quad \bar{V} = 1 - \sqrt{2c}$$

Following this strategy, each customer earns an expected surplus of $1 - \sqrt{2c} - p$.

Customers are willing to search only if $p \leq 1 - \sqrt{2c}$. The monopolist maximizes its profit by choosing $p = 1 - \sqrt{2c}$. Since all customers purchase one item at this price, the firm's profit in the absence of any investment is:

$$(3) \quad \Pi_M^\emptyset = 1 - \sqrt{2c}$$

For the remainder of the paper, we assume $c < \frac{1}{2}$ so that positive profits are attainable in the absence of any investment.

S-Learning. The monopolist can reduce customers' search costs by investing in technology that solicits information from consumers at the time of purchase.^{vi} We assume the implementation cost to be T_{SL} . Here, the firm offers customers the opportunity to identify their preferred item without searching product-by-product. The firm learns a customer's tastes by having the customer answer a series of questions. Let the costs to a customer of completing the required tasks equal E . For a customer to be willing to provide the solicited information rather than searching on her own, we must have:

$$(4) \quad E \leq \sqrt{2c}$$

Throughout this paper, we assume condition (4) is met so that *S-Learning* has potential value.

Furthermore, consistent with Dewan et al. (2003), we assume the collected information allows the

firm to identify the customer's most preferred product. With an infinite number of choices from the value distribution, the firm can find a product that is valued at 1. Thus, the customer earns a net surplus of $1 - E - p$. The profit-maximizing price is: $p = 1 - E$. The firm earns a profit of:

$$(5) \quad \Pi_M^{SL} = 1 - E - T_{SL}$$

O-Learning. A second approach to learning is based on data from the historical observation of customers. The firm uses information it has collected from past shopping behavior (e.g., past purchases of related and unrelated items, browsing behavior) to construct a customer profile that can be used to identify a customer's most preferred product. We assume this learning approach can be implemented at a cost T_{OL} without any associated per-customer costs. For example, the process may be entirely automated, thus requiring an sophisticated infrastructure, but approximately zero user-specific costs.

Since customers have heterogeneous tastes, recommendations based on *O-Learning* are valid only if the firm has an established relationship with the particular user. We assume such a relationship exists with a proportion α_M of the population. Furthermore, in line with our assumptions regarding the accuracy of predictions made on the basis of *S-Learning*, we assume that predictions based on these historical relationships are perfect, i.e., they allow the firm to find a product of value 1 for each of these customers. Since *O-Learning* requires no effort on the part of customers, a customer's surplus from purchasing this recommended item equals $1 - p$.

The firm maximizes profit on this segment by choosing $p = 1$ and thus earns a profit of:^{vii}

$$(6) \quad \Pi_M^{OL} = \alpha_M - T_{OL}$$

Choosing the Optimal Learning Approach. The monopolist chooses *O-Learning*, *S-Learning*, or No Learning depending on which action yields the highest profit. The optimal investment is given by Proposition 1:

PROPOSITION 1 (Monopoly): *A monopolist's optimal learning technology depends on the relative costs of these two technologies, the size of its current customer base, and the effort required of customers when S-Learning is employed. For instance, O-Learning is more likely to be optimal the larger is the investment cost and customer effort of S-Learning, the larger is the size of the firm's current customer base, and the smaller is the investment cost of O-Learning. In particular,*

- (a) *No Learning is optimal if $T_{SL} \geq \sqrt{2c} - E$ and $T_{OL} \geq \sqrt{2c} + \alpha_M - 1 \equiv \hat{H}_M$*
- (b) *S-Learning is optimal if $T_{SL} < \sqrt{2c} - E$ and $T_{SL} \leq 1 - \alpha_M + T_{OL} - E$*
- (c) *O-Learning is optimal if $T_{OL} < \sqrt{2c} + \alpha_M - 1$ and $T_{SL} > 1 - \alpha_M + T_{OL} - E$*

Figure 3(a) illustrates these results. A monopolist employs *S-Learning* if the implementation cost of *S-Learning* is sufficiently cheap, the implementation cost of *O-Learning* is sufficiently expensive and the search costs in the absence of either learning approach are sufficiently high. Furthermore, *S-Learning* is less likely to be employed as the size of the firm's database, α_M , increases (and thus *O-Learning* is more valuable) and as the effort required during the solicitation procedure (E) increases (and thus *S-Learning* is less valuable).

Figure 3 here please

An interesting result is that a monopolist does not always invest in *O-Learning* when it is socially optimal to do so. To see this, note that under *No Learning* and *S-Learning* a monopolist is able to extract the entire surplus that is created. For example, when a monopolist has *S-Learning*, it offers this tool to all customers, as would be efficient to do, and sets a price equal to customers' common expected net value. However, the monopoly outcome does not match the welfare-maximizing outcome when *O-Learning* is available. Recall that in this case a monopolist only sells to customers with whom it has a historical relationship. However, the social optimum prescribes that the remainder of the population would also purchase the product, but through traditional search. This does not occur under a monopoly because high prices ($p = 1$) discourage purchases by this segment. Specifically, the social optimal would create a total surplus of

$W_{SO}^H = \alpha_M + (1 - \alpha_M)(1 - \sqrt{2c}) - T_{OL}$ which exceeds the profit a monopolist earns (see Equation 6).

Thus, there are parameters such that the monopolist invests in *S-Learning* or No Learning, whereas the efficient outcome would utilize *O-Learning*.

Competition

Now suppose a firm faces competition. In particular, there are two firms, *A* and *B*, both of which sell J items, where J again approaches infinity. Customers purchase at most one item in total. All item valuations are drawn from $U[0,1]$. Thus, a randomly selected item from firm *A* has the same expected value as a randomly chosen item from firm *B*. However, to account for the differences in databases that firms have, we assume that firm *B* has an historical relationship with a proportion α_B of the population, whereas firm *A* has previously interacted with the remaining $1 - \alpha_B$ proportion of the population. This is obviously an oversimplification, but a similar assumption is made in other papers such as Klemperer (1995). We employ this assumption to highlight the ability of *O-Learning* to mitigate price competition. However, the benefit from reduced rivalry would be lessened if databases partially overlap or if there are customers that are in neither database.

Learning technologies may not be easily adjusted over short time horizons (at least not as easily as prices can be). Thus, consistent with Syam et al. (2005), we model these decisions in a sequential manner.^{viii} First, firm *A* decides upon its learning approach ($D_A = \text{SL}, \text{OL}, \text{or } \emptyset$). Second, firm *B* chooses its learning approach ($D_B = \text{SL}, \text{OL}, \text{or } \emptyset$).^{ix} Third, firms *A* and *B* choose prices (p_A, p_B) independently and simultaneously. We solve the game recursively. First, we derive the pricing equilibrium given the learning approaches. Then, we consider firm *B*'s best response. Finally, we study firm *A*'s initial investment choice, which allows us to identify the Sub-game Perfect Nash Equilibrium to the full game.

Pricing. Prices depend on which, if any, learning approach the firms have invested in. For each possible permutation of (D_A, D_B) , we derive the equilibrium prices. This analysis is reported in the appendix (and summarized in Table A1). For illustration purposes, we examine a few cases in the text:

(i) $D_A = D_B = \emptyset$. Since an expected value of a draw from either firm is identical for all customers (regardless of the number and realizations of previous draws), all customers will buy from the firm with the lowest price provided that $p_i \leq 1 - \sqrt{2c}$. This is standard Bertrand competition with the well known unique pricing equilibrium at $p_A = p_B = 0$.

(ii) $D_A = D_B = SL$. Again, all customers view the firms as perfect substitutes, i.e., purchasing from firm A yields an expected surplus of $1 - E - p_A$ and purchasing from firm B yields an expected surplus of $1 - E - p_B$. The unique pricing equilibrium is $p_A = p_B = 0$. Thus, neither firm is able to recover its sunk investment in *S-Learning*.

(iii) $D_A = D_B = OL$. Here, customers view the two firms' product offerings quite differently since only the firm with historical information is able to offer personalized recommendations and thus allow customers to avoid costly search. For example, a segment of size α_B earns an expected surplus of $1 - p_B$ if it purchases from firm B and an expected surplus of $1 - \sqrt{2c} - p_A$ if it purchases from firm A . There is a pure-strategy equilibrium at the prices $p_A = p_B = 1$ as long as condition (7) is met:

$$(7) \quad c > \frac{(\text{Max}[\alpha_B, 1 - \alpha_B])^2}{2}$$

Condition (7) ensures that it is not profitable for either firm to try to poach each other's customers. Since we are interested in environments where search costs are sizeable enough to warrant

investment in learning technologies, we assume condition (7) holds throughout the remainder of the paper.

In the cases where only one firm invests in *O-Learning*, the pricing equilibrium involves mixed strategies. The appendix details the derivation of these equilibria. However, it is relevant to note here that the analysis relies on an additional assumption:

$$(8) \quad E > 1 - \frac{1}{2 - \text{Min}[\alpha_B, 1 - \alpha_B]}$$

Thus, our results only apply when *S-Learning* imposes non-trivial costs on customers.

The Reaction Decision. Now, we consider the second stage of the game: firm *B*'s investment decision given D_A . Firm *B* chooses D_B to maximize its profit, anticipating what the pricing equilibrium will be given D_A and D_B . The appendix contains the detailed analysis. In the text, we summarize our results using Proposition 2 and Figures 3(b)-(d):

PROPOSITION 2 (Competitive Response): *A firm's optimal learning technology is strongly influenced by the learning technology adopted by its rival. In particular,*

- (a) *Because O-Learning can reduce price rivalry, a firm that faces a rival has a greater incentive to invest in O-Learning than a monopoly does.*
- (b) *If a rival uses S-Learning, a firm has no incentive to also invest in S-Learning and has less incentive to invest in O-Learning (relative to the case where the rival uses no learning technology).*
- (c) *If a rival uses O-Learning, a firm has less incentive to invest in either O-Learning or S-Learning (relative to the case where the rival uses no learning technology).*

The presence of a rival and that rival's ability to offer personalized recommendations dramatically affects what learning approach will be optimal for a firm. Consider the case where firm *A* does not have the ability to personalize recommendations. Figure 3(b) illustrates the optimal response by firm *B*. The key characteristic to note is that, holding the availability of historical information constant, a firm that faces a rival is more likely to invest in *O-Learning* than it would in the absence of competition. If neither firm personalizes, price competition is intense

since the firms' product offerings are undifferentiated. Investing in *O-Learning* not only allows a firm to reduce customers' search costs, it also provides a credible mechanism to mitigate competition by creating distinct segments of the population. This advantage exists even though perfect segmentation cannot be achieved, i.e., the resulting pricing equilibrium involves mixed strategies in which firm *A* sells to *B*'s historical customer base with some positive probability. Thus, *O-Learning* reduces the degree of rivalry but does not eliminate it altogether. This benefit does not apply to investments in *S-Learning*. Although *S-Learning* also reduces customers' search costs and thus allows firm *A* to sustain higher prices, all customers view the trade-off between purchasing from firm *A* and purchasing from firm *B* in the same way. Therefore, no finer segmentation of the population occurs.

Figure 3(c) illustrates the optimal response by firm *B* if firm *A* possesses *S-Learning*. The most glaring observation is that firm *B* should never respond by also investing in *S-Learning*. Such an investment would create fierce competition that drives prices down to marginal cost. Firm *B* would strictly prefer not to make any investment (and thus avoid the investment cost T_{SL}). Furthermore, *A*'s investment in *S-Learning* also discourages *B* from investing in *O-Learning* (since $\hat{H}_S < \hat{H}_\emptyset$). The implication of this result is quite interesting. Firm *A* is very formidable if it has *S-Learning*. Often, firm *B*'s best response is to not invest in either learning approach, thus allowing firm *A* to maintain a competitive advantage. It is only prudent for firm *B* to invest in *O-Learning* if this learning approach is cheap and its own historical base is sufficiently large.

Figure 3(d) illustrates the optimal response by firm *B* if firm *A* invests *O-Learning*. A critical feature of this graph is that the size of the "Invest in *S-Learning*" region is smaller in Figure 3(d) than in Figure 3(b). Choosing *S-Learning* will require smaller T_{SL} for the case when the firm faces a rival who has *O-Learning*. Furthermore, such an investment by a rival also expands the *No Learning* region downward since $\hat{H}_\emptyset > \hat{H}_H$. Thus, firm *A*'s investment in *O-*

Learning reduces, but does not eliminate, B 's incentive to invest in *S-Learning*. Furthermore, firm B is less likely to invest in *O-Learning* if firm A has already made such an investment (rather than if it had done nothing). The reason for this is as follows. Firm A 's investment in *O-Learning*, mitigates cutthroat competition by creating two distinct segments of customers. A second investment in *O-Learning* by firm B could crystallize this segmentation, i.e., move the market from a mixed strategy equilibrium in which leakage between segments may occur to a pure-strategy equilibrium in which each firm only sells to its historical base at their reservation value. But, the benefit of this second investment is less than the value from the initial investment. Thus, when such investments involve substantial costs, firm B may “settle” for imperfect segmentation.

Modeling Extensions. The basic model assumes that each firm can invest in at most one learning technology. In this subsection, we extend the model to allow a firm to invest in both technologies if it so chooses. We also allow each firm to customize prices based on their previous interactions (or lack thereof) with the firm. In particular, each firm can charge one price to consumers who have an historical relationship with them (i.e., is a member of α_i) and another price to new customers (i.e., customers who belong to $\alpha_{.i}$). Proposition 3 summarizes the impact of this modeling extension:

PROPOSITION 3 (Multiple Learning Technologies and Price Customization): *If it is possible for a firm to customize prices and to invest in both *O-Learning* and *S-Learning*, then it is never optimal to respond to a rival that employs *S-Learning* by using a strategy that uses *S-Learning*. In particular,*

- (a) *If a rival uses *S-Learning* only, then it is not optimal to respond by investing in *S-Learning* only or by investing in both *S-Learning* and *O-Learning*.*
- (b) *If a rival uses both *S-Learning* and *O-Learning*, then it is not optimal to respond by investing in *S-Learning* only or by investing in both *S-Learning* and *O-Learning*.*

Proposition 3 helps generalize our finding that *S-Learning* by a rival can eliminate a firm's incentive to also invest in *S-Learning*. If firm A uses any learning technology (*S-Learning*, *O-Learning*, or both), an investment in *S-Learning* by firm B simply intensifies competition for firm

A's customers. Since firm *B* can never offer a strictly lower cost search option to customers for whom it does not have historical information, firm *B* could only induce customers to switch if it offered a sufficiently lower price to firm *A*'s customers. But, firm *A* has an incentive to lower its prices enough to prevent such switching. Thus, an investment in *S-Learning* (which is aimed primarily at attracting new customers), in equilibrium, is not successful at expanding one's market share, and therefore, such an investment is not warranted.

EMPIRICAL EVIDENCE

We now present a cross-sectional test of our main analytical results using a survey of purchasing managers in different firms. We select purchasing managers for four reasons. First, purchasing managers, acting as business customers, are knowledgeable about their suppliers' personalization effort and achievement (c.f., Heide and John 1990). Second, purchasing managers are likely to be well informed about alternative suppliers of a product, being responsible for the purchase of the particular product or product line. This enables us to construct measures of *S-Learning* and *O-Learning* for rival suppliers of the same product. Third, we avoid the problem of self-reporting bias since asking about the learning approach used by a firm directly from its sales manager will expose the survey to common method variance (Bertrand and Mullainathan 2001). Fourth, choosing purchasing managers to reflect on their suppliers' strategies also avoids any potential discrepancy between the beliefs of the sales managers with the actual selling approaches employed by their frontline sales representatives. Also, there is ample evidence in the marketing literature suggesting that some firms are unable to assess their own efforts in discovering customer specific preferences (Gatignon and Xuereb 1997; Voss and Voss 2000).

Purchasing managers from a variety of industries involving manufacturing, retailing, material, and equipment industries were contacted by email and directed to complete an online survey via a website link. They were asked to provide information about a focal (or primary)

supplier of a specific product they were responsible for purchasing and a rival, i.e., an alternate (or secondary) supplier of the same product. We received a total of 189 responses with information on 398 suppliers. After accounting for the undeliverable email addresses, the response rate was 10.4%. Overall, we received information on a focal and a rival supplier of 182 products. For each supplier, the respondent assessed both the extent of personalization achieved (effectiveness in delivery of personalized products) and the amount of *S-Learning* employed (efforts made by a supplier to solicit information directly from the buyer) using a five-point scale.^x In Table 1, we provide the descriptive statistics of the different measures used in the survey.

Table 1 here please

A main prediction of our model is that one firm's choice depends on the learning approach adopted by its competitor. Particularly, our analytical model suggests that an *S-Learning* focal supplier eliminates a rival supplier's use of *S-Learning*. In contrast, an *O-Learning* focal supplier does not eliminate a rival's use of *O-Learning*. Together these two results imply:

HYPOTHESIS: *The difference in S-Learning employed by an S-Learning supplier and its rival is likely to be larger than the difference in O-Learning employed by an O-Learning supplier and its rival.*

Note the use of *O-Learning* by a supplier is, by definition, not observed by the respondents. They only observe the extent of personalization achieved and the amount of information directly solicited by the supplier from them (i.e., use of *S-Learning*). However, if the extent of personalization achieved is high and the amount of solicited information is low, then the personalization must be achieved through use of *O-Learning*. Accordingly, we use the extent of solicited information by a supplier firm at the time of transaction to apportion the personalization achieved by the firm on account of *S-Learning* and *O-Learning*.^{xi} Next, we label focal suppliers as *S-Learning* or *O-Learning* types depending on their extent of use of a specific learning approach. Accordingly, we sort the focal suppliers in terms of their level of *S-Learning* employed. We label

the top 50 focal suppliers as *S-Learning* type. Likewise we sort the focal suppliers in terms of their level of *O-Learning* employed and label the top 50 as *O-Learning* type. For *S-Learning* focal suppliers, we compute the difference in *S-Learning* with their respective rival, ΔS_S (where the subscript denotes the focal supplier type in terms of the learning approach). For *O-Learning* focal suppliers, we compute the difference in *O-Learning* with their respective rival, ΔO_O . Based on our hypothesis, we expect $\Delta S_S > \Delta O_O$.

Table 2 reports the results of a 2-sample t-test that strongly supports our hypothesis. The mean difference between ΔS_S and ΔO_O is 0.65 and significantly positive ($p < .001$). To further validate our result, we use three analyses. First, we use alternate cutoffs to label *S-Learning* and *O-Learning* suppliers. We use top 25 focal suppliers (i.e., in terms of either *S-Learning* or *O-Learning*), top 75, and top 100. For each of these three alternate categorizations, we compute ΔS_S and ΔO_O . In each of these cases, we find significant support of our hypotheses (see Table 2). Second, we compute the difference in *S-Learning* employed by an *O-Learning* supplier and its respective rival, ΔS_O . Likewise we compute the difference in *O-Learning* employed by an *S-Learning* supplier and its respective rival, ΔO_S . Unlike the earlier results, we find that there is no significant difference between ΔS_O and ΔO_S . This is true for all the 4 alternate characterizations – top 25, top 50, top 75, and top 100. This shows that the significant positive difference between ΔS_S and ΔO_O is not driven by measurement scales. Third, irrespective of our earlier categorization of a supplier as *S-Learning* or *O-Learning*, for each of the 182 products we calculate (i) the absolute difference in *S-Learning* between the focal supplier and the rival, i.e., $\Delta S = |S_{foc} - S_{alt}|$, and (ii) the absolute difference in *O-Learning* between the focal supplier and the rival, i.e., $\Delta O = |O_{foc} - O_{alt}|$. If $\mathbf{z} = \Delta S - \Delta O$, across all products, we expect $\mathbf{z} > 0$. Consistent with this conjecture, this difference score was significantly positive ($\bar{z} = .29$, $p < .001$).

Table 2 here please

Prior to this current research, one would have expected competing firms to choose their learning approach solely based on customer-specific and product category-specific characteristics. For example, Simonson (2005) suggests that the optimal personalization depends on the degree to which customers' preferences are stable over time, customers' skill at conveying these preferences, and customers' ability to recognize a personalized product that suits their preferences. Gilmore and Pine (1997) suggest that a firm's learning approach will depend on customer's willingness to participate in the learning process and the firm's technological sophistication. If such customer and category characteristics are the sole drivers of the choice of a specific learning approach, then one would expect that firms in the same market would employ similar learning approaches, i.e., rivals of *S-Learning* firms are more likely to use *S-Learning* while rivals of *O-Learning* firms are more likely to use *O-Learning*. In stark contrast, we find significant empirical evidence that suggests rivals of *S-Learning* firms are less likely to use *S-Learning* while rivals of *O-Learning* firms are more likely to use *O-Learning*. This finding supports our contention that beyond customer and product category factors, competition plays an important role in shaping the optimal learning approach employed by a firm.

Next, we explore the strategic consequence of our finding. In other words, can a firm choose a learning approach so as to influence the subsequent choice of its rival?

STRATEGIC INVESTMENTS IN LEARNING TECHNOLOGIES

Notice that in the modeling section, we characterized the outcome of the second and third stages of the game. Now, we now consider the first stage, i.e., firm *A*'s investment choice. In particular, firm *A* chooses $D_A = SL, OL, \text{ or } \emptyset$ to maximize its profit, anticipating *B*'s best response, $D_B(D_A)$. We identify several "strategies" that may be optimal for firm *B* depending on the particular set of parameters. We should be clear that we are not attempting to describe all

possible equilibrium outcomes. For example, we will not examine “non-strategic” parameter configurations. By “non-strategic”, we mean those situations where firm A 's investment choice does not affect B 's investment decision in the second stage. For instance, if T_{SL} is sufficiently large and T_{OL} is sufficiently small, firm B will invest in *O-Learning* regardless of firm A 's decision in the first stage (see Figures 3(b)-(d)) and thus B 's action does not influence A 's subsequent choice.^{xii}

Table 3 presents six equilibria to the full game in which strategic factors influence investment choices. The second column lists firm B 's best response to each of the possible investments choices by firm A . The last column lists firm A 's investment choice in the Sub-game Perfect Nash Equilibrium. For example, in Strategy #6, firm B will choose *S-Learning* if A chooses no learning, no learning if A chooses *S-Learning*, and *O-Learning* if B chooses *O-Learning*. Anticipating these responses by B , firm A maximizes its profit by selecting *O-Learning*. In the appendix, we outline the conditions required for each of these respective equilibria hold. In the text, we provide numerical examples to illustrate these six strategies and to serve as verification that the prescribed parameter spaces are non-empty.

Table 3 here please

Strategy#1: “Take Your Pick”: Firm A uses its first-mover advantage to secure the most cost efficient learning approach for itself.

Example: $\alpha_b = .5$, $c = .2$, $E = .5$, $T_{SL} = .05$, $T_{OL} < .3$. Under these parameters, both learning approaches are relatively cheap. Firm B will invest in whichever approach firm A does not choose. Thus, the final outcome will involve one firm investing in *S-Learning* and the other investing in *O-Learning*. By moving first, firm A can choose whether to be the *S-Learning* player or the *O-Learning* player.^{xiii} In this particular example, the *S-Learning* firm earns a net profit of .2 and the

O-Learning firm earns a net profit of $.375 - T_{OL}$. Thus, firm A chooses $D_A = SL$ if $T_{OL} > .175$ and chooses *O-Learning* if $T_{OL} < .175$.

Strategy#2: “Beat ‘em to the Punch”: Firm A takes a dominant industry position by adopting *S-Learning* before firm B does.

Example: $\alpha_B = .5, c = .2, E = .5, T_{SL} < .066, T_{OL} > .342$. For these parameters, *S-Learning* is viable (for one firm). The firm that adopts *S-Learning* earns a positive net profit while the non-adopter earns zero profit. Firm A, by moving first, secures positive profit for itself and, in the process, puts firm B at a severe competitive disadvantage.

Strategy#3: “Beat ‘em to the Punch with a Helping hand”: Firm A prevents an attack by firm B by adopting *O-Learning*, which actually improves the profit of Firm B.

Example: $\alpha_B = .35, c = .2, E = .5, T_{SL} = .1, T_{OL} = .4$. Because of the skewed amount of historical data ($\alpha_B < .5$), *O-Learning* is viable for firm A but not for firm B. If firm A does not invest in either learning approach, firm B will choose *S-Learning* and take the dominant position in the industry. Firm A can prevent this investment by choosing either type of learning approach. It chooses the one that yields the highest profit for itself, which is *O-Learning* for these parameters. This choice has a positive spillover effect on Firm B because it relaxes competition. Firm B may even receive a disproportionate share of the net benefit, e.g., with these particular parameters, equilibrium profit (net of investment costs) is .095 for firm A and .129 for firm B.

Strategy#4: “Be My Guest”: Firm A foregoes the opportunity to adopt *O-Learning*, leaving it for firm B to adopt.

Example: $\alpha_B = .5, c = .2, E = .5, T_{SL} = 1, T_{OL} = .35$. *S-Learning* is prohibitively costly, but *O-Learning* is of moderate expense – low enough to be warranted by one firm but high enough to dissuade a second firm from investing in it. Thus, either firm A or firm B will invest in *O-Learning*. In a somewhat counterintuitive move, A’s best strategy is to not invest in *O-Learning* and let firm B “be its guest” to this learning approach. This accommodating move is advantageous

because it allows A to avoid paying T_{OL} while still benefiting from the segmentation that arises when firm B employs O -Learning.

Strategy#5: “Ride My Coattails”: Firm A takes a leading industry position by adopting O -Learning first, but this action allows Firm B to be a free-rider.

Example: $\alpha_B = .35$, $c = .2$, $E = .5$, $T_{SL} = 1$, $T_{OL} = .25$. Similar to “Be My Guest”, S -Learning is prohibitively costly and O -Learning is only viable for a single firm. But, for these parameters, firm A achieves higher profit by investing in O -Learning (rather than leaving this learning approach to its rival). Firm B benefits from the ensuing reduction in price rivalry without making an investment of its own. Although this may seem like an “expected” result (relative to “Be My Guest”), we should stress that this strategy is only optimal for a more stylized setting. For example, it requires a skewed amount of historical data, i.e., $\alpha_B < .5$, and a lower range of T_{OL} .

Strategy#6: “Follow Me”: Firm A guides the industry to a cooperative outcome by adopting O -Learning first.

Example: $\alpha_B = .5$, $c = .4$, $E = .5$, $T_{SL} = .3$, $T_{OL} = .4$. Firm A essentially has two choices: 1) play aggressively by choosing $D_A = SL$ and thus discouraging any investment (and associated profit) by B ; or 2) play accommodating by choosing $D_A = OL$ and thus inducing firm B to also invest in O -Learning. For these specific parameters, the second choice leads to higher profits for firm A . A 's decision has a dramatic impact on the evolution of the market. The “Follow Me” strategy guides the market to a conciliatory outcome in which each firm earns a profit of .1; whereas choosing $D_A = SL$, which would be optimal for $T_{SL} < .294$, leads the market down a path of fierce competition in which firm B earns zero profit.

CONCLUDING REMARKS

This paper explores how competitive factors affect what approach (if any) a firm should use to learn about customer preferences. In particular, our analysis of a monopoly firm's choice of the learning approach reveals:

- Investing in a particular learning approach is worthwhile only if search costs are sufficiently large so that the savings from reduced search outweigh the costs of implementing the approach.
- *O-Learning* will be more preferred as the size of a firm's historical base increases (so the required information can be collected for a greater fraction of the population) and as the customers' effort requirement under *S-Learning* rises.
- A monopolist is biased away from *O-Learning*, i.e., there are cases where efficiency would have the firm invest in *O-Learning*, but the monopolist would either make no investment or invest in *S-Learning* instead.

Allowing for competition, we find that learning approaches already adopted by one's rivals have an important impact upon one's current decisions. Specifically, we find that:

- Total market revenue is highest if both firms invest in *O-Learning*.
- A firm facing a rival has a stronger incentive than a monopolist does to invest in *O-Learning*.
- A firm facing a rival who has invested in *S-Learning* has no incentive to invest in that same learning approach and also has less incentive to invest in *O-Learning*.
- A firm facing a rival who has invested in *O-Learning* may still benefit from investing in either learning processes (but is less likely to make either investment relative to a firm who faces a rival who does not personalize its products).

Finally, we find empirical support consistent with our theory. Specifically, we find that firms are more likely to share a similar degree *O-Learning* rather than a similar degree of *S-Learning* when they personalize their products for their customers.

Overall, our theoretical and empirical results suggest that *S-Learning* deters a rival's incentive to invest in *S-Learning* significantly more than *O-Learning* deters a rival's incentive to

invest in *O-Learning*. Consequently, there is an opportunity to behave strategically when deciding upon what learning approach to pursue (e.g., developing more efficient ways to utilize click-stream data or developing more efficient surveys to solicit customer preferences). By investing (or not investing) in a particular learning approach, a firm can induce a desired response from a rival. For example, in the “Follow Me” strategy that we identify, by investing in *O-Learning*, a firm may be able to induce its rival to also invest in *O-Learning*, and thus minimize the amount of price rivalry in the market. Furthermore, the choice of the learning approach could also be made for defensive or offensive purposes in order to secure a long-term competitive advantage (either because rivals will then forego any investment of their own – as in “Beat ‘em to the Punch” and “Beat ‘em to the Punch with a Helping Hand” – or because rivals will then adopt the less cost-efficient learning approach – as in “Take Your Pick”). Finally, not investing in any learning approach may also be a viable strategy. For example, as *O-Learning* becomes feasible for an industry, a firm may want to “play dumb”, e.g., signal that it is unprepared or unwilling to adopt this new approach in an effort to lure rivals into undergoing the costly investment instead (as in “Be My Guest”).

There are number of potential directions for future research. First, a greater degree of dynamics could be added to the current model. For example, a firm that employs *S-Learning* may be able to reduce the required effort (E) in future interactions, as less information needs to be solicited (Brady et al. 2000). Furthermore, if a firm can use both forms of learning then it may be possible to follow evolving strategies, e.g., attract new customers with an *S-Learning* approach and then utilize an *O-Learning* approach once sufficient history with those customers has been accumulated.

Second, many other issues ignored in the current model may affect the incentive to adopt *O-Learning* or *S-Learning* technologies. For instance, privacy issues may arise under *O-Learning*.

Thus, it would be important to consider whether customers are willing to have historical information collected and used. Implementation through a third party may be a potential solution (Hagel and Singer 1999). Also, personalized recommendations based on *O-Learning* and *S-Learning* are not likely to be perfectly accurate. Thus, an important consideration may be which learning approach allows the firm to better match the preferences of their customers. Finally, customers may prefer to be more involved in the process, rather than being simply told what their preferred product is (Huffman and Kahn 1998). In practice, it seems likely that customers need some assurance that the recommended product is indeed more preferred to the available alternatives.

The empirical study presented in this paper should be viewed as preliminary. As such, there are many directions left for future research. First, in the present survey, we are unable to track the dynamics of learning approaches. It would be very interesting to use field experiments and/or longitudinal data to track the impact of firms' investment decisions on subsequent prices, profit and/or learning approaches of rivals. Second, the survey collected a limited amount of data. It would be interesting for future research to look for commonalities and asymmetries that arise across industries. In particular, the model predicts that the cost of each type of learning approach, the effort required from customers under *S-Learning*, and a firm's historical market share will impact which, if any, learning approach a firm will adopt. A more extensive data set may be able to test for these hypothesized relationships and uncover new ones.

REFERENCES

- Ansari, A., C.F. Mela. 2003. E-customization. *J. of Mar. Res.* **40**(May) 131-145.
- Bertrand, M. and S. Mullainathan. 2001. Do people mean what they say? Implications for subjective survey data. *The Amer. Econ. Rev.* **91**(2) 67-72.
- Bohlmann, J.D., P.N. Golder, and D. Mitra. 2002. Deconstructing the pioneer's advantage: Examining vintage effects and consumer valuations of quality and variety. *Man. Sci.* **48**(9) 1175-1195.
- Brady, D., K. Kerwin, D. Welch, L. Lee, R. Hof. 2000. Customizing for the masses. *Bus. Week* March **20**(3673) 130B.
- Chen, Y., G. Iyer. 2002. Research note: Consumer addressability and customized pricing. *Marketing Sci.* **21**(2) 197-208.
- Chen Y., C. Narasimhan, J. Zhang 2001. Individual marketing with imperfect targetability. *Marketing Sci.* **20**(1) 23-41.
- Chuang, J. C.-I., M.A. Sirbu. 1999. Optimal bundling strategy for digital information goods: network delivery of articles and subscriptions. *Information Economics & Policy* **11**(2) 147-176.
- Court, A. 2005. Hotels try to make your stay personal. *USA Today* October 25, Money 09b.
- Dana, Jr. J.D. 1999. Equilibrium price dispersion under demand uncertainty: The roles of costly capacity and market structure," *The RAND J. of Econ.* **30**(4) 632-660.
- DeFigueiredo, J.M., M.K. Kyle. 2006. Surviving the gales of creative destruction: The determinants of product turnover. *Strategic Man. J.* **27**(3) 241-264.
- Dellaert, B.G.C., S. Stremersch. 2005. Marketing mass-customized products: Striking a balance between utility and complexity. *J. of Mar. Res.* **42** 219-227.
- Dewan, R., B. Jing, A. Seidmann. 2003. Product customization and price competition on the Internet. *Man. Sci.* **49**(8) 1055-1070.
- Fruchter, G.E., Z.J. Zhang. 2004. Dynamic targeted promotions: A customer retention and acquisition perspective. *J. of Ser. Res.* **7**(1) 3-19.
- Fudenberg, D., J.M. Villas-Boas. 2006. Behavior-Based Price Discrimination and Customer Recognition. *Handbook on Economics and Information Systems*. Volume 1 (T.J. Hendershott, Ed.), Elsevier Science 377-436..

- Gatignon, H., J.-M. Xuereb. 1997. Strategic orientation of the firm and new product performance. *J. of Mar. Res.* **34**(1) 77-90.
- Gilmore, J.H. II, B.J. Pine. 1997. The four faces of mass customization. *Har. Bus. Rev.* **75**(1) 91-102.
- Golder, P.N. 2000. Historical method in marketing research with new evidence on long-term market share stability. *J. of Mar. Res.* **37**(2) 156-172.
- Green, P.E., V. Srinivasan. 1990. Conjoint analysis in marketing: New development with implications for research and practice. *J. of Mar.* **54**(4) 3-19.
- Hagel, J., M. Singer. 1999. Net worth: shaping markets when consumers make the rules. Harvard Business School Press, Boston.
- Hall R. 1983. *Zero Inventories*. Dow-Jones Irwin, Homewood, IL.
- Heide, Jan B. and George John 1990. Alliances in industrial purchasing: The determinants of Joint Action in Buyer-Supplier Relationships. *Journal of Marketing Research* 27 (4), 24-26..
- Huffman, C., B. Kahn. 1998. Variety for sale: Mass customization or mass confusion. *J. of Ret.* **74**(4) 491-513.
- Johnson, E.J., W.W. Moe, P.S. Fader, S. Bellman, G.L. Lohse. 2004. On the depth and dynamics of online search behavior. *Management Sci.* **50**(3) 299-308.
- Lee, H.L., C.S. Tang. 1998. Variability reduction through operations reversal. *Man. Sci.* **44**(2) 162-172.
- Lippman, S.A, J.J. McCall. 1976. The economic of job search. *Econ. Inquiry* **14**(2) 155-189.
- McCall, J.C. 1965. The economics of information and optimal stopping rules. *J. of Bus.* (July) 300-317.
- McGahan, A.M., P. Ghemawat. 1994. Competition to retain customers. *Mar. Sci.* **13**(2) 165-176.
- Moser, K. 2006. How they know what you like before you do. *Chr. Sci Mon.* **98**(57, Feb. 16) 13-16.
- Murthi, B.P.S.; S. Sarkar. 2003. The role of the management sciences in research on personalization. *Man. Sci.* **49**(10) 1344-1362.
- Peppers, D. M. Rogers. 1997. *Enterprise One to One: Tools for Competing in the Interactive Age*. Doubleday, New York.
- Pine. B.J. II; A.C. Boynton. 1993. Making mass customization work. *Harvard Bus. Rev.* **71**(5) 108-118.

- Raghu, T.S., P.K. Kannan, H.R. Rao, A.B. Whinston. 2001. Dynamic profiling of consumers for customized offerings over the Internet. A model and analysis. *Decision Support Sys.* **32**(2) 117-134.
- Randall, T., C. Terwiesch, K.T. Ulrich. 2005. Principles for user design of customized products. *Cal. Man. Rev.* **47**(4) 68-85.
- Randall, T., C. Terwiesch, K.T. Ulrich. 2007. User design of customized products. *Mar Sci.* **26**(2) 268-280.
- Reinartz, W., J.S. Thomas, V. Kumar. 2005. Balancing acquisition and retention resources to maximize customer profitability. *J. of Mar.* **69**(January) 63-79.
- Resnick, P., H.R. Varian. 1997. Recommender systems. *Comm. ACM* **40**(3) 56-58.
- Rust, R.T. P.K. Kannan, N. Peng. 2002. The customer economics of Internet privacy. *J. Acad. of Mar. Sci.* **30**(4) 455-464.
- Sellandurai, R.S. 2004. Mass customization in operations management: Oxymoron or reality. *Omega* **32**(4) 295-300.
- Shaffer, G., Z.J. Zhang. 1995. Competitive coupon targeting. *Mar. Sci.* **14** 395-416.
- Shaffer, G., Z.J. Zhang. 2002. Competitive one-to-one promotions. *Man. Sci.* **48**(9) 1143-1160.
- Shugan, S.M. 2002. Editorial: Marketing Science, Models, Monopoly Models, and Why We Need Them. *Mar. Sci.* **21**(3) 223-228.
- Shugan, S.M. 2007. It's the Findings, Stupid, Not the Assumptions. *Mar. Sci.* **26**(4) 449-459.
- Silvestro, R., L. Fitzgerald, R. Johnston, C. Voss. 1992. Towards a classification of service processes. *Inter. J. of Ser. Industry Man.* **3**(3) 62-75.
- Simonson, I. 2005. Determinants of customers' responses to customized offers: Conceptual framework and research propositions. *J. of Marketing* **69** 32-45.
- Sismeiro, C., R.E. Bucklin. 2004. Modeling purchase behavior at an E-commerce web site: A task-completion approach. *J. of Mar. Research* **41**(3) 306-323.
- Srikumar, K., B. Bhasker. 2004. Personalized product selection in Internet business. *J. of Elec. Com. Res.* **5**(4) 216-227.
- Swaminathan, J.M., S.R. Tayer. 1998. Managing broader product lines through delayed differentiation using vanilla boxes. *Man. Sci.* **44**(12.2) S161-S172.
- Syam, N.B., B. Dellaert. 2002. Consumer-producer interaction: A strategic analysis of the market for customized products. *Review of Mar. Sci.* **1**(1).

- Syam, N. B., R. Ruan, J.D. Hess. 2005. Customized products: A competitive analysis. *Mar. Sci.* **24**(4) 569-584.
- Syam, N.B., J.D. Hess. 2006. Acquisition versus retention: Competitive customer relationship management. University of Houston, March 21.
- Syam, N.B., N. Kumar. 2006. On customized goods, standard goods and competition. *Mar. Sci.* **25**(5) 525-537.
- Toubia, O., D.I. Simester, J.R. Hauser, E. Dahan. 2003. Fast polyhedral adaptive conjoint estimation. *Mar. Sci.* 22(3) 273-303
- Ulrich, K.T., S. Pearson. 1998. Assessing the importance of design through product archaeology. *Man. Sci.* **44**(3) 352-369.
- Vargo, S.L., R.F. Lusch 2004. Evolving to a new dominant logic for marketing. *J. of Mar.* 68(1) 1-17.
- Voss, G.B., Z.G. Voss. 2000. Strategic orientation and firm performance in an artistic environment. *J. of Mar.* **64**(1) 67-83.
- Zipkin, P. 2001. The limits of mass customization. *MIT Sloan Man. Rev.* **42**(3) 81-87.
- Zucker, L.G., M.R. Darby, J.S. Armstrong. 2002. Commercializing knowledge: University science, knowledge capture, and firm performance in biotechnology. *Man. Sci.* **48**(1) 138-153.

TABLE 1
DESCRIPTIVE STATISTICS OF SURVEY DATA

<i>Variable</i> ¹	<i>Description</i>	<i>N</i>	<i>Mean</i>	<i>StDev</i>	<i>Max</i>	<i>Min</i>
<i>C_{foc}</i>	Focal supplier's level of personalization learning	216	3.28	1.06	5	1
<i>C_{alt}</i>	Alternative supplier's level of personalization learning	182	2.91	1.04	5	1
<i>O_{foc}</i>	Focal supplier's <i>O-Learning</i>	216	1.13	0.64	4	0
<i>O_{alt}</i>	Alternative supplier's <i>O-Learning</i>	182	1.16	0.46	2.50	0
<i>S_{foc}</i>	Focal supplier's <i>S-Learning</i>	216	2.16	1.21	5	0.20
<i>S_{alt}</i>	Alternative supplier's <i>S-Learning</i>	182	1.75	1.10	5	0.20

¹All variables are computed on a scale of 1 to 5.

TABLE 2**LEARNING APPROACH DIFFERENCES WITH RIVAL
FOR S-LEARNING AND O-LEARNING FOCAL SUPPLIERS**

Classification Criterion	Average ΔS_S	Average ΔO_O	$\Delta S_S - \Delta O_O$	Average ΔS_O	Average ΔO_S	$\Delta S_O - \Delta O_S$
Top 25	1.40	0.62	0.78***	0.38	0.55	-0.17 ^{NS}
Top 50	1.20	0.55	0.65***	0.52	0.45	0.07 ^{NS}
Top 75	1.09	0.45	0.64***	0.54	0.48	0.06 ^{NS}
Top 100	0.99	0.39	0.60***	0.50	0.47	0.03 ^{NS}

*** p<0.001

^{NS} p>0.1

TABLE 3

SIX STRATEGIES FOR INVESTING IN LEARNING APPROACHES

<i>Strategy</i>	<i>Firm B's Best Response</i> $(D_B(\emptyset), D_B(SL), D_B(OL))$	<i>Firm A's Action</i>
#1: Take Your Pick	(SL, OL, SL)	SL if T_{OL} is high, OL if T_{OL} is low
#2: Beat 'em to the Punch	(SL, \emptyset , \emptyset or SL)	SL
#3: Beat 'em to the Punch with a Helping Hand	(SL, \emptyset , \emptyset)	OL
#4: Be My Guest	(OL, \emptyset or OL, \emptyset or OL)	\emptyset
#5: Ride My Coattails	(OL, \emptyset , \emptyset)	OL
#6: Follow Me	(SL, \emptyset , OL)	OL

FIGURE 1

ELEMENTS OF THE PERSONALIZATION PROCESS AND REPRESENTATIVE RESEARCH

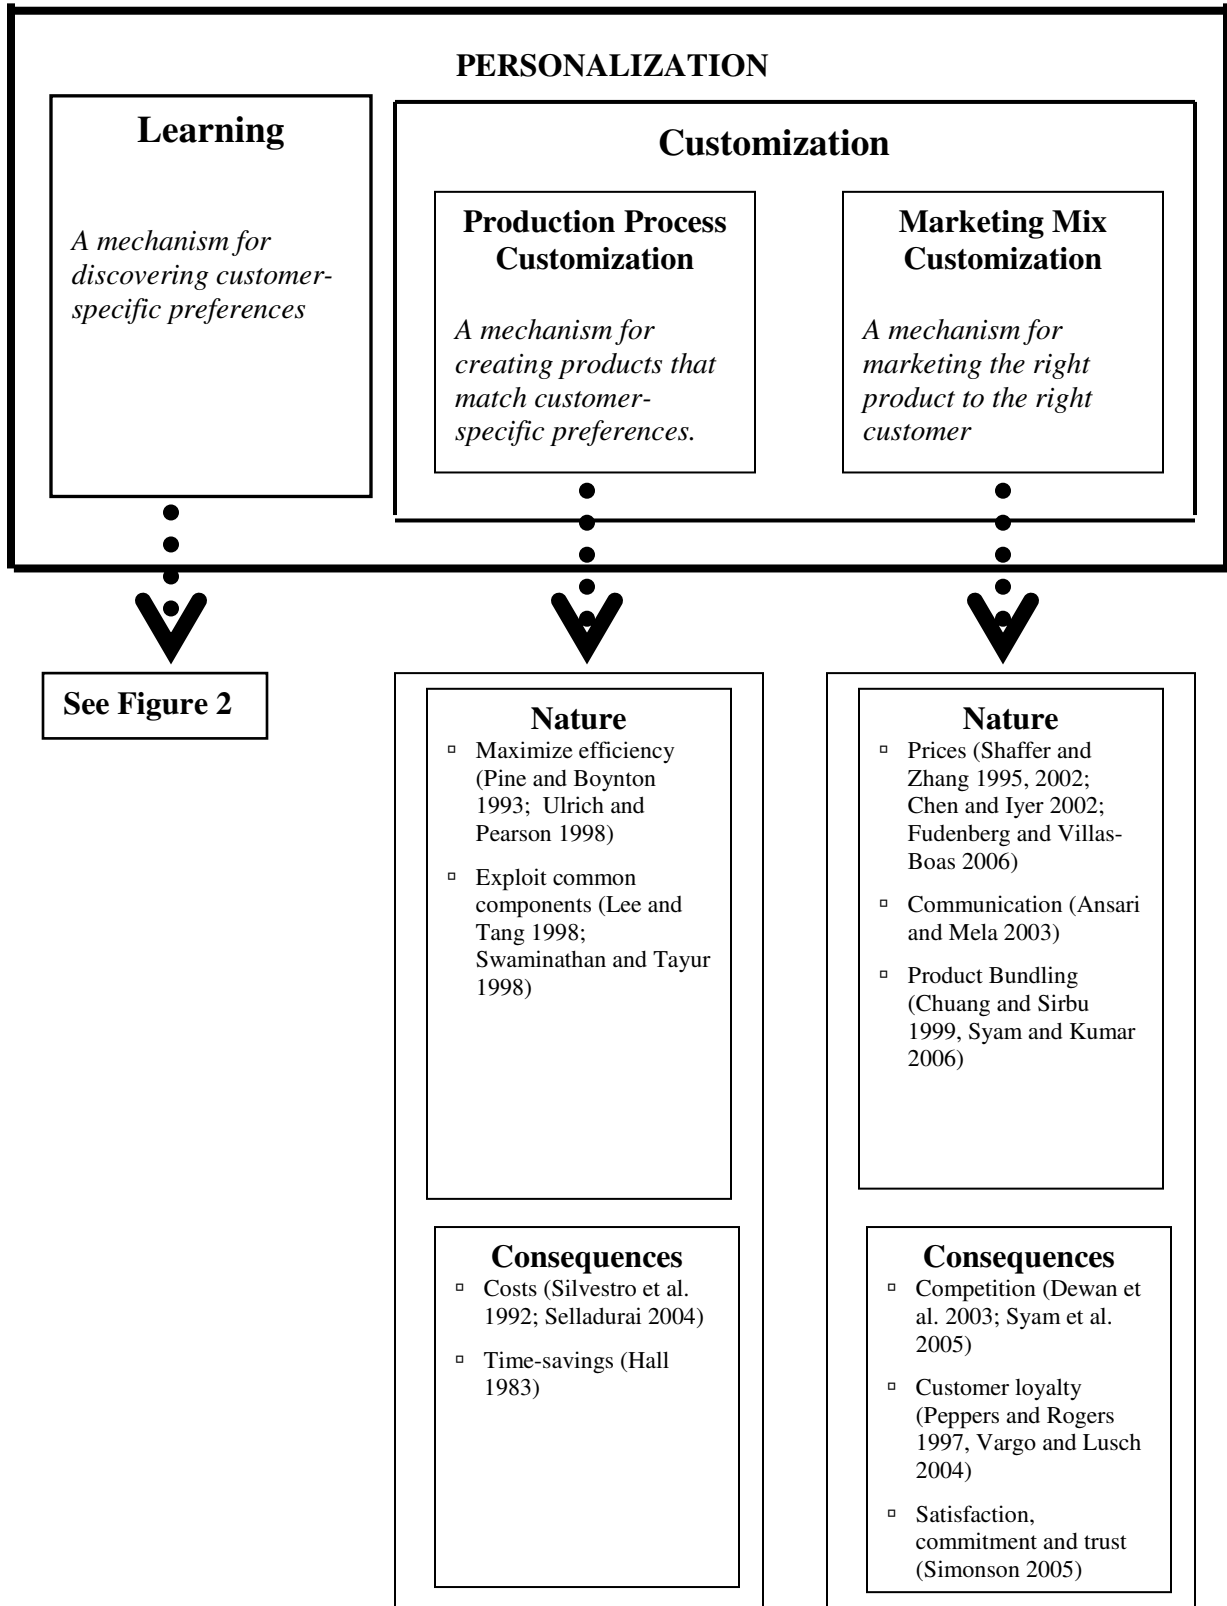


FIGURE 2

LITERATURE RELATED TO THE LEARNING STAGE OF PERSONALIZATION

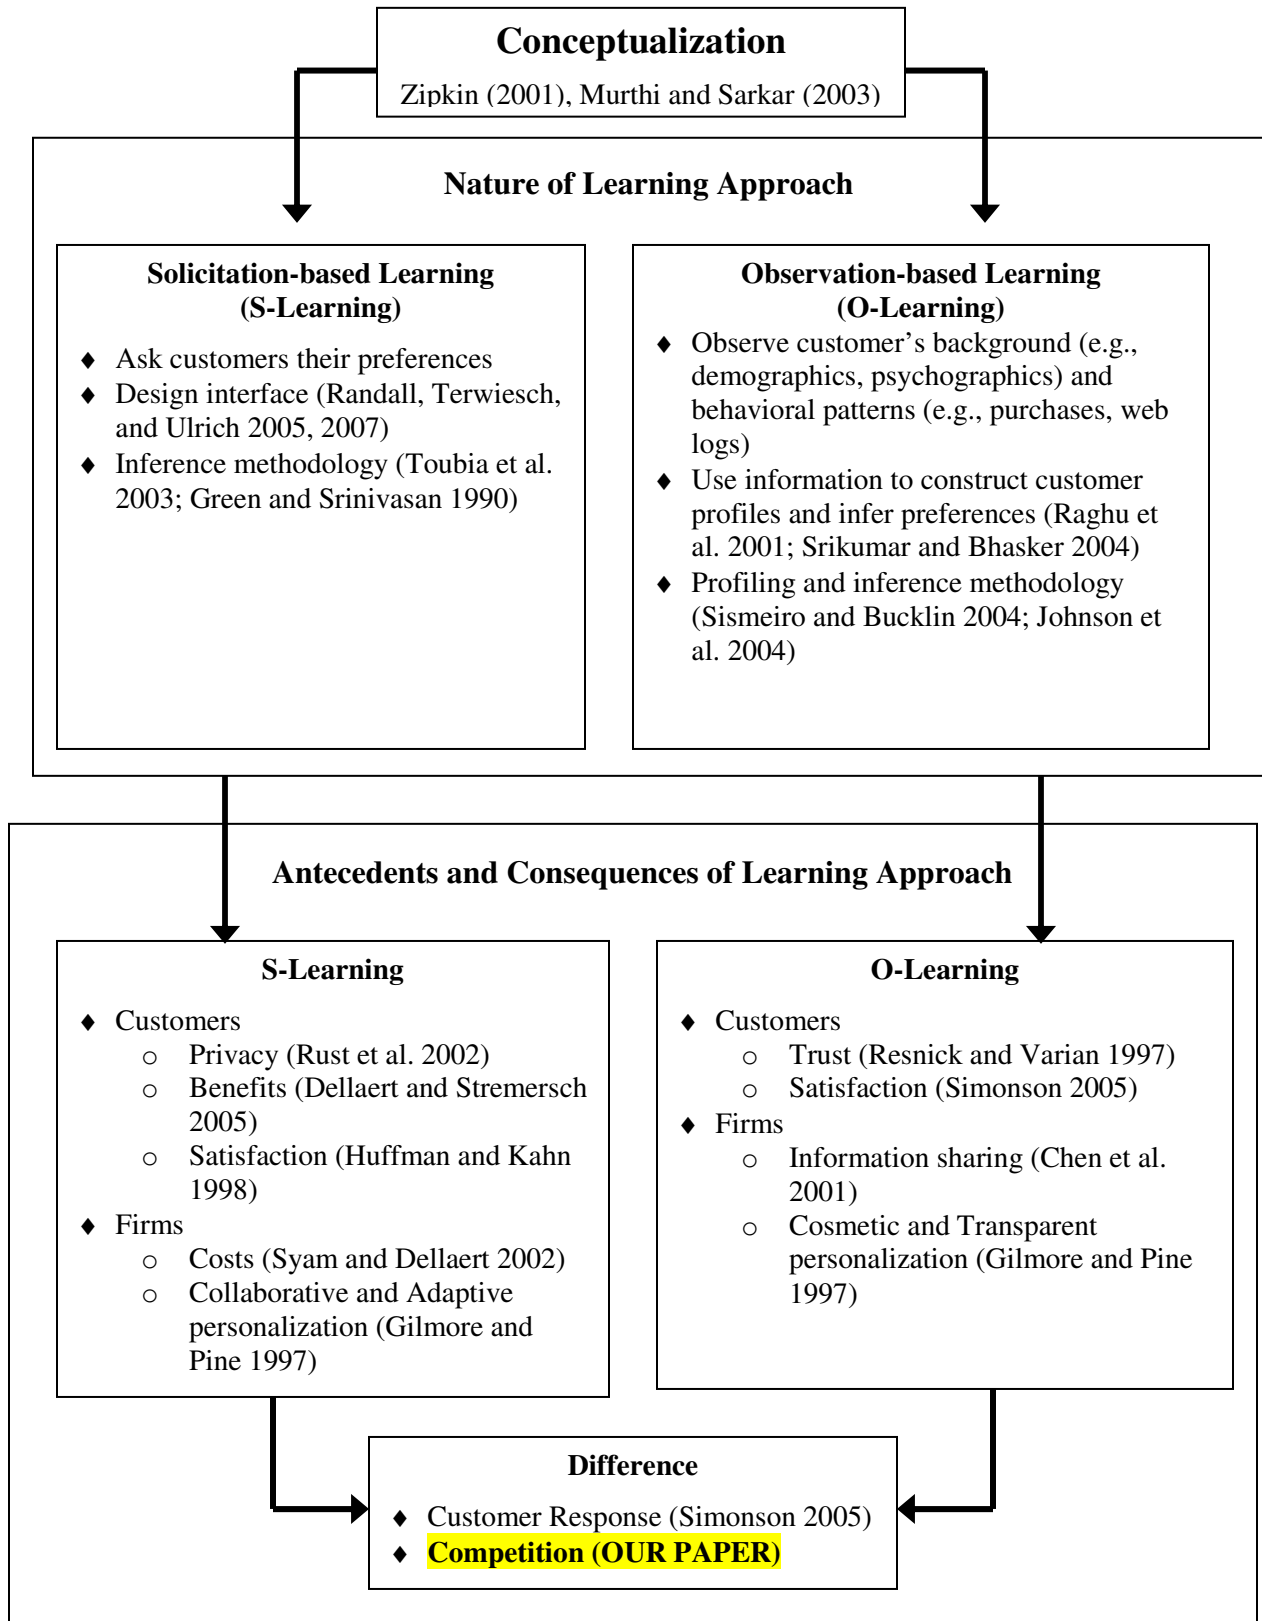
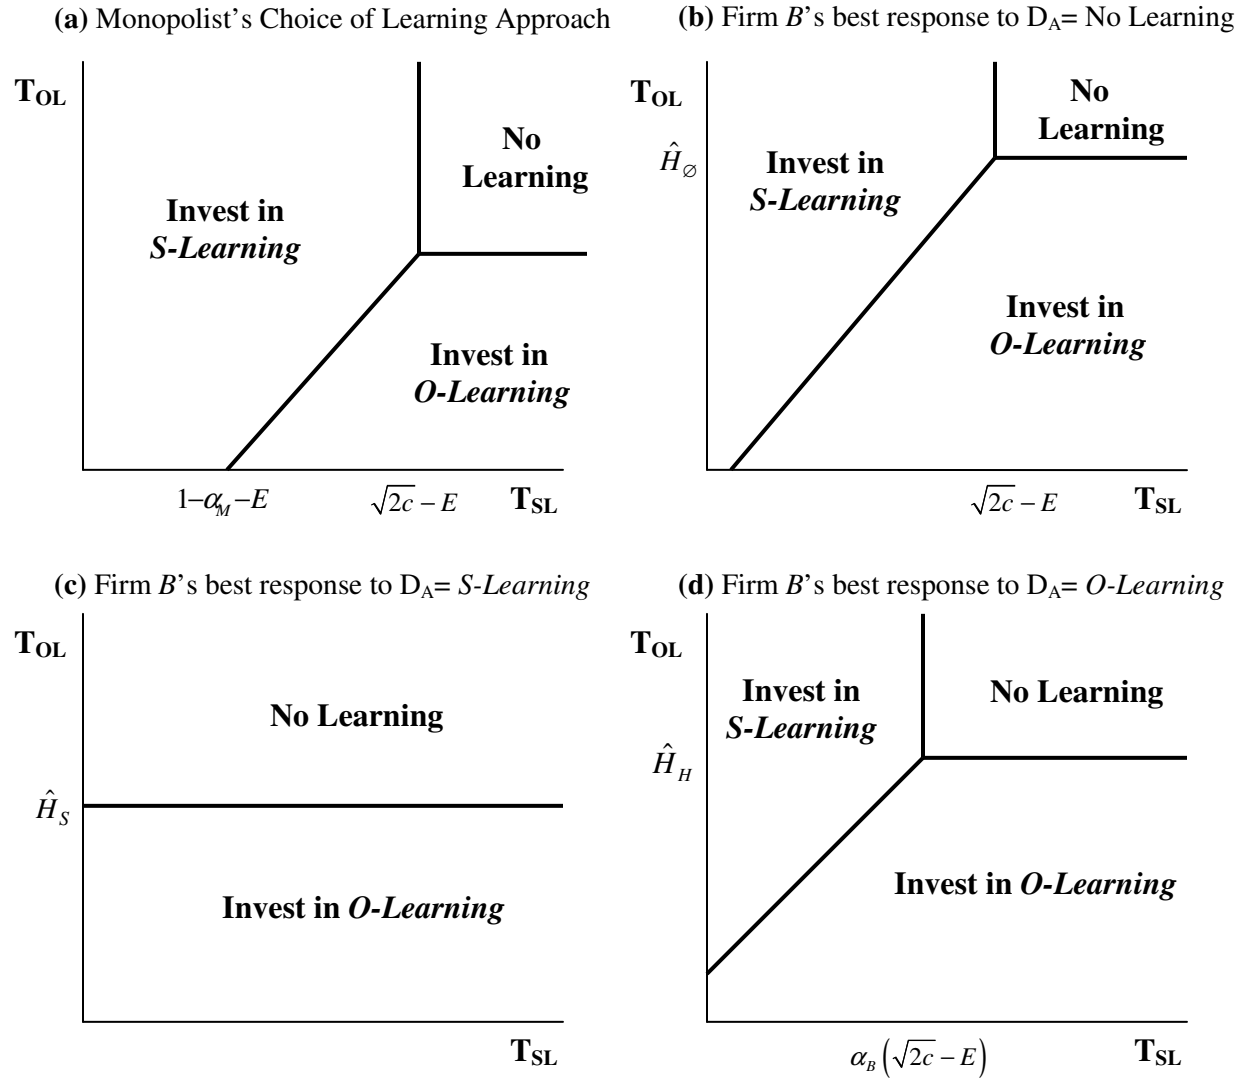


FIGURE 3

OPTIMAL LEARNING APPROACH OF A FIRM



Appendix

Proof of Proposition 1

The monopolist chooses the No Learning option if $\Pi_M^\varnothing \geq \Pi_M^{SL}$ and $\Pi_M^\varnothing \geq \Pi_M^{OL}$, where Π_M^\varnothing , Π_M^{SL} , and Π_M^{OL} are given in equations (3), (5), and (6), respectively. These conditions simplify to the expressions given in Proposition 1(a). Similarly, the monopolist chooses the *S-Learning* option if $\Pi_M^{SL} > \Pi_M^\varnothing$ and $\Pi_M^{SL} \geq \Pi_M^{OL}$, and the monopolist chooses the *O-Learning* option if $\Pi_M^{OL} > \Pi_M^\varnothing$ and $\Pi_M^{OL} > \Pi_M^{SL}$. The simplified expressions for these conditions are reported in parts (b) and (c) of Proposition 1.

Competitive Response—Derivation of the Pricing Equilibrium

The pricing equilibria are reported in Table A1 and are derived below.

Table A1 The Pricing Equilibrium and Expected Profit with Competition

		Firm A		
		<i>No Learning</i>	<i>S-Learning</i>	<i>O-Learning</i>
Firm B	<i>No Learning</i>	$p_A = p_B = 0$ $\Pi_A = \Pi_B = 0$	$p_A = \sqrt{2c} - E$ $p_B = 0$ $\Pi_A = \sqrt{2c} - E - T_{SL}$ $\Pi_B = 0$	Mixed Strategy Equilibrium in prices $\Pi_A = (1 - \alpha_B)(1 - (1 - \alpha_B)[1 - \sqrt{2c}]) - T_{OL}$ $\Pi_B = \alpha_B(1 - \sqrt{2c})$
	<i>S-Learning</i>	$p_A = 0$; $p_B = \sqrt{2c} - E$ $\Pi_A = 0$; $\Pi_B = \sqrt{2c} - E - T_{SL}$	$p_A = p_B = 0$ $\Pi_A = \Pi_B = -T_{SL}$	Mixed Strategy Equilibrium in prices $\Pi_A = (1 - \alpha_B)(1 - (1 - \alpha_B)(1 - E)) - T_{OL}$ $\Pi_B = \alpha_B(1 - E) - T_{SL}$
	<i>O-Learning</i>	Mixed Strategy Equilibrium in prices $\Pi_A = (1 - \alpha_B)(1 - \sqrt{2c})$ $\Pi_B = \alpha_B(1 - \alpha_B[1 - \sqrt{2c}]) - T_{OL}$	Mixed Strategy Equilibrium in prices $\Pi_A = (1 - \alpha_B)(1 - E) - T_{SL}$ $\Pi_B = \alpha_B(1 - \alpha_B(1 - E)) - T_{OL}$	$p_A = p_B = 1$ $\Pi_A = 1 - \alpha_B - T_{OL}$ $\Pi_B = \alpha_B - T_{OL}$

$D_A = D_B = \varnothing$: A customer earns a surplus of $1 - \sqrt{2c} - p_A$ by purchasing from firm A and a surplus of $1 - \sqrt{2c} - p_B$ by purchasing from firm B. Assuming ties split demand equally, the demand faced by A is:

$$D(p_A) = \begin{cases} 0 & \text{if } p_A > p_B \text{ or } p_A > 1 - \sqrt{2c} \\ 1/2 & \text{if } p_A = p_B \text{ and } p_A \leq 1 - \sqrt{2c} \\ 1 & \text{if } p_A < p_B \text{ and } p_A \leq 1 - \sqrt{2c} \end{cases} \quad (\text{A1})$$

Firm A's profit is $p_A D(p_A)$. Thus, A's best response is $p_A(p_B) = p_B - \varepsilon$ where ε is an arbitrary small number. As is well-known, the equilibrium of this duopoly game converges to $p_A = p_B = 0$ and each firm earns zero profit.

$D_A = D_B = SL$: A customer earns of surplus of $1 - E - p_A$ by purchasing from firm A and a surplus of $1 - E - p_B$ by purchasing from firm B. Thus, firm A faces demand:

$$D(p_A) = \begin{cases} 0 & \text{if } p_A > p_B \text{ or } p_A > 1 - E \\ \frac{1}{2} & \text{if } p_A = p_B \text{ and } p_A \leq 1 - E \\ 1 & \text{if } p_A < p_B \text{ and } p_A \leq 1 - E \end{cases} \quad (\text{A2})$$

A's best response is $p_A(p_B) = p_B - \varepsilon$. Again, the equilibrium of this duopoly game converges to $p_A = p_B = 0$. Each firm earns a profit of $-T_{SL}$.

$D_A = D_B = OL$: Customers for which firm A has historical information, a segment of size $1 - \alpha_B$, earn a surplus of $1 - p_A$ if they purchase from firm A and a surplus of $1 - \sqrt{2c} - p_B$ if they purchase from B. For the remaining segment (of size α_B) surplus equals $1 - \sqrt{2c} - p_A$ if they purchase from A and $1 - p_B$ if they purchase from B. In the case of ties, we assume the customer stays with the firm with which they have had a previous relationship. Firm A' demand is:

$$D(p_A) = \begin{cases} 0 & \text{if } p_A > p_B + \sqrt{2c} \text{ or } p_A > 1 \\ 1 - \alpha_B & \text{if } p_B - \sqrt{2c} \leq p_A \leq p_B + \sqrt{2c} \\ 1 & \text{if } p_A < p_B - \sqrt{2c} \end{cases} \quad (\text{A3})$$

The pure-strategy equilibrium is $p_A = p_B = 1$ if condition (7) is met. To see this, note that firm A earns a profit of $1 - \alpha_B$ and firm B earns a profit of α_B in this pricing outcome. For this to be an equilibrium neither firm can earn higher profit by deviating from these prices. From (A3), it is obvious that the best deviation for each firm would be to the price of $1 - \sqrt{2c} - \varepsilon$, which would result in a profit of $1 - \sqrt{2c} - \varepsilon$. Such a deviation is not profitable for firm A if $1 - \alpha_B \geq 1 - \sqrt{2c} - \varepsilon$ and a deviation is not profitable for firm B as long as $\alpha_B \geq 1 - \sqrt{2c} - \varepsilon$.

$D_A = SL, D_B = \emptyset$: Customers obtain a surplus of $1 - E - p_A$ by purchasing from firm A and a surplus of $1 - \sqrt{2c} - p_B$ by purchasing from B. Assuming ties are settled by buying from the firm with the personalized offer, firm A's demand is:

$$D(p_A) = \begin{cases} 0 & \text{if } p_A > p_B + \sqrt{2c} - E \text{ or } p_A > 1 \\ 1 & \text{if } p_A \leq p_B + \sqrt{2c} - E \end{cases} \quad (\text{A4})$$

Each firm maximizes profit by pricing just low enough to attract the entire market, i.e., A's best response is $p_A(p_B) = p_B + \sqrt{2c} - E$, B's best response is $p_B(p_A) = p_A - \sqrt{2c} + E$, with the constraint that firms never choose negative prices. This constraint is first to bind firm B. Thus, the equilibrium price is for firm B to choose $p_B = 0$, to which firm A responds with $p_A = \sqrt{2c} - E$. All sales go to firm A. Thus,

$$\Pi_A = \sqrt{2c} - E - T_{SL} \text{ and } \Pi_B = 0.$$

$D_A = OL, D_B = \emptyset$: The customers for which firm A has historical information, a segment of size $1 - \alpha_B$, earn a surplus of $1 - p_A$ if they purchase from firm A and a surplus of $1 - \sqrt{2c} - p_B$ if they purchase from B. For the remaining segment of size α_B , surplus equals $1 - \sqrt{2c} - p_A$ if they purchase from A and $1 - \sqrt{2c} - p_B$ if they purchase from B. In the case of ties, we assume the customer stays with the firm with which they have had a previous relationship. Thus, demand faced by firm A equals:

$$D(p_A) = \begin{cases} 0 & \text{if } p_A > p_B + \sqrt{2c} \text{ or } p_A > 1 \\ 1 - \alpha_B & \text{if } p_B \leq p_A \leq p_B + \sqrt{2c} \\ 1 & \text{if } p_A < p_B \end{cases} \quad (\text{A5})$$

On the other hand, the demand faced by firm B is:

$$D(p_B) = \begin{cases} 0 & \text{if } p_B > p_A \text{ or } p_B > 1 - \sqrt{2c} \end{cases} \quad (\text{A6})$$

$$D(p_B) = \begin{cases} \alpha_B & \text{if } p_A - \sqrt{2c} \leq p_B \leq p_A \\ 1 & \text{if } p_B < p_A - \sqrt{2c} \end{cases}$$

A pure-strategy equilibrium does not exist in this scenario. To see this, note that firm A 's best response to p_B is either $p_B + \sqrt{2c} - \varepsilon$ or $p_B - \varepsilon$. In either case, firm B could cause a discrete jump in its demand by cutting its price by 2ε . One might expect an equilibrium at $p_B = 0$. Such a price would induce A to choose $p_A = \sqrt{2c} - \varepsilon$. But, firm B could make positive profits by setting $p_B = \sqrt{2c} - 2\varepsilon$, thus demonstrating that $p_B = 0$ is not an equilibrium.

Thus, we must look for a mixed-strategy equilibrium. Suppose p_A is drawn from the interval $[p_L^A, p_H^A]$ and continuously distributed according to the density function $f(p)$, where $p_H^A > p_L^A$. Firm B 's prices are distributed according to the density function $g(p)$ for the interval $[p_L^B, p_H^B]$ where $p_H^B > p_L^B$. Furthermore, we allow for the possibility that there is a mass point at $\overline{p_B}$. Finally, assume that the following condition holds:

$$p_L^A \geq 1 - \sqrt{2c} \quad (\text{A7})$$

This condition drastically reduces the complexity of the problem. After calculating the exact specification of p_L^A , we will be able to show that condition (7) guarantees that (A7) is satisfied.

If (A7) is met, firm A will never sale to customers for whom it does not have historical information. Furthermore, personalized recommendations lead to sales only if $p_A \leq p_B + \sqrt{2c}$. One immediate result is that:

$$p_L^B = p_L^A - \sqrt{2c} \quad (\text{A8})$$

It is not optimal for Firm B to put any weight on $p_B < p_L^A - \sqrt{2c}$ since all prices strictly below $p_L^A - \sqrt{2c}$ result in one unit of sales (with a probability of one). Thus, profit is strictly increasing in p_B until this threshold is reached. Furthermore, it is not optimal for Firm A to put any weight on $p_A < p_L^B + \sqrt{2c}$ since all prices strictly below $p_L^B + \sqrt{2c}$ result in $1 - \alpha_B$ of sales.

The expected profit for each firm is given by the following two equations:

$$E[\Pi_A] = p_A (1 - \alpha_B) (1 - G(p_A - \sqrt{2c})) - T_{OL} \quad (\text{A9})$$

$$E[\Pi_B] = p_B \left[\alpha_B + (1 - \alpha_B) (1 - F(p_B + \sqrt{2c})) \right] \quad (\text{A10})$$

where $F(\cdot)$ and $G(\cdot)$ are the cumulative density functions of $f(\cdot)$ and $g(\cdot)$ respectively. All prices in the support must yield the same expected profit for a mixed strategy equilibrium to exist. In particular, we know that $p_A = p_L^A$ and $p_A = p_H^A$ must yield the same profit:

$p_L^A (1 - \alpha_B) (1 - G(p_L^A - \sqrt{2c})) - T_{OL} = p_L^A (1 - \alpha_B) (1 - G(p_L^A - \sqrt{2c})) - T_{OL}$. Furthermore, from (A8) we know that $G(p_L^A - \sqrt{2c}) = 0$. Substituting this into the preceding equality and canceling:

$$p_L^A = p_L^B (1 - G(p_L^B - \sqrt{2c})) \quad (\text{A11})$$

Note that this implies $G(p_L^B - \sqrt{2c}) < 1$. In other words, there must be a mass point above the interval $[p_L^B, p_H^B]$, denoted by $\overline{p_B}$, for which firm B does not sell to A 's historical customers with any positive probability. To maximize profit, given that sales will only be made to A 's non-historical customers, firm B must choose $\overline{p_B} = 1 - \sqrt{2c}$. This earns a profit of:

$$\Pi_B = \alpha_B (1 - \sqrt{2c}) \quad (\text{A12})$$

All other p_B in the relevant range must earn the same profit. Specifically, at $p_B = p_L^B$, firm B will make one unit of sales and will earn the profit given in (A12). Thus, we have:

$$p_L^B = \alpha_B (1 - \sqrt{2c}) \quad (\text{A13})$$

Combining (A8) and (A13) we see that $p_L^A = \alpha_B (1 - \sqrt{2c}) + \sqrt{2c}$. Substituting this equality into equation (A9) and recalling that $G(p_L^A - \sqrt{2c}) = 0$, we see that firm A earns an expected profit of:

$$E[\Pi_A] = (1 - \alpha_B) \left(1 - (1 - \alpha_B) \left[1 - \sqrt{2c} \right] \right) - T_{OL} \quad (\text{A14})$$

Finally, recall that throughout this derivation we assumed (A7) was satisfied. Substituting in our result that $p_L^A = \alpha_B (1 - \sqrt{2c}) + \sqrt{2c}$, (A7) reduces to the condition that $c \geq \frac{(\alpha_B)^2}{2(1 + \alpha_B)^2}$. It is easy to see that this is a weaker restriction than condition (7) since $\alpha_B > 0$. Thus, condition (A7) is met over the entire relevant range of parameters.

$D_A = OL, D_B = SL$: This scenario is very similar to the preceding case. In effect, the expected search cost associated with purchasing from firm B has fallen from $\sqrt{2c}$ to E . Thus, the analysis proceeds exactly as above except that every $\sqrt{2c}$ that appears in equations (A5)-(A14) and the associated text should be replaced with an E . However, an important point to note is that condition (A7) now becomes $p_L^A \geq 1 - E$ where $p_L^A = \alpha_B (1 - E) + E$. Thus, the analysis assumes condition (8) is met.

Proof of Proposition 2 (Competitive Response)

In this section, we consider the second stage of the game: firm B 's investment decision given D_A . Firm B chooses D_B to maximize its profit, anticipating what the pricing equilibrium will be given D_A and D_B . Table A1 provides the expected profit, net of investment costs, for each possibility. Figures 3(b)-(d) illustrate the learning approach that maximizes firm B 's profit given firm A 's choice.

Derivation of Figure 3(b): If $D_A = \emptyset$, firm B maximizes its profit by choosing *No Learning* if both conditions (A15) and (A16) are met:

$$T_{SL} \geq \sqrt{2c} - E \quad (\text{A15})$$

$$T_{OL} \geq \alpha_B \left(1 - \alpha_B \left[1 - \sqrt{2c} \right] \right) \equiv \hat{H}_{\emptyset} \quad (\text{A16})$$

S-Learning is most profitable if both conditions (A17) and (A18) are met:

$$T_{SL} < \sqrt{2c} - E \quad (\text{A17})$$

$$T_{SL} \leq T_{OL} + \sqrt{2c} - E - \alpha_B \left(1 - \alpha_B \left[1 - \sqrt{2c} \right] \right) \quad (\text{A18})$$

O-Learning is most profitable if both conditions (A19) and (A20) are met:

$$T_{OL} < \hat{H}_{\emptyset} \quad (\text{A19})$$

$$T_{SL} > T_{OL} + \sqrt{2c} - E - \alpha_B \left(1 - \alpha_B \left[1 - \sqrt{2c} \right] \right) \quad (\text{A20})$$

Now we verify Proposition 2(a), i.e., holding the availability of historical information constant, a firm that faces a rival is more likely to invest in *O-Learning* than it would in the absence of competition. To see this, note that $\hat{H}_{\emptyset} > \hat{H}_M$ (as defined in Proposition 1(a) and equation (A16), respectively) when $\alpha_M = \alpha_B$.

Derivation of Figure 3(c)—Proposition 2(b): If $D_A = SL$, it is never optimal for firm B to also invest in *S-Learning*. Instead, *O-Learning* is most profitable for B only if condition (A21) is met:

$$T_{OL} < \alpha_B \left(1 - \alpha_B \left[1 - E \right] \right) \equiv \hat{H}_S \quad (\text{A21})$$

If this inequality is violated, it is optimal for firm B not to invest in either learning approach.

To see that A 's investment in S -Learning also discourages B from investing in O -Learning, we show that $\hat{H}_S < \hat{H}_\emptyset$:

$$\hat{H}_\emptyset - \hat{H}_S = \alpha_B^2 (\sqrt{2c} - E) > 0 \quad (\text{A22})$$

Derivation of Figure 3(d)—Proposition 2(c): If $D_A = \text{OL}$, it is optimal for firm B to choose *No Learning* if both conditions (A23) and (A24) are met:

$$T_{SL} \geq \alpha_B (\sqrt{2c} - E) \quad (\text{A23})$$

$$T_{OL} \geq \alpha_B \sqrt{2c} \equiv \hat{H}_H \quad (\text{A24})$$

S -Learning is most profitable if both conditions (A25) and (A26) are met:

$$T_{SL} < \alpha_B (\sqrt{2c} - E) \quad (\text{A25})$$

$$T_{SL} \leq T_{OL} - \alpha_B E \quad (\text{A26})$$

O -Learning is most profitable if both conditions (A27) and (A28) are met:

$$T_{OL} < \hat{H}_H \quad (\text{A27})$$

$$T_{SL} > T_{OL} - \alpha_B E \quad (\text{A28})$$

To see that the size of the “Invest in S -Learning” region is smaller in Figure 3(d) than in Figure 3(b), compare conditions (A17) and (A25). To see that S -Learning by a rival also expands the *No Learning* region downward, we calculate: $\hat{H}_\emptyset - \hat{H}_H = \alpha_B (1 - \alpha_B) (1 - \sqrt{2c}) > 0$.

Proof of Proposition 3 (Multiple Learning Technologies and Price Customization)

Here, we allow firms to invest in both technologies if they so choose and also to customize prices, i.e., charge different prices depending on whether a customer is new to the firm or is part of its existing database. To prove the results in Proposition 3, we only need to examine the cases where a rival (Firm A) uses S -Learning and when a rival uses both S -Learning and O -Learning. Table A2 gives the equilibrium prices and expected profit for these two cases, where p_A^N is the price charged by Firm A to new customers (i.e., customers that are in segment α_B), p_A^O is the price charged by Firm A to its old customers (i.e., customers that are in segment α_A), p_B^N is the price charged by Firm B to new customers (i.e., customers that are in segment α_A), and p_B^O is the price charged by Firm B to its old customers (i.e., customers that are in segment α_B).

Table A2 The Pricing Equilibrium and Expected Profit with Multiple Learning Technologies

		Firm A	
		<i>S-Learning</i>	<i>S-Learning and O-Learning</i>
Firm B	No Learning	$p_A^N = p_A^O = \sqrt{2c} - E$ $p_B^N = p_B^O = 0$ $\Pi_A = \sqrt{2c} - E - T_{SL}$ $\Pi_B = 0$	$p_A^N = \sqrt{2c} - E$ $p_A^O = \sqrt{2c}$ $p_B^N = p_B^O = 0$ $\Pi_A = (1 - \alpha_B)\sqrt{2c} + \alpha_B(\sqrt{2c} - E) - T_{SL} - T_{OL}$ $\Pi_B = 0$
	<i>S-Learning</i>	$p_A^N = p_A^O = 0$ $p_B^N = p_B^O = 0$ $\Pi_A = \Pi_B = -T_{SL}$	$p_A^N = 0$ $p_A^O = \sqrt{2c} - E$ $p_B^N = p_B^O = 0$ $\Pi_A = (1 - \alpha_B)(\sqrt{2c} - E) - T_{SL} - T_{OL}$ $\Pi_B = -T_{SL}$
	<i>O-Learning</i>	$p_A^N = 0$ $p_A^O = \sqrt{2c} - E$ $p_B^N = 0$ $p_B^O = \sqrt{2c} - E$ $\Pi_A = (1 - \alpha_B)(\sqrt{2c} - E) - T_{SL}$ $\Pi_B = \alpha_B(\sqrt{2c} - E) - T_{OL}$	$p_A^N = 0$ $p_A^O = \sqrt{2c}$ $p_B^N = 0$ $p_B^O = \sqrt{2c} - E$ $\Pi_A = (1 - \alpha_B)\sqrt{2c} - T_{SL} - T_{OL}$ $\Pi_B = \alpha_B(\sqrt{2c} - E) - T_{OL}$
	<i>S-Learning and O-Learning</i>	$p_A^N = p_A^O = 0$ $p_B^N = 0$ $p_B^O = \sqrt{2c} - E$ $\Pi_A = -T_{SL}$ $\Pi_B = \alpha_B(\sqrt{2c} - E) - T_{SL} - T_{OL}$	$p_A^N = 0$ $p_A^O = \sqrt{2c} - E$ $p_B^N = 0$ $p_B^O = \sqrt{2c} - E$ $\Pi_A = (1 - \alpha_B)(\sqrt{2c} - E) - T_{SL} - T_{OL}$ $\Pi_B = \alpha_B(\sqrt{2c} - E) - T_{SL} - T_{OL}$

To see that Proposition 3(a) holds, notice that if Firm A uses *S-Learning*, then Firm B earns more profit from choosing *O-Learning* rather than both *S-Learning* and *O-Learning* for all $T_{SL} > 0$. And, if Firm A uses *S-Learning*, then Firm B earns more profit from choosing No Learning rather than choosing *S-Learning*. Thus, neither *S-Learning* nor both *S-Learning* and *O-Learning* can be the optimal response to *S-Learning*.

To see that Proposition 3(b) holds, notice that if Firm A uses both *S-Learning* and *O-Learning*, then Firm B earns more profit from choosing *O-Learning* rather than both *S-Learning* and *O-Learning*, and Firm B earns more profit from choosing No Learning rather than choosing *S-Learning* for all $T_{SL} > 0$. Thus, neither *S-Learning* nor both *S-Learning* and *O-Learning* can be the optimal response to the combination of both *S-Learning* and *O-Learning*.

Strategic Investments in Learning Technologies

Here, we provide the conditions under which the six strategies identified in the text are subgame-perfect. The numerical examples in the text demonstrate that the prescribed parameter space is non-empty for each of these strategies. In addition to the conditions listed below, all examples also satisfy conditions (4), (7), and (8).

“Take Your Pick”

This strategy requires $D_B(D_A = OL) = SL$ and $D_B(D_A = SL) = OL$. Thus, we need conditions (A21), (A25), and (A26) to hold. Table A1 reports firm A’s profit under these two scenarios. Firm A chooses $D_A = OL$ if $(1 - \alpha_B)(1 - (1 - \alpha_B)(1 - E)) - T_{OL} > (1 - \alpha_B)(1 - E) - T_{SL}$ or:

$$T_{SL} > T_{OL} + (1 - \alpha_B)[(1 - \alpha_B)(1 - E) - E] \quad (\text{A29})$$

If condition (A29) is not met, then $D_A = SL$.

“Beat ‘em to the Punch”

This strategy is valid when $D_B(D_A = SL) = \emptyset$, $D_B(D_A = OL) \neq OL$ and $D_B(D_A = \emptyset) = SL$. Thus, we need conditions (A17) and (A18) to hold, condition (A21) to be violated, and either condition (A27) or condition (A28) to be violated. Firm A’s profit is maximized at $D_A = SL$ only if $\sqrt{2c} - E - T_{SL} > (1 - \alpha_B)(1 - (1 - \alpha_B)(1 - E)) - T_{OL}$.

“Beat ‘em to the Punch (with a Helping hand)”

This strategy is valid if $D_B(D_A = SL) = \emptyset$, $D_B(D_A = OL) = \emptyset$ and $D_B(D_A = \emptyset) = SL$. Thus, we need condition (A21) to be violated and conditions (A17), (A18), (A23), and (A24) to hold. Firm A’s profit is maximized at $D_A = OL$ only if $\sqrt{2c} - E - T_{SL} < (1 - \alpha_B)(1 - (1 - \alpha_B)[1 - \sqrt{2c}]) - T_{OL}$.

“Be My Guest”

This strategy requires $D_B(D_A = \emptyset) = OL$ and $D_B(D_A = OL) \neq SL$. Thus, we need conditions (A19) and (A20) to hold and either condition (A25) or (A26) to be violated. Firm A’s profit is maximized at $D_A = \emptyset$ only if $(1 - \alpha_B)(1 - \sqrt{2c}) > (1 - \alpha_B)(1 - (1 - \alpha_B)[1 - \sqrt{2c}]) - T_{OL}$ or:

$$T_{OL} > (1 - \alpha_B)(\sqrt{2c}(2 - \alpha_B) + \alpha_B - 1) \quad (\text{A30})$$

“Ride My Coattails”

This strategy is valid if $D_B(D_A = OL) = \emptyset$ and $D_B(D_A = \emptyset) = OL$, i.e., conditions (A19), (A20), (A23), and (A24) hold. Firm A chooses $D_A = OL$ only if condition (A30) is violated. These conditions can be met only if $\alpha_B < 1/2$. To see this, note that satisfying (A24) and violating (A30) simultaneously requires $\alpha_B \sqrt{2c} < T_{OL} < (1 - \alpha_B)(\sqrt{2c}(2 - \alpha_B) + \alpha_B - 1)$, which is non-empty only if:

$$(1 - \alpha_B)(\sqrt{2c}(2 - \alpha_B) + \alpha_B - 1) - \alpha_B \sqrt{2c} > 0 \quad (\text{A31})$$

This relationship is violated at $\alpha_B = 1/2$ since at this parameter value the LHS of (A31) becomes:

$-\frac{1 - \sqrt{2c}}{4} < 0$. Furthermore, taking the derivative of (A31) and evaluating at $\alpha_B = 1/2$ we find:

$$\frac{\partial LHS(A31)}{\partial \alpha_B} \Big|_{\alpha_B = 1/2} = 1 - 3\sqrt{2c} \quad (\text{A32})$$

(A32) is negative for all c that satisfy (8). Therefore, condition (A31) cannot be satisfied for $\alpha_B \geq 1/2$.

“Follow Me”

This strategy requires $D_B(D_A = SL) = \emptyset$ and $D_B(D_A = OL) = OL$. Thus, we need condition (A21) to be violated and conditions (A27), and (A28) to be met. Firm A selects $D_A = OL$ over $D_A = SL$ only if $1 - \alpha_B - T_{OL} > \sqrt{2c} - E - T_{SL}$. We also need to verify that firm A would not want to choose $D_A = \emptyset$. For the example in the text this is easily verified since $D_B(D_A = \emptyset) = SL$ and thus this option would result in $\Pi_A = 0$.

ⁱ It should be noted that previous research does not explicitly distinguish between the two learning approaches and their competitive effects. For instance, Ansari and Mela (2003), p. 132, use examples of both *O-Learning* – a company personalizing a website “based on revealed preferences data” – and *S-Learning* – allowing “users to self-customize the site” – to define on-site personalization.

ⁱⁱ There is a related stream of literature which studies the trade-off between efforts to increase customer retention and efforts to enhance customer acquisition. Specifically, Reinartz, Thomas and Kumar (2005) examine how to achieve the right balance between these marketing efforts. Others explore how these decisions depend upon market share (Fruchter and Zhang 2004, McGahan and Ghemawat 1994) and competitive forces (Syam and Hess 2006). Our paper adds to this literature by studying how different learning approaches influence the ability and willingness of a firm to pursue customer retention versus acquisition.

ⁱⁱⁱ Note that at that time Yahoo! had a market share exceeding 55%

^{iv} For example, reflect.com offers a product line of lipsticks with varying colors and finishes that are priced at a common level, and *Lands’ End* offers personalized jeans with all variants priced at \$54 (Syam et al. 2005).

^v Technically, this specification assumes search is conducted with recall. However, this is not critical since no customer chooses to consume a previously rejected option (see Lippman and McCall 1976).

^{vi} This model specification is based on the idea that *S-Learning* requires a significant up-front investment in the appropriate infrastructure, interfaces, and algorithms. This is consistent with Dewan et al. (2003) who model the cost of “gathering and processing information” as a fixed cost.

^{vii} If α_M is sufficiently small, the firm would prefer a price of $1 - \sqrt{2c}$ (instead of $p = 1$) in order to sell to all customers. But, in this case, the *O-Learning* increases the firm’s costs without enhancing revenue and thus cannot be the optimal strategy.

^{viii} However, it is straight-forward to extend the model to a setting where firms choose their learning approaches simultaneously. But, such an environment is less interesting since the strategic investments in learning technologies (to be discussed in the next section) could not motivate investment decisions in such a setting because a rival firm would not be able to observe the focal firm’s choice before making its own decision.

^{ix} This sequential structure is consistent with the argument in Dewan et al (2003) that firms differ in “their readiness to adopt new technologies.”

^x The complete survey instrument is available with the authors on request.

^{xi} For example, if the total personalization is reported as 3 (on a scale of 5) and the extent of solicited information is reported at 2 (on a scale of 5), then *S-Learning* of the supplier firm is calculated as 2/5 of 3, i.e., 1.2. The rest of the personalization, i.e., $(3 - 1.2) = 1.8$, is attributed to *O-Learning*.

^{xii} Such non-strategic decisions have already been analyzed in the preceding section.

^{xiii} Strategy #1 is similar to the result in Syam and Hess (2006) that the first mover can benefit by forcing a rival to adopt the less profitable strategy. However, in the current model, depending on the parameters, either learning approach may be the preferred one. And, one is assured that the rival will choose the opposite strategy only under a subset of market parameters.