Title
Essays on Technology and Data Analytics in Operations Management

Author
Wang, Yiwei

Publication Date
2021

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Essays on Technology and Data Analytics in Operations Management

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Management

by

Yiwei Wang

Dissertation Committee:
Professor Shuya Yin, Chair
Professor L. Robin Keller
Professor Lauren X. Lu

2021

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

Taking this opportunity, I would like to acknowledge many people who have encouraged, mentored, and helped me throughout my Ph.D. study and entire student life. This dissertation could not be written without their continuous support.

I would like to express the deepest gratitude to my dissertation committee chair Professor Shuya Yin. Prof. Yin has provided me with invaluable guidance and support during my entire Ph.D. study. She has also encouraged me to follow my research interests. She is a great advisor, and I am very proud to be her student. Without her persistent help, this dissertation would not have been possible.

I would also like to extend my sincerest appreciation for Professor Lauren Xiaoyuan Lu, a fantastic co-author, and a remarkable mentor. Professor Lu has sparked my research interest in the field of data-driven and empirical operations management. I am inspired by her enthusiasm every time she talks about an exciting research finding. She is knowledgeable, responsive, and efficient in every meeting. Doing research with her is indeed one of my most memorable memories.

I am indebted to my committee member, Professor Robin Keller, for her kindness in sharing thoughts and comments on my research ideas. She is altruistic and always available to students when they need help. She is absolutely a role model for me and many young scholars to follow, not only in terms of research but also in treating other people.

I am thankful to Prof. Luyi Gui for her guidance and assistance during the first year of my Ph.D. study. I also want to thank Prof. Vidyanand Choudhary, Prof. Jeannette Song, Prof. Rick So, Prof. John Turner, Prof. Mingdi Xin, Prof. Vibhanshu Abhishek, Prof. Yuhong He and Prof. Carlton Scott for their advice and support. Moreover, I want to express my special thanks to Prof. Max Zuo-jun Shen, who encouraged me to get into the field of Operations Management when I was an undergraduate student at UC Berkeley.

I am grateful to my undergraduate classmate, serial entrepreneur Michael Chen Wang, for sharing many insightful thoughts about the retail sector with me. My trip to the headquarter of Y-Closet in the summer of 2018 was a monumental moment in my Ph.D. study. I am indebted to Pengcheng Shi and Henry Shi at AI List Capital and CEO Jay Bean at Freshlime.com for the opportunity to implement my research in a real-world setting.

Last, I thank all my friends and my family for their unconditional support. My father, a professor at Zhejiang University, encouraged me to pursue a Ph.D. after my undergraduate study. My mother, who was also my first math teacher, stood on my side throughout the seven-year journey. Weixuan Chen, thank you for your support and love. To all of you, I dedicate this dissertation.

CURRICULUM VITAE

Yiwei Wang

EDUCATION

Doctor of Philosophy in Business Administration **2021**
University of California – Irvine *Irvine, CA*

Bachelor of Science in Operations Research **2013**
University of California – Berkeley *Berkeley, CA*

RESEARCH INTEREST

Retail Operations and Analytics

Online and Field Experiments

Empirical Operations Management

Applied Economic Modeling

AWARDS AND HONORS

Long Marilyn US-China Research Fellowship 2018

Doctoral Student Fellowship, UC Irvine 2014

Alpha Pi Mu Honors Society, Berkeley Chapter 2013

SERVICES

Session Chair: POMS (2021); INFORMS (2021)

Ad-hoc Reviewer: Management Science

Member: INFORMS, MSOM, POMS

ABSTRACT OF THE DISSERTATION

Essays on Technology and Data Analytics in Operations Management

By

Yiwei Wang

Doctor of Philosophy in Management

University of California, Irvine, June 2021

Professor Shuya Yin, Chair

Motivated by recent advances in technologies and big data analytics, this dissertation consists of three essays on operations management. The first essay investigates the firm's technology choice under product quality uncertainty. We propose a three-stage game to study how a firm should choose between a mature technology and an innovative technology in upgrading an existing product, where the level of quality improvement is jointly determined by two attributes – one is determined by the firm, and the other is outside of the firm's control. We characterize conditions under which each technology is more likely to be adopted.

The second essay empirically investigates how customer email engagement affects the profitability of subscription-based service providers. We analyze the outcome of a field experiment conducted by a large U.S. car wash chain, which offers tiered subscription services to consumers and employs an RFID-based technology to track subscriber service events. We apply survival analysis and difference-in-differences methods to estimate the effects of email engagement on subscribers' retention and service consumption. We find that a one-month engagement with two emails separated by a half-month interval increased the likelihood of subscriber retention by 7.4% five months after the experiment started and decreased the subscriber churn odds by 26.3% for the entire five-month duration. Meanwhile, we find that the same engagement increased a subscriber's per-period service consumption by 8.8%. Our

study highlights that email engagement is a double-edged sword—it increases both customer retention and service consumption, and it may decrease profitability when the increased operating cost to serve retained customers outweighs the benefit of customer retention.

The third essay empirically examines the impact of curated box retailing, i.e., shipments of retailer-selected products seeking to surprise and delight customers at regular intervals. We conducted a field experiment to analyze curated box retailing's impact on a leading retailer. We randomly selected 580 customers to receive curated boxes for two consecutive months and post-treatment observation for another seven months. Each box contained exactly six items for product sampling and purchase. We find that monthly dispatch of one curated box for two months substantially increased overall product sales in all retail channels and caused positive cross-channel demand spillovers to the online and home try-on sales. At the same time, curated box retailing led to a reduction in excessive product sampling and returns in the home try-on channel. Our research provides implications for retailers in adopting and optimizing the curated box retailing strategy.

Chapter 1

Product Innovation and Technology Choice under Quality Uncertainty

1.1 Introduction

Consider a firm that currently sells a product in the market. Given the much-shortened product life cycle and rapid product upgrades, the firm evaluates some options to update the product. Product updates can be on various dimensions, including expanding functionality, performance stability, cost structure, quality, consumer convenience, etc. In this paper, we focus on the improvement of the product's overall quality. Take the battery industry as an example. For a battery product, one of the key determinants in its overall quality is the underlying technology choice. For instance, a major battery manufacturer currently faces a

This chapter is in conjunction with Prof. Shuya Yin and Prof. Vidyanand Choudhary at University of California - Irvine. This chapter is word-for-word the same as the working paper "Product Innovation and Technology Choice under Quality Uncertainty" by Yiwei Wang, Vidyanand Choudhary and Shuya Yin, prepared for journal review at *Journal of Management Information Systems*.

difficult situation in whether to use a mature or innovative technology to improve its existing battery product's quality. In the past, this manufacturer has supplied the existing product – Lithium-ion (Li-ion) batteries mainly for consumer electronics (e.g., laptop computers and mobile phones). Those standard Li-ion battery products have been well tested and can provide a consistent and stable product quality. Given the booming electric vehicle market, the firm has decided to upgrade its current product to make high-capacity batteries suitable for the automobile sector. However, the Li-ion battery usually only has limited cycle life if packed in large quantities (Boston Consulting Group, 2010). On the horizon is an innovative battery technology that is less constrained by the fixed cycle life and may lead to a much greater improvement in the overall capacity of the product (Hu, 2016). Nevertheless, this new technology can be risky, as its exact nature is less well understood and could lead to a worse situation regarding the battery cycle life (and subsequently its overall quality). Indeed, product failure frequently occurs when a new technology is involved. The recent incidents of faulty new batteries catching fire stand as extreme cases for the above statement (BBC, 2017).

Given the trade-offs mentioned earlier, it could be quite challenging and complicated for firms to assess the costs and benefits of product improvement due to the uncertainty involved in innovative technologies. This is especially so when multiple attributes measure the product's overall quality. Consider the battery example again. A battery's quality is determined mainly by two complementary attributes: *capacity* and *cycle life*. Capacity is the maximum amount of energy stored in the battery. Firms can usually control and determine the level of the battery capacity before the R&D process (Boston Consulting Group, 2010). Cycle life is the total number of charge/discharge cycles the battery can perform. This attribute is generally outside of firms' control (Boston Consulting Group, 2010). To this end, firms may partially leverage the controllable capacity attribute (not completely, as firms cannot control the level of cycle life) to influence the resulting overall product quality. The exogenous cycle life parameter, at the same time, might be deterministic or uncertain, depending on the type of technology used. Under the mature technology, firms have prior experience with previous product developments. Hence, they treat the resulting product's cycle life as a constant. However, under the innovative technology, firms have rather limited prior knowledge of product development, despite its high potential in terms of cycle life. Therefore, firms anticipate the resulting product's cycle life to be an uncertain parameter *a priori*. There are other examples in practice that fit into the setting described here. When mobile phones first came out, the technology behind the product was innovative. Relative to the landline phones, mobile phones' mobility was much improved, but phone calls' quality was much poorer due to the cellular network's poor infrastructure. Similarly, when the Blu-ray discs started to replace the DVDs, they had high potentials on storage capacity but may not be as compatible as the old version of the product, which imposed some uncertainty on its adoption by consumers.

In this paper, we propose a stylized three-stage economic model to understand the critical trade-offs involved in technology choice in the process of improving an existing product. Specifically, we aim to understand:

1. Given a technology, what should the firm's optimal strategy be, in terms of the level of quality improvement, pricing, and the product offering strategy?

2. Under what conditions is a mature or an innovative technology more likely to be adopted?

3. How would the firm's technology decision be influenced by factors such as the cost structure in using each technology and the risk factor in the innovative technology?

To reflect the motivating example from the battery industry and address the questions mentioned above, we model a firm that considers improving the quality of its existing product. There are two potential technologies available to improve the product quality: one is relatively mature, and the quality improvement is predictable; the other is innovative, and the quality improvement carries some risk. In particular, we assume that the overall product quality is the multiplication of two attributes: the firm endogenously determines one, and the other is exogenously affected by some external factors. To make the model description more intuitive, we borrow the terminology used in the battery example to reflect the two attributes: *capacity* is an endogenous decision variable, and *cycle life* is modeled as an exogenous parameter, which could be either deterministic or uncertain. The two attributes can certainly be defined entirely differently from capacity and cycle life in the context of a different product type with different attributes.

The firm makes decisions in three stages. In the first stage, the firm selects which technology to use. In the second stage, the firm decides on the capacity improvement of the existing product. Note that the capacity decision is associated with a fixed investment cost in this stage, irrespective of whether or not the improved product will indeed be adopted in the subsequent stage. If the mature technology is adopted, the firm will set the price for the improved product in the third stage. If the innovative technology is adopted, then at the beginning of the third stage, the uncertainty in the product cycle life is realized, which can either be high quality under improvement success or low quality under improvement failure. A unique feature under this technology is that the firm may strategically offer the improved product or the existing product depending on the realized value of cycle life and set the selling price accordingly. As a result, in the third stage, there are three product offering strategies that may possibly be optimal:

- (II) strategy: Improved product is offered regardless of the realized value of cycle life.
- (IE) strategy: Improved product is offered under a high value of cycle life and the existing product is offered under a low value of cycle life.
- (EE) strategy: Existing product is offered regardless of the realized value of cycle life.

Note that the (EI) strategy where the existing product is offered under high cycle life while the improved product is offered under low cycle life is never optimal for the firm. The (EE) strategy, offering the existing product regardless of the cycle life realization, happens when high marginal product cost deters the firm's product improvement. After all decisions are made, the market demand realizes, and sales occur. The firm then engages in the production and incurs the marginal production cost. Finally, the profit is generated. Analysis of our model framework demonstrates the following main findings:

(1) The firm only uses mature technology to improve the capacity when the unit production cost is low. However, when the innovative technology is used, the firm may improve the existing product's capacity even if the unit production cost is high, provided that the cycle life is sufficiently high when the improvement is successful.

(2) Under the mature technology, the firm always offers the improved product. On the contrary, under innovative technology, the firm has more flexibility in leveraging the product offering strategy. Thus, the firm uses the (II) strategy, i.e., always offers the improved product if both the marginal cost and the uncertain cycle life fluctuation are sufficiently low. Otherwise, the firm offers the improved product only under high cycle life and offers the existing product under low cycle life, i.e., the (IE) strategy.

(3) We now consider the firm's preference over the two technologies. Let us first assume that the expected cycle life under the innovative technology is equal to or higher than the cycle life under the mature technology. Under this circumstance, the innovative technology has an advantage on expected cycle life, which compensates for the risk caused by its uncertainty. Interestingly, we find that the fixed cost and the production cost play opposite roles on the firm's preference over the two technologies. A higher fixed development cost makes the mature technology more likely to be adopted, while a higher unit production cost makes the innovative technology more preferred by the firm. This finding demonstrates the importance of decomposing the cost structure involved in product improvement and provides insights into policy-making when an incentive program is in consideration to motivate the adoption

of a particular technology.

(4) Finally, if the mature technology is advantageous, relative to the innovative one, in terms of *both* the lack of uncertainty *and* the expected cycle life, the impact of the production cost on the firm's technology preference is not monotonic any more. Specifically, when the production cost is low, increasing the production cost benefits the mature technology. As the production cost reaches a certain threshold, a further increase in the production cost benefits the innovative technology. When the production cost is low, the firm always offers the improved product regardless of the underlying technology. This places the innovative technology at a disadvantageous position as the innovation uncertainty can no longer be compensated by the flexible product offering strategy.

The rest of this paper is organized as follows. In §1.2, we review related literature. In §1.3, we describe the model and specify the objective functions under both technologies. In §1.4, we carry out the analysis for each technology option, and examine the firm's technological preference in the base model where the (expected) cycle life is the same under both technologies. In §1.5, we show that the main results obtained in the base model remain robust after we relax the equal cycle life assumption. We state conclusions, managerial insights, and future research directions in §1.6.

1.2 Literature Review

The first stream of literature that is related to our paper is about product/quality choice. A focus in this literature is on how firms can design multiple products of different qualities to segment consumers with different levels of willingness-to-pay for quality. The early work can be traced back to, e.g., Mussa and Rosen (1978) and Moorthy and Png (1992). Since then, many factors that can potentially influence the product/quality design have been explored, including the product characteristics such as information goods (Bhargava and Choudhary, 2001, 2008), product cannibalization (Desai, 2001), the structure of distribution channels (Xu, 2009), and production technologies (Thatcher and Oliver, 2001; Thatcher and Pingry, 2004). The product/quality design is also studied in operations management through product upgrades/improvement or new product introduction. The issues usually focus on whether or not the existing version of the product should be made available together with the new and upgraded version (Levinthal and Purohit, 1989; Liang et al., 2014; Lim and Tang, 2006), on the firms' pricing strategies (Dhebar, 1994; Kornish, 2001), or on the impact of product characteristics (Bala and Carr, 2009) and used goods markets (Yin et al., 2010). A common feature of the majority of papers in this literature stream is that they model the product quality as a deterministic attribute.

There are a few papers that consider uncertain product quality. For example, like our paper, Feldman et al. (2019) model the overall quality of an experience good as a mix of an endogenous attribute (determined by the firm) and an exogenous and uncertain attribute (determined by nature). In their paper, the firm's product design is perfectly observable, but the end consumers' overall quality is also affected by ex-ante quality uncertainty that always exists. Their focus is to understand how social learning affects the firm's product design. A key differences of our paper is that, we propose a three-stage model which implicitly considers the technology choice decisions underlying the quality improvement process. So the critical trade-off that the firm faces is about the risks and benefits from adopting the innovative technology. A similar choice between the mature and the innovative technologies has been studied in Krishnan and Bhattacharya (2002). In their paper, if the firm decides to explore the innovative technology, it can learn about this technology's uncertainty over time and decides on when to adopt the new technology. They find that the demand function, product development cost, and the development cycle's length jointly affect the optimal technology decision. Our paper, however, is fundamentally different in terms of the research questions, modeling of uncertainty, and the decisions involved, i.e., pricing, demand, and quality. In

Krishnan and Bhattacharya (2002), the uncertainty lies in the timing of the adoption of new technology, and whether the technology is viable or not is exogenous. In our setting, the uncertainty is associated with a part of quality (cycle life) choice, and the other part of the quality choice, the capacity, is endogenous. Also, the product offering is determined by overall quality, pricing, and market demand. All of those factors play a role in how the decision-maker reacts. Moreover, we focus on the decomposition of different costs, which further leads to heterogeneous effects on the optimal technology choice, whereas their paper does not capture the effect of marginal cost.

Furthermore, our paper is related to the operations management for innovative technologies. For instance, Chen (2001) analyzes the impact of a green product attribute on the firm's new product development. Chen et al. (2013) examine the firm's optimal product-line design and production decisions under the novel vertical co-product technology. Lim et al. (2014) study the impact of range and resale anxieties associated with different business practices in the electric vehicle industry. Both Avci et al. (2014) and Mak et al. (2013) study various operational decisions associated with the battery swapping business model. Wang et al. (2013) consider the choice between a conventional and sustainable technology in making capacity decisions over time, without implicitly modeling the firm's quality decision. To abstract away from the details in a specific industry, in this paper, we use a stylized economic model to comprehensively understand the firm's technology choice problem while internalizing the quality and pricing decisions.

1.3 Model Framework

A firm sells a product in the market and considers upgrading it. The quality of the product is affected by two main attributes. The firm can endogenously determine one attribute while external factors determine the other attribute. There are two potential technologies that the firm can consider in improving the quality of the existing product. They are different due to the nature of the external factors. Specifically, for one technology, those factors can be well studied and predicted. So, the improvement of this technology on product quality can lead to predictable and stable performance. As a result, it is named as the "mature", or "deterministic" technology. However, for the other technology, its improvement on product quality exerts some uncertainty in some of the external factors, resulting in a random outcome, namely, the "innovative" or "uncertain" technology. Indeed, many new technologies in their early development stages exhibit some uncertainty in their performance, even though they may present advantages on other dimensions. Consider the battery industry as an example, where the two main attributes associated with the overall battery quality are capacity and cycle life. The firm can set the battery capacity as a decision variable while the cycle life is often exogenous to the firm's decision. If a well-established battery improvement technology is adopted, then the battery is likely to have a predictable performance on cycle life; if an innovative technology is adopted, then the resulting battery cycle life may be stochastic.

In the rest of this section, we propose a stylized model to understand the firm's technology choice decision. We first describe the sequence of events (see Figure 1.1) and then characterize the demand and profit functions. We will borrow the terms from the battery industry in our model setup to ease the terminologies. "capacity" is used to represent the attribute in quality improvement controlled by the firm, and "cycle life" represents the attribute that is exogenous to the firm. Moreover, these terms measure how much *improvement* on capacity and cycle life that the new technology can make to the existing product. We assume that the multiplication of the two attributes measures the overall quality improvement on the existing product due to their interconnected influence on product quality.

Following Figure 1.1, the firm decides on whether to improve the product quality of an existing product, which it always does unless the related costs are sufficiently high. The second decision that the firm makes is to choose which technology to use to improve its existing product. Depending on this strategic decision, some operational decisions will follow, including the capacity level in quality improvement and the selling price, etc. We describe the sequence of events separately under a chosen technology.

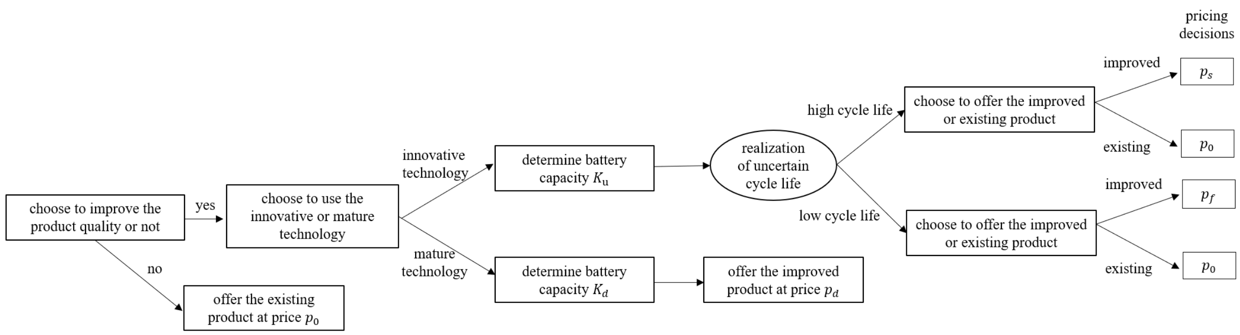

Figure 1.1: Sequence of Events

1.3.1 Sequence of Events

- *Deterministic/Mature Technology.* The sequence of events in this case follows the bottom branch in Figure 1.1 and it is quite straightforward. The firm invests on a capacity level, K_d, which leads to a development cost, and a corresponding retail price, p_d, for the product sold in the market, where subscript "d" represents the deterministic/mature technology. Note that the overall quality improvement in this case is deterministic, since it is measured by $K_d T_d$, the multiplication of capacity K_d and deterministic cycle life T_d. Accordingly, the market responds and leads to a sales quantity or demand.

- *Uncertain/Innovative Technology.* The sequence of events under this technology follows the top branch of Figure 1.1 and it is more complicated due to the uncertain nature of cycle life, measured by $\widetilde{T}$. For simplicity, we assume that $\widetilde{T}$ follows a two-point distribution, i.e., a high or a low status, measured by $T_u + \theta$ (which implies a situation of successful quality improvement) or $T_u - \theta$ (which implies a situation of failed quality improvement), respectively, with equal probabilities, where $\theta \in [0, T_u]$.

(1) First, the firm decides to invest on the improvement level on capacity, K_u. This decision is assumed to be made before the uncertainty of cycle life is revealed. This assumption can be justified that, in general, capacity decisions are mostly strategic and long-term that need to be decided ahead of time. The amount of investment associated with capacity improvement or the development cost is incurred when this decision is made.

(2) Second, the value of cycle life is revealed, which can be either high or low.

(3) Third, contingent on the chosen capacity level and the high or low cycle life, the firm can decide whether to offer the improved or the existing product and its corresponding retail price. Specifically, under high cycle life, $\widetilde{T} = T_u + \theta$, the firm sets p_s when the improved product is sold, and p_0 when the existing product is sold. The subsequent model analysis indicates that the improved product will always be offered if the firm discovers that the cycle life is high, as long as the marginal production cost is not too high, which is not surprising. Under low cycle life, $\widetilde{T} = T_u - \theta$, the firm sets p_f or p_0. We assume that the pricing decision is made after realization of cycle life. This assumption helps the model tractability and it can be justified since pricing is usually an operational decision that can be made rather quickly.

(4) Lastly, the market responds and demand is realized for the product sold in the market.

Following the sequence of events under the two technologies, there is one key differentiation between the two technologies that are worth noting. In the deterministic case, it is never optimal for the firm to initially set a positive capacity level $K_d > 0$ and later decide to offer the original product (since this is dominated by setting $K_d = 0$ in the beginning to save the development cost). However, in the uncertain case, depending on the realization of the uncertain cycle life, it might benefit the firm to offer the existing product later (but not the improved product) even if the firm previously set a positive capacity improvement level, $K_u > 0$, with associated development cost. This may happen when the realized value of cycle life turns out to be low. Finally, we also assume that $T_d \in [T_u - \theta, T_u + \theta]$ to avoid non-interesting cases where one technology is always dominant. For ease of reference, a summary of notation used in the paper is presented in the Appendix.

1.3.2 Demand Characterization

Before we characterize the firm's profit function, it is important to understand how consumers react to the quality improvement (if any) and the retail price. Here, we adopt a commonly used demand model where the overall demand is determined as follows:

$$Demand = a + overall\ \ quality\ \ improvement - b \cdot retail\ \ price, \tag{1.1}$$

where parameter a is positive and measures the baseline market demand if the firm offers the existing product at zero price and parameter b measures how demand is sensitive to the price change, which can be normalized to 1. For mathematical tractability, $b = 1$ will be assumed for the rest of the paper. Note however that relaxing this assumption would not change major results of our analysis. The actual demand function is specific to which technology is adopted and the subsequent decisions made in the process:

- Under the mature technology, we have overall quality improvement expressed as $K_d T_d$ and retail price as p_d. So, following equation (1.1), the overall demand can be rewritten as

$$Demand = a + K_d T_d - p_d. \tag{1.2}$$

It is clear that the existing product is sold if the firm sets $K_d = 0$.

- Under the innovative technology, the overall quality improvement and retail price are contingent on the realization of random cycle life and on whether the firm chooses to offer the improved or the existing product.

$$\textit{Demand} = \begin{cases} a + K_u(T_u + \theta) - p_s & \text{if } \widetilde{T} = T_u + \theta \;\; \textit{and the improved product is sold}, \\ a + K_u(T_u - \theta) - p_f & \text{if } \widetilde{T} = T_u - \theta \;\; \textit{and the improved product is sold}, \\ a - p_0 & \text{if } \textit{the existing product is sold}, \end{cases} \tag{1.3}$$

where p_0 represents the base retail price when there is no quality improvement (or it is the selling price of the existing product).

Recall that in the base model, we assume that the two technologies have an equal cycle life (in expectation), i.e., $T_d = T_u$. This further leads to the observation that the deterministic model becomes a special case of the uncertain model with $\theta = 0$. This can also be easily verified from the demand functions in (1.2) and (1.3). For clarity, in §1.4, we will still carry out the analysis of the base model under the mature technology in addition to the analysis in the model with uncertainty. Note that in the extended model considered in §1.5 where $T_d \neq T_u$, the deterministic model is no longer a special case of the uncertain model. With the demand function, we can derive the firm's profit function and its optimization problem(s) involved in the decision process.

1.3.3 Firm's Profit Function

Following the structure of demand characterization, we present the firm's profit function and its optimization problems separately in the deterministic and uncertain cases.

Under the *mature technology*, the firm sets the capacity level and the selling price. Since there is no further information gained or no other events occurring between the two decisions, mathematically, it is equivalent for the firm to set both K_d and p_d at the same time to maximize the following profit function:

$$
\begin{aligned}
\max_{K_d, p_d} \quad & \Pi_d = \overbrace{(a + K_d T_d - p_d)}^{\text{market demand}} \cdot \overbrace{(p_d - cK_d)}^{\text{unit profit}} - \overbrace{dK_d^2 T_d^2}^{\text{fixed cost}} \\
\text{s.t.} \quad & K_d \geq 0 \;\; and \;\; cK_d \leq p_d \leq a + K_d T_d.
\end{aligned}
\tag{1.4}
$$

The constraints are imposed to ensure the non-negativity of decisions, demand and profit margins.

There are two sets of cost information that need to be discussed. The first cost is the additional marginal cost in producing a unit of an improved product. The marginal production cost is assumed to be proportionally increasing in capacity improvement, as producing a product with higher quality requires more resources, t. So, the marginal cost is measured by cK_d. We also assume that the battery cycle life does not contribute to the product's marginal production cost. For example, the battery cycle life generally depends on the chemical properties but not the amount of material used, thus not contributing to the marginal cost during production (The Economist, 2010).

In addition to the marginal production cost, the second type of cost is the development cost associated with the level of overall quality improvement, which is characterized by $dK_d^2 T_d^2$, where d is a positive development cost coefficient. This form of cost can be considered fixed since it is not related to the number of units produced. The quadratic function implies that the benefit from higher values of quality is increasingly costly. This function form has been widely adopted in literature, see, e.g, Thatcher and Oliver (2001), Barua et al. (1991). The increasing difficulty in improving quality has also been widely documented in practice UPI news (2017). Note that there is no fixed cost if the firm offers the existing product because, by definition, both attributes K_d and T_d represent levels of quality improvement, which are zero for the existing product.

Under the *innovative technology*, according to the sequence of events presented earlier, the firm's decisions can be framed as a three-stage game where the standard backward induction approach is used to solve the optimal decisions. Specifically, we can consider the technology decision to be the first stage decision, the capacity decision made before uncertain cycle life as the second stage decision, and the retail pricing decision made after its realization as the third stage decision. In solving the problem, we apply backward induction. We first solve the third stage pricing problem (knowing the capacity level and the realized value of cycle life). Then, we move to solve the second stage capacity decision, anticipating the best response functions (in terms of pricing) in the third stage and the expected cycle life. The detail is presented below.

- In the third stage, the capacity level K_u is known, and the value of cycle life is realized. If the value of cycle life is high, i.e., $\widetilde{T} = T_u + \theta$, *and* the firm decides to sell the improved product, then the firm sets p_s to maximize the sales profit (since the fixed costs occurred in the second stage is sunk):

$$\begin{aligned} \max_{p_s} \quad & \Pi_s^{stage\ 3} = \overbrace{[a + K_u(T_u + \theta) - p_s]}^{\text{market demand}} \cdot \overbrace{(p_s - cK_u)}^{\text{unit profit}} \\ \text{s.t.} \quad & cK_u \leq p_s \leq a + K_u(T_u + \theta), \end{aligned} \tag{1.5}$$

  where the constraints are applied to guarantee non-negative demand and profit margins. Similarly, if the value of cycle life is low, i.e., $\widetilde{T} = T_u - \theta$, *and* the firm decides to offer the improved product, then the firm sets p_f to maximize its profit:

$$\begin{aligned} \max_{p_f} \quad & \Pi_f^{stage\ 3} = \overbrace{[a + K_u(T_u - \theta) - p_f]}^{\text{market demand}} \cdot \overbrace{(p_f - cK_u)}^{\text{unit profit}} \\ \text{s.t.} \quad & cK_u \leq p_f \leq a + K_u(T_u - \theta). \end{aligned} \tag{1.6}$$

On the other hand, if the firm chooses to offer the existing product in the case (regardless of the realization of cycle life), then the firm sets p_0 to maximize $\Pi_0 = (a - p_0)p_0$, which is concave in p_0 with a global optimal $p_0^* = \frac{a}{2}$ and a corresponding optimal profit $\Pi_0^* = \frac{a^2}{4}$.

- In the second stage, anticipating the best response pricing functions, $p_s^*(K_u)$, $p_f^*(K_u)$ and $p_0^* = \frac{a}{2}$, and in expectation of the uncertain cycle life, the firm sets the capacity level, K_u, to maximize its expected total profit given the innovative technology is chosen:

$$\begin{aligned} \max_{K_u} \quad & E(\Pi) = \frac{1}{2}\max\{\overbrace{[a + K_u(T_u + \theta) - p_s^*(K_u)] \cdot [p_s^*(K_u) - cK_u]}^{\text{offer improved product given high cycle life}}, \overbrace{\frac{a^2}{4}}^{\text{offer existing product}}\} \\ & + \frac{1}{2}\max\{\overbrace{[a + K_u(T_u - \theta) - p_f^*(K_u)] \cdot [p_f^*(K_u) - cK_u]}^{\text{offer improved product given low cycle life}}, \overbrace{\frac{a^2}{4}}^{\text{offer existing product}}\} - \overbrace{dK_u^2({T_u}^2 + \theta^2)}^{\text{fixed cost}} \\ \text{s.t.} \quad & K_u \geq 0. \end{aligned} \tag{1.7}$$

Note that the fixed development expense is captured by the term $d\mathrm{E}[K_u^2\widetilde{T}^2] = dK_u^2\mathrm{E}[\widetilde{T}^2]$, or $dK_u^2({T_u}^2 + \theta^2)$. The quadratic fixed cost is consistent with that in the deterministic case. In the uncertain case, it is now also in expectation of cycle life since it is incurred before its realization.

1.4 Analysis

In this section, we will individually solve the base model's optimal decisions for the mature and innovative technologies and then compare the firm's optimal profits under the two technologies to understand the firm's best technological choice. We start the analysis of the deterministic case in §1.4.1.

1.4.1 Model Analysis under Mature Technology

Recall that the firm's problem is presented in equation (1.4) above. Note from our model analysis that the fixed development cost coefficient, d, needs to be relatively high so that the firm's optimal level of improvement on capacity is not unbounded (or infinite). Specifically, in this deterministic model, we assume $d > d_0$, where $d_0 = \frac{(T_d-c)^2}{4T_d^2}$ when $0 < c \leq T_d$ and $d_0 = 0$ when $c \geq T_d$, to ensure the optimal improvement on capacity and the firm's optimal profit is finite. Accordingly, it is straightforward to show that the firm's profit function is jointly concave in (K_d, p_d) with a unique global optimal solution. The following proposition summarizes the optimal solution to the firm's problem under mature technology. All the technical proofs are presented in the Appendix.

Proposition 1 (**Optimal Capacity and Pricing Decisions under the Mature Technology**). *Under the mature technology, the firm's optimal capacity level and retail price are as follows:*

- *The firm will serve an improved product to the market if the marginal production cost of quality improvement is relatively low, i.e., when $0 \leq c \leq T_d$. The optimal capacity and price are:*

$$K_d^* = \frac{a(T_d - c)}{4dT_d^2 - (T_d - c)^2} \ \textit{and} \ p_d^* = \frac{a(-c^2 + cT_d + 2dT_d^2)}{4dT_d^2 - (T_d - c)^2}.$$

  Accordingly, the firm's optimal profit is:

$$\Pi_d^* = \frac{a^2 dT_d^2}{4dT_d^2 - (T_d - c)^2}.$$

- *Otherwise, when the marginal cost is high, i.e., $c \geq T_d$, the firm will serve the existing product without improving its quality. So, $K_d^* = 0$, $p_d^* = p_0^* = \frac{a}{2}$ and $\Pi_d^* = \Pi_0^* = \frac{a^2}{4}$.*

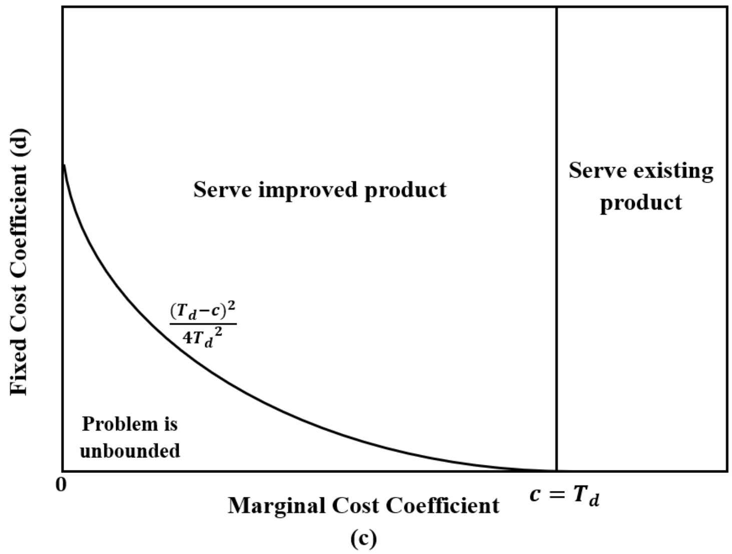

Figure 1.2: Optimal Quality/Product Decision under the Mature Technology

A visual presentation of Proposition 1 is also given in Figure 1.2 on a (c, d) plane. Following this figure, we can make some observations. First of all, the optimal level of capacity improvement is getting smaller as either the marginal cost, c, or the fixed development cost coefficient, d, increases. In other words, higher costs deter quality improvement. Second, as the cycle life, T_d, increases, the firm is more willing to exert effort to improve the existing product's capacity level. According to equation (1.2), since T_d serves a multiplier role to K_d in the demand function, the cycle life helps the firm gain more demand when the firm improves the capacity level.

1.4.2 Model Analysis under the Innovative Technology

Since the sequence of events, in this case, is more involved, as discussed in §1.3, the analysis is also more complicated. We start with the firm's optimal pricing problems in the third stage and then backtrack to the second stage's capacity problem. In the third stage, given the level of improvement on capacity, K_u, the state-contingent optimal pricing decision is summarized in Lemma 1.

Lemma 1 (**Optimal Product Offering and Its Pricing under the Innovative Technology**). *Under the innovative technology, given capacity level K_u, the optimal type of product offered and its retail price are dependent on the realized value of cycle life $\widetilde{T}$ as follows:*

- *If the marginal cost is low, i.e., $0 \leq c \leq T_u - \theta$, the firm offers the improved product regardless of the realized state of cycle life. The corresponding price is $p_s^* = \frac{1}{2}[a + K_u(c + T_u + \theta)]$ if cycle life is high or it is $p_f^* = \frac{1}{2}[a + K_u(c + T_u - \theta)]$ if cycle life is low.*

- *If the marginal cost is medium, i.e., $T_u - \theta \leq c \leq T_u + \theta$, the firm offers the improved product if cycle life is high at $p_s^* = \frac{1}{2}[a + K_u(c + T_u + \theta)]$ and offers the existing product if cycle life is low at $p_0^* = \frac{a}{2}$.*

- *If the marginal cost is high, i.e., $c \geq T_u + \theta$, the firm offers the existing product regardless of the realized cycle life at $p_0^* = \frac{a}{2}$.*

The insight behind Lemma 1 is as follows. Let us first understand how the firm can make a non-negative profit by selling an improved product. First of all, both the sales amount, $a + (K_u)(realized\ state\ of\ cycle\ life) - (b)(retail\ price)$, and the marginal profit, $retail\ price - cK_u$, need to be non-negative. When the marginal cost is on the high side, there is no feasible price that generates a positive profit. So, offering the improved product cannot lead to a non-negative profit (even without considering the sunk fixed development cost incurred). Thus the only option is for the firm to continue to offer the existing product.

Similarly, even if the marginal cost is only moderate, but if the realized state of cycle life is low, the firm faces the same dilemma: the retail price needs to be sufficiently low to generate positive sales, but it also needs to be high enough to cover the production cost. In that situation, the firm will again forgo the improved product and keep the existing product.

Understanding these insights has some significant influence on how the firm would set its level of improvement on the existing product's capacity in the second stage.

In the second stage, knowing the distribution of the uncertain cycle life and anticipating the best reactions in the third stage, the firm will now set the optimal improvement on capacity level. Similar to the deterministic model, we first impose a lower bound on the coefficient in the fixed development cost, i.e., $d > d_1$, in order to ensure the optimal level of improvement on capacity and the firm's optimal profit will be bounded (or finite).[1]

Proposition 2 (**Optimal Capacity Improvement under the Innovative Technology**). *Under the innovative technology, the firm's optimal level of capacity improvement is characterized below (see Figure 1.3).*

- *Region (a) with a low marginal cost, i.e., $0 \leq c \leq T_u - \theta$: In stage 2, the firm's optimal capacity improvement level is:*

$$K_u^* = K_a = \frac{a(T_u - c)}{4d(T_u^2 + \theta^2) - (T_u - c)^2 - \theta^2}.$$

In stage 3, the firm offers the improved product regardless of the state of cycle life. Its profit is:

$$\Pi_{II}^* = \frac{a^2\left(\theta^2 - 4d\left(\theta^2 + T_u^2\right)\right)}{4\left(c^2 - 2cT_u - (4d-1)\left(\theta^2 + T_u^2\right)\right)}.$$

- *Region (b) with a medium marginal cost, i.e., $T_u - \theta \leq c \leq T_u + \theta$: In stage 2, the firm's optimal capacity improvement level is:*

$$K_u^* = K_b = \frac{a(T_u + \theta - c)}{8d\left(T_u^2 + \theta^2\right) - (T_u + \theta - c)^2}.$$

In stage 3, the firm offers the improved product if cycle life is high while it offers the

[1]Note that the cut-off value d_1 takes different functions depending on the value of c. Specifically, if $0 \leq c \leq T_u - \theta$, $d_1 = I_a = \frac{(T_u-c)^2+\theta^2}{4(T_u^2+\theta^2)}$; if $T_u - \theta \leq c \leq T_u + \theta$, $d_1 = I_b = \frac{(T_u+\theta-c)^2}{8(T_u^2+\theta^2)}$, and if $c \geq T_u + \theta$, $d_1 = 0$. Further, $I_a = I_b$ when $c = T_u - \theta$ and $I_b = 0$ when $c = T_u + \theta$.

existing product if cycle life is low, namely, the (IE) strategy. Its profit is:

$$\Pi^*_{IE} = \frac{a^2\left(c^2 - 2c(\theta + T_u) + (1-16d)\theta^2 + (1-16d)T_u^2 + 2\theta T_u\right)}{8\left(c^2 - 2c(\theta + T_u) + (1-8d)\theta^2 + (1-8d)T_u^2 + 2\theta T_u\right)}.$$

- *Region (c) with a high marginal cost, i.e., $c \geq T_u + \theta$, it is too costly for the firm to improve product quality. So, $K_u^* = 0$, and the firm serves the existing product to the market to get a profit of $\frac{a^2}{4}$.*

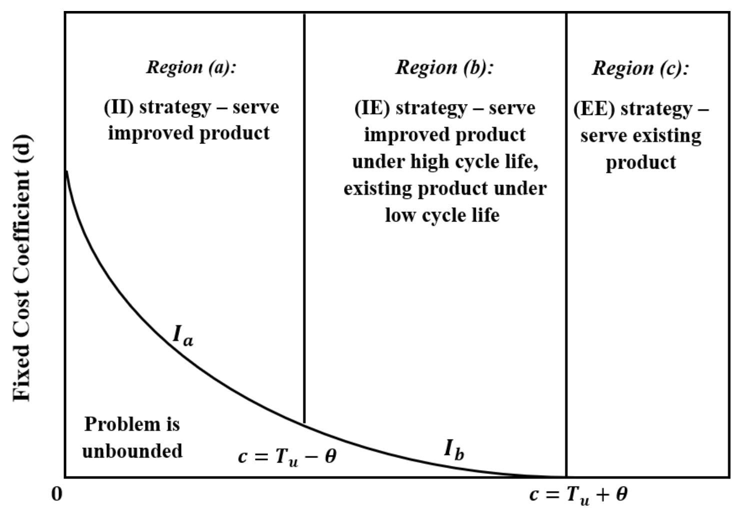

Figure 1.3: Optimal Quality/Product Decision under the Innovative Technology

By viewing Lemma 1 and Figure 1.3, we can conclude that all the three scenarios stated in Lemma 1, in terms of the product offering strategy, can be optimal depending on the model parameters. Proposition 2 further characterizes the corresponding level of capacity improvement, and the optimal profit in each scenario. There are a number of points that are worth noting about this result.

First of all, by following the expression of the optimal level of capacity improvement, K_u^*, it is straightforward to show that K_u^* always decreases in the marginal production cost, c, or in the fixed cost coefficient, d. This is quite intuitive but important to understand the three regions presented in the above result.

Let us start with region (c), where the marginal cost is high, $c \geq T_u + \theta$. The outcome in this region is quite intuitive since it is too costly to produce the improved product. So, it is natural for the firm to choose not to improve the existing product, i.e., the (EE) strategy. As the marginal cost parameter, c decreases, we are moving to region (b), where the firm is now more willing to invest in capacity improvement in the second stage. However, such investment is only for the scenario where innovation results in high cycle life. Indeed, a high cycle life would serve as a multiplier in front of the improved capacity (as indicated in the demand function in equation (1.3)), which leads to high demand and benefit for the firm from offering the improved product. If cycle life is low, then according to Lemma 1, the firm cannot find any price to guarantee positive sales profit if it sells the improved product. In this case, the firm should forgo the improved product and offer the existing product instead. This strategy is also known as the (IE) strategy. Finally, note that the existence of region (b) requires θ to be strictly positive. That is, $\theta > 0$. If $\theta = 0$, then it is clear that this region is gone and the model degenerates to a deterministic model where Figure 2 is also reduced to Figure 1 under the mature technology.

Now, let us consider region (a) where the marginal production cost is low, $0 \leq c \leq T_u - \theta$. The low marginal cost enables the firm to offer the improved product regardless of whether cycle life being high or low, i.e., adopt the (II) strategy. So, the firm will always invest on a higher level of capacity improvement in the second stage (comparing with what is offered in region (b) at the same fixed cost coefficient d), and then offer the improved product in the third stage regardless of cycle life realizations.

Recall that the firm ultimately needs to choose between the mature and the innovative technologies. The main difference between the two technologies is that the innovative one exerts uncertainty on cycle life. So, to solve the firm's technological choice problem, it is important to understand how the uncertainty in cycle life affects the firm's optimal decisions and profit. According to the distribution of cycle life, i.e., $\widetilde{T} = T_u + \theta$ or $\widetilde{T} = T_u - \theta$

with equal probabilities, parameter θ measures the volatility of cycle life. Proposition 2 and Figure 1.3 immediately lead to the following result in terms of the effect of θ on the firm's optimal product offering.

Proposition 3 (**Effect of Uncertainty on the Firm's Optimal Product Offering Strategy**). *The firm adopts the (IE) strategy if the level of cycle life uncertainty is high, i.e., $\theta \geq |T_u - c|$; otherwise, the firm adopts the (II) strategy if marginal cost is low, i.e., $c \leq T_u$ and uses the (EE) strategy if marginal cost is high, i.e., $c \geq T_u$, regardless of the realization of the cycle life. See also Figure 1.4.*

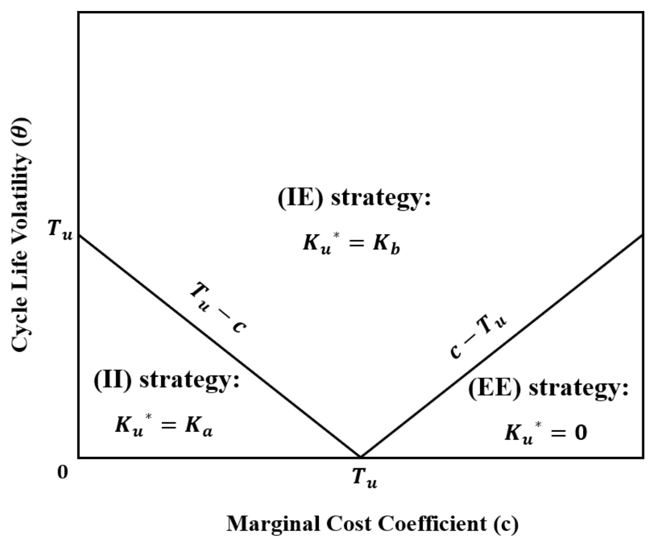

Figure 1.4: Effect of Uncertainty on the Optimal Product Offering Strategy

Here is an explanation of why the firm uses the (IE) strategy when θ is high. Recall that the overall quality improvement is measured by $K_u(T_u+\theta)$ under high cycle life and by $K_u(T_u-\theta)$ under low cycle life, where the realized value of cycle life serves as a multiplier factor that can influence the benefit of the improvement on the capacity level. When θ is sufficiently high, it motivates the firm to increase the level of capacity improvement in the second stage since the firm knows that the potential benefit of capacity improvement is amplified under high cycle life in the third stage and that the existing product can be a backup if cycle life turns out to to be relatively low. The effect of θ on the firm's optimal product offering strategy helps us interpret its effect on the firm's optimal profit in the following result.

Proposition 4 (The Effect of Uncertainty on the Firm's Optimal Profit). *Under the innovative technology, the firm's optimal profit increases in θ, the volatility of the cycle life, except when fixed cost carries sufficient weight, i.e., $d \geq \bar{d} = \frac{1}{4}$, and the level of cycle life uncertainty is small, i.e., $\theta \leq T_u - c$, in which case the firm's profit decreases in θ.*

The explanation of the effect of θ on the firm's profit in Proposition 4 is based on the effect of θ on the firm's product offering strategy in Proposition 3. Let us first consider the case when $c \geq T_u$. Proposition 3 indicates that the firm offers the existing product when θ is low, for either state of the cycle life, and it offers the (IE) strategy when θ is high. Clearly, in the low θ case, the firm's profit is independent of θ. In the high θ case, the firm always benefits from an increase in θ. This is because the firm offers the improved product only if cycle life turns out to be high where the quality improvement increases in θ due to $\widetilde{T} = T_u + \theta$. Note that if cycle life is low, θ does not affect the firm's profit since the existing product is offered in this scenario.

Now let us consider the case when $c \leq T_u$. Proposition 3 indicates that the firm offers the improved product when θ is low, regardless of the realization of the cycle life, and it offers the (IE) strategy when θ is high. For the high θ case, we can follow a similar logic that we used previously when $c \geq T_u$ to conclude that the firm's profit increases in θ. For the low θ case, the firm sells the improved product regardless of cycle life realization. In this case, the effect of an increase in θ is a double-edged sword. On the one hand, if cycle life is high ($\widetilde{T} = T_u + \theta$), higher θ leads to higher demand and benefits the firm more. On the other hand, if cycle life is low ($\widetilde{T} = T_u - \theta$), it is the opposite: higher θ leads to lower demand and hence harms the firm more. Note that in this case, low d leads to high improvement on capacity, K_u, and moreover, $K_u^* = K_a$ increases in θ when $d \leq \frac{1}{4}$. So, when d is low since cycle life serves as a multiplier to improvement on product capacity (see the demand function in equation (1.2)), as θ increases, its benefit on demand due to high cycle life outweighs its harm due to low cycle life. The firm benefits from an increase in θ when d is low, and the development cost is also small. However, as d increases, K_u becomes lower. Since K_u^* is also decreasing in θ (when $d \geq \frac{1}{4}$), its multiplier effect is reduced. With high development cost, the harm of an increase in θ on the firm's profit is higher than its benefit. As a result, the firm's profit decreases in θ. In the next subsection, it will become evident that this understanding of θ on the firm's profit plays a significant role in the firm's technology choice in the first stage.

1.4.3 Firm's Technology Choice Decisions under Equal Expected Cycle Life

We are now ready to analyze the firm's first stage technological choice decisions by comparing the optimal profits under both technologies. To make the comparison valid, we focus on regions where the optimal capacity under both technologies are bounded (i.e., $d > \max(d_0, d_1)$). In the base model, since we assume an equal cycle life for both technologies, i.e., $T_d = T_u$, neither technology has (dis)advantage on this attribute. Consequently, the firm's technological choice is largely influenced by the effect of cycle life uncertainty θ on the firm's optimal profit presented in Proposition 4. Recall from Propositions 1 and 2 that when the marginal cost is too high, $c \geq T_u + \theta$, the firm will not improve the product quality under either the deterministic or the innovative technology and the existing product will be served. In this case, the technological choice is irrelevant. So, we focus on the case where $c \leq T_u + \theta$ in the following proposition.

Proposition 5 (**Firm's Technology Choice Decisions for Product Improvement**). *The firm's preference over the two technologies is as follows (see also Figure 1.5):*

1. *For a sufficiently low marginal cost, i.e.,* $0 \leq c < \bar{c}$*, where* $\bar{c} < T_u$ $(= T_d)$*, there exists*

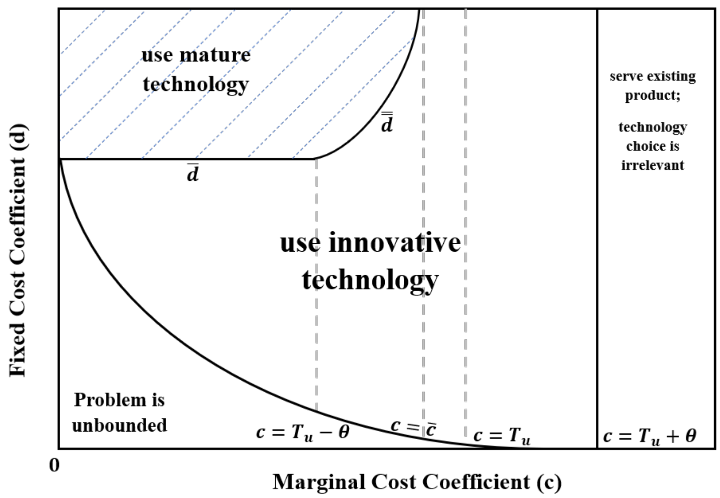

Figure 1.5: Technology Choice Decisions under Equal Expected Cycle Life ($T_d = T_u$)

a threshold value, $\tilde{d}$, such that the firm prefers the innovative technology if the fixed cost is below this threshold, i.e., $d \leq \tilde{d}$, and prefers the mature technology otherwise.[2] *Moreover, $\tilde{d}$ is (weakly) increasing in c.*

2. For a sufficiently high marginal cost, i.e., $\bar{c} \leq c \leq T_u + \theta$, the firm prefers the innovative technology.

Proposition 5 directly indicates that the mature technology is favorable when the fixed quality improvement cost coefficient, d, is high and the marginal cost, c, is low. This further implies the opposite impact of the marginal and the fixed costs on the firm's technology choice decisions, where high fixed costs encourage the mature technology while high marginal costs favor the innovative technology (if the fixed cost is not too high). The explanation for this result stems from the product offering strategy adopted in the uncertain model (presented in Figure 1.3) and the effect of uncertainty on the firm's profit.

[2] Note that $\tilde{d} = \bar{d} = \frac{1}{4}$ when $c \leq T_u - \theta$ and $\tilde{d} = \bar{\bar{d}} = \frac{(c-T_u)^2(T_u+\theta-c)^2}{4\theta^2(4c^2-8cT_u+3T_u^2)+8\theta T_u^2(c-T_u)+12T_u^2(c-T_u)^2}$ when $T_u - \theta \leq c \leq \bar{c}$, where $\bar{\bar{d}}$ always increases in c in this range, and that $\bar{d} = \bar{\bar{d}}$ when $c = T_u - \theta$. Note also that $\bar{c}$ is the left root of the quadratic function, $4\theta^2(4c^2 - 8cT_u + 3T_u^2) + 8\theta T_u^2(c - T_u) + 12T_u^2(c - T_u)^2$, the denominator of $\bar{\bar{d}}$.

We first consider the case when the marginal cost is low, i.e., when $c \leq T_u - \theta$. Under the innovative technology, according to Figure 1.3, this case falls in region (a) and the firm always provides the improved product regardless of the realization of cycle life. Under the mature technology, according to Figure 1.2, the firm also offers the improved product. Indeed, the deterministic model is a special case of the uncertain model with $\theta = 0$. In the uncertain model, according to Proposition 4 and the discussion thereafter, we observe that an increase in θ may benefit or harm the firm's profit, depending on whether the realized state of cycle life is high or low. The discussion there further indicates that for a sufficiently low fixed cost coefficient, d, i.e., $d \leq \tilde{d}$ $(= \bar{d} = \frac{1}{4})$, the benefit of a higher θ outperforms its harm which leads to an overall benefit to the firm. Hence, the firm favors the innovative technology with $\theta > 0$. On the other hand, when $d \geq \tilde{d}$, it is the opposite, and a higher θ results in net harm to the firm and the firm would prefer the mature technology with $\theta = 0$.

Next, we consider the case when the marginal cost is moderate, i.e., when $T_u - \theta \leq c \leq T_u$ $(= T_d)$. Note that this case requires $\theta > 0$. Under the innovative technology, according to Figure 1.3, this case falls in region (b) and the firm adopts the (IE) product strategy, which implies that the improved product is offered if cycle life turns out to be high while the existing product is sold if cycle life is low. Under the mature technology, Figure 1.2 indicates that the firm always offers the improved product. In this case, essentially, we compare an uncertain model with $\theta > 0$ that offers an (improved/existing) product and a deterministic model that offers an improved product. Our analysis shows that, similar to the low marginal cost case, there exists a threshold value, $\tilde{d} = \bar{\bar{d}}$, such that the innovative technology is favored if $d \leq \tilde{d}$ while the mature technology is preferred otherwise.

We explain this result based on the logic behind the firm's choice on the capacity level. In the uncertain model, the capacity level, K_u, determined in the second stage, has two direct impacts on the firm's profit. One is through the fixed development cost incurred in the second stage, $dK_u^2(T_u^2 + \theta^2)$. The other is through consumer demand (or utility) in the third stage. Due to the (IE) product strategy, its impact on the consumer's demand (or utility) is measured by $K_u(T_u+\theta)$ only when cycle life is high. Note that the existing product is offered when cycle life is low, and hence K_u is not relevant. This allows the firm to improve the product quality much more significantly in the uncertain model than that in the deterministic model as long as the fixed cost coefficient, d, is not too high. Consequently, it leads to a much higher increase in consumer demand and in the firm's profit than that in the deterministic model. Hence, innovative technology is preferred. However, as d increases to be relatively high, i.e., when $d \geq \bar{\bar{d}}$, the improvement on K_u in the uncertain model is constrained due to the high fixed cost, which further limits the increase in consumer demand and the firm's profit. This makes the uncertain model lose its advantage in outperforming the deterministic model. As a result, the firm would adopt mature technology instead.

Moreover, in the case when $T_u - \theta \leq c \leq T_u$ $(= T_d)$, our analysis shows that the threshold value for the fixed cost coefficient, $\bar{\bar{d}}$, exists only when $T_u - \theta \leq c < \bar{c}$ $(< T_u)$. Otherwise, when $c \geq \bar{c}$, the innovative technology always dominates the mature one regardless of the level of the fixed cost coefficient. Also, we show that $\bar{\bar{d}}$ increases in c in its relevant range. This implies that a higher marginal cost makes the innovative technology more favorable. Recall that in the scenario when $\bar{\bar{d}}$ is applicable, the (IE) product offering strategy is adopted under the uncertain model, and an improved product is offered under the deterministic model. As the marginal cost c increases, under the innovative technology, the firm can still be quite aggressive in setting a high level of capacity improvement. This is because the improved product is offered only when the cycle life turns out to be high (and the existing product is offered when the cycle life is low so that the high marginal cost does not apply to the high level of high capacity improvement). However, under the mature technology, the improved product is always offered, which makes the capacity improvement more significantly restricted as the marginal production cost increases, relative to that in the uncertain case. Consequently, as marginal cost c increases, the innovative technology is more likely to dominate the mature technology.

The effect of high c on the firm's technological choice continues to the case when c increases beyond T_u. That is, when $T_u \leq c \leq T_u + \theta$, the innovative technology always dominates the mature one. Note that in this case, the firm will not improve the product quality at all under the mature technology (see Figure 1.2), which is clearly dominated by the uncertain case, because no improvement is always an option in the uncertain model.

Finally, the effect of cycle life uncertainty θ on the firm's technology preference is summarized in the proposition below. This result immediately follows Proposition 4.

Proposition 6 (The Effect of Uncertainty on the Firm's Technology Preference). *As the cycle life uncertainty, θ, increases, it is more likely for the firm to adopt the innovative technology.*

An intuitive explanation for above result is that the firm has more flexibility in the product offering strategy under the innovative technology, relative to the case under the mature technology, and that an increase in θ further enhances this flexibility. Note from Proposition 3 that, as θ increases, the firm is more likely to use the (IE) product offering strategy. That is, the improved product is offered only under the high state of cycle life (where an increase in θ contributes to a higher overall quality improvement) while the existing product is offered under the low state of cycle life (where the value of θ becomes irrelevant). An alternative approach to understand this effect is from Proposition 4. Note from this proposition that the firm's profit under the innovative technology is always increasing in θ, except when $d \geq \bar{d} = \frac{1}{4}$ and $c \leq T_u - \theta$. Putting aside the exceptional area, since the firm's profit increases in θ under the innovative technology and its profit is irrelevant to θ under the mature technology, it is apparent that the innovative technology is more likely to be preferred as θ increases. Now, consider the exceptional case where the firm's profit decreases in θ. Following from Figure 1.5 that in this region, the mature technology always dominates. As θ increases, this region becomes smaller, which implies that the dominance of the mature technology becomes less likely.

1.5 The Model with Unequal Cycle Life

So far, we have assumed that the two technologies have an equal (expected) cycle life ($T_d = T_u$). As a result, the difference between the two technologies is about whether or not there is uncertainty in cycle life and the focus is about how the uncertainty, together with the cost structure, affects the firm's decisions. When we extend the model to relax the assumption on equal cycle life, the firm's technological choice will be impacted not only by the innovation uncertainty, but also by the level of the cycle life. Similar to the base model, we focus on the model parameter set where the optimal capacity improvement is upper bounded. That is, we assume that $d \geq \max(d_0, d_1)$. Recall from the setup of the base model that $T_d \in [T_u - \theta, T_u + \theta]$. Otherwise, one technology will always dominate the other. Also, similar to the base model, when marginal cost carries too much weight, $c \geq T_u + \theta$, the firm will not improve the product quality under either technology. So, we focus on the interesting case when $c \leq T_u + \theta$. The optimal decisions are presented in the following proposition and it is also displayed in Figure 1.6 on a (T_d, c) plane.

Proposition 7 (**Firm's Technology Choice Decisions under Unequal Expected Cycle Life**)**.** *There exist cut-off functions, c' and c'', where $c'' \leq c'$, such that the mature technology is always preferred when marginal cost is low, i.e, $c \leq c''$ and the innovative technology is always favored when marginal cost is high, i.e., $c \geq c'$. For medium values of marginal cost, i.e., $c'' \leq c \leq c'$, there exists a threshold value, d', such that the firm prefers the innovative technology if the fixed cost is low, i.e., $d \leq d'$ and prefers the mature technology otherwise.*[3]

[3] Note that when $0 \leq c \leq T_u - \theta$, $d' = d_a = \frac{-X_{1a}}{\theta^2(c-T_d)^2}$, $c' = c_1$, $c'' = c_3$; when $T_u - \theta \leq c < T_u + \theta$, $d' = d_b = \frac{-X_{1b}}{(c-T_d)^2(-c+\theta+T_u)^2}$, $c' = c_2$, $c'' = c_4$. Note also that $X_{1a} = 4c(T_d - T_u)(c(T_d + T_u) - 2T_dT_u) - 4\theta^2(c-T_d)^2$, $X_{1b} = c^2\left(T_d^2 - 4\left(\theta^2 + T_u^2\right)\right) - 2cT_d\left(T_d(\theta + T_u) - 4\left(\theta^2 + T_u^2\right)\right) + T_d^2\left(-3\theta^2 - 3T_u^2 + 2\theta T_u\right)$; $c_1 =$

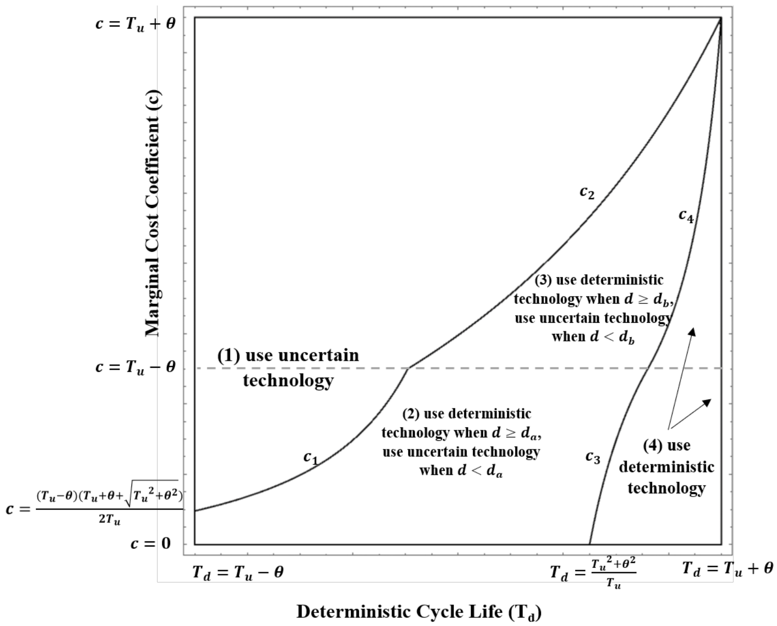

Figure 1.6: Technology Choice Decisions under Unequal Expected Cycle Life ($T_u \neq T_d$)

This result has a number of implications. First of all, in region (1) of Figure 1.6 where the marginal cost c is sufficiently high, i.e., $c \geq c'$, the firm would always adopt the innovative technology. The positive effect of high marginal costs on the innovative technology is quite consistent with the observation in the base model. However, different from the base model, if the mature technology gains significant advantage in terms of its cycle life, i.e., when $T_d \geq \frac{T_u^2+\theta^2}{T_u}$, this technology will outperform the innovative one given that the marginal cost is low enough when $c \leq c''$. See region (4) in Figure 1.6. This is not surprising since high cycle life T_d works in favor of the mature technology.

In region (2) and (3) where the marginal cost is moderate when $c'' \leq c \leq c'$, we show that there exists a threshold value for the fixed cost coefficient, d', below which the innovative technology is preferred, and the mature technology is favored otherwise. The effect of the fixed cost on the firm's technology choice decisions in the extended model with unequal cycle

---

$\frac{T_d\left(\theta^2+(T_u-T_d)\left(\sqrt{\theta^2+T_u^2}+T_u\right)\right)}{\theta^2-T_d^2+T_u^2}$, $c_2 = \frac{T_d\left(-2\sqrt{\theta^2+T_u^2}(\theta-T_d+T_u)+T_d(\theta+T_u)-4\left(\theta^2+T_u^2\right)\right)}{T_d^2-4(\theta^2+T_u^2)}$, $c_3 = \frac{2T_d\left(-\theta^2+T_dT_u-T_u^2\right)}{-\theta^2+T_d^2-T_u^2}$, $c_4 = \frac{T_d\left(-\sqrt{2}\sqrt{\theta^2+T_u^2}(\theta-T_d+T_u)+T_d(\theta+T_u)-2\left(\theta^2+T_u^2\right)\right)}{T_d^2-2(\theta^2+T_u^2)}$.

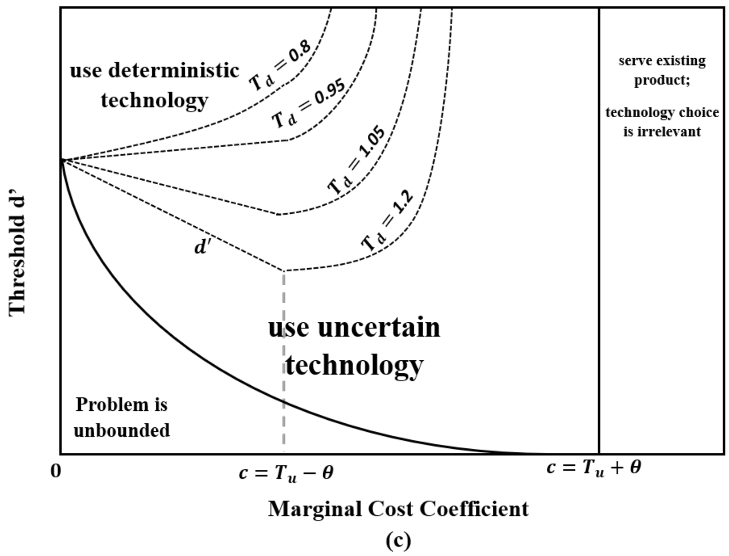

Figure 1.7: Numerical Plot: $T_u = 1$ and $\theta = 0.5$

life is quite consistent with its effect in the base model. What is interesting here is the effect of the marginal cost c. Note that d' is a function of c and other model parameters (T_d, T_u, θ) and it could take different forms depending on the values of these parameters. Figure 1.7 presents a numerical plot showing how (c, T_d) affects the threshold d'.

Following Figure 1.7, we can make a number of observations. First, if $c \geq T_u - \theta$, the threshold d' always strictly increases in c. This implies that a higher marginal cost makes the innovative technology more likely to be dominant. This is aligned with the effect of c when $c \geq c'$ in which case the innovative technology always dominates. However, if c is on the low side, $c \leq T_u - \theta$, the impact of c on d' is not obvious and is actually contingent on the relationship between the two technologies' cycle life, T_d and T_u. Specifically, an increase in c works in favor of a technology that has an advantage in terms of cycle life. For example, when $T_d > T_u$, the mature technology has a strictly higher cycle life. We observe that d' always strictly decreases in c, which implies that a higher c makes the mature technology more likely to be dominant. Recall from Figure 1.3 that when $c \leq T_u - \theta$ $(\leq T_u < T_d)$, the firm will improve the capacity level and always offer the improved product under both the mature and the innovative technologies. In order for the innovative technology to outperform the mature one, the firm needs to be able to set a much higher capacity level under the innovative technology than that under the deterministic model. However, with the same product offering strategy under both technologies, as c increases, the firm has less room to set a much higher capacity improvement level under the innovative technology in consideration of the risk of low cycle life. Together with the higher (expected) cycle life under the deterministic technology, it will be harder for the innovative technology to dominate the mature technology. In other words, the mature technology is more likely to dominate as c increases.

1.6 Conclusions, Managerial Insights, and Future Research

It is well documented that innovative technology plays a crucial role in firms' new product development process. There have been studies in the literature about over-time learning of the uncertain nature of disruptive technology. However, how the associated cost structure and the innovation uncertainty due to external factors affect firms' internal decisions such as product quality improvement, pricing, and product offering flexibility are less well studied in Operations Management. In this paper, we propose a scenario where a firm considers improving the overall quality of its existing product and faces two technology choices: one technology is more mature or deterministic, and the other one is more innovative, but at a price of development uncertainty. The firm can choose how much to improve the main quality attribute, which only partially determines the overall product quality improvement since some exogenous factors also determine the overall quality. Irrespective of which technology to adopt, there is an upfront fixed development cost associated with overall improvement. Subsequently, contingent on the product offered and demand realization, there is a unit variable production cost linked with improvement. Under the mature technology, there is no uncertainty so that the firm can predict the decisions and profit at the beginning of the decision process. Under the innovative technology, the firm needs to commit to a level of investment to improve the main quality attribute (i.e., battery capacity) before observing the uncertain state's realization. It can, however, adjust the product offering strategy, pricing, and manufacturing decisions, contingent on the realized state.

Our analysis demonstrates some interesting findings:

1. Even though the innovative technology presents uncertainty in the outcome, it provides the firm more flexibility in the product offering. Specifically, the firm can choose to forgo the improved product and offer the existing product instead of the realized state of uncertainty that turns out to be low. This phenomenon happens when the variance in the uncertainty is high.

2. It is essential to decompose the fixed development cost and the variable production cost in adopting a technology since they may affect the innovative and mature technologies differently. For example, as long as the firm exercises the flexible product offering strategy under the innovative technology, i.e., the (IE) strategy is in use, a higher fixed cost plays in favor of the mature technology. In comparison, a higher variable cost plays in favor of the innovative technology.

3. If an innovative technology can provide a high potential in increasing overall product quality when the state of uncertainty is in a good situation, then the firm is more likely to try out the innovative technology and makes it a dominant choice.

These results shed light on several managerial insights. First of all, our theoretical result on the different effects of the fixed and variable costs can provide some guidelines for policy-makers in designing governmental incentive programs to promote sustainable technologies, which are usually associated with a high level of innovation. For example, if it is desirable to encourage firms to adopt a risky yet potentially high-performing technology (e.g., novel battery technology, carbon sequestration, solar farm, etc.), the policymakers might focus on providing incentives to compensate firms' high fixed development cost without focusing on the variable cost side. Indeed, these financial incentive programs have been in place for sectors such as renewable energy technologies. Available incentives include rebates programs, tax credits, and breaks, etc., which are used to help the companies mitigate the burden due to high initial fixed costs. On the other hand, innovative technology might become more mature over time. As a result, the demand for the product under this technology would increase over time, leading to the economy of scale. Hence, we can anticipate the marginal production cost to decrease. This cost reduction may dampen the firm's incentive to carry out further innovations.

Second, our analysis indicates that in considering the innovative technology, if the firm can easily switch back to offer the existing product if the uncertain event turns out to be an unfavorable outcome, then the firm should focus more on the potential benefit of the uncertain event when it turns out to be good. The higher the potential benefit is, the more the firm should consider the innovative technology. The flexibility of the product offering strategy (without a high switching cost) is critical for innovative technology to outperform mature technology.

There are many valuable future research directions. First, a direct extension of our paper can be a competition model. There could be multiple firms who need to either simultaneously or sequentially determine which technology to adopt, given the fundamental trade-offs considered in this paper. The outcome will shed light on each firm's own product decisions and carry implications for social welfare since both firms' technology choices jointly affect the overall innovation level. Another way to incorporate competition is to consider versioning at the product level. There is again one monopolist firm in the market, but the firm may decide whether or not to keep selling the existing product when an improved product becomes available for sales in the market. Second, in the current paper, we assumed that the uncertain event in terms of product cycle life has two states under the innovative technology. Since the high state of the random event plays a more important role than the low state, it is worthwhile to explore the situation where we still have two states but with different levels of increase and decrease on the base level. This might separate the high potential and the overall variance in the uncertain event under the innovative technology. One may also follow the literature and consider learning of the uncertainty under the innovative technology over time and update its distribution belief. However, together with the internal decisions on capacity, pricing, and product offering, the model might be analytically challenging. Finally, previous literature has studied the impact of product obsolescence on the environment when the firm upgrades its existing product. It will be interesting to extend our current setting to analyze the environmental impact of product updates while considering innovation uncertainty and different technology choices.

1.7 Appendix

Summary of Notation

For ease of reference, we list below the basic notation used in Chapter 1.

- a: baseline market demand parameter;
- b: price sensitivity coefficient;
- T_d: cycle life under the mature technology;
- $\widetilde{T}$: cycle life under the innovative technology; a random variable with two possible states, high at $T_u + \theta$ and low at $T_u - \theta$;
- T_u: expected cycle life under the innovative technology;
- θ: cycle life volatility, where $\theta \in [0, T_u]$;
- K_d: capacity improvement under the mature technology;
- K_u: capacity improvement under the innovative technology;
- p_d: product sales price if mature technology is used to improve product;
- p_s: product sales price under innovative improvement when cycle life is high;
- p_f: product sales price under the innovative improvement when cycle life is low;
- p_0: product sales price if the existing product is offered;
- c: unit production cost to achieve product capacity improvement;
- d: fixed development cost coefficient for overall product quality improvement;
- Π_0: firm's profit when the existing product is offered;
- Π_d: firm's profit when the mature technology is adopted;
- Π_u: firm's overall profit when the innovative technology is adopted;
- Π_s^{stage3}: firm's stage 3 profit with the innovative technology and high cycle life;
- Π_f^{stage3}: firm's stage 3 profit with the innovative technology and low cycle life;
- Π_{IE}: firm's stage 2 expected profit when the innovative technology is adopted and the (IE) product offering strategy is used;
- Π_{II}: firm's stage 2 expected profit when the innovative technology is adopted and the firm serves the improved product in stage 3 regardless of the cycle life realization.

Proof of Proposition 1

We use backward induction to solve this optimization problem.

Stage 3 pricing problem: Given K_d determined in stage 2, the firm solves the following pricing problem:

$$
\begin{aligned}
\max_{p_d} \quad & \Pi_d = (a + K_d T_d - p_d) \cdot (p_d - cK_d) - dK_d^2 T_d^2 \\
\text{s.t.} \quad & p_d \geq 0, \;\; p_d - cK_d \geq 0, \;\; \textit{and} \;\; a + K_d T_d - p_d \geq 0,
\end{aligned}
\tag{1.8}
$$

where the constraints guarantee the non-negativity of the profit margin and demand. Given the value K_d, it is straightforward to show that the firm's profit is concave in its selling price. Hence, if $K_d(c - T_d) \leq a$, the unconstrained optimal $p_d^* = \frac{1}{2}(a + cK_d + K_d T_d)$ is also the global optimal solution and the corresponding profit is $\frac{1}{4}(a + K_d(T_d - c))^2 - dK_d^2 T_d^2$; otherwise (when $K_d(c - T_d) \geq a$), the firm will set the price high enough to generate zero sales and the firm's overall profit is $-dK_d^2 T_d^2$.

Stage 2 capacity problem: Knowing the firm's stage 3 pricing strategy, the firm sets its capacity decision in stage 2 to maximize its profit. Following from the above analysis, there are two options for the value of K_d that lead to different profit functions for the firm, i.e., option (a) $K_d(c - T_d) \leq a$ where the firm's profit function is $\frac{1}{4}(a + K_d(T_d - c))^2 - dK_d^2 T_d^2$, or option (b) $K_d(c - T_d) \geq a$ where the firm's profit function is $-dK_d^2 T_d^2$. Option (b) is always dominated. Hence, we focus on the analysis of option (a), which leads to the following four scenarios depending on the model parameters.

- When $c \leq T_d$ and $d \leq \frac{(T_d - c)^2}{4T_d^2}$, the firm's problem is unbounded. That is, K_d is set to be infinite which leads to an infinite profit.

- When $c \leq T_d$ and $d > \frac{(T_d - c)^2}{4T_d^2}$, the objective function is concave with the optimal capacity $K_d^* = \frac{a(T_d - c)}{4dT_d^2 - (T_d - c)^2}$. Accordingly, $p_d^* = \frac{-a(c^2 - cT_d - 2dT_d^2)}{4dT_d^2 - (T_d - c)^2}$, and $\Pi_d^* = \frac{a^2 dT_d^2}{4dT_d^2 - (T_d - c)^2}$.

- When $c \geq T_d$ and $d \leq \frac{(T_d-c)^2}{4T_d^2}$, the firm's profit is convex and decreasing in K_d. Hence, $K_d^* = 0$.

- When $c \geq T_d$ and $d > \frac{(T_d-c)^2}{4T_d^2}$, the firm's profit is convex. Note from the pricing problem that $K_d(c - T_d) \leq a$. We have $\Pi_d(K_d = 0) = 0$ and $\Pi_d(K_d = \frac{a}{(c-T_d)}) = -\frac{a^2 d T_d^2}{(c-T_d)^2} \leq 0$. Hence, the firm's profit is maximized at $K_d^* = 0$.

Note that when both variable and fixed cost coefficients are sufficiently low, the firm's problem is unbounded, which leads to an infinite profit, and hence is not interesting. So, for the rest of the paper, we assume that either $d \geq \frac{(T_d-c)^2}{4T_d^2}$ or $c \geq T_d$. □

Proof of Lemma 1

Given K_u determined in stage 2, in stage 3, there are two steps needed in order to characterize the firm's product offering strategy. Given a realized value of cycle life, the firm can choose to offer the improved product or the existing product. So, we first need to solve for the firm's corresponding optimal pricing decision when an existing product is offered and the pricing decision when the improved product is offered. Second, we compare the firm's profits under the existing and the improved products to configure the firm's optimal product offering strategy.

Let us start with the case when the realized cycle life is high, that is, $\widetilde{T} = T_u + \theta$. The case with a low value of cycle life can be analyzed similarly. Indeed, the analysis is also very similar to that of the firm's stage 2 problem under the mature technology.

If the improved product is offered, then the firm solves the following maximization problem:

$$\begin{aligned} \max_{p_s} \quad & \Pi_s^{stage\ 3} = [a + K_u(T_u + \theta) - p_s] \cdot (p_s - cK_u) \\ \text{s.t.} \quad & p_s \geq 0, \;\; p_s - cK_u \geq 0, \;\; and \;\; a + K_u(T_u + \theta) - p_s \geq 0. \end{aligned} \tag{1.9}$$

Given the value of K_u, it is straightforward to show that, if $K_u(c-T_u-\theta) \leq a$, the constrained optimal price is $p_s^* = \frac{1}{2}[a + K_u(T_u + \theta + c)]$; otherwise (when $K_u(c - T_u - \theta) > a$), the firm sets a price high enough to generate zero sales and profit.

If the existing product is offered, equivalently, $K_u = 0$. This reduces the firm's profit function to be $[a - p_0]p_0$ which is concave in p_0 with its optimal value at $p_0^* = \frac{a}{2}$ and profit at $\frac{a^2}{4}$.

Comparing the firm's profits under the improved product and under the existing product, we conclude that, if $c \geq T_u + \theta$, serving the existing product is a dominant strategy which yields the firm's profit as $(\frac{a^2}{4})$. If $c \leq T_u + \theta$, serving the improved product always dominates.

Let us now consider the case when the realized value of cycle life is low, that is, $\widetilde{T} = T_u - \theta$. The analysis is very similar to the case when the cycle life is high. So, if the improved product is offered, the firm solves the following problem:

$$\begin{aligned} \max_{p_f} \quad & \Pi_f^{stage\ 3} = [a + K_u(T_u - \theta) - p_f] \cdot (p_f - cK_u) \\ \text{s.t.} \quad & p_f \geq 0, \;\; p_f - cK_u \geq 0, \;\; and \;\; a + K_u(T_u - \theta) - p_f \geq 0. \end{aligned} \tag{1.10}$$

At optimality, if $K_u(c-T_u+\theta) \leq a$, the constrained optimal price $p_f^* = \frac{1}{2}[a+K_u(T_u-\theta+c)]$; otherwise (when $K_u(c-T_u+\theta) > a$), the firm sets a price high enough to generate zero sales and profit in this stage. If the existing product is offered, the analysis is done previously, the optimal price is $p_0^* = \frac{a}{2}$ and the profit is $\frac{a^2}{4}$. Comparison of the firm's profits under the improved product and the existing product leads to the following conclusion: If $c \geq T_u - \theta$, serving the existing product is the best strategy. Otherwise, when $c \leq T_u - \theta$, serving the improved product is the optimal strategy.

Combining the above cases of the high and low cycle life, we summarize the optimal pricing decisions under the innovative technology in three regions of the model parameters as follows:

- Region (a): For $c \leq T_u - \theta$, the firm offers the improved product at $p_s^* = \frac{1}{2}[a + K_u(c+$

$T_u + \theta)]$ if cycle life is high, and at $p_f^* = \frac{1}{2}[a + K_u(c + T_u - \theta)]$ if cycle life is low. This is the (II) strategy where the improved product is always offered regardless of the realization of cycle life.

- Region (b): For $T_u - \theta < c \leq T_u + \theta$, the firm offers the improved product if cycle life is high at $p_s^* = \frac{1}{2}[a + K_u(c + T_u + \theta)]$ and offers the existing product if cycle life is low. This is the (IE) strategy where the improved product is offered when the cycle life is high while the existing product is offered when the cycle life is low.

- Region (c): For $c \geq T_u + \theta$, the firm offers the existing product regardless of the realized value of cycle life, i.e., the (EE) strategy. In this case, the optimal price is $p_0^* = \frac{a}{2}$ and the optimal profit is $\Pi_{\mathrm{EE}} = \frac{a^2}{4}$.

□

Proof of Proposition 2

Given the fact that the product offering and pricing strategy in stage 3 is sensitive to the model parameters as described in the three regions in the proof of Lemma 1, we now analyze the firm's optimal level of capacity improvement in stage 2 in each of the three regions separately.

Region (a): In this region, the (II) strategy is used in stage 3. Accordingly, the firm's overall profit across both periods can be written as

$$\Pi_{\mathrm{II}} = \frac{1}{8}(a + K_u(T_u + \theta - c))^2 + \frac{1}{8}(a + K_u(T_u - \theta - c))^2 - dK_u^2\left(T_u^2 + \theta^2\right),$$

where the firm sets $K_u \geq 0$ to maximize the above profit function. This function is bounded only when the fixed cost coefficient d is sufficiently high, that is, when $d \geq I_a = \frac{(T_u - c)^2 + \theta^2}{4(T_u^2 + \theta^2)}$. Otherwise, the firm would set K_u to be infinitely high and generate infinite amount of profit, which is not an interesting case. So, we assume the condition $d \geq I_a = \frac{(T_u-c)^2+\theta^2}{4(T_u^2+\theta^2)}$ when $c \leq T_u - \theta$. Under these conditions, the firm's profit is concave in K_u with a unique optimal solution at $K_u^* = K_\mathrm{a} = \frac{a(T_u-c)}{-(T_u-c)^2-\theta^2+4d(T_u^2+\theta^2)}$.

Region (b): In this region, the (IE) strategy is used in stage 3. Accordingly, the firm's overall profit across both periods can be written as

$$\Pi_{\mathrm{IE}} = \frac{1}{8}(a + K_u(T_u + \theta - c))^2 + \frac{a^2}{8} - dK_u^2\left(T_u^2 + \theta^2\right),$$

where the firm sets $K_u \geq 0$ to maximize the above profit function. This function is bounded only when the fixed cost coefficient d is again sufficiently high, that is, when $d \geq I_b = \frac{(T_u+\theta-c)^2}{8(T_u^2+\theta^2)}$. Otherwise, the firm would set K_u to be infinitely high and generate infinite amount of profit, which is not an interesting case. So, we assume the condition $d \geq I_b = \frac{(T_u+\theta-c)^2}{8(T_u^2+\theta^2)}$ when $T_u - \theta \leq c \leq T_u + \theta$. Under these conditions, the firm's profit is concave in K_u with a unique optimal solution at $K_u^* = K_\mathrm{b} = \frac{a(T_u+\theta-c)}{-(T_u+\theta-c)^2+8d(\theta^2+T_u^2)}$.

Region (c): In this region, the firm knows that the existing product will be offered in stage 3 regardless of the realized cycle life. Hence, there is no use to invest on the product capacity improvement in stage 2 and $K_u^* = 0$.

Taking into account all cases, the optimal solutions are as follows:

$$K_u^* = \begin{cases} K_\mathrm{a} & \text{if } c \leq T_u - \theta \\ K_\mathrm{b} & \text{if } T_u - \theta \leq c \leq T_u + \theta \\ 0 & \text{if } c \geq T_u + \theta \end{cases} \tag{1.11}$$

Accordingly, the optimal profit is:

$$\Pi_u^* = \begin{cases} \frac{a^2\left(\theta^2-4d\left(\theta^2+T_u^2\right)\right)}{4(c^2-2cT_u-(4d-1)(\theta^2+T_u^2))} & \text{if } c \leq T_u - \theta \\ \frac{a^2\left(c^2-2c(\theta+T_u)+(1-16d)\theta^2+(1-16d)T_u^2+2\theta T_u\right)}{8(c^2-2c(\theta+T_u)+(1-8d)\theta^2+(1-8d)T_u^2+2\theta T_u)} & \text{if } T_u - \theta \leq c \leq T_u + \theta \\ \frac{a^2}{4} & \text{if } c \geq T_u + \theta \end{cases} \tag{1.12}$$

□

Proof of Proposition 4

We show the effect of θ on the firm's optimal profit under the innovative technology. According to Proposition 2, the firm's profit expression depends on the model parameters.

(1) When $c \geq T_u$, we consider two cases in terms of θ:

- If $\theta \leq c - T_u$, the firm offers the existing product. The optimal profit is always $\frac{a^2}{4}$, which is not affected by θ.
- If $\theta \geq c - T_u$, the firm offers the (IE) strategy where the first order derivative of the firm's profit with respect to θ is $\frac{\partial \Pi_{\mathrm{IE}}^*}{\partial \theta} = \frac{2a^2 d(T_u+\theta-c)(c\theta+T_u(T_u-\theta))}{(c^2-2c(\theta+T_u)+(1-8d)\theta^2+(1-8d)T_u^2+2\theta T_u)^2}$, which is always non-negative and the profit is increasing in θ.

In summary, when $c \geq T_u$, the firm's optimal profit is first independent of and then increases as θ increases.

(2) When $c \leq T_u$, we again consider two cases in terms of θ:

- If $\theta \leq T_u - c$, the firm offers the (II) strategy and the optimal profit is $\Pi_{\mathrm{II}}^* = \frac{a^2\left(\theta^2-4d\left(\theta^2+T_u^2\right)\right)}{4(c^2-2cT_u-(4d-1)(\theta^2+T_u^2))}$. Taking the first order derivative of Π_{II}^* with respect to θ, we have $\frac{\partial \Pi_{\mathrm{II}}^*}{\partial \theta} = -\frac{a^2(4d-1)\theta(c-T_u)^2}{2(c^2-2cT_u-(4d-1)(\theta^2+T_u^2))^2}$. It is easily observed that the effect of θ on profit depends on d. Specifically, when $d \geq 0.25$, the first order derivative is non-positive, indicating that the firm's optimal profit decreases as θ increases; when $d \leq 0.25$, the first order derivative is non-negative, and the firm's optimal profit increases in θ.

- If $\theta \geq T_u - c$, again, the firm offers the (IE) strategy where the first order derivative of the firm's profit with respect to θ is $\frac{\partial \Pi^*_{\mathrm{IE}}}{\partial \theta} = \frac{2a^2 d(T_u+\theta-c)(c\theta+T_u(T_u-\theta))}{(c^2-2c(\theta+T_u)+(1-8d)\theta^2+(1-8d)T_u^2+2\theta T_u)^2}$, which is always non-negative and the profit is increasing in θ.

In summary, when $c \leq T_u$, if $d \leq 0.25$, the firm's optimal profit always increases in θ; if $d \geq 0.25$, the firm's optimal profit first decreases and then increases in θ with the turning point at $\theta = T_u - c$. $\square$

Proof of Proposition 5

In the base model, we have $T_d = T_u$. So, we use T_u to replace T_d in this proof for ease of exposition. To guarantee bounded solution and profit under both technologies, we assume that $d \geq \max(d_0, d_1)$. Due to the fact that the optimal solution under the innovative technology is model parameter dependent, we carry out the comparison following the three regions defined previously.

Region (a) where $c \leq T_u - \theta$: In this region, the firm improves the product capacity in stage 2 and then always offers the improved product regardless of realized cycle life under the innovative technology. In this case, the model under the mature technology is a special case of the model under the innovative technology when $\theta = 0$. Hence, the comparison between the two technologies essentially reduces to examine how the firm's profit behaves in terms of θ under the innovative technology, which has been shown in Proposition 4. Following from this proposition, we conclude that the firm prefers the mature technology when $d \geq 0.25$ (and $\theta \leq T_u - c$ which is implied by this region) and prefers the innovative technology otherwise.

Region (b1) when $T_u - \theta \leq c \leq T_u$: In this region, under the innovative technology, the firm adopts the (IE) product offering strategy in stage 3. Taking the difference of the optimal profits under the two technologies, we have:

$$\Pi_{\mathrm{IE}}^{*}-\Pi_{\mathrm{d}}^{*}=\frac{a^2[(c-T_u)^2(-c+\theta+T_u)^2-4dG(c)]}{8[4dT_u^2-(T_u-c)^2][c^2-2c(\theta+T_u)+(1-8d)\theta^2+(1-8d)T_u^2+2\theta T_u]}, \quad (1.13)$$

where $G(c) = [c^2(4\theta^2+3T_u^2)+2cT_u(-4\theta^2-3T_u^2+\theta T_u)+T_u^2(3\theta^2+3T_u^2-2\theta T_u)]$. Note that the denominator of the difference function is always positive due to the assumption that $d \geq \max(d_0, d_1)$. The numerator is a linear function of d with coefficient $G(c)$ that is quadratic and convex in c. Note that $G(c = T_u-\theta) = 4\theta^2$ and $G(c = T_u) = -T_u^2\theta^2$. So, the profit difference changes its sign exactly once at $\bar{\bar{d}}$ if $T_u-\theta < c \leq \bar{c}$, and is always positive regardless of d if $\bar{c} < c \leq T_u$, where $\bar{c} = \frac{T_u\left(-2\theta\left(\sqrt{\theta^2+T_u^2}-2\theta\right)+3T_u^2-\theta T_u\right)}{4\theta^2+3T_u^2}$ and $\bar{\bar{d}} = \frac{(c-T_u)^2(-c+\theta+T_u)^2}{4\theta^2(4c^2-8cT_u+3T_u^2)+8\theta T_u^2(c-T_u)+12T_u^2(c-T_u)^2}$. Hence, if $c \leq \bar{c}$, the firm's profit difference changes its sign and prefers the innovative technology when $d \leq \bar{\bar{d}}$ and the mature technology otherwise. If $c \geq \bar{c}$, the profit difference is always positive and it is optimal to always use the innovative technology.

Finally, we prove that $\bar{\bar{d}}$ increases in c on its feasible interval, $c \in [T_u-\theta, \bar{c})$. As defined earlier, $\bar{c} = \frac{T_u\left(-2\theta\left(\sqrt{\theta^2+T_u^2}-2\theta\right)+3T_u^2-\theta T_u\right)}{4\theta^2+3T_u^2}$ is the left root of the denominator of $\bar{\bar{d}}$, which is renamed as $g(c) = 4\theta^2\,(4c^2-8cT_u+3T_u^2)+8\theta T_u^2(c-T_u)+12T_u^2(c-T_u)^2$. We now take the first order derivative of $\bar{\bar{d}}$ with respect to c which leads to $\frac{\partial\bar{\bar{d}}}{\partial c} = \frac{(T_u-c)(T_u+\theta-c)h(c)}{2g^2(c)}$, where $h(c) = (4\theta^2+3T_u^2)c^3+(3\theta T_u^2-12\theta^2T_u-9T_u^3)c^2+(-6\theta T_u^3+9\theta^2T_u^2+9T_u^4)c+\theta^3T_u^2-\theta^2T_u^3+3\theta T_u^4-3T_u^5$. Note that $\frac{(T_u-c)(T_u+\theta-c)}{2g^2(c)} \geq 0$. So, in order to show that $\bar{\bar{d}}$ increases in c, we only need to show that $h(c) \geq 0$ for any c in the given range. Note that $h(c)$ is a cubic function of c with $\frac{\partial h(c)}{\partial c} = (12\theta^2+9T_u^2)c^2+(6\theta T_u^2-24\theta^2T_u-18T_u^3)c-6\theta T_u^3+9\theta^2T_u^2+9T_u^4$, which is quadratic and convex in c and its left root coincides with $\bar{c}$, the upper bound of the feasible range of c. Hence, we have $\frac{\partial h(c)}{\partial c} \geq 0$, and accordingly, $h(c)$ increases in c and $h(c)$ reaches its lowest value when $c = T_u-\theta$. Since $h(c = T_u-\theta) = 4\theta^3(T_u-\theta)(T_u+\theta) \geq 0$, this implies that $h(c) \geq 0$. Consequently, $\frac{\partial\bar{\bar{d}}}{\partial c} \geq 0$ and $\bar{\bar{d}}$ increases in c in the given range.

Region (b2) when $T_u \leq c \leq T_u + \theta$, the firm could not profitably serve the improved product in stage 3 under mature technology so will not improve at all in stage 2. However, under the innovative technology, the firm uses the (IE) strategy and its profit is always higher than that under the no improvement strategy and hence higher than that under the mature technology.

Region (c) when $c \geq T_u + \theta$, the firm could not profitably serve the improved product in stage 3 under either technology. So, technological preference is irrelevant in this case.

$\square$

Proof of Proposition 7

Now we consider the case when $T_d \neq T_u$. Note that if $T_d \geq T_u + \theta$, then it is apparent that the mature technology always dominates. So, here we assume that $T_d \leq T_u + \theta$. Similarly, if $T_d \leq T_u - \theta$, then it is apparent that the innovative technology always dominates. So, here we assume that $T_d \geq T_u - \theta$. Note also that when $c \geq T_d$, the firm will not at all improve the capacity level of the existing product under the mature technology, which is not an interesting case to consider. So, in this proof, we focus on the case where $c \leq T_d$ and $T_u - \theta \leq T_d \leq T_u + \theta$. Given this, we separately consider the following three cases.

Case (1) when $T_d \leq T_u$ and $c \leq T_u - \theta$ which implies that $c \leq T_d$: The profit difference between the two technologies can be expressed as follows:

$$\Pi^*_{\mathrm{II}} - \Pi^*_{\mathrm{d}} = \frac{dX_{1a} + \theta^2(c - T_d)^2}{4\left(c^2 - 2cT_d + (1-4d)T_d^2\right)\left(c^2 - 2cT_u - (4d-1)\left(\theta^2 + T_u^2\right)\right)}, \tag{1.14}$$

where $X_{1a} = 4c(T_d - T_u)(c(T_d + T_u) - 2T_dT_u) - 4\theta^2(c - T_d)^2$. Both terms in the denominator are linearly decreasing in d. Given $d \geq \max(d_0, d_1)$, both terms are negative, resulting with a positive denominator. The numerator is linear in d. The coefficient of d may be positive or negative, depending on X_{1a}. Given X_{1a} is a quadratic, concave function in terms of c, negative when $c = 0$, positive when $c = T_d$, it is easy to show that for $c \in [0, T_u - \theta]$ there exists one solution to $X_{1a} = 0$, which we define as $c_1 = \frac{T_d\left(\theta^2+(T_u-T_d)\left(\sqrt{\theta^2+T_u^2}+T_u\right)\right)}{\theta^2-T_d^2+T_u^2}$. Specifically, if $c \geq c_1$, the profit difference is always positive, hence the firm always uses the innovative technology; otherwise if $c \leq c_1$, mature technology is preferred if $d \geq d'$, and innovative technology is preferred if $d \leq d'$, where $d' = d_a = \frac{-X_{1a}}{\theta^2(c-T_d)^2}$.

Case (2) when $T_d \leq T_u$ and $T_u - \theta \leq c \leq T_d$: We also compare the profit difference at optimality, which is:

$$\Pi_{\mathrm{IE}}^* - \Pi_{\mathrm{d}}^* = \frac{dX_{1b} + (c-T_d)^2(-c+\theta+T_u)^2}{8\left(c^2 - 2cT_d + (1-4d)T_d^2\right)\left(c^2 - 2c(\theta+T_u) + (1-8d)(\theta^2+T_u^2) + 2\theta T_u\right)}, \tag{1.15}$$

where $X_{1b} = c^2\left(T_d^2 - 4\left(\theta^2+T_u^2\right)\right) - 2cT_d\left(T_d(\theta+T_u) - 4\left(\theta^2+T_u^2\right)\right) + T_d^2\left(-3\theta^2 - 3T_u^2 + 2\theta T_u\right)$. The denominator is positive for reasons similar as before. Depending on the sign of X_{1b}, the numerator is either always positive, or changes sign exactly once as d increases. Given X_{1b} is a quadratic, concave function in terms of c, and positive when $c = T_d$, it is easy to show that for $c \in [0, T_u - \theta]$ there exists one solution to $X_{1b} = 0$, which we define as $c_2 = \frac{T_d\left(-2\sqrt{\theta^2+T_u^2}(\theta-T_d+T_u)+T_d(\theta+T_u)-4\left(\theta^2+T_u^2\right)\right)}{T_d^2-4(\theta^2+T_u^2)}$. In this case, if $c \geq c_2$, the profit difference is always positive, hence the firm always uses the innovative technology; otherwise if $c \leq c_2$, mature technology is preferred if $d \geq d'$, and innovative technology is preferred if $d \leq d'$, where $d' = d_b = \frac{-X_{1b}}{(c-T_d)^2(-c+\theta+T_u)^2}$.

Case (3) when $T_d \geq T_u$ and $c \leq T_u - \theta$: The specific expression of the profit difference is the same as Equation (1.14). Given $T_d \geq T_u$, it is easily observed that X_{1a} is always negative, which means the mature technology is preferred when d is sufficiently large. Since $d \geq \max(d_0, d_1)$, if $d_0 \geq d_1$, then when $d = d_0$, the optimal profit under the innovative technology is finite, which is smaller than the mature profit as it approaches infinity as d approaches d_0. So the mature technology is preferred regardless of d if $d_0 \geq d_1$, or equivalently, if $X_{2a} \leq 0$, where $X_{2a} = \theta^2(2T_d - c) + (T_d - T_u)(c(T_d + T_u) - 2T_dT_u)$. Otherwise, if $X_{2a} \geq 0$ then when $d \leq d' = d_a$, the innovative technology is preferred, and when $d \geq d' = d_a$, the mature technology is preferred. X_{2a} is a linear function of c, hence $c \geq c_3$, then there exists such d', otherwise if $c \leq c_3$, the mature technology is always preferred, where $c_3 = \frac{2T_d\left(-\theta^2+T_dT_u-T_u^2\right)}{-\theta^2+T_d^2-T_u^2}$.

Case (4) when $T_d \geq T_u$ and $T_u - \theta \leq c \leq T_d$: In this case, the firm can either use the mature technology to improve the existing product, or use the innovative technology and adopt the (IE) strategy. According to the analysis of Equation (1.15), if $c \geq c_2$, the profit difference is always positive, hence the firm always prefers the innovative technology; if $c \leq c_2$, the mature technology is preferred if d is sufficiently large. We note that $d \geq \max(d_0, d_1)$ holds, but now the expression of d_1 is different. Similar as before, if $d_0 \geq d_1$, or equivalently, if $X_{2b} \leq 0$ then the mature technology dominates for any d, where $X_{2b} = c^2\left(T_d^2 - 2\left(\theta^2 + T_u^2\right)\right) - 2cT_d\left(T_d(\theta + T_u) - 2\left(\theta^2 + T_u^2\right)\right) - T_d^2(T_u - \theta)^2$. Otherwise, if $X_{2b} \geq 0$, then the innovative technology is preferred if $d \leq d' = d_b$, and the mature technology is preferred if $d \geq d' = d_b$. Given X_{2b} is a quadratic, concave function in terms of c, and positive when $c = T_d$, then there exists one root $c_4 = \frac{T_d\left(-\sqrt{2}\sqrt{\theta^2+T_u^2}(\theta-T_d+T_u)+T_d(\theta+T_u)-2\left(\theta^2+T_u^2\right)\right)}{T_d^2-2(\theta^2+T_u^2)}$, such that if $c \geq c_4$ there exists such d', and otherwise if $c \leq c_4$ the mature technology dominates for any d. $\square$

Chapter 2

Does Customer Email Engagement Improve Profitability? Evidence from a Field Experiment of a Subscription-based Service Provider

2.1 Introduction

Subscription, defined as a business model in which customers pay a recurring fee at regular intervals, is an increasingly common way for consumers to buy access to products and services, e.g., health club, curated subscription box service, meal plan, car wash, etc. This business model has experienced tremendous growth over the past decade, especially in the consumer service and retail sectors. According to a recent report (SUBTA, 2019), by

This chapter is in conjunction with Prof. Lauren Lu at Dartmouth College, and Mr. Pengcheng Shi at AI List Capital. This chapter is word-for-word the same as the working paper "Does Customer Email Engagement Improve Profitability – Evidence from a Field Experiment of a Subscription-based Service Provider" by Yiwei Wang, Lauren Lu, and Pengcheng Shi, currently under second round review at *M&SOM* practice-based research competition.

2023, 75% of direct-to-consumer retailers will incorporate a subscription model into their businesses. Moreover, the largest subscription-based service providers had more than $2.6 billion in sales revenue in 2016, which significantly increased from $57 million in 2011 (Chen et al., 2018). From a firm's perspective, running a subscription model has many benefits, including ensuring a consistent and predictable revenue stream, facilitating personalized interactions, as well as improving customer lifetime value and profitability. From a customer's perspective, using a subscription-based service improves convenience and potentially saves money if she uses the service frequently.

Despite all these benefits, a key challenge for subscription-based companies is customer attrition. Indeed, roughly 40% of subscribers churn within six months of initial enrollment (Chen et al., 2018). To tackle this issue, subscription-based companies employ many strategies to boost customer retention. For instance, they can offer recently churned customers special promotions to motivate them to re-subscribe or maintain an interactive online review platform to facilitate more direct communications from and to customers. The predominant and most cost-effective method to improve customer retention is sending emails to customers at regular intervals, which is a key digital engagement strategy in practice (Data & Marketing Association, 2015). Engagement emails serve the purposes of both providing information on company activities and reminding customers about their subscribed services. According to a 2015 survey of firms in consumer retailing and service sectors, 81% of respondents contacted their customers more than twice a month via email in 2014, and 9 out of 10 companies declared the strategy of email engagement to be of "great strategic importance" to them (Data & Marketing Association, 2015).

As emails and other digital communications have been growing explosively in the last two decades, consumers are now constantly bombarded with marketing emails and text messages, and the effect of email engagement on customer retention has become elusive. According to Data & Marketing Association (2015), 75% of customers resent a brand after receiving excessive engagement emails from the company. Even if email engagement does increase customer retention, its net impact on firm profitability is unclear because engaged customers may increase their service consumption, thereby causing service providers' operating cost to increase. For many consumer services, the marginal cost of providing additional services is substantial compared to the relatively low subscription fee paid by individual customers. The existing literature on email engagement has primarily focused on its benefit of increasing customer retention but ignores its associated operating cost to serve retained customers. Therefore, it is not clear whether customer email engagement will improve the profitability of subscription-based service providers.

To fill these gaps in our understanding of email engagement, we seek to answer the following research questions: *(1) How does email engagement affect service subscribers' retention and service consumption? (2) How should service providers optimize their email engagement strategies to maximize profitability?* In answering these research questions, we analyze the outcome of a field experiment conducted by our partnering company, a large American car wash chain, which offers tiered subscription services to consumers and employs an RFID-based technology to track subscriber service events. This company owns 130 car wash branches in 16 U.S. states and serves over 168,000 service subscribers nationwide. We apply survival analysis and difference-in-differences methods to examine the effects of email engagement on subscriber retention and service consumption. The experiment adopts a longitudinal design with email engagement for one month and post-treatment observation for another four months. Our dataset is unique for analyzing subscriber behaviors because it contains granular, time-stamped service transaction data collected using RFID devices attached to each subscriber's car. Note that, unlike online platform or e-commerce settings where real-time tracking of individual customer transactions has become widely available, 92.8% of service transactions in the U.S. still happen in brick-and-mortar facilities (U.S. Department of Commerce, 2017), where granular-level data collection is challenging or even infeasible. Our paper presents several interesting and relevant findings. First, we observe from the field experiment that a one-month engagement with two emails separated by a half-month interval increased the likelihood of subscriber retention by 7.3% five months after the experiment started and decreased subscriber churn odds by 26.3% for the entire five-month duration. Second, we find that the same treatment increased a subscriber's per-period service consumption by 8.8%. Third, we identify two behavioral mechanisms associated with email engagement on service consumption: (1) The engagement emails acted as a reminder to subscribers and increased their service consumption immediately after they received emails, but the reminder effect decayed within the short term and exhibited fatigue after the second email; (2) The engagement emails led to habit formation of increased service consumption in the long term after the engagement stops. In sum, these results suggest whether email engagement improves profitability in subscription-based service settings depends on the relative magnitudes of the engagement effects on subscriber retention and service consumption. This finding stands in sharp contrast to the existing literature on customer engagement, which uses customer retention as the primary outcome measure and mostly finds that email engagement is always beneficial.

Building on our empirical findings, we conduct a data-driven analysis to find the optimal email engagement strategy by computing customer lifetime value and the operating cost of serving subscribers. We find that email engagement increases profit when deployed on all top-level subscribers and mid-level subscribers who infrequently utilized service but decreases profit when deployed on all basic-level subscribers and mid-level subscribers who frequently utilized service. Therefore, we recommend that the company use a selective strategy by sending engagement emails to only profitable subscribers. Our counterfactual analysis estimates that the firm can increase its profit by 13.9% if it adopts this selective engagement strategy. To conclude, our study highlights that email engagement is a double-edged sword for service providers—it increases both subscriber retention and service consumption, and it may decrease profitability when the increased operating cost to serve retained subscribers outweighs the benefit of subscriber retention. Subscription-based service providers need to adopt a data-driven approach to optimize their email engagement strategies.

2.2 Literature Review

Our work is related to four streams of literature. First, our work naturally falls within the literature of consumer behaviors in response to customer engagement strategies. In various industries, companies have employed various customer engagement strategies, such as emails, customer training, etc., to increase customer retention. Field experiments have been conducted to evaluate the effectiveness of these engagement strategies. For example, Du et al. (2020) implement different engagement strategies to study the effect of text message reminders on the loan repayment rates on a peer-to-peer lending website. Karlan et al. (2016) provide empirical evidence to show that text reminders increase deposits among microfinance customers. Similar to our paper, Calzolari and Nardotto (2017), Charness and Gneezy (2009), and Retana et al. (2016) conduct field experiments to analyze the effect of engagement on customer retention or service consumption. Calzolari and Nardotto (2017) find that sending emails increases service consumptions in a health club. Charness and Gneezy (2009) study the post-intervention effects of paying people to attend the gym. They find that providing financial incentives is effective in the formation of healthy habits. Neither of the two papers, however, investigates the effect of emails on customer retention despite their subscription settings. Retana et al. (2016) document that doing one-shot new customer training can effectively increase short-term customer retention for pay-per-use cloud computing services. However, their analysis does not provide evidence on how customer engagement affects service consumption, a key metric that will affect a service provider's operating cost and profitability. With survival analysis and difference-in-differences analysis, our paper is the first to jointly examine the effects of email engagement on the retention of subscribers and their service consumption behavior. We identify key behavioral mechanisms behind the email engagement. First, our study explores the patterns of decay and fatigue of email engagement's reminder effect that increases customers' service consumption immediately after they receive emails. Second, we find evidence that email engagement leads to long-term habit formation of increased service consumption. While the existing literature primarily focuses on the benefit of customer engagement, e.g., increasing revenue through customer retention, it ignores the costs associated with increased service consumption. Our paper conducts heterogeneous analyses over two important customer characteristics, i.e., the frequency of service consumption and the level of service subscription, in order to evaluate our industry collaborator's email engagement strategy. We recommend that the company use a selective strategy by sending engagement emails to only profitable subscribers.

The second literature our work contributes to is subscription-based service operations. Operations management researchers have developed various models to study how to manage subscription-based service operations. For example, Belavina et al. (2017) study the differences in the operational and environmental implications between a subscription model and a pay-per-use model for online grocery delivery. Both Cachon and Feldman (2011) and Randhawa and Kumar (2008) compare the profitability between subscription and pay-per-use models while considering service congestion costs. Danaher (2002) investigates the optimal subscription pricing structure for different cell phone plans. Subscription models have also been examined in information systems under the topic of bundling (Bakos and Brynjolfsson, 1999). Unlike informational goods subscription, however, consumer service subscription is usually associated with substantial marginal operating costs. In this paper, we make the first attempt to use data generated from a longitudinal field experiment and use a data-driven approach to assess the trade-off between the benefits of email engagement in improving customer retention and the increased operating costs caused by higher service consumption of engaged customers.

Third, our work contributes to the growing data-driven, practice-based research in operations management. This literature has analyzed a wide range of operational issues in the real world, such as inventory management (Caro and Gallien, 2010), pricing (Caro and Gallien, 2012; Fisher et al., 2018; Ferreira et al., 2016), information provision (Cui et al., 2019; Han et al., 2020), and product life cycle (Hu et al., 2019). Within service operations, there have been data-driven works studying delivery service (Cui et al., 2020c,b), on-demand service (Bai et al., 2019; Cui et al., 2020a), education (Zhang et al., 2017), etc. Our study is an industry-academia collaboration, and our research findings have a direct practical impact on our partnering company. Our data-driven analysis yields an optimal engagement strategy that can potentially increase our partnering company's profit by 13.9%, which demonstrates the real-world relevance of this research.

Fourth, our work is tangentially related to the literature that studies how firms can use innovative technologies to track and study consumer behavior. Many novel technologies such as RFID, WIFI-based tracking, and mobile targeting have recently been adopted to study operations management problems in specific industries such as healthcare (Staats et al., 2017), brick-and-mortar retailing (Ghose et al., 2019; Hui et al., 2013), and e-commerce (Zhang et al., 2019). We complement this literature by showing that large-scale deployment of RFID stickers in a physical setting (specifically, a car wash) is a cost-effective, convenient method to enable a granular analysis of customer service consumption behaviors. Novel data collection technologies, such as RFID stickers, are essential because unlike settings such as an online platform where tracking of customer transactions is common, most service transactions in the U.S. happen in settings where customer tracking and data collection are still not feasible.

2.3 Experiment Setting and Hypothesis Development

2.3.1 Experiment Setting

We analyze a field experiment executed at a large U.S. car wash chain. This company operates over 130 drive-through car wash branches located in 16 states and has a customer base of over 168,000 individuals. This company mainly operates using a subscription model that offers customers a fixed monthly fee to access uncapped services at any branch operated by this company. An innovative aspect of this company's operations is the use of RFID (i.e., radio-frequency identification) devices. Specifically, each subscriber is required to attach an RFID sticker underneath her car windshield. The sticker will be immediately destroyed if removed from the vehicle, thus is not transferable among subscribers. With this novel data-collection device, the company's computer system can track each service event's time and location and then link it to the subscriber's service and renewal history. By partnering with this company, we had obtained a data set which contains all records of consumer activities (e.g., service events, subscription purchases, and subscription status changes) either since January 1, 2016, or since the date the company acquired a local branch, whichever is later.

Our experiment aims to examine the effects of email engagement on customer behaviors. The company intended to achieve two purposes with email engagement: (1) to deliver information about recent activities at the company; (2) to remind customers about their service subscriptions. In this experiment, all engagement communications were sent through the email channel. Although we cannot provide the actual content of these emails due to the company's restriction, all emails' format is identical. In general, an engagement email contains three components: (1) the company's logo; (2) a car-wash-related picture; (3) correspondence with subscribers. The specific wordings of such correspondence were generated randomly by the centralized customer relationship management (CRM) software operated by a SaaS (i.e., Software as a Service) company FreshLime, which we also collaborated with for this project. Despite slight variations in wordings, all engagement emails achieve the two purposes mentioned above. The CRM software randomizes the assignment of engagement emails to participants. Therefore, the experiment's treatment process was independent of branch characteristics, e.g., our experiment does not suffer from confounding issues that different branches might apply treatments differently or that treatments might be based on subscribers' pre-treatment characteristics and transaction history. Furthermore, before the experiment started, customers did not know whether or when they would receive engagement emails. Figure 2.1 illustrates the timeline of the experiment. Because the participants

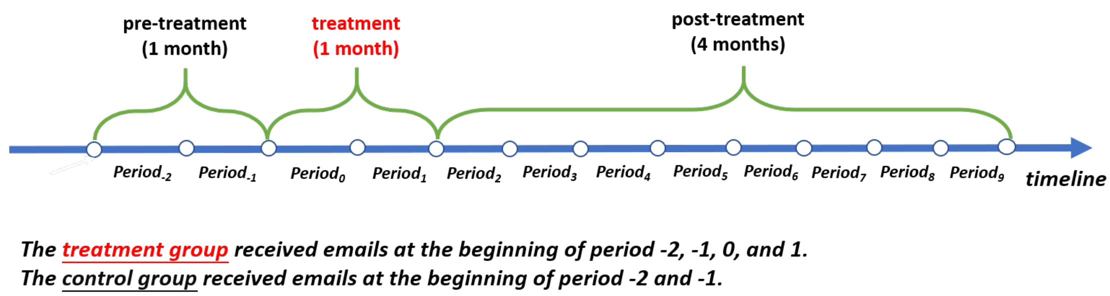

Figure 2.1: Timeline of the Experiment

received each engagement email precisely 15 days after the previous one, we define a 15-day time bucket to be one period. The entire duration of the experiment spanned a total of 12 periods or 6 months. The first month of the experiment includes period -2 and period -1, which are the pre-treatment periods, during which all experiment participants received emails at the beginning of periods -2 and -1. The second month of the experiment includes period 0 and period 1, which are the treatment periods, during which the treatment group received two additional emails at the beginning of periods 0 and 1. The control group no longer received any email during the treatment periods. From period 2 to period 9, which are the post-treatment periods, neither the treatment group nor the control group received any additional emails, but subscriber renewal and service consumption were tracked during this 4-month post-treatment period. We shall point out that the experiment's starting dates for all participants were staggered, ranging from December 1, 2018, to April 30, 2019, followed

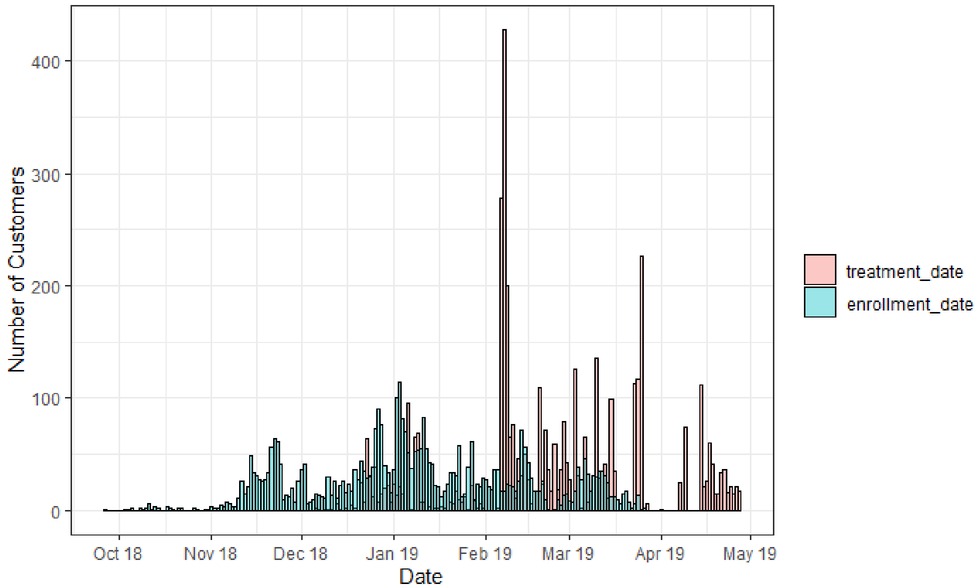

Figure 2.2: The Distribution of Customer Enrollment and Treatment Dates

by an observation period ending on October 31, 2019. To facilitate our analysis, we create a total of 22 half-month time buckets for the entire 11-month duration of this experiment. The actual experiment starting date for a participant could fall on any date during a half-month time bucket. Figure 2.2 shows the dates when the customers entered the system and received the treatment.

2.3.2 Hypothesis Development

Given the experiment setting, we now develop five testable hypotheses to study the effects of email engagements on consumer behavior.

First, as shown in Retana et al. (2016), in the context of information technology, customer engagement activities such as customer training can increase satisfaction and loyalty because the customers will better match expectations with specific features of the cloud computing service, resulting in improved customer experiences. The emotional connections developed through the engagement can also increase customer switching costs. As a result, the authors find that customer engagement effectively increases service retention. In our setting, we examine the impact of customer engagement (i.e., emails) on subscription-based service providers. So, it is plausible that similar effects would appear, and we hypothesize that email engagement would increase subscriber retention.

Hypothesis 1 (Subscriber Retention). *Email engagement increases subscriber retention.*

Second, as shown in §2.2, the literature on customer engagement finds that email engagement can effectively increase customers' service consumption in a wide range of consumer service settings (e.g., banking, health club, vaccine shots, etc.). It is reasonable to hypothesize that email engagement would have a similar impact in our setting, and therefore the treated customers' level of service consumption would increase.

Hypothesis 2 (Subscriber Consumption). *Email engagement increases subscribers' service consumption.*

Why does email engagement lead to increased service consumption? The most adopted theory is the so-called the "reminder effect." According to this theory, an engagement email serves as a behavioral stimulus, which increases the engaged customers' attention to get service (Karlan et al., 2016; Calzolari and Nardotto, 2017). If engagement emails act as a reminder, then it is reasonable to expect that this reminder effect weakens within a short period. As documented in a considerable number of psychological studies (Rubin, 1974; Baddeley, 2007), a memory stimulus (e.g., a reminder) causes temporary peaking of a subject's attention, which would decrease over time unless another stimulus arrives.

Hypothesis 3 (Reminder Effect: Decay). *Subscribers' service consumption immediately increases after receiving engagement emails. However, this effect decays even in the short term.*

In the field experiment, two treatment emails were sent to treated customers at a 15-day interval. This design allows us to test how customers respond to sequential email engagements.

Researchers find when individuals become accustomed to recurring stimuli, they become desensitized and less responsive to future stimuli they receive (Calzolari and Nardotto, 2017; Boksem and Tops, 2008). As a result, the first stimulus's effect is likely more significant than the subsequent ones. So, we conjecture that the first treatment email's effect on subscribers' consumption would be more significant than the second treatment email in our setting.

Hypothesis 4 (Reminder Effect: Fatigue). *The first treatment email's positive effect on subscribers' service consumption is larger than the second treatment email.*

In our setting, email engagement lasted for two months (including both pre-treatment and treatment periods). According to the literature, after customers maintain a high level of service consumption for a considerable long period, habituation may happen (Charness and Gneezy, 2009). As a result, the increased service consumption would persist after email engagement stopped. However, the positive effect on service consumption may weaken over time since the formed habit might eventually be unlearned in the long term without continued email engagement.

Hypothesis 5 (Habit Formation). *The positive effect of email engagement on subscribers' service consumption persists even after engagement termination, but this effect may weaken in the long term.*

2.4 Data Description

2.4.1 Sample Construction

To systematically examine the effects of email engagement on subscriber retention and service consumption, our industry collaborator conducted a field experiment that started on December 1, 2018 and finished on October 31, 2019. During this experimental period, 4,393 new customers were selected to participate in an email engagement program initiated by the company. To be eligible to participate, customers must first enroll in the company's subscription program online and get an RFID sticker to place on their windshield at any of its local branches. Then, a fixed recurring monthly fee will be deducted from a subscriber's credit card on file until she cancels the subscription. The participants' average subscription tenure at the start of the experiment was less than half a month, while the maximum subscription tenure was two months, which means these were newly enrolled customers.

Furthermore, each customer must enroll in one of the three subscription programs: basic-level, mid-level, and top-level. Basic-level subscribers pay a monthly fee between \$10 and \$20; mid-level subscribers pay a monthly fee between \$20 and \$30; top-level subscribers pay a monthly fee between \$30 and \$40. Note that the service options offered by the company are identical throughout the entire car wash chain, while the monthly subscription fee may differ slightly across branches. Among the participants, we exclude the following subscribers from our analysis: (1) those who canceled subscriptions before the treatment started, (2) promotion subscribers who paid no subscription fees. This exclusion reduces our final sample to 4,077 customers, among whom 1,435 (35.2%) customers chose the basic-level program, 1,763 (43.2%) chose the mid-level program, and 879 (21.6%) chose the top-level program. Table 2.1 provides the definition and summary statistics of key variables used in our study.

Among all experiment participants, 1,626 (40%) were assigned to the control group, and 2,451 (60%) were assigned to the treatment group. As a balance check, we confirm that both the treatment and control groups were comparable along with a set of important pre-treatment variables, including the number of distinct branches visited before treatment, the subscription tenure before treatment, and service consumption in period -2 and -1. A non-parametric test indicates that all p-values for comparing these variables between the treatment and control groups are above 0.277, thus statistically insignificant. Table 2.2 shows the result of this balance check. Since the treatment began, only a small fraction of

Variable	Definition	Obs.	Mean	SD
Basic-level Subscribers				
Treat	1 if a subscriber received treatment, and 0 otherwise	1435	0.53	0.50
Distinct Branches Visited	Number of distinct branches visited before treatment	1435	1.34	0.63
Subscription Tenure	Total months of subscription prior to treatment	1435	1.52	0.40
Total Prior Consumption	Service consumption (visits) in period$_{-2}$ and period$_{-1}$	1435	2.87	2.33
Monthly Subscription Fee	Monthly fee paid by each subscriber	1435	14.90	1.53
Mid-level Subscribers				
Treat	1 if a subscriber received treatment, and 0 otherwise	1763	0.64	0.48
Distinct Branches Visited	Number of distinct branches visited before treatment	1763	1.37	0.68
Subscription Tenure	Total months of subscription prior to treatment	1763	1.46	0.38
Total Prior Consumption	Service consumption (visits) in period$_{-2}$ and period$_{-1}$	1763	3.15	2.73
Monthly Subscription Fee	Monthly fee paid by each subscriber	1763	24.59	3.08
Top-level Subscribers				
Treat	1 if a subscriber received treatment, and 0 otherwise	879	0.64	0.48
Distinct Branches Visited	Number of distinct branches visited before treatment	879	1.35	0.66
Subscription Tenure	Total months of subscription prior to treatment	879	1.46	0.39
Total Prior Consumption	Service consumption (visits) in period$_{-2}$ and period$_{-1}$	879	3.54	3.40
Monthly Subscription Fee	Monthly fee paid by each subscriber	879	35.66	2.39

Table 2.1: Summary Statistics

	Control	Treatment	p-value
Observations	1,626	2,451	
Distinct Branches Visited Before Treatment	1.36	1.35	0.659
	(0.66)	(0.66)	
Subscription Tenure (Month) Before Treatment	1.48	1.48	0.558
	(0.34)	(0.43)	
Consumption in Period$_{-2}$	1.07	1.03	0.337
	(1.36)	(1.29)	
Consumption in Period$_{-1}$	2.05	2.11	0.277
	(2.11)	(1.95)	

Table 2.2: Balance Check

customers ($< 4\%$) upgraded their subscription levels, and we find no evidence to prove that service upgrade behavior correlates with the treatment applied in our experiment.

2.4.2 Dependent Variables

We observe each customer's service consumption and subscription renewals for the entire 6-month duration of the experiment, as illustrated in Figure 2.1. Two main dependent variables we consider are subscriber retention and service consumption.

The subscriber retention variable captures whether a participant was renewing her service subscription by the end of each period. Given our data set, identification of a customer churn event is straightforward. Namely, if a subscriber has not renewed the service in any given month, she is marked as churned by the end of that month. The specific date of churn is then defined to be the subscriber's membership expiration date, which is one month after the date of her last subscription renewal event. Only a small fraction of subscribers ($< 2\%$) marked as churned would resume subscription at a later date during the experiment, and we find no evidence that such behavior is systematically associated with the assignment of treatment. For this group of subscribers, we identify the churn period to be the period after the first churn event occurs. Using alternative identification approaches (e.g., the period after the last churn event) or removing these customers does not qualitatively change our results. In our sample, 30.5% of participants had churned by the end of the experiment, i.e., period 9.

The service consumption variable represents the total number of visits each subscriber made in any period. It is worth noting that email engagement may affect a subscriber's service consumption through the direct effect of increasing her per-period consumption or indirectly by reducing her churn rate. Our data set's unique feature is that if a customer has no consumption in a period, we can clearly identify whether it is due to her subscription cancellation or service inactivity. Thus, in this paper, we focus on the direct effect of email engagement on a subscriber's service consumption.

2.4.3 Key Control Variables

For our survival analysis, we include two sets of control variables. The first set of variables is related to pre-treatment subscriber characteristics, including subscription tenure, subscription expense, distinct branches visited prior to the treatment date, consumption in period -1, consumption in period -2, and the subscription level (i.e., basic-level, mid-level or top-level). We denote this set of variables by X_i.

The second set of control variables is related to the timing of treatment for each participant. Specifically, we employ 12 dummy variables $StartingPeriod_{ni}$ to represent the starting time of the experiment ranging from December 1, 2018, to April 30, 2019, whose value is equal to 1 if the email engagement for subscriber i started in period n. We use two binary dummy variables $Email1Weekday_i$ and $Email2Weekday_i$ to control for whether subscriber i received the two engagement emails on a weekday or a weekend during the treatment duration. In our sample, all experiment participants survived through the pre-treatment periods (i.e., period -2 and -1), so it is unnecessary to add control variables for the pre-trend of customer retention behavior. Similar time-related control variables have been previously used for survival analysis in the literature (see, e.g., Retana et al. (2016)).

For our service consumption analysis, we will leverage our data set's panel structure by including subscriber fixed effects to control for potential pre-treatment heterogeneity at the subscriber level. With subscriber fixed effects, it is unnecessary to include other time-invariant control variables at the subscriber level. To capture the branch-level consumption heterogeneity in different periods, we include a set of time-variant control variables, denoted by Z_{jt}, which includes aggregate service consumption by subscribers and pay-per-use customers at branch j in period t. We control for time-variant service qualities and system congestion levels at different branches by including these control variables.

2.5 Empirical Methods

Our empirical strategy employs two main methods. First, we use survival analysis to investigate the effect of email engagement on subscriber retention. Second, we adopt the difference-in-differences method to analyze the effect of email engagement on service consumption.

2.5.1 Subscriber Retention

We employ linear probability and probit models to estimate the effect of email engagement on subscriber retention. The dependent variable is $Survival_i$, indicating whether a subscriber "survived" by the end of the experiment, i.e., period 9. The specifications for the linear probability and the probit models are given as follows:

$$Survival_i = \alpha_0 + \alpha_1 Treat_i + \alpha_2 X_i + \alpha_3 Time_i + \epsilon_i, \tag{2.1}$$

$$Pr(Survival_i) = \Phi(\alpha_0 + \alpha_1 Treat_i + \alpha_2 X_i + \alpha_3 Time_i + \epsilon_i). \tag{2.2}$$

$Treat_i$ is a binary variable, which equals 1 if subscriber i received an additional month of email engagement, and 0 otherwise. α_i captures the treatment effect, which will be positive if email engagement increases subscriber retention. X_i is the set of control variables that capture pre-treatment subscriber characteristics, and $Time_i$ is the set of control variables related to the timing of the experiment, both of which are described in §2.4.3.

Besides linear probability and probit models, we employ the logit hazard model to analyze email engagement's effect on subscriber churn. The logit hazard model has been widely adopted in longitudinal data analysis (Singer et al., 2003). To apply this model, we use a binary outcome variable $Churn_{it}$ to indicate how likely subscriber i will churn in period t. For subscriber i, these variables equal 0 for all periods before the churn event occurs, equal 1 for the period when the churn event occurs, and equal null for periods after the churn event. Our model is specified as follows:

$$\log\left[\frac{p(Churn_{it})=1}{1-p(Churn_{it})=1}\right]=\sum_{t=0}^{9}\beta_t D_t+\alpha_1 Treat_i+\alpha_2 X_i+\alpha_3 Time_i+\epsilon_i. \tag{2.3}$$

This logit hazard model uses a logit function to "link" all explanatory variables on the right-hand side of this equation to the outcome variable $Churn_{it}$. Thus, the term on the left-hand side of this equation represents the log-hazard odds of the churn event. D_t is the indicator variable that equals 1 for period t, and 0 otherwise. β_t represents the underlying baseline hazard that all subscribers are subject to in period t. The treatment effect is captured by the coefficient α_1, which will take a negative value if the treatment has a positive effect on churn reduction. With this model, $e^{\alpha_1}-1$ corresponds to the percentage change of the hazard odds (i.e., the probability of churn over the probability of being retained) for the treatment group. X_i and $Time_i$ are the sets of control variables identical to those used in our linear probability and probit models, and ϵ_i is the error term. It should also be noted that the logit hazard model is similar to the regular logit model, and therefore we can use the standard maximum likelihood method to estimate it.

2.5.2 Service Consumption

To estimate the effect of email engagement on service consumption, we apply the difference-in-differences (DID) model with count data and conduct a Poisson regression. The DID model has been widely used in economics for policy evaluations (see, e.g., Duflo (2001)) and recently adopted in the empirical service operations literature (Cui et al., 2020b). In our model, the dependent variable C_{ijt} is the total number of service visits of subscriber i

made at branch j in period t. The count data model (i.e., Poisson) is appropriate for our setting as the number of visits only takes non-negative integer values. Recall from §2.4.2, we focus on the effect of email engagement on service consumption conditional on a subscriber being retained. To exclude the churn effect on service consumption, we remove observations of churned customers during their post-churn periods. To account for the possible issue of over-dispersion of zero entries, we use the robust variance-covariance matrix for our Poisson maximum likelihood estimator (Retana et al., 2016). Our baseline DID model is specified as:

$$\log\left(E[C_{ijt}]\right) = \beta_0 + \beta_1 Treat_i \times Post_{it} + Post_{it} + Z_{jt} + \mu_i + \theta_t + \epsilon_{ijt}. \tag{2.4}$$

Our observation unit is a service transaction of a subscriber in a branch. Here i denotes each subscriber; j denotes a specific car wash branch; t denotes the period number; μ_i denotes the subscriber fixed effect; θ_t captures the period fixed effect; ϵ_{ijt} is the error term. $Treat_i$ is a binary variable that equals 1 for the treatment group and 0 for the control group. $Post_{it}$ equals 0 for the first two periods (i.e., period -2 and -1), and 1 for all periods (including two treatment periods and eight post-treatment periods, i.e., period 0 through period 9) after the treatment started. We incorporate the term $Post_{it}$ in our specification to account for subscriber-specific time trends. For the control group, $Post_{it}$ is also well defined even for subscribers in the control group because we know which periods are their pre-treatment periods and thus find their corresponding "post-treatment" periods. The coefficient β_1 captures the treatment effect of the two engagement emails sent during the treatment period on subscribers' service consumption.

Note that the dependent variable is a count variable, and the Poisson regression model specifies the *log* of the expected count as a function of the predictive variables (Wooldridge 2010). So, the coefficient β_1 can be interpreted as follows: with email engagement, the *log* of expected service consumption increases by β_1. In other words, given email engagement, the percentage change in the expected service consumption is $e^{\beta_1} - 1$. In our model, we observe all participants for exactly 12 periods (i.e., two pre-treatment periods, two treatment periods and eight post-treatment periods). The matrix Z_{jt} contains two vectors, which capture the total consumption for all subscribers or pay-per-use customers at each branch j in period t. Note that we can only obtain the consumption information for pay-per-use customers from the Point-of-Sale (POS) data because these customers were not equipped with the RFID tracker. The POS data were aggregated at the branch-period level instead of the subscriber-period level. These time-variant covariates control for system congestion and service quality at each facility. In our setting, the relationship between branch and subscriber is not hierarchical because each subscriber may visit multiple branches. Therefore, subscriber is the highest level for clustering, and the standard errors are clustered at the subscriber level.

To explore the decay and fatigue patterns of engagement emails' reminder effect on service consumption (i.e., Hypothesis 3 and 4), we conduct a second DID analysis focusing on a short 37-day time-frame (i.e., 7 days before and 30 days after the treatment). For this analysis, the panel data is constructed at the daily level, and we label the date any subscriber received her first treatment email as Day 0 (for the control group we also label the matching treatment date, despite no email was dispatched). Because of the daily level panel structure, for this regression, our dependent variable is a binary variable $Service_{it}$, indicating whether a subscriber i had service or not on day t. We do not use the number of visits as our dependent variable, as it is unlikely that customers will get two car washes within a day. With a binary dependent variable, we adopt the following logistic regression for our DID analysis to study the effect of email engagement on the probability of service consumption at the subscriber-day level:

$$
\begin{aligned}
Pr(Service_{it}) = {} & \beta_0 + \beta_1 Treat_i \times Day_{-1,it} + \beta_2 Treat_i \times Day_{0,it} + \beta_3 Treat_i \times Day_{1\sim7,it} \\
& + \beta_4 Treat_i \times Day_{8\sim14,it} + \beta_5 Treat_i \times Day_{15,it} + \beta_4 Treat_i \times Day_{16\sim22,it} \\
& + \beta_7 Treat_i \times Day_{23\sim29,it} + \alpha_1 X_i + \alpha_2 Time_i + Day_{-1,it} + Day_{0,it} \\
& + Day_{1\sim7,it} + Day_{8\sim14,it} + Day_{15,it} + Day_{16\sim22,it} + Day_{23\sim29,it} \\
& + Day_t + \epsilon_{it}.
\end{aligned} \tag{2.5}
$$

$Treat_i$ is a binary variable, which equals 1 if subscriber i is in the treatment group. $Day_{0,it}$, $Day_{1,it}$ etc., are dummy variables equal to 1 for the corresponding time bucket and 0 otherwise. β_2 and β_5 capture the effect of email engagement on Day 0 and Day 15 when the first and second treatment emails were dispatched. β_3 and β_4 capture the effect of email engagement on the average daily probability to get service during the two weeks following the dispatch of the first treatment emails (i.e., Day 1 to Day 7 and Day 8 to Day 14); β_6 and β_7 capture the effect of email engagement on the average daily probability to get service during the two weeks following the dispatch of the second treatment email (i.e., Day 16 to Day 22 and Day 23 to Day 29). X_i is the set of control variables that capture pre-treatment subscriber characteristics, and $Time_i$ is the set of control variables related to the experiment's timing, both of which are described in §2.4.3. Day_t is the day fixed effect (except for Day -1, Day 0, and Day 15). ϵ_{it} is the error term.

Finally, to investigate the mechanism of habit formation over post-treatment periods (i.e., Hypothesis 5), we conduct additional regression analyses to explore the dynamic, long-term effects of email engagement on service consumption. Recall that each period in our experiment contains 15 days. To have an intuitive interpretation for the long-term effects, we study the dynamics at the month level, i.e., two consecutive periods are referred to as a month, although these months do not have to coincide with the calendar months. We run the regression analysis given by the following model:

$$\begin{aligned}\log\left(E[C_{ijt}]\right) = {} & \beta_0 + \beta_1 Treat_i \times DuringTreatMonth0_{it} + \beta_2 Treat_i \times PostTreatMonth1,2_{it} \\ & + \beta_3 Treat_i \times PostTreatMonth3,4_{it} + DuringTreatMonth0_{it} + PostTreatMonth1,2_{it} \\ & + PostTreatMonth3,4_{it} + Z_{jt} + \mu_i + \theta_t + \epsilon_{ijt}.\end{aligned} \tag{2.6}$$

where C_{ijt} represents the number of services visits subscriber i received at branch j in period t. $DuringTreatMonth0_{it}$, $PostTreatMonth1,2_{it}$ and $PostTreatMonth3,4_{it}$ are dummy variables that capture dynamics of the treatment. Specifically, $DuringTreatMonth0_{it}$ equals 1 for month 0 (i.e., periods 0-1), when the treatment was being applied, and equals 0 otherwise; $PostTreatMonth1,2_{it}$ equals 1 for post-treatment months 1 and 2 (i.e., periods 2-5), and equals 0 otherwise; $PostTreatMonth3,4_{it}$ equals 1 for post-treatment months 3 and 4 (i.e., periods 6-9), and equals 0 otherwise. To this end, the coefficients β_1, β_2 and β_3 capture the treatment effect for associated time buckets. Similar DID specifications have previously been adopted to study the long-term effect of sudden removal and restoration of high-quality delivery options for an e-commerce retail platform (Cui et al., 2020b).

2.5.3 Identification

This section discusses potential issues related to identifying the causal relationship between email engagement and consumer behaviors. Causal inference has been a notoriously difficult empirical question due to endogeneity problems such as self-selection and unobserved heterogeneity (Rubin, 1974). Field experiments, however, provide a clean way to identify causal effects, overcoming potential confounding factors that may result in biased estimation of the actual treatment effect. In this research, we exploit a controlled experiment setting, where the treatment of interest is the two emails an experiment subject received in periods 0 and 1. In the following, we will show that this exogenous intervention is sufficient to allow us to identify the causal effect of email engagement in our context. First, in our experiment, the treatment application process was managed by a centralized CRM software system. In particular, new subscribers from all car wash branches were placed in a first-come, first-serve (FCFS) queue and then randomly assigned to either the treatment group or the control group. According to (Rubin, 1974), proper randomization helps establish the comparability of treatment and control groups. We validate the randomization process's effectiveness by conducting a balance check between the treatment group and the control group. The results are shown in Table 2.2. According to Table 2.2 , there is no statistically significant difference between the treatment group and the control group, for pre-treatment subscriber characteristics such as distinct branches visited, subscription tenure, and pre-treatment consumptions.

Moreover, before the treatment commenced, all experiment participants were unaware of the total number of emails they would receive, nor did they know the frequency of those emails. Therefore, our setting is free from the self-selection bias, i.e., when subjects select themselves into a group, resulting in a biased sample.

For this experiment, we track when each subscriber received an engagement email. However, we do not observe whether she opened the email. So, our analysis is focused on the notion of intention to treat (Rubin, 1974) by studying all customers who received engagement emails. The efficacy of intention to treat is of primary interest to our industry collaborator, as a service provider can only control intention to treat but not directly control whether engagement emails are viewed. Hence, the issue of whether subscribers viewed those emails is beyond the scope of this paper.

One potential concern in our experiment is unobserved inter-temporal and cross-sectional heterogeneity, which may arise because each participant received engagement emails in different time periods and consumed services at different branches. These differences can potentially correlate with the outcomes and yield a biased estimate of the average treatment effect. For all consumption analyses, we address this issue by using a panel data difference-in-differences approach, which identifies the causal effect by relying on the within-subscriber variation across time. We control for pre-treatment heterogeneity of subscriber characteristics by including the subscriber fixed effects. Besides, we include period fixed effects to account for inter-temporal heterogeneity. Finally, we incorporate key time-variant service consumption variables aggregated at the branch level to account for system congestion and service quality at each branch over time. Controlling for these time-variant covariates can effectively mitigate the problem of unobserved heterogeneity. For the survival analysis and the daily level consumption analysis, we include a set of variables to control for heterogeneous subscriber characteristics and each subscriber's experiment starting time, as described in §2.4.3.

It should be noted that our industry collaborator designed and conducted this field experiment without stratified sampling. So, the fractions of customers enrolled in different subscription levels are not equal between the treatment and control groups. As an additional robustness check, we use matching algorithms to create 1,620 treatment-control pairs from the full sample and re-estimate the treatment effect using the matched sample. To do so, we first perform an exact match between a pair of control and treatment customers according to their service level and the treatment starting period. Then we employ the nearest distance matching algorithm with Mahalanobis distance to create "equivalent" treatment-control pairs. That is, for each treatment-control pair ij, the following expression is minimized:

$$D_{ij} = \sqrt{(X_i - X_j)' S^{-1} (X_i - X_j)}. \tag{2.7}$$

Vector X contains a customer's pre-treatment characteristics, including subscription tenure at the time of treatment, the number of unique branches visited, and pre-treatment service consumption. S^{-1} is the covariance matrix between each of customer i and j's pre-treatment characteristics. After matching, for both the treatment and control groups, the percentage of customers in each service level is identical, and the total number of customers who started the experiment in each period is identical. We use this matched sample to conduct a robustness check. Estimation results obtained using the matched sample are quantitatively similar to those obtained using the full sample.

2.6 Empirical Results

We present the estimation results for the effects of email engagement on subscriber retention in §2.6.1 and the effects on service consumption in §2.6.2. §2.6.3 explores the mechanisms through which email engagement influences consumer behavior.

2.6.1 The Effect of Email Engagement on Subscriber Retention

We first visually inspect the effect of email engagement on subscriber retention. Figure 2.3 shows the cumulative survival probability for the treatment and control groups by the end of each period. At the beginning of periods -2 and -1, the first and second engagement emails were sent to all participants in the experiment, and the corresponding survival probabilities for both the treatment and control groups were 100%, as our sample is constructed to include only participants who survived at least through the treatment starting date. At the beginning of periods 0 and 1, two engagement emails were sent to the treatment group but not the control group. Figure 2.3 shows that immediately after the treatment commenced, the survival probabilities for both groups started to decline, but that of the treatment group constantly stayed above that of the control group, suggesting a positive subscriber retention effect due to the treatment. To quantify email engagement's effect on subscriber retention,

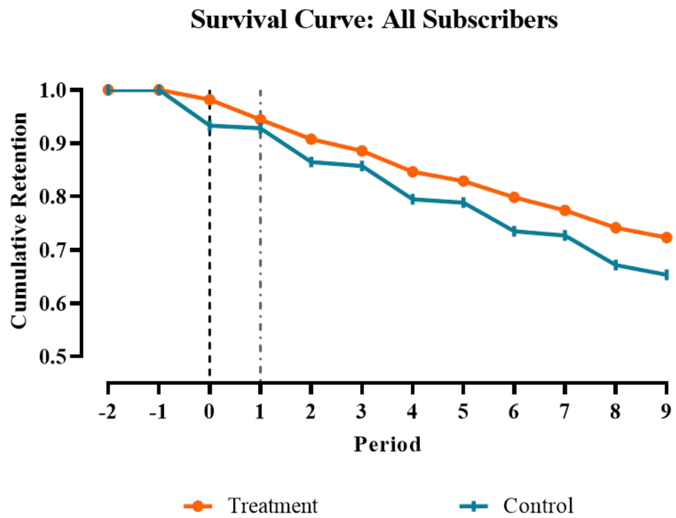

Figure 2.3: Subscriber Survival Curve

we estimate both a linear probability and a probit model. Table 2.3 Columns 1-3 reports the estimation results. To help interpret the magnitude of the coefficients, we report the corresponding marginal effect in the bottom section of the table. Results from both regressions are qualitatively and quantitatively similar and statistically significant. Column (1) shows that under the linear probability model, email engagement increased the likelihood of subscriber survival through period 9 by 7.4%. The probit model yields a similar estimation, as shown in Column (2). The estimated effect of email engagement on the subscriber churn hazard rate is presented in Column (3). The coefficient of $Treat_i$ is -0.305 and statistically significant, indicating that the treatment group had a lower churn rate. Moreover, this coefficient translates to a 26.3% reduction in the hazard odd, which indicates that the ratio of the probability of a subscriber canceling her service to the probability of the subscriber retaining her service is reduced by 26.3% by period 9, which is five months after the treatment started. Altogether, these results imply that email engagement effectively increased subscriber retention (or decreased customer churn). These results support Hypothesis 1.

In addition to the main regression analysis, we explore the heterogeneity in the effects of

Samples	Full Sample			Infrequent Users			Frequent Users		
Dependent Variable	Survival LPM	Survival Probit	Churn Hazard	Survival LPM	Survival Probit	Churn Hazard	Survival LPM	Survival Probit	Churn Hazard
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treat	0.074***	0.213***	-0.305***	0.061**	0.167**	-0.230**	0.091***	0.278***	-0.436***
	(0.016)	(0.047)	(0.066)	(0.023)	(0.064)	(0.087)	(0.023)	(0.072)	(0.103)
Marginal Effect	0.074***	0.073***		0.061**	0.059**		0.091***	0.091***	
Δ Hazard Odds			-26.3%***			-20.5%**			-35.3%***
Subscriber Characteristics	Y	Y	Y	Y	Y	Y	Y	Y	Y
Time Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	4,077	4,077	4,077	2,186	2,186	2,186	1,891	1,891	1,891

Notes. †p<0.1, *p<0.05, **p<0.01, ***p<0.001. LPM stands for linear probability model.

Table 2.3: The Effect of Email Engagement on Subscriber Retention: Full Sample, Frequent and Infrequent Users

Samples	Basic-level Subscribers			Mid-level Subscribers			Top-level Subscribers		
Dependent Variable	Survival LPM	Survival Probit	Churn Hazard	Survival LPM	Survival Probit	Churn Hazard	Survival LPM	Survival Probit	Churn Hazard
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treat	0.064**	0.199*	-0.290*	0.046†	0.135†	-0.184†	0.145***	0.387***	-0.522***
	(0.026)	(0.078)	(0.116)	(0.025)	(0.071)	(0.100)	(0.036)	(0.096)	(0.122)
Marginal Effect	0.064**	0.064*		0.046†	0.049†		0.145***	0.146***	
Δ Hazard Odds			-25.2%*			-16.8%†			-40.7%***
Subscriber Characteristics	Y	Y	Y	Y	Y	Y	Y	Y	Y
Time Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	1,435	1,435	1,435	1,763	1,763	1,763	879	879	879

Notes. †p<0.1, *p<0.05, **p<0.01, ***p<0.001. LPM stands for linear probability model.

Table 2.4: The Effect of Email Engagement on Subscriber Retention: Subscription Levels

email engagement using two important customer characteristics: subscription level and pre-treatment consumption frequency. These are two key dimensions that capture customer preference and behavioral patterns in our setting. A customer's subscription level reflects her self-selected preference of service level, while a customer's consumption frequency reveals her actual pattern of service consumption.

The heterogeneous analysis results for pre-treatment consumption frequency are reported in Table 2.3 Columns 4-9. In the field experiment, all participants were new subscribers, and the maximum subscription tenure before treatment was two months. Hence, we use the median of the total pre-treatment consumption to classify customers into two groups: infrequent users (i.e., less than or equal to four visits) and frequent users (more than four visits). We then conduct linear probability, probit, and logit hazard models for each subsample. According to Columns 6-9, email engagement reduced the hazard odds by 20.5% for infrequent users and 35.3% for frequent users. Both estimates are statistically significant. Therefore, email engagement has a stronger retention effect on frequent users than on infrequent users.

The heterogeneous analysis results for subscription levels are reported in Table 2.4. According to Columns 3, 6, and 9, email engagement reduced the hazard odds by 25.2%, 16.8%, and 40.7% for basic-level, mid-level, and top-level subscribers, respectively. All estimates are statistically significant. In sum, email engagement has the strongest retention effect on top-level subscribers and the weakest retention effect on mid-level subscribers.

2.6.2 The Effect of Email Engagement on Service Consumption

In this section, we examine the effect of email engagement on subscribers' service consumption. Figure 2.4 shows the average service consumption for the treatment and control groups during each period. In the pre-treatment periods (i.e., periods -2 and -1), the pre-trends of the treatment and control groups are almost identical to each other, which supports the parallel pre-trend assumption of our DID specification. During and after the treatment periods (i.e., periods 0 and 1), the control group's service consumption declined much more quickly than the treatment group. Moreover, this effect was persistent and lasted through the end of period 9.

After inspecting Figure 2.4, we turn to the DID regression results of service consumption using the full sample, reported in Table 2.5 Column (1). Note that there were 4,077 subscribers in the full sample. However, 55 were dropped from the DID analysis because they did not obtain any service in any period despite paying the subscription fee. This reduces the total number of subscribers to 4,022. Each subscriber might visit more than one car wash

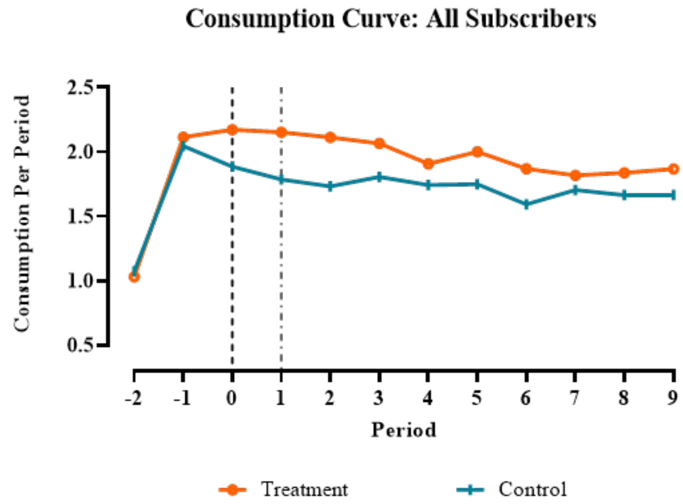

Figure 2.4: Service Consumption Curve

branch, so our regression is conducted at the subscriber-branch level. Also, we construct our panel data to be unbalanced as each subscriber-branch pair has an observation for a period only if the subscriber renews subscription through that period. Column (1) shows that the treatment group's consumption increase is positive and statistically significant with a magnitude of 8.8% ($= e^{0.084} - 1$). This result supports Hypothesis 2.

Samples	Full Sample	Infrequent Users	Frequent Users	Basic-Level Subscribers	Mid-Level Subscribers	Top-Level Subscribers
	(1)	(2)	(3)	(4)	(5)	(6)
Treat × Post	0.084***	0.135***	0.076**	0.117**	0.084**	0.065
	(0.025)	(0.043)	(0.030)	(0.038)	(0.037)	(0.061)
Branch Characteristics	Y	Y	Y	Y	Y	Y
Subscriber Characteristics	Y	Y	Y	Y	Y	Y
Time Fixed Effects	Y	Y	Y	Y	Y	Y
Observations	84,267	38,820	45,447	30,970	36,471	16,826
Log-Likelihood	-94,629	-33,558	-60,899	-34,132	-39,780	-20,666

Notes. †p<0.1, *p<0.05, **p<0.01, ***p<0.001. Standard errors are given in parentheses.

Table 2.5: The Effect of Email Engagement on Service Consumption

We next report the heterogeneous treatment effects for frequent and infrequent users. Columns 2-3 of Table 2.5 represent the treatment effect on service consumption for infrequent and frequent users. For infrequent users, email engagement increased their consumption by 14.5%, or 0.35 visits in absolute terms. For frequent users, email engagement increased their consumption by 7.9%, or 0.66 visits in absolute terms. The results of both regressions are statistically significant.

Finally, we conduct the heterogeneous analysis of service consumption for different subscription service levels. For basic-level, mid-level and top-level subscribers, email engagement increased their consumption by 12.4%, 8.8% and 6.7%, respectively. The estimation results are statistically significant for basic-level, mid-level subscribers. Taken together, the increase of service consumption is the smallest for top-level subscribers, medium for mid-level subscribers, and the largest for basic-level subscribers.

2.6.3 Evidence on Mechanisms

This section explores behavioral mechanisms through which email engagement led to the observed increase in both subscriber retention and consumption. Specifically, we present evidence on the decay and fatigue patterns of email engagement's reminder effect, and investigate the mechanism of long-term habit formation caused by email engagement.

The Reminder Effect of Email Engagement: Decay and Fatigue

To explore the decay and fatigue patterns of email engagement's reminder effect, we conduct a DID analysis (with a 37-day time-frame) on service consumption according to Equation 2.5. The estimations results are presented in Table 2.6. We first discuss the decay of the reminder effect. As expected, there is no statistically significant difference in the probability of obtaining service ($\beta_1 = -0.048$) between the treatment and control group the day before treatment (Day -1). On Day 0, when the first treatment email was sent, the treatment effect is significant and positive ($\beta_2 = 0.373$), which translates to a 37.3% increase in the daily consumption probability. From Day 1 to Day 7, the treatment effect is still positive and significant ($\beta_3 = 0.144$) but decreased to a 14.4% increase in the daily consumption probability. The estimated DID coefficient from Day 8 to Day 14 ($\beta_4 = 0.149$) is quantitatively close to the estimation for the week before. This decreasing pattern of the positive reminder effect within two weeks of receiving engagement emails supports Hypothesis 3. We next discuss

Dependent Variable	Average Daily Probability to Get Service
Treat × Day -1	-0.048
	(0.107)
Treat × Day 0	0.373***
	(0.114)
Treat × Day 1-7	0.144**
	(0.055)
Treat × Day 8-14	0.149**
	(0.055)
Treat × Day 15	0.116
	(0.117)
Treat × Day 16-22	$0.093^{\dagger}$
	(0.056)
Treat × Day 23-29	$0.106^{\dagger}$
	(0.059)
Subscriber Characteristics	Y
Time Fixed Effects	Y
Observations	150,849
Log-Likelihood	-51,952

Notes. †p<0.1, *p<0.05, **p<0.01, ***p<0.001.
Standard errors are given in parentheses. The omitted category is day -7 to -2.

Table 2.6: The Decay and Fatigue of the Reminder Effect of Email Engagement

the fatigue of the reminder effect due to repeated email engagements. On Day 15, the second treatment email was sent. For that day, the treatment effect was positive but insignificant ($\beta_5 = 0.116$). For Day 16 to Day 22, a significant increase in the daily probability to get service emerged again ($\beta_6 = 0.093$), although at a much smaller magnitude than that for Day 1 to Day 7. For Day 23 to Day 29, the estimated DID coefficient was significant at ($\beta_7 = 0.107$). If we compare the estimations of DID coefficients of the first and second treatment email, we observe that the increase in consumption probability is stronger for the first treatment email than that for the second treatment email in any time bucket (treatment day, the first week after treatment, the second week after treatment) throughout the 37-day time-frame. These observations support Hypothesis 4 that the reminder effect exhibits a pattern of fatigue.

Habit Formation Induced by Email Engagement

To investigate subscribers' long-term habit formation induced by email engagement, we conduct a dynamic DID analysis presented in Equation 2.6. Table 2.7 Columns (1) reports the regression results using the Poisson DID specification for the full sample. The increase in service consumption is positive and statistically significant in all treatment and post-treatment periods. At the same time, we also observe a decline of the treatment effect over the long run. Specifically, for month 0, the treatment coefficient takes the value of 0.116, which translates to a 12.3% ($= e^{0.116} - 1$) increase in service consumption. For month 1 and 2, we observe an 8.3% ($= e^{0.080} - 1$) increase in service consumption; for months 3 and 4, we see a 6.3% ($= e^{0.061} - 1$) increase in service consumption. These results support Hypothesis 5, i.e., the persistent post-treatment increase in service consumption indicates that email engagement induced habit formation in the treatment group. However, the weakening of such effect over time implies that email engagement, once stopped, failed to induce increased service consumption habits permanently.

Samples	Full Sample	Infrequent Users	Frequent Users	Basic Level	Mid Level	Top Level	Warm Weather	Cold Weather
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treat × $\text{DuringTreatMonth}_0$	0.116***	0.120**	0.125***	0.135***	0.097*	0.137*	0.093*	0.113***
	(0.025)	(0.042)	(0.031)	(0.039)	(0.037)	(0.060)	(0.046)	(0.031)
Treat × $\text{PostTreatMonth}_{1-2}$	0.080**	0.138*	0.073*	0.104*	0.091*	0.066	0.070	0.083*
	(0.031)	(0.058)	(0.034)	(0.044)	(0.042)	(0.082)	(0.062)	(0.034)
Treat × $\text{PostTreatMonth}_{3-4}$	0.061*	0.108*	0.059	0.097*	0.053	0.056	0.018	$0.073^{\dagger}$
	(0.030)	(0.046)	(0.038)	(0.047)	(0.046)	(0.073)	(0.052)	(0.037)
Branch Characteristics	Y	Y	Y	Y	Y	Y	Y	Y
Subscriber Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y
Time Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y
Observations	84,267	38,820	45,447	30,970	36,471	16,826	29,161	54,833
Log-Likelihood	-94,222	-33,507	-60,456	-34,010	-39,608	-20,546	-31,563	-62,152

Notes. $^{\dagger}p<0.1$, $^{*}p<0.05$, $^{**}p<0.01$, $^{***}p<0.001$. Standard errors are given in parentheses. For the heterogeneous analysis on weather, we do not have the location information of 13 subscribers.

Table 2.7: Long-Term Dynamic Effects of Email Engagement on Service Consumption

Do treated customers differ in terms of the degree to which they form consumption habits? To answer this question, we conduct three additional heterogeneous analyses to explore heterogeneous habit among different customers. We first conduct the dynamic DID (i.e., Equation 2.6) on infrequent and infrequent users. Table 2.7 Columns 2-3 summarize the results of our estimation. The estimated coefficients of treatment are significant and quantitatively similar for both infrequent users (i.e., 12.7%) and frequent users (i.e., 13.3%) during the treatment month 0. During post-treatment months 1 and 2, the estimated treatment effects are significant, at 14.8% for infrequent users and 7.6% for frequent users. During post-treatment months 3 and 4, the estimated treatment effect is significant (i.e., 11.4%) for infrequent users but insignificant for frequent users, at a level of 6.1%. The faster reduction of estimated coefficients for frequent users indicates that the level of consumption habit formation is weaker for frequent users than for infrequent users.

Next, we conduct similar regressions for customers in different subscription levels. Table 2.7 Columns 4-6 report the results. Among all three subscription levels, the decay of treatment effect is strongest for top-level subscribers, followed by mid-level subscribers, and weakest for basic-level subscribers. In other words, habit formation is strongest for basic-level subscribers, followed by the mid-level subscribers, and weakest for the top-level subscribers. In particular, for the top-level subscribers, the estimated treatment effects become insignificant immediately after the experiment ended.

Finally, because our treatment takes place during winter time, we analyze the effect of weather on the formation of consumption habits. To do that, we associate each customer with the state that she is in and use the median of the average temperature of these states to classify customers into two groups: those living in warm or cold weather states. We conduct the dynamic regression model (i.e., Equation 2.6) with respect to each group. Table 2.7 Columns 7-8 report our estimations. We observe significant treatment effects during all treatment/post-treatment months for customers living in cold weather states; however, the treatment effect is only significant during the treatment month 0 for customers living in warm weather states. Therefore, habituation of increased consumption is more pronounced for customers living in cold weather states than those living in warm weather states.

2.7 Data-Driven Email Engagement Strategies

In this section, we seek to answer the research question: does email engagement always improve profitability? To answer this question, in §2.7.1, we first conduct a 2×3 (i.e., consumption frequency $\times$ subscription level) estimation of the treatment effect of email engagement on subscriber retention and service consumption. Next, in §2.7.2, §2.7.3 and §2.7.4, we use the estimation results to conduct a data-driven analysis to optimize the email engagement strategy.

2.7.1 The Cost-Benefit Trade-off of Email Engagement

Our paper is motivated by the industry practice of subscription-based operations, where customers pay a fixed monthly subscription fee to enjoy often times unlimited service consumption. A service provider's revenue depends on the total number of subscribers, while its operating cost depends on all active subscribers' aggregate service consumption. Therefore, there exists an important trade-off in email engagement. That is, when the firm engages its subscribers, its revenue increases due to an increased retention rate; meanwhile, it incurs a higher operating cost due to increased service consumption of retained subscribers. Therefore, email engagement increases both the revenue and the cost, and the net effect on profitability is not immediately apparent without quantifying it from data.

To proceed, we adopt the notion of customer lifetime value, which is defined as the predicted total subscription revenue generated by a customer over her entire life as a subscriber. In

Samples	Basic-Level Infrequent Users			Mid-Level Infrequent Users			Top-Level Infrequent Users		
Dependent Variable	Survival LPM	Survival Probit	Churn Hazard	Survival LPM	Survival Probit	Churn Hazard	Survival LPM	Survival Probit	Churn Hazard
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treat	0.012	0.029	-0.036	0.026	0.074	-0.102	0.177***	0.469***	-0.619***
	(0.039)	(0.116)	(0.164)	(0.035)	(0.098)	(0.134)	(0.053)	(0.139)	(0.173)
Marginal Effect	0.012	0.010		0.026	0.026		0.177***	0.179***	
Δ Hazard Odds			-3.5%			-9.7%			-46.2%***
Subscriber Characteristics	Y	Y	Y	Y	Y	Y	Y	Y	Y
Time Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	795	795	795	917	917	917	419	419	419
Samples	Basic-Level Frequent Users			Mid-Level Frequent Users			Top-Level Frequent Users		
Dependent Variable	Survival LPM	Survival Probit	Churn Hazard	Survival LPM	Survival Probit	Churn Hazard	Survival LPM	Survival Probit	Churn Hazard
	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)
Treat	0.074†	0.265†	-0.396	0.067†	0.204†	-0.326*	0.136**	0.377**	-0.532***
	(0.042)	(0.154)	(0.245)	(0.036)	(0.109)	(0.157)	(0.050)	(0.137)	(0.182)
Marginal Effect	0.074†	0.072†		0.067†	0.067†		0.136**	0.137**	
Δ Hazard Odds			-32.7%			-27.8%*			-41.3%***
Subscriber Characteristics	Y	Y	Y	Y	Y	Y	Y	Y	Y
Time Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	621	621	621	821	821	821	449	449	449

Notes. †p<0.1, *p<0.05, **p<0.01, ***p<0.001. LPM stands for linear probability model.

Table 2.8: Consumption Frequency $\times$ Subscription Level Heterogeneous Analysis on Retention

Samples	Basic-level Infrequent	Mid-level Infrequent	Top-level Infrequent	Basic-level Frequent	Mid-level Frequent	Top-level Frequent
	(1)	(2)	(3)	(4)	(5)	(6)
Treat $\times$ DuringTreatMonth$_0$	0.115†	0.129*	0.153	0.147**	0.099*	0.139*
	(0.070)	(0.063)	(0.106)	(0.047)	(0.046)	(0.070)
Treat $\times$ PostTreatMonth$_{1-2}$	0.092	0.145*	0.242	0.107†	0.098†	0.012
	(0.067)	(0.061)	(0.214)	(0.055)	(0.054)	(0.074)
Treat $\times$ PostTreatMonth$_{3-4}$	0.088	0.094	0.197	0.100†	0.072	0.017
	(0.071)	(0.065)	(0.136)	(0.061)	(0.060)	(0.086)
Branch Characteristics	Y	Y	Y	Y	Y	Y
Subscriber Fixed Effects	Y	Y	Y	Y	Y	Y
Time Fixed Effects	Y	Y	Y	Y	Y	Y
Observations	14,767	17,119	6,934	16,203	19,352	9,892
Log-Likelihood	-12,933	-14,266	-6,286	-20,996	-25,226	-14,178

Notes. †p<0.1, *p<0.05, **p<0.01, ***p<0.001.

Table 2.9: Consumption Frequency $\times$ Subscription Level Heterogeneous Analysis on Consumption

other words, we will estimate the long-term effect of email engagement on subscriber retention and service consumption over infinite periods. To do that, we first create six sub-samples from the full sample by assigning each customer into one of six groups: whether the customer is an infrequent or frequent user, and whether she is enrolled in the basic-level, mid-level, or top-level subscription service. We then conduct regressions to analyze the treatment effect on both retention (i.e., Equation 2.1 – 2.3) and consumption (i.e., Equation 2.6) for each group. Our regression estimations are reported in Tables 2.8 and 2.9. We use the estimated consumption increase during the treatment month for numerical calibration, as email engagement is assumed to be repeated on a per-period basis. Technically, to accurately estimate attrition, we need a sufficiently long observation period so that enough customer churns occur. Our study chooses to use five months (entire experiment duration) to estimate the hazard rate, where email engagement stops after one month. Because firms would continuously engage customers with emails in reality, our estimate of the retention benefit of email engagement is conservative.

2.7.2 The Benefit of Increased Subscriber Retention

To compute customer lifetime value, we first estimate the baseline hazard rate for each customer group. In our sample, the average per-period hazard rate is $h_{ib} = 3.95\%$ for basic-level, infrequent users, $h_{fb} = 2.89\%$ for basic-level frequent users; $h_{im} = 4.29\%$ for mid-level infrequent users, $h_{fm} = 3.71\%$ for mid-level frequent users and $h_{it} = 6.94\%$ for top-level infrequent users, $h_{ft} = 5.49\%$ for top-level frequent users in the control group. Given the estimates of the heterogeneous engagement effects on hazard reduction in Table 2.8, the per-period hazard rates with email engagement are $h^{'}_{ib} = 3.81\%$, $h^{'}_{fb} = 1.94\%$, $h^{'}_{im} = 3.87\%$, $h^{'}_{fm} = 2.68\%$, $h^{'}_{it} = 3.73\%$ and $h^{'}_{ft} = 3.22\%$. According to Table 2.1, the monthly fees paid by basic-level, mid-level, and top-level subscribers are \$14.90, \$24.60, and \$35.70, respectively. These values translate to per-period revenue $R_b = \$7.50$, $R_m = \$12.30$, and $R_t = \$17.90$, respectively. For a customer with a churn hazard rate h and per-period revenue R, we follow the literature (e.g., Fader and Hardie (2007)) and calculate the total revenue generated over his lifetime as:

$$\sum_{i=0}^{\infty} R(1-h)^i = \frac{R}{h}. \tag{2.8}$$

Then the firm's total revenue increase due to increased retention is given by:

$$TR = \frac{R}{h'} - \frac{R}{h}. \tag{2.9}$$

2.7.3 The Cost of Increased Service Consumption

To estimate the cost of increased service consumption, we communicated with the car wash chain to estimate a set of parameters related to its operating cost. In our context, the operating cost mainly includes electricity cost ($0.5/wash), natural gas cost ($0.12/wash), water cost ($0.16/wash), chemicals cost ($0.43/wash for basic-level service, $0.64/wash for mid-level and top-level services), possible repair and maintenance of machinery ($0.47/wash), labor and administration cost ($1.8/wash for basic-level service, $2.04/wash for mid-level service and $2.22 for top-level service). This amounts to a total of $c_b = \$3.48$ for basic-level service, $c_m = \$3.93$ for mid-level service and $c_t = \$4.11$ for top-level service per wash. These estimates are consistent with survey results reported by the industry newsletter (Auto Laundry News, 2016).

Finally, we estimate the average service consumption for each subscription level to calculate the total operations costs. From our data, we calculate that the average per-period consumption is $q_{ib} = 1.12$, $q_{im} = 1.09$, and $q_{it} = 1.11$ for basic-level, mid-level, and top-level infrequent users, and $q_{fb} = 2.40$, $q_{fm} = 2.38$, $q_{ft} = 2.85$ for basic-level, mid-level, and top-level frequent users, across all periods after the treatment started. According to the results in Table 2.9, with treatment, the service consumption for infrequent users are $q'_{ib} = 1.26$, $q'_{im} = 1.24$, and $q'_{it} = 1.29$; with treatment, the service consumption for frequent users are

$q_{ib}^{'} = 2.78$, $q_{im}^{'} = 2.63$, and $q_{it}^{'} = 3.27$. We then calculate the effect of email engagement on the operating cost for a subscriber's lifetime as follows:

$$TC = c\sum_{i=0}^{\infty} q^{'}(1-h^{'})^{i} - c\sum_{i=0}^{\infty} q(1-h)^{i} = c(\frac{q}{h'} - \frac{q}{h}). \tag{2.10}$$

where c is the operations cost of each consumption; q is the consumption without email engagement in period i, and $q^{'}$ is the consumption with email engagement in period i; h is the estimated per-period hazard rate without email engagement, and $h^{'}$ is the estimated per-period hazard rate with email engagement.

2.7.4 Optimizing the Email Engagement Strategy

If we assume the cost of deploying email engagement is negligible (i.e., \$0.0001 as in our collaborator's case), then the net profit of email engagement is given by $Profit = TR - TC$. With this formula, we can numerically estimate the net profit of email engagement for each subscription level. Table 2.10 summarizes our estimates. As it turns out, deploying email engagement on top-level infrequent users is most beneficial, which can lead to a profit improvement of \$145.6 for each subscriber. Deploying email engagement on top-level frequent users can lead to a profit improvement of \$25.8. For mid-level infrequent users, the net benefit is much weaker but still positive at \$5.0. For mid-level frequent users and all basic-level users, email engagement is counterproductive, as the revenue improvement is offset by a much greater increase of the operating cost due to subscribers' increased service consumption. Strikingly, email engagement on mid-level frequent users and basic-level infrequent users yields a net reduction of \$6.1 and \$9.4 on profit, respectively. Moreover, basic-level frequent users currently contribute negative profit (= -\$29.5) even without email engagement. For this group, conducting email engagement yields a reduction of \$82.6 on profit, resulting in a total net profit of -\$112.1 per subscriber.

Samples	Basic-Level Frequent	Basic-Level Infrequent	Mid-Level Frequent	Mid-Level Infrequent	Top-Level Frequent	Top-Level Infrequent
Δ Profit	-82.6	-9.4	-6.1	5.0	25.8	145.6
95% Confidence Intervals	$[-229.4, -12.8]$	$[-43.8, 42.4]$	$[-52.2, 67.1]$	$[-52.1, 80.2]$	$[-59.5, 161.0]$	$[24.2, 318.2]$

Table 2.10: Estimated Financial Impact of Email Engagement

To conclude, we have two policy recommendations. First, the car wash chain should only target its email engagement program at all top-level service subscribers, and mid-level service subscribers who infrequently utilize service. According to our data, the total fraction of basic-level, mid-level, and top-level service subscribers are 30.9%, 40.0%, and 29.1%, respectively. Compared to no email engagement, deploying email engagement on all subscribers would result in a profit improvement of 10.8%; selective email engagement would result in a profit improvement of 24.7%. Consequently, by adopting our recommendation of a selective engagement strategy, the car wash chain can increase its profit by 13.9%. Second, we recommend that the car wash chain adjust its pricing scheme or set a service consumption limit to cut loss on basic-level subscribers. If we assume the car wash chain breaks even on serving this group of subscribers (i.e., zero profit), it can obtain an additional profit increase of 4.1%.

2.8 Conclusion

Leveraging a field experiment conducted by a U.S. car wash chain, our study is the first to jointly quantify the causal effect of email engagement on subscriber retention and service consumption for subscription-based services. We observe that the reminder effect of email engagement exhibits patterns of decay and fatigue, and that due to subscribers' habit formation, email engagement still influences their behavior in the long term after engagement termination. Our analysis indicates that email engagement is a double-edged sword that increases both the retention and service consumption of subscribers. From a service provider's perspective, a higher retention rate of subscribers increases its revenue; at the same time, additional service consumption increases its operating costs. Therefore, email engagement must be implemented with caution. We use empirical estimations from the field experiment to calibrate a data-driven model to optimize the engagement strategy for heterogeneous subscriber groups. We find that the car wash chain can increase profit by 13.9% if it adopts a selective engagement strategy. Such a strategy can be conveniently implemented at this car wash chain's 130 branches operating in 16 states via easy reprogramming of the CRM software managed by FreshLime, a Software-as-a-Service company we collaborate with for this project.

More generally, our work is relevant to all subscription businesses where the fulfillment of a physical product or service delivery incurs a substantial marginal operating cost. We hope that our work inspires other companies in the subscription space to re-examine their current email engagement policies and to conduct appropriate cost and benefit analyses. For instance, many online retailers have started to offer subscription box services to their customers (e.g., clothes, jewelry, toy, etc.). Under this business model, a firm sells the product access instead of product ownership to its subscribers for a fixed monthly fee. We note that our paper's findings also apply to the setting of product subscription, where marginal operating costs (i.e., transportation, inventory, and labor costs) are substantial.

Our research demonstrates that combining empirical methods (e.g., field experiments) and personalized data collection technologies (e.g., RFID devices) can enable researchers to investigate interesting consumer behavioral problems in the service sector where activities occur in brick-and-mortar facilities. Personalized data collection devices can allow service providers to overcome customer tracking barriers and gather granular-level customer data, which opens up opportunities for data-driven analytical research. We shall note that in our experiment, the duration of email engagement is one month, whereas in practice, email engagement can be made indefinitely until customers unsubscribe from the engagement email service. This implies that our field experiment probably captures the lower bound of the effect of email engagement, particularly in terms of habit formation. On the other hand, repeated emails might cause customers to become insensitive to or even annoyed by them and eventually unsubscribe from the email service. Moving forward, a promising future research direction is to implement email engagement experiments which allow different engagement duration and frequencies. It is likely there exists a nonlinear relationship between the effect of email engagement and engagement duration or frequencies. Finding a profit-maximizing email engagement strategy will be an interesting problem to investigate. Another limitation of our study is that no demographic information is available for the subscribers, as demographic tracking is not easily achieved in brick-and-mortar settings, even with RFID sensors. However, if additional demographic information is available, we can fine-tune the current analysis by segmenting subscribers based on demographic characteristics. This will allow for a more granular analysis and design of email engagement strategies.

Additional Tables and Figures

Service Options	Basic-Level Service	Mid-Level Service	Top-Level Service
Self-Serve Vaccums	Y	Y	Y
Wheel Cleaning	Y	Y	Y
Spot-Free Rinse	Y	Y	Y
Power Air Dry	Y	Y	Y
Triple Conditioner		Y	Y
Rain Repellent		Y	Y
Tire Shine		Y	Y
Hot Wax			Y
Liquid Glaze			Y
Bug Prep			Y

Table 2.11: All Service Options Included in Different Subscription Service Levels

Dependent Variable	Survival LPM	Survival Probit	Churn Hazard
	(1)	(2)	(3)
Treat	0.079***	0.227***	-0.333***
	(0.018)	(0.052)	(0.074)
Marginal Effect	0.079***	0.078***	
Δ Hazard Odds			-28.3%***
Subscriber Characteristics	Y	Y	Y
Time Fixed Effects	Y	Y	Y
Observations	3,240	3,240	3,240

Notes. †p<0.1, *p<0.05, **p<0.01, ***p<0.001. LPM stands for Linear Probability Model.

Table 2.12: The Effect of Email Engagement on Retention Using Matched Sample

Dependent Variable	Matched Sample Consumption
Treat × Post	0.093***
	(0.026)
Branch Characteristics	Y
Subscriber Fixed Effects	Y
Time Fixed Effects	Y
Observations	66,902
Log-Likelihood	-73,512

Notes. †p<0.1, *p<0.05, **p<0.01, ***p<0.001.

Table 2.13: The Effect of Email Engagement on Consumption Using Matched Sample

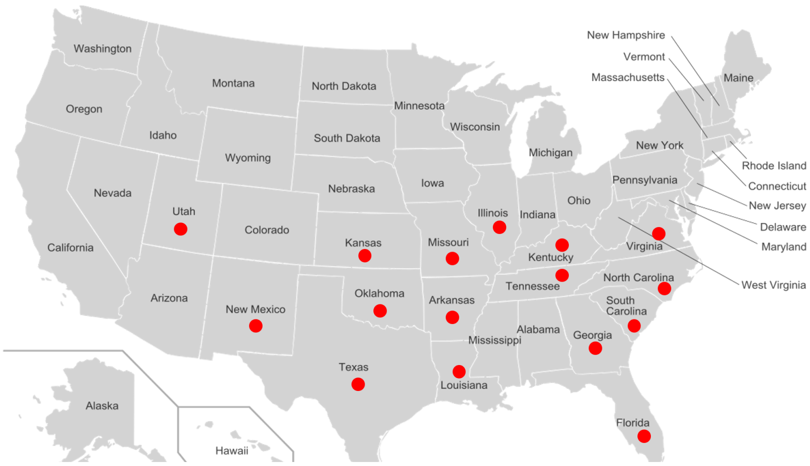

Figure 2.5: The Car Wash Company's Branch Locations

Figure 2.6: A Sample Engagement Email

Chapter 3

The Value of Curated Box Retailing — Evidence from a Field Experiment

3.1 Introduction

Curated box retailing – shipments of retailer-selected products seeking to surprise and delight customers at regular intervals, has become one of the most popular retailing trends in recent years. Some examples of product categories offered in curated boxes include apparel and accessories (Stitch Fix and Y-Closet), books (Bookroo), and cosmetic items (Birchbox and Ipsy). This business model grew over 890% from 2014 to 2018, and there were over 3,500 different kinds of boxes nationwide as of 2019 (Chen et al., 2018; Andonova et al., 2021). Since late 2019, the sudden outbreak of COVID-19 pandemic and subsequent government interventions, e.g., the declaration of national emergency and state-level "stay-at-home" or-

This chapter is in conjunction with Prof. Lauren Lu at Dartmouth College. This chapter is based on working paper "The Value of Curated Box Retailing – Evidence from a Field Experiment" by Yiwei Wang and Lauren Lu, prepared for journal review at *Management Science*.

ders, have fundamentally changed the retailing landscape. The closings of brick-and-mortar stores and showrooms forced customers to seek alternative ways to shop, which accelerated the boom of curated box retailing. According to a recent report (Forbes, 2020), out of 1,000 people surveyed, over 20% have used at least one curated box service; meanwhile, nearly all product categories sold through curated boxes have experienced double-digit growth since the beginning of COVID-19.

Besides stand-alone curated box providers, this retailing format's strong growth has also attracted a growing number of established retailers, e.g., Walmart, Target, Nordstrom, Amazon, etc. It is estimated that by 2023, 75% of direct-to-customer retailers will have incorporated a curated box channel to complement their existing retailing channels (Forbes, 2020). Compared with other novel retail strategies (e.g., showrooms and zero inventory stores), curated box retailing is a flexible and affordable way to execute two key functions of a retail channel: (1) product fulfillment; (2) information provision. For product fulfillment, curated box is a fast, convenient method for customers to buy products directly within the box (i.e., in-box sales). For information provision, curated box helps customers obtain non-digital product information through touch and feel, which reduces sales friction when they see similar products in the future. Moreover, curated box delivers assortment recommendations (i.e., exciting and novel items) that serve as guidance for customers to explore and buy unfamiliar products. Because of these advantages, an established retailer generally expects a curated box channel to amplify positive customer behaviors, e.g., acquiring prospective customers, increasing product sales quantity and variety, and decreasing excessive product sampling in the retailer's existing retailing channels.

Against the backdrop of omnichannel product offerings with minimal switching costs, customers may seek product information through the curated box channel, only to buy the actual product from a different seller (Brynjolfsson et al., 2013). It is thus unclear whether adopting a curated box retailing strategy can result in the expected outcomes. To our best knowledge, there is no extant empirical research on the impact of curated box retailing on an established retailer despite its popularity in various industries. A related stream of literature has examined the impact of channel integration, e.g., opening a new store or showroom, on customer demand in the retailer's existing retailing channels, yet the results have been inconclusive: setting up a new channel may not affect demand (Pauwels and Neslin, 2015; Avery et al., 2012), may cannibalize demand (Wang and Goldfarb, 2017; Brynjolfsson et al., 2009), but may also boost demand (Bell et al., 2018b, 2020; Zhang et al., 2019).

We provide the first empirical study to uncover the impact of introducing curated box retailing on an established retailer. Our objective is two-fold: (1) to document any positive customer behaviors induced by the curated box; (2) to identify behavioral mechanisms which drive the observed effects. To do that, we collaborate with a leading Chinese fashion retailer Y-Closet. This firm sells apparel and accessories to over 320,000 female customers in 80 cities through an online channel and a home try-on channel. Our partnership allows us to design and implement a longitudinal experimental study by randomly selecting 580 Y-Closet customers to receive monthly curated box delivery for two consecutive months and post-treatment observation for another seven months. Each box contained exactly six items for product sampling and purchase (i.e., called in-box sales) within a six-day sampling period. Granular, time-stamped consumer transaction records (i.e., in-box sales, online sales, home try-on sales, and samplings) were tracked for an entire year (i.e., three pre-treatment months, two treatment months, and seven post-treatment months). We then combine Mahalanobis matching algorithm and difference-in-differences methods to examine the causal effect of curated box retailing on customer outcomes.

Our paper providers several interesting and relevant findings. We find that the in-box sales conversion rate of curated boxes was 5.76%. The sales conversion rate was higher for the first (6.12%) than the second box (5.40%); higher for accessories (10.64%) than clothes (5.66%). Moreover, we find that monthly dispatch of one curated box for two months substantially increased overall product sales in all retail channels, and caused positive cross-channel demand spillovers to the online and home try-on sales. At the same time, curated box retailing led to a reduction in excessive product sampling and and an increase of the sales conversion rate in the home try-on channel. Also, curated boxes had a more substantial positive impact on the online sales of accessories than clothes because of fit uncertainty. We identify fundamental mechanisms associated with the curated box on product purchase and sampling. First, curated boxes impact customer behaviors through information provision. Specifically, the provision of product recommendations encourages customers' exploration and purchase of unfamiliar items and makes product search more efficient via home try-on. The provision of product tactile information causes a stronger pre-purchase uncertainty reduction for online sales than home try-on sales. Second, curated boxes lead to increased product consumption and exploration habits in the long run, such that removing the curated box channel increases sales in the remaining retail channels.

In sum, we provide the first experimental study on the curated box, which is a novel retail strategy that delivers substantial demand and operational benefits to established retailers. The exact and numerous practical benefits induced by the curated box are further decomposed into constituent parts, i.e., customers' likelihood to purchase a product, purchase quantity, and variety in the online and home try-on channels, and customers' likely to sample a product, sample quantity, and variety via home try-on, which we have examined systematically in this paper. Our research provides important implications for industry practitioners in adopting and optimizing the curated box retail strategy.

3.2 Literature Review

First, our paper contributes to the literature that studies innovative distribution approaches in retail operations management. Traditionally, retail operations management centered around brick-and-mortar stores, used both as the channel to fulfill demand and to deliver product information (Caro et al., 2020). With the rise of e-commerce, online retailers have started to reinvent their distribution channels, which disrupted the traditional store-centered retail model. This disruption has led to numerous business models that brought innovations to the two core functions of a retail channel, i.e., information and fulfillment (Bell et al., 2018b). On the one hand, traditional brick-and-mortar retailers are offering options such as "research online, pickup in-store" (Gao and Su, 2017a; Gallino and Moreno, 2014) to communicate product information online, or even transforming their physical stores into zero inventory stores (Bell et al., 2018a) to increase the operational efficiency in demand fulfillment. On the other hand, online-first retailers have adopted novel approaches to facilitate customer-product interactions. Extant literature has shown that innovative retailing strategies, such as showrooms (Bell et al., 2018b; Gao and Su, 2017b), pop-up stores (Zhang et al., 2019), and virtual fitting software (Gallino and Moreno, 2018) can deliver substantial demand and operational benefits, e.g., increasing customer engagement and decreasing excessive product return. We complement this literature stream by providing the first examination of an innovative yet increasingly popular retailing format – curated box retailing. We find monthly distribution of firm-curated products can substantially increase customers' positive behaviors, such as increased purchase quantity and variety, decreased product sampling, and improved sales conversion rate in the home try-on channel.

Second, this paper complements a literature stream that studies product sampling. Early works in this literature either focus on the sampling of physical products in a store (Heiman et al., 2001; Marks and Kamins, 1988; Bawa and Shoemaker, 2004), or sampling of information goods in an online environment (Niculescu and Wu, 2014; Chellappa and Shivendu, 2005; Cheng and Liu, 2012). Similar to our paper, a few recent papers have investigated product sampling in an omnichannel setting. In particular, Lin et al. (2019) empirically examine the relationship between offline product sampling and online reviews. Han et al. (2020) study the impact of warehouse-based free sample distribution on subsequent customer activities on an e-commerce platform. However, existing studies have exclusively focused on the information function of product sampling, where free samples are essentially a marketing tool to enhance brand awareness. In our setting, customers can choose to buy or return any product from the curated box by the end of the sampling period. To that end, curated boxes act as a retail channel that simultaneously addresses product fulfillment and information delivery. We find that in-box sales contributed to a significant fraction of overall product sales after the curated box channel was established.

Finally, our paper also contributes to the growing data-driven, practice-based research in operations management. This literature stream has examined various facets of retail operations management, such as pricing (Caro and Gallien, 2012; Chen et al., 2016; Ferreira et al., 2016; Fisher et al., 2018), inventory (Caro and Gallien, 2012; Boada-Collado and Martínez-de Albéniz, 2020; Calvo et al., 2020), assortment decisions (Bernstein and Martínez-de Albéniz, 2017; Bertsimas et al., 2018; Aouad and Segev, 2020), demand forecasting (Huang and Van Mieghem, 2014; Martínez-de Albéniz et al., 2020), and delivery (Cui et al., 2020b,c). These works utilized granular, consumer behavior data tracked in real-time, and extrapolated meaningful patterns from the data to provide actionable implications. In this research, we implemented a controlled quasi-experiment, and tracked time-stamped customer activities for an entire year. We leverage this granular data-set to conduct causal inference on channel-specified customer behaviors induced by curated boxes. Our estimated results are practically relevant, as they shed light on how industry practitioners can design or optimize their curated box delivery strategies.

Our main contributions to the literature are as follows. First, we undertake, to our knowledge, the first examination of curated box, a novel retail strategy that jointly addresses the two essential functions of a retail channel: fulfillment and information provision. The curated box is flexible, cost-effective, and efficient compared to showrooms or physical stores, which are much more operationally complex and long-lasting (Bell et al., 2018b). In fact, our industry collaborator achieved significant supply-side and demand-side benefits in this experiment with only two box deliveries. Second, while many previous empirical works on channel innovation relied on natural shocks to study aggregate, market-level outcomes, our paper leverages disaggregate, real-time transaction records of customer activities in curated box, home try-on and online channels to uncover fundamental drivers of the observed effects on customer outcomes separately for each channel. Third, we cleanly designed and implemented a controlled quasi-experiment that used physical products as the intervention. Unlike virtual interventions, e.g., text-messages (Zhang et al., 2019; Sun et al., 2019a), emails (Wang et al., 2020), software (Gallino and Moreno, 2018) or algorithms (Feldman et al., 2018), physical products convey both digital and non-digital information, which helps elucidate unidentified behavioral mechanisms. Hence, our paper makes a methodological contribution to the operations management literature that utilizes experimental tools in empirical research, see, e.g., Terwiesch (2019) and references therein. Fourth, we explore a setting where the firm simultaneously operates multiple retailing channels that face inherent discrepancies in their ability to deliver information over many key consumer preference dimensions, e.g., whether the tactile product information is provided, whether the size information is precise, whether product recommendations are available, etc. We identify previously unexplored mechanisms regarding the information function of each channel due to the presence of curated box retailing.

3.3 Institutional Background and Theoretical Motivation

3.3.1 Institutional Background

This paper studies the impact of curated box retailing on demand and the corresponding customer outcomes. We do that by collaborating with Y-Closet, a leading fashion retailer in China. Founded in 2015, Y-Closet has over 320,000 customers and received funding from major venture capital firms such as Softbank and Sequoia Capital (Crunchbase, 2021). Before 2018, Y-Closet sold apparel and accessories to female customers in over 80 Chinese cities exclusively through two channels: (1) an online channel; (2) a home try-on channel.

Primary Category	Secondary Category
Clothes	Coat, Down-suit, Dress, Formal Dress, Hoodie, Jacket, Jumpsuit, Pants, Shirt, Shorts, Skirt, Suit, Suit-Dress, Sweater, T-shirt, Top, Trench Coat,Vest
Accessories	Decorations, Bag, Bracelet, Earring, Non-prescription Glasses, Hat, Jewelry, E-commerce Items, Necklace, Ring, Scarf, Waist Seal, Watch

Notes. Clothes are associated with multiple sizes. Accessories are associated with one standard size.

Table 3.1: Product Categories Sold by Y-Closet

Under the online channel, customers can only browse information (e.g., description, pictures, and product reviews) and place orders via mobile app/website without physical product inspections before purchase. Y-Closet, however, specializes in apparel and accessories, where most of the product categories are known to have significant non-digital attributes (see Table 3.1 for detailed product categories sold by Y-Closet). Thus, the customers' inability to touch and feel products with these non-digital attributes can cause significant uncertainty for online sales (Bell et al., 2018b).

Customers who use the home try-on channel, relative to online shoppers, encounter less product uncertainty as they can physically sample the product before purchase. The process of the home try-on channel is as follows. First, customers browse product information online and select to receive a subset of products to sample at home for a limited time. When the time is up, customers can purchase any desired product and return the rest to the firm. If a customer buys a product that she previously sampled at any future point, that order is classified as a home try-on purchase.

In March 2018, Y-Closet decided to launch a curated box channel called "VIPLOOK" to complement its existing retail channels. This new channel targeted the firm's current customers with monthly deliveries of curated products. From Y-Closet's perspective, "VIPLOOK" offers two advantages. First, the curated box channel provides the customers an opportunity to touch, feel, try and experience a greater set of firm-selected products. Second, the curated box gives customers an expert opinion to search for the right product. These benefits will likely affect the customers' future shopping trajectory in terms of in-box sales, online sales, and home try-on activities. We formalize these conjectures as six testable hypotheses in the subsequent section.

3.3.2 Cross-Channel Demand Spillover Induced by Curated Boxes

As the customers could not buy products through the curated box channel before its establishment, it is evident that in-box sales would be higher after the curated box channel became available. What is not clear is the impact of curated box retailing on the retailer's existing channels. Previous studies have shown that customer-product interactions through an experience-centric retail channel can result in beneficial outcomes, e.g., enhance brand awareness and reduce pre-purchase friction. In particular, Bell et al. (2018b, 2020) and Zhang et al. (2019) find that the establishment of a retail channel increases customer demand in the online channel, which is termed the "demand spillover" effect. Specifically, Bell et al. (2018b) find that the opening of showrooms generates additional demand in the same region through the retailer's online channel. Zhang et al. (2019) leverage a randomized field experiment to show that opening a temporary pop-up store increases demand at retailers that sell related products on the same e-commerce platform yet did not participate in the pop-up store event. Bell et al. (2020) find that customers who have had experience at a "zero-inventory store" tend to deepen their understanding of an online retailer's products and order more quantities and a larger variety of products. In our case, curated box retailing is an innovative fulfillment channel which provides opportunities for customers to experience an additional set of curated products. So, it is plausible that similar effects would appear, to the extent that introducing the curated box channel would lead to positive demand spillover to the retailer's existing channels.

Hypothesis 1 (**Cross-Channel Demand Spillover**). *Curated box increases customer demand in the firm's existing retail channels.*

3.3.3 The Effects of Curated Boxes' Information Provision on Customer Behaviors

Why does the curated box lead to positive cross-channel demand spillover? The most adopted theory suggests that a novel retail channel can provide additional product information that is difficult to obtain via the retailer's existing channels, affecting customers' future shopping activities (Bell et al., 2020; Gallino and Moreno, 2018). In our setting, the curated box mainly provides information about product recommendations, and tactile product information via touch and feel. We first explore the curated box's provision of product recommendations. As curated boxes are intended to positively surprise the customers, the product assortment inside each box can act as a "billboard", which influences customers' purchase decisions by helping them find obscure items that they would not have otherwise known about (Brynjolfsson et al., 2010). If this is the case, it is reasonable that curated boxes lead to customers' product exploration, resulting in more categories of products being purchased overall, and in the firm's existing channels.

Hypothesis 2 (**Product Recommendation – Exploration**). *Curated box increases product categories purchased overall, and in the firm's existing channels.*

Besides encouraging product exploration, expert recommendations have been shown to deliver substantial operational benefits. Bell et al. (2018b, 2020) find that after acquiring product taste and fit in a showroom, customers tend to decrease excessive sampling and returns, hence become more efficient in product search. So, the customers would likely apply the recommendations obtained from the curated box to simplify product search through the retailer's other retail channels. We conjecture that the treatment group would sample less but purchase more products via home try-on. This implies that the conversion rate improves (i.e., ratio of items purchased to items sampled) in the home try-on channel.

Hypothesis 3 (**Product Recommendation – Improved Sales Conversion**). *Curated box reduces product sampling, but increases product purchase in the home try-on channel.*

Secondly, the curated box channel can delivery additional tactile information. That is, curated boxes offer opportunities for the customers to physically sample many products, which reduces the pre-purchase friction for similar products (Zhang et al., 2019) in the future. Consequently, curated boxes likely have a more substantial influence on customers' subsequent purchase activities in the online channel than in the home try-on channel. This channel-specific difference arises because product touch, feel and experience should be more beneficial for online shopping, where the lack of physical product information can cause significant sales friction (Bell et al., 2018b).

Hypothesis 4 (**Tactile Information**). *Curated box's positive effects on product sales are stronger for online sales than for home try-on sales.*

3.3.4 The Effects of Fit Uncertainty on Customer Behaviors Induced by Curated Boxes

We also explore curated boxes' heterogeneous effects on demand spillovers due to "fit uncertainty". In this research, our industry collaborator sells female fashion products. One unique feature that distinguishes fashion products from other consumer products is the size dimension. Indeed, customers face an additional risk of "fit uncertainty" when they are unsure about the correct product size, even though they might like the product design or color (Li et al., 2020; Gallino and Moreno, 2018). In our setting, this risk mainly pertains to the online channel. We further note that the products sold by Y-Closet can be classified into two categories: (1) clothes, which have multiple sizes; (2) accessories, which have one standard size. We conjecture that the level of fit uncertainty would be a much stronger impediment to the sales of clothes than accessories in the online channel, and curated boxes' positive effects on product sales are likely stronger for accessories than clothes.

Hypothesis 5 (**Fit Uncertainty**). *Curated box's positive effects on product sales are stronger for accessories than clothes in the online channel.*

3.3.5 Customer Habit Formation Induced by Curated Boxes

Apart from information provision and fit uncertainty, we examine customers' long-term habit formation induced by curated boxes. Our experiment's unique feature involves both the establishment and termination of a curated box channel. As a result, it is possible to separately analyze curated boxes' effects in treatment months (i.e., when the curated box channel was in use) and post-treatment months (i.e., after the curated box channel was suspended). Consumer behavioral research shows that a considerable period of customer engagement will lead to the formation of long-term habits (Wang et al., 2020). So, it is plausible that customers would form increased consumption habits by the end of the treatment period. In other words, the curated box's positive effects on customer behaviors would likely persist during the post-treatment period even without additional box deliveries. Moreover, we conjecture the curated box channel's termination causes product sales to increase in the remaining channels.

Hypothesis 6 (Long-Term Habit Formation). *Curated box's positive effects on overall sales demand persist even after the termination of box delivery. Moreover, removal of the curated box channel induces purchase demand spillovers to the online and home try-on channels.*

3.4 Experiment Setting and Data Description

3.4.1 Experiment Setting

Our field experiment, designed to be a small-scale pilot study of the "VIPLOOK" program, quantifies the "VIPLOOK" program's potential impact on Y-Closet. Starting one of the months between April 1, 2018, and October 31, 2018, 580 existing Y-Closet customers were randomly selected to receive the curated box for two consecutive months[1]. Each box contained exactly six fashion products. Figure 3.1 shows the timeline of this experiment. Before month 1, no experiment participants received any curated box. All participants received the first box in month 1 and the second one in month 2. The curated box channel was suspended at the end of month 2, and the participants no longer received any boxes after month 2. The entire experiment duration spanned 12 months (i.e., 3 pre-treatment months, 2 treatment months, and 7 post-treatment months), and customer activities (i.e., in-box purchase, online purchase, home try-on sampling, and purchase) were tracked during this period. Overall, the 580 customers sampled 24,495 items and bought 1,103 items in the home try-on chan-

[1]Our field experiment was conducted before the COVID-19 pandemic in 2020-2021.

nel. They purchased 1,943 items in the online channel and purchased 394 items in curated boxes over the entire 12-month period. We note that the products contained in each box

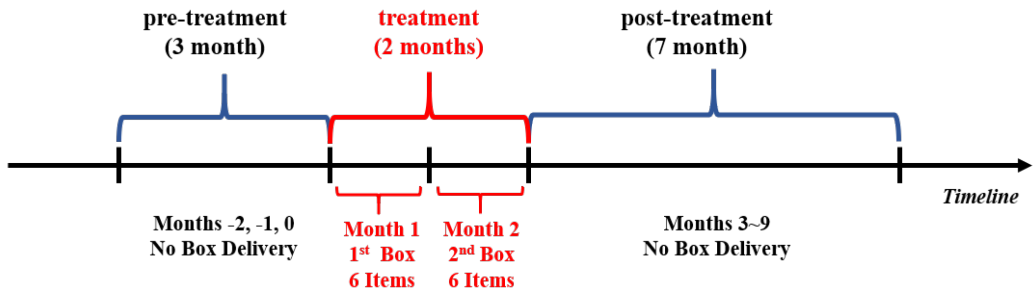

Figure 3.1: Timeline of the Experiment

were all newly arrived items chosen by the same product team in the Beijing headquarters of Y-Closet. Therefore, the experiment's treatment process did not suffer from confounding issues that different individuals might apply treatments differently. Also, selected customers did not know whether or what they would receive before the experiment.

Inside each box, customers would find a detailed explanation of the "VIPLOOK" program, a price list, and policies associated with this service in addition to six products. We note that each selected customer received a different assortment of products. However, personalization was not involved in the product curation process. In other words, the specific content of a curated box was not linked to the recipient's past consumption history. The fashion buyer team prepared each box such that it contained at least one complete outfit (i.e., one top and one bottom) and other matching pieces or accessories (e.g., bags, rings, bracelets, etc.), providing more options for "mix-and-match." The packaged curated boxes were then randomly assigned to the customers based on the default size. Since the products were intended to be sampled but not used, customers were not allowed to remove the tag from a product. Moreover, Y-Closet mandated a six-day sampling period, at the end of which customers must decide whether to buy any product from the box before returning the rest to Y-Closet. There was no penalty for not buying anything from the box. For our experiment, the box curation service, e.g., two-way transportation, product organization, packaging, etc., was free of charge, but the customers must pay the same prices as in other retail channels (i.e., online and home try-on) to buy any product they want to keep.

We note that the experiment's starting months were staggered, ranging from April 1, 2018 to October 31, 2018, followed by a seven-month observation period, with the last customer group's observation ending on July 31, 2019. The actual experiment starting month for a customer could fall between April 2018 and October 2018. To this end, the issue of unobservable inter-temporal heterogeneity may arise because customers could receive treatments in different months. These differences could be correlated with outcome variables, yielding a biased estimate of the average treatment effect. To overcome this issue, we follow the literature (Cui et al., 2020b; Wang et al., 2020) and use a panel data difference-in-differences (DID) approach, which identifies the causal effects by relying on the within-individual variation across time. We will further discuss our identification strategy in §3.5.2.

3.4.2 Sample Construction

We first construct a customer-level data-set from the field. Recall, we randomly selected 580 customers to receive curated boxes for two consecutive months. We collect detailed in-box sales data (i.e., purchased or not, purchase quantity and value) from each customer for both curated boxes. At the same time, we append the above data-set with the time-stamped customer-level transaction data, e.g., online sales, home try-on sampling, and home try-on sales, for the entire experiment period (i.e., 3 pre-treatment months, 2 treatment months, and 7 post-treatment months).

Based on the data-set, we leverage the volume of our raw data to conduct an exact matching procedure. We match each treated customer to all other customers in her city, who had the same participation in online purchase, home try-on sampling, and home try-on purchase activities (denoted by binary variables) in each of the three months before the

	Control	Treatment	p-value
Observations	580	580	
Monthly Online Sales			
Purchased or Not (Fraction)	0.07	0.07	1
	(0.25)	(0.25)	
Items Purchased	0.39	0.25	0.545
	(5.51)	(1.67)	
Categories Purchased	0.15	0.15	0.975
	(1.02)	(0.85)	
Value of Products Purchased (RMB)	58.03	48.36	0.769
	(727.18)	(319.32)	
Monthly Home Try-on Sampling			
Sampled or Not (Fraction)	0.29	0.29	1
	(0.46)	(0.46)	
Items Sampled	5.14	5.24	0.886
	(12.05)	(12.08)	
Categories Sampled	2.03	1.96	0.790
	(3.94)	(3.98)	
Value of Products Sampled (RMB)	1410.66	1369.29	0.853
	(4051.33)	(3534.76)	
Monthly Home Try-On Sales			
Purchased or Not (Fraction)	0.13	0.13	1
	(0.34)	(0.34)	
Items Purchased	0.26	0.26	0.948
	(0.86)	(0.93)	
Categories Purchased	0.22	0.22	0.932
	(0.66)	(0.72)	
Value of Products Purchased (RMB)	43.33	41.47	0.856
	(168.45)	(179.53)	
Consumer Demographics			
Consumer Tenure (Days)	257.62	265.91	0.321
	(142.69)	(141.65)	
Female (%)	100	100	1
City			1

Notes. Customer behavior variables are averaged over the three months before treatment. The treatment-control pairs are exactly matched on the "city" variable and must come from the same city, i.e., see Table 3.3 for details.

Table 3.2: Summary Statistics and Balance Check

experiment. Next, we use a one-to-one matching algorithm with Mahalanobis distance to create 580 "equivalent" treatment-control pairs. That is, for each treatment-control pair ij, the following expression is minimized:

$$Distance_{ij} = \sqrt{(X_i - X_j)' S^{-1} (X_i - X_j)}. \tag{3.1}$$

where vector X contains a customer's demographic information and a wide range of pre-treatment characteristics, including customer tenure at the time of treatment, total categories, total number, and total value of products purchased through the online or home try-on channel, and sampled through the home try-on channel in each of the three months before

City	Control	Treatment	City	Control	Treatment
Baoding	1 (0.2)	1 (0.2)	Nanjing	43 (7.4)	43 (7.4)
Baotou	1 (0.2)	1 (0.2)	Nanning	4 (0.7)	4 (0.7)
Beijing	49 (8.4)	49 (8.4)	Nanping	1 (0.2)	1 (0.2)
Changde	1 (0.2)	1 (0.2)	Ningbo	3 (0.5)	3 (0.5)
Changsha	25 (4.3)	25 (4.3)	Ningde	1 (0.2)	1 (0.2)
Changzhou	1 (0.2)	1 (0.2)	Panjin	2 (0.3)	2 (0.3)
Chengdu	59 (10.2)	59 (10.2)	Putian	1 (0.2)	1 (0.2)
Chenzhou	1 (0.2)	1 (0.2)	Qiannan Dist.	1 (0.2)	1 (0.2)
Chongqing	22 (3.8)	22 (3.8)	Qingdao	5 (0.9)	5 (0.9)
Dalian	3 (0.5)	3 (0.5)	Qitaihe	1 (0.2)	1 (0.2)
Daqing	1 (0.2)	1 (0.2)	Quanzhou	4 (0.7)	4 (0.7)
Dongguan	7 (1.2)	7 (1.2)	Quzhou	1 (0.2)	1 (0.2)
Foshan	3 (0.5)	3 (0.5)	Shanghai	54 (9.3)	54 (9.3)
Fuzhou	7 (1.2)	7 (1.2)	Shaoxing	1 (0.2)	1 (0.2)
Ganzhou	1 (0.2)	1 (0.2)	Shenyang	3 (0.5)	3 (0.5)
Guangzhou	22 (3.8)	22 (3.8)	Shenzhen	15 (2.6)	15 (2.6)
Guilin	1 (0.2)	1 (0.2)	Shuangyashan	1 (0.2)	1 (0.2)
Guiyang	3 (0.5)	3 (0.5)	Shuozhou	1 (0.2)	1 (0.2)
Haikou	1 (0.2)	1 (0.2)	Suzhou	22 (3.8)	22 (3.8)
Handan	2 (0.3)	2 (0.3)	Taiyuan	3 (0.5)	3 (0.5)
Hangzhou	15 (2.6)	15 (2.6)	Taizhou	1 (0.2)	1 (0.2)
Hefei	2 (0.3)	2 (0.3)	Tianjin	36 (6.2)	36 (6.2)
Hengyang	3 (0.5)	3 (0.5)	Tianmen	1 (0.2)	1 (0.2)
Huaihua	1 (0.2)	1 (0.2)	Weifang	1 (0.2)	1 (0.2)
Huainan	2 (0.3)	2 (0.3)	Wenzhou	11 (1.9)	11 (1.9)
Huanggang	1 (0.2)	1 (0.2)	Wuhan	56 (9.7)	56 (9.7)
Huzhou	1 (0.2)	1 (0.2)	Wuzhou	1 (0.2)	1 (0.2)
Ji An	1 (0.2)	1 (0.2)	Xi An	18 (3.1)	18 (3.1)
Jiangmen	1 (0.2)	1 (0.2)	Xiamen	4 (0.7)	4 (0.7)
Jinhua	2 (0.3)	2 (0.3)	Xianyang	1 (0.2)	1 (0.2)
Langfang	2 (0.3)	2 (0.3)	Xing'an	1 (0.2)	1 (0.2)
Leshan	1 (0.2)	1 (0.2)	Xining	2 (0.3)	2 (0.3)
Liaocheng	1 (0.2)	1 (0.2)	Ya An	1 (0.2)	1 (0.2)
Lijiang	1 (0.2)	1 (0.2)	Yantai	1 (0.2)	1 (0.2)
Linfen	2 (0.3)	2 (0.3)	Yinchuan	3 (0.5)	3 (0.5)
Linyi	1 (0.2)	1 (0.2)	Zhengzhou	19 (3.3)	19 (3.3)
Longyan	1 (0.2)	1 (0.2)	Zhenjiang	1 (0.2)	1 (0.2)
Maanshan	1 (0.2)	1 (0.2)	Zhumadian	1 (0.2)	1 (0.2)
Meishan	1 (0.2)	1 (0.2)	Zhongshan	2 (0.3)	2 (0.3)
Mianyang	1 (0.2)	1 (0.2)			

Table 3.3: Balance Check of the City Variable

treatment. S^{-1} is the covariance matrix between each of customer i and j's pre-treatment characteristics. As a balance check, we summarize a set of important pre-treatment variables of the treatment and control group in Table 3.2 and 3.3. A non-parametric test indicates that all p-values for comparing these variables between the treatment and control groups are above 0.321. Therefore, we confirm no statistical difference between the two groups after matching these variables, and the two groups have extremely similar tendencies towards their shopping behaviors. We report some key statistics of our sample: Among all 1,160 customers, 7% placed at least one online purchase order, 29% of customers sampled at least one product in the home try-on channel, among whom 45% purchased at least one product after sampling. On a monthly average, the customers purchased 0.32 products through the online retailing channel, sampled 5.19 products through the home try-on channel, and purchased 0.26 products through the home try-on channel before the experiment started.

3.4.3 Dependent Variables

We observe each customer's activities in both the online and home try-on channels for the 12-month experiment duration as illustrated in Figure 3.1. These activities can be classified into three categories, i.e., online purchase, home try-on sampling, and home try-on purchase. To test Hypotheses 1-6, we use the following dependent variables: ordered or not, order quantity, and variety, separately for each customer activity. These variables have been widely used to capture consumption behaviors in the retail operations literature (Bell et al., 2020; Caro and Gallien, 2010). To facilitate a multi-period panel difference-in-differences analysis, we follow the literature and choose customer-month as our analysis unit.

Ordered or not is a binary variable indicating whether a customer placed an order during a month for a specific activity. For online or home try-on sales, this variable denotes whether a customer purchased any product in a month through a particular retail channel. For product sampling via home try-on, this variable captures whether a customer sampled any product during a particular month.

Order quantity and order variety are count variables that capture the total number of products or categories of products ordered by a customer within a month for each customer activity. If a customer incurs no transaction activity within a month, her order variety is zero. Y-Closet carries a total of 32 sub-categories of products, which can be broadly classified into clothes (19 sub-categories) and accessories (13 sub-categories) as presented in Table 3.1.

3.4.4 Key Control Variables

Given our data-set's panel structure, we include the customer fixed effects to control for pre-treatment heterogeneity at a granular customer level. With customer fixed effects, it is unnecessary to include other time-invariant control variables at the customer level. Furthermore, we include month fixed effects to control temporal differences of outcomes across all customers. To capture the time-related heterogeneity in market characteristics, we implement four time-variant control variables, denoted by the matrix X_{jt}. The first two variables capture the aggregate number of sampling and purchase activities for each city j in month t. The last two variables capture the total number of samplers and buyers in each city j in month t. These variables account for temporal factors such as fashion trends and weather that may simultaneously affect the treated and control customers who live in the same geographical area.

3.5 Empirical Methods

3.5.1 Difference-in-Differences Estimation

We apply the DID model with two-way fixed effects to estimate curated box's impact on customer outcomes in different retailing channels. The unit of observation in the analysis is a customer in a specific month. For the binary variable ordered or not, we conduct the logistic regression for our DID analysis, where the specific model is given by Equation 3.2:

$$Pr(Y_{ijt}) = \beta_0 + \beta_1 Treat_i \times Post_{it} + Post_{it} + X_{jt} + \mu_i + \theta_t + \epsilon_{ijt}. \tag{3.2}$$

Y_{ijt} represents whether the customer i, from city j, placed an order in month t, for different customer activities. More specifically, we conduct five regressions to examine whether a customer had purchase through any retail channel, whether a customer had purchase through the retailer's online or home try-on channels, whether a customer had online purchase, whether a customer sampled any product via home try-on, and whether the customer had any purchase after home try-on, respectively. $Treat_i$ is a dummy variable equal to 1 if the customer i received the curated boxes and 0 otherwise; $Post_{it}$ equals 0 for the first three months before treatment (months -2 , -1 and 0), and 1 for all months (including 2 treatment months and 7 post-treatment months) after the treatment started (months 1 - 9). We incorporate the term $Post_{it}$ in our specification to account for customer-specific time trends. The matrix X_{jt} includes four vectors corresponding to the aggregate sampling and purchase transactions, and the total number of samplers and buyers for the city j in month t. Together, these variables control for market-level (i.e., city) heterogeneity, such as fashion trends and weather across different months. μ_i captures customer i's fixed effect, which controls for time-invariant unobserved factors. θ_t captures the month fixed effect which controls for unobserved time-invariant factors. ϵ_{ijt} is the error term. The DID coefficient β_1 estimates the treatment effect of the two curated boxes sent during the treatment period on customers' probability of order placement for different activities. The standard errors are clustered by customers.

Next, we study the impact of curated boxes on order quantity and variety for product sales and sampling in different channels. Recall that the order quantity and order variety variables capture different customer behaviors. Order quantity measures the total number of products ordered, whereas order variety captures the total number of unique product categories ordered by each customer. For instance, for a customer who purchased four different dresses, the order quantity variable equals four but the order variety variable equals one. A Poisson model is appropriate for this setting as both quantity and categories take non-negative integer values. This model has also been commonly used in the literature to model the demand process in fashion retailing settings (see, e.g., Kalyanam et al. (2007); Caro and Gallien (2010); Li et al. (2020)). Our DID specification is given by Equation 3.3:

$$ln[E(Z_{ijt})] = \beta_0 + \beta_1 Treat_i \times Post_{it} + Post_{it} + X_{jt} + \mu_i + \theta_t + \epsilon_{ijt}. \tag{3.3}$$

Z_{ijt} represents the quantity (variety) of orders placed for a specific activity in month t. Similar to Equation 3.2, we conduct five different regressions for order quantity, and five for order variety. $Treat_i$ is a dummy variable which equals 1 for the treated customer; $Post_{it}$ is a dummy variable for the corresponding time buckets of the treatment. X_{jt} captures the market characteristics, including aggregate demand and number of customers for city j in month t. μ_i captures customer i's fixed effect; θ_t captures the month fixed effect; ϵ_{ijt} is the error term. β_1 is our main DID coefficient, which estimates the impact of curated box delivery on order quantity (variety). Note that the Poisson regression model specifies the log of the expected count as a function of the dependent variable (Wooldridge, 2010). In this sense, the coefficient β_1 may be interpreted as follows: with curated box delivery, the log of expected order quantity (variety) will increase by β_1. In other words, given curated box delivery, the percentage change in the expected order quantity (variety) is $e^{\beta_1} - 1$.

To investigate the curated box's provision of fit information (i.e., Hypothesis 5), we conduct a DID analysis focusing on customer purchases towards accessories and clothes. We adopt the following specification:

$$\begin{split} ln[E(Z_{ijmt})] = \beta_0 + \beta_1 Treat_i \times Post_{it} + \beta_2 Treat_i \times Post_{it} \times A_m + Post_{it} \times A_m \\ + Treat_i \times A_m + Post_{it} + A_m + X_{jt} + \mu_i + \theta_t + \epsilon_{ijt}. \end{split} \tag{3.4}$$

The dependent variable Z_{ijmt} captures the number of accessories and clothes sold to customer i, from city j, in month t in the online channel. The dummy variable A_m indicates whether the purchase order was for accessories ($A_m = 1$) or clothes ($A_m = 0$). So, β_1 captures the treatment effect of curated box delivery on clothes, and β_2 captures the differential effect of curated box delivery between clothes and accessories. Hence, $\beta_1 + \beta_2$ captures the treatment effect of curated box delivery on accessories. All other variables in Equation 3.4 remain the same as in Equation 3.3, and we cluster the standard errors by customers.

Finally, to test the curated box's effect on habit formation (i.e., Hypothesis 6), we follow Cui et al. (2020b) and Wang et al. (2020) to combine pre-treatment, treatment, and post-treatment periods in one dynamic DID regression. We adopt a Poisson model with the following specification:

$$\begin{aligned} ln[E(Z_{ijt})] = \beta_0 + \beta_1 Treat_i \times TreatmentPeriod_{it} + \beta_2 Treat_i \times PostTreatmentPeriod_{it} \\ + TreatmentPeriod_{it} + PostTreatmentPeriod_{it} + X_{jt} + \mu_i + \theta_t + \epsilon_{ijt}. \end{aligned} \tag{3.5}$$

The dependent variable Z_{ijt} represents the quantity (variety) of product sales overall, in the online channel, and in the home try-on channel. X_i, μ_i, and θ_t are the sets of control variables identical to those used in Equation 3.2 and 3.3, and ϵ_{ijt} is the error term. The dummy variable $TreatmentPeriod_{it}$ equals 1 if t belongs to the treatment period (i.e., months 1 - 2) and 0 otherwise. The dummy variable $PostTreatmentPeriod_{it}$ equals 1 if t belongs to the post-treatment period (i.e., months 3 - 9) and 0 otherwise. The coefficient β_1 estimates the treatment effect of curated boxes by comparing treatment and pre-treatment months, and β_2 estimates the treatment effect across the post-treatment and pre-treatment months.

3.5.2 Identification

This section presents some potential issues related to identifying the causal relationship between curated box retailing and customer outcomes. Causal inference has been a difficult empirical question because of endogeneity issues such as self-selection or unobserved heterogeneity (Rubin, 1974). In this research, we exploit an experimental setting, where our treatment is the two curated boxes a customer received in months 1 and 2. We will show that this exogenous intervention is sufficient to overcome confounding factors and identify the curated box delivery's causal effects.

First, in our research, the intervention involved physical products. Unlike virtual interventions, which were more commonly used, e.g., text message, email, or software, products can convey both digital and non-digital information, which allows us to uncover previously unidentified behavioral mechanisms. The design of our intervention, however, renders additional challenges for a clean identification strategy. In our setting, the intervention, although not randomized to a subset of treated customers, was maintained exogenous to the fullest extent possible. First, our field experiment was conducted as an internal pilot study for the "VIPLOOK" curated box program, without any form of advertising. Before the experiment started, all experiment participants were unaware of whether they would receive the curated box or know the specific content inside each box. Thus, the experiment participants' exposure to curated boxes was exogenously determined. In other words, our setting is free from the self-selection bias, i.e., when subjects pro-actively select themselves to receive the curated boxes, resulting in a biased sample.

Despite the fact that each customer received a different product assortment, our treatment application process was centrally managed by the same product team in the Beijing headquarters of Y-Closet. This design helps to mitigate the concern that different agents may apply treatments differently. For each box, the product team jointly decided on the product assortment that was not based on individual transaction history. Packaged boxes were then randomly assigned to the treated customers based on the default size. To further ensure that all treated customers were exposed to the same treatment condition, all curated products were in brand new condition, and the experiment participants were not allowed to remove the tag from the product. Besides, each customer had precisely six days to sample the products inside each curated box. Finally, all boxes were distributed using the same delivery service operated by SF express, known for its reliability and punctuality (Cui et al., 2020b).

One potential concern that may arise in our setting is unobserved inter-temporal and cross-sectional heterogeneity. In our setting, the treated customers were in different regions, had different consumption habits, and received curated boxes in different months. These differences can correlate with the outcomes and yield a biased estimate of the average treatment effect. Hence, a direct comparison between the treated customers and the average Y-Closet customers who did not receive any curated box may be subject to bias and result in unparallel pre-treatment trends between the treatment and control groups.

Following the approaches documented in the literature (Bell et al., 2020; Cui et al., 2020b), we seek to apply a series of econometric approaches (exact matching, Mahalanobis matching, and panel DID) to construct a comparable control group. We first exactly match each treated customer to all customers who live in her city (see Table 3.3). This approach addresses the concern of time-invariant geographical heterogeneity. Next, we use the nearest distance matching with Mahalanobis distance to create 580 "equivalent" treatment-control pairs, taking into account a wide range of customer characteristics, including customer tenure, categories, quantity, and the total value of products purchased and sampled through the online and home try-on channels. We validate the matching process's effectiveness by conducting a balance check between the treatment and control groups. According to Table 3.2, there is no statistically significant difference between the treatment group and the control group for all pre-treatment customers characteristics mentioned above.

We subsequently use a panel DID approach to identify the causal effect. We control for pre-treatment heterogeneity of customer characteristics by both the customer fixed effects and month fixed effects. Including the two-way fixed effects strengthens the validity of our identification. Finally, we incorporate key time-variant variables aggregated at the city level to account for market-level heterogeneity such as fashion trends and weather over time. Controlling for these variables can effectively mitigate the problem of unobserved heterogeneity.

3.6 Empirical Results

In §3.6.1, we present the descriptive analysis of in-box sales from the curated boxes. We explore the curated box's impact on product purchase activities in §3.6.2. §3.6.3 explores the mechanisms through which curated box retailing influences consumer behavior. §3.6.4 shows the heterogeneous treatment effects induced by fit uncertainty.

3.6.1 Exploring In-Box Sales and Its Impact on Overall Sales

In this section, we explore the curated box channel's product fulfillment function, and the impact of in-box product sales on the retailer's overall product sales. We first present the descriptive analysis of in-box sales associated with curated boxes. Table 3.4 shows the number of products sent, purchased, and the conversion rate of curated boxes. Out of the 6,960 items distributed, 401 were adopted. The overall conversion rate is 5.76%. Comparing the two curated boxes' conversion rates, the first box's conversion rate of 6.12% is slightly higher than the second box's conversion rate of 5.40%. The conversion rate of accessories is 10.6%, which is much higher than the clothes' conversion rate of 5.7%. This difference is likely due to that accessories are usually priced lower than clothes.

Next, we explore in-box product sales' impact on overall sales (i.e., online sales, home try-on sales and curated box sales). Table 3.5 Column (1) shows that compared to controls, the treated customers' probability of purchasing increased by 324% ($= e^{1.445} - 1$), which is statistically significant. Columns (2) and (3) show that the increases in order quantity and variety are positive and statistically significant at 566% for order quantity and 253% for order variety for the treatment group. In sum, curated boxes led to a substantial increase of

	Month 1	Month 2	Total
Total Items Sent	3480	3480	6960
Total Items Purchased	213	188	401
Conversion Rate	6.12%	5.40%	5.76%
Total Clothes Sent	3432	3387	6819
Total Clothes Purchased	208	178	386
Conversion Rate (Clothes)	6.06%	5.26%	5.66%
Total Accessories Sent	48	93	141
Total Accessories Purchased	5	10	15
Conversion Rate (Accessories)	10.42%	10.75%	10.64%

Table 3.4: In-Box Sales Conversion Rates of Curated Boxes

overall product sales.

Activity	Online, Home Try-On and Curated Box Sales			Online and Home Try-On Sales		
Dependent Variable	Purchased or Not	Quantity	Variety	Purchased or Not	Quantity	Variety
	(1)	(2)	(3)	(4)	(5)	(6)
Treat × Post	1.445***	1.896***	1.260***	0.624**	1.593***	0.718***
	(0.188)	(0.102)	(0.125)	(0.197)	(0.103)	(0.129)
Δ Percentage	324%***	566%***	253%***	86.6%**	392%***	105%***
Market Characteristics	Y	Y	Y	Y	Y	Y
Customer Fixed Effects	Y	Y	Y	Y	Y	Y
Month Fixed Effects	Y	Y	Y	Y	Y	Y
Log-Likelihood	-1,252	-3,693	-2,320	-1,252	-3,166	-1,819
Observations	5,376	5,376	5,376	3,984	3,984	3,984

Notes. †$p<0.1$, *$p<0.05$, **$p<0.01$, ***$p<0.001$. Standard errors are given in parenthesis.

Table 3.5: The Effect of Curated Boxes on Overall Product Sales

What fraction of overall product sales can be attributed to in-box sales? To analyze this issue, we note that in-box sales were zero in all pre-treatment, treatment, and post-treatment periods for the control group since these customers did not receive any curated box at all. Hence, performing a DID analysis on in-box sales is infeasible as such analysis relies on variations of the customer outcome variable over time (Wooldridge, 2010). Alternatively, we utilize the curated box's differential treatment effects between overall sales (i.e., curated box, online and home try-on sales) and sales in the firm's existing channels (i.e., online and home try-on sales) to measure the in-box sales' contribution to the treatment effect. Specifically, the percentage increase (i.e., 324%) in Table 3.5 Columns (1) denotes the curated box's effects on product sales across all three channels. The percentage increase (i.e., 86.6%) in

Table 3.5 Column (4) captures the curated box's effects on online and home try-on sales. The net difference (i.e., 237.4% = 324% - 86.6%) represents the percentage increase in the treated customers' probability of purchasing due to the availability of the in-box sales option. Similarly, Table 3.5 Columns 2-3 and 5-6 show that curated boxes increased treated customers' order quantity by 174% and increased their order variety by 148%, given the possibility to purchase from within the curated box. Consequently, treated customers were more likely to place a purchase order and bought more categories of products through the curated box channel than through the firm's existing channels. In other words, the in-box sales function enhanced customers' exploration and adoption of unfamiliar products. Since the curated box channel was in use for two months, but the other retailing channels were always available for nine months, in-box sales contributed to a significant fraction of overall product sales.

3.6.2 Curated Boxes' Effects on Online and Home Try-On Sales

We have shown that curated boxes led to an increase of overall product sales. However, since overall sales include all products sold in the curated box, online, and home try-on channels, we do not know if curated boxes led to demand spillover or cannibalization in the online and home try-on channels. The positive and statistically significant DID coefficients in Table 3.5 Columns 4-6 indicate that combining online and home try-on sales, the treated customers' probability of making a purchase, order quantity, and order variety increased by 87%, 392%, and 105%. These results demonstrate that curated box positively induced all purchase behaviors in the firm's existing channels. In subsequent paragraphs, we separately explore the curated box's effects on customer behaviors in each channel.

First, Table 3.6 presents the estimation results for curated box's impact on online sales. The estimation results indicate that curated boxes had significant and positive effects on products

Dependent Variable	Purchased or Not	Online Sales Quantity	Online Sales Variety
	(1)	(2)	(3)
Treat × Post	0.976***	2.295***	0.993***
	(0.302)	(0.139)	(0.197)
Δ Percentage	165%***	892%***	170%***
Market Characteristics	Y	Y	Y
Customer Fixed Effects	Y	Y	Y
Month Fixed Effects	Y	Y	Y
Log-Likelihood	-463	-1,902	-889
Observations	1,932	1,932	1,932

Notes. †p<0.1, *p<0.05, **p<0.01, ***p<0.001. Standard errors are given in parenthesis.

Table 3.6: The Effect of Curated Boxes on Online Sales

sold through the online channel. Specifically, treated customers increased their probability of purchasing, order quantity, and variety by 165%, 892%, 170%, respectively. Hence, curated box retailing had a strong and positive impact on customer purchase activities in the online channel.

Activity	Home Try-On Sales			Home Try-On Sampling		
Dependent Variable	Purchased or Not	Quantity	Variety	Sampled or Not	Quantity	Variety
	(1)	(2)	(3)	(4)	(5)	(6)
Treat × Post	0.305	0.426*	0.402*	-0.290†	-0.152***	-0.089†
	(0.255)	(0.167)	(0.179)	(0.163)	(0.033)	(0.052)
Δ Percentage	35.7%	53.1%*	49.5%*	-25.1%†	-14.1%***	-8.5%†
Market Characteristics	Y	Y	Y	Y	Y	Y
Customer Fixed Effects	Y	Y	Y	Y	Y	Y
Month Fixed Effects	Y	Y	Y	Y	Y	Y
Log-Likelihood	-577	-1,038	-928	-1,553	-15,982	-7,308
Observations	2,640	2,640	2,640	5,388	5,496	5,496

Notes. †p<0.1, *p<0.05, **p<0.01, ***p<0.001. 9 customers are excluded from the analysis in Column (4) for having positive values in all months. Standard errors are given in parenthesis.

Table 3.7: The Effect of Curated Boxes on Home Try-On Sampling and Sales

Next, we present the estimation results for the curated box's influence on product sales and sampling in the home try-on channel. The estimation results are presented in Table 3.7. Columns 1-3 demonstrate that the percentage increases in the customers' probability of purchasing, order quantity, and order variety through the home try-on channel are 35.7%, 53.1%, 49.5%.

Combining all results above, we find strong evidence that curated box retailing not only amplified the overall customer purchase activities, but also led to spillover of increased customer purchase behaviors in the online and home try-on channels, which supports Hypothesis 1.

3.6.3 Evidence on Mechanisms

The Impact of Curated Boxes' Information Provision on Customer Behaviors

First, Hypothesis 2 states that curated boxes can provide product recommendations which encourage the customers purchase more categories of products. Table 3.5 Columns 3 and 6 show that customers' purchase variety increased 253% across all channels, and increased 105% in the retailer's existing retail channels. Thus, we find strong evidence to support Hypothesis 2. The fact that curated boxes induced an increase of order variety in the online and home try-on channels indicates that the customers *proactively* engaged in product exploration after being passively exposed to more product categories through curated boxes.

Besides enhancing product exploration, Hypothesis 3 states that the provision of product recommendations can improve the sales conversion rate in the home try-on channel. Table 3.7 Columns 4-6 show that customers' probability of product sampling, sampled quantity and variety, decreased by 25.2%, 14.1%, and 8.5% in the home try-on channel. All estimates are statistically significant. Overall, these results indicate that curated boxes led to a reduction in customers' sampling activities. Given curated boxes also led to amplified purchase activities in the home try-on channel, we find that the customers sampled less but purchased more products via home try-on because of the curated boxes. This result implies that the purchase conversion rate improves, and the return rate of sampled products (i.e., $= 1 -$ conversion rate) decreases in the home try-on channel, which supports Hypothesis 3. Considering that curated boxes also induced the probability of purchasing to increase in the online channel, provision of product recommendations causes a reduction of the customers' dependence on home try-on, and an increase in search efficiency. Since operating a home try-on program is costly, our finding has important operational implications. Based on our conversation with Y-Closet, a home try-on program is associated with: (1) inventory costs due to product wear and tear; (2) labor costs due to product curation and cleaning; (3) two-way transportation costs. All aforementioned costs critically depend on the return rate of sampled products. Hence, the product recommendations provided by curated boxes give rise to beneficial operational spillovers, which significantly mitigates the issue of product returns and lowers the corresponding operational costs.

Next, Hypothesis 4 states that curated boxes can provide tactile product information, which would benefit online sales more than home try-on sales. To find empirical evidence for Hypothesis 4, we compare curated boxes' positive effects on product sales in the online channel, reported by the estimated coefficients in Table 3.6 Columns 1-3, and in the home try-on channel, reported by the estimated coefficients in Table 3.7 Columns 1-3. We observe that all the estimated DID coefficients for online sales regarding the probability of purchasing, order quantity, and order variety are more significant for online sales than for home try-on sales. These results support Hypothesis 4.

Habit Formation Induced by Curated Boxes

To investigate the mechanism of long-term habit formation induced by curated boxes, we conduct the dynamic DID analysis presented in Equation 3.5. Table 3.8 Columns (1) and (2) report the regression results for overall sales. The increases in order quantity and variety are positive and statistically significant in all treatment and post-treatment periods. Thus, we find evidence to prove Hypothesis 6, i.e., the significant post-treatment increase in product sales can be interpreted by the formation of product consumption and exploration habits for the treatment group.

Interestingly, we also observe amplification, rather than decay, of the curated box's positive

Activity	Overall Sales		Online Sales		Home Try-On Sales	
Dependent Variable	Quantity	Variety	Quantity	Variety	Quantity	Variety
	(1)	(2)	(3)	(4)	(5)	(6)
Treat × TreatmentPeriod	2.445***	2.065***	2.155***	1.035***	0.258	0.171
	(0.144)	(0.174)	(0.189)	(0.252)	(0.325)	(0.334)
Treat × PostTreatmentPeriod	1.642***	0.731***	2.347***	0.956***	0.460**	0.450*
	(0.109)	(0.137)	(0.147)	(0.215)	(0.174)	(0.187)
Δ Percentage in TreatmentPeriod	1,053%***	689%***	763%***	182%***	29.4%	18.6%
Δ Percentage in PostTreatmentPeriod	417%***	108%***	945%***	160%***	58.4%***	58.6%***
Market Characteristics	Y	Y	Y	Y	Y	Y
Customer Fixed Effects	Y	Y	Y	Y	Y	Y
Month Fixed Effects	Y	Y	Y	Y	Y	Y
Log-Likelihood	-3,672	-2,284	-1,893	-887	-1,021	-916
Observations	5,376	5,376	1,932	1,932	2,640	2,640

Notes. $^{\dagger}p<0.1$, $^{*}p<0.05$, $^{**}p<0.01$, $^{***}p<0.001$. Overall sales include products sold via the online, home try-on, and curated box channels. Standard errors are given in parenthesis.

Table 3.8: Long-Term Dynamic Effects of Curated Boxes on Product Sales

effects on online sales and home try-on sales after the treatment termination. The estimated percentage increase in order quantity is 763% in the treatment period and 945% in the post-treatment period for the online channel. In terms of order variety for online sales, the estimated coefficients are similar in treatment (i.e, 182%) and post-treatment periods (i.e., 160%). This phenomenon is even more pronounced for home try-on sales. For the home try-on channel, the estimated treatment effects on order quantity and variety are small and insignificant during the treatment period, which suggests that curated boxes did not significantly impact product sales in the home try-on channel when curated box retailing was available. The estimated coefficients are significantly higher for the post-treatment period, resulting in a percentage increase of 58.4% for sales quantity and 56.8% for sales variety. Both estimations are significant. These results strongly support Hypothesis 6, indicating positive demand spillover from the curated box channel to the online and home try-on channels after the termination of the curated box channel.

3.6.4 The Effects of Fit Uncertainty

Apart from information provision, we also conduct a heterogeneous analysis related to fit uncertainty. Specifically, we conduct the DID analysis presented in Equation 3.4. Table 3.9 reports the estimation results. The two-way interaction term $Treat \times Post$ represents curated boxes' effects on clothes, and the three-way interaction term $Treat \times Post \times Accessories$ represents the differential treatment effects between clothes and accessories. To measure the treatment effects of curated boxes on accessories, we perform an F-test on $Treat \times Post + Treat \times Post \times Accessories$. We report the corresponding treatment effect for accessories sales in the bottom section of the table. According to the results, curated boxes led to a more substantial increase of accessories sales than clothes sales in the online channel, which supports Hypothesis 5.

Dependent Variable	Online Sales Quantity (1)
Treat × Post	2.231***
	(0.142)
Treat × Post × Accessories	1.259†
	(0.646)
Treatment Effect for Accessories	3.490***
Market Characteristics	Y
Customer Fixed Effects	Y
Month Fixed Effects	Y
Log-Likelihood	-2,093
Observations	3,864

Notes. †p<0.1, *p<0.05, **p<0.01, ***p<0.001. Standard errors are given in parenthesis.

Table 3.9: The Effect of Curated Boxes on the Online Sales of Clothes and Accessories

3.7 Conclusion

Curated box retailing is an efficient and effective strategy to address two key retailing objectives regarding product fulfillment and information provision. We partnered with a leading fashion retailer, Y-Closet, to conduct the first longitudinal experiment to analyze the curated box retailing's impact on an established retailer. Our analysis shows that the overall in-box sales conversion rate of the curated box was 5.76%. The sales conversion rate was higher for the first than the second box and higher for accessories than clothes. Meanwhile, distributing one curated box each month for two months can induce substantial demand spillover to the existing retail channels. The observed demand spillover can be further decomposed into increased probability of purchasing, sales quantity and variety in the online and home try-on channels. Finally, we find curated boxes caused a more substantially increase of the online sales of accessories than clothes due to fit uncertainty. We also explore underlying mechanisms through which a curated box affects customer behaviors. First, we find that the curated box provides important recommendations and tactile product information. Specifically, provision of product recommendations induces customer exploration of unfamiliar items, and increases sales conversion rate via home try-on. Provision of product tactile information causes a stronger increase of online sales than home try-on sales. Second, we find evidence that curated box retailing causes customer habituation of product consumption and exploration in the long run, and removing the curated box channel increases demand in the firm's existing retail channels.

Our research rationalizes the increasing examples of established retailers incorporating the curated box retailing strategy to complement their current businesses. One potential future research direction is to optimally design the content of the curated box based on historical in-box sales information. Second, since not every customer who receives the curated box will engage in in-box purchase, how to target the correct set of customers based-off each customer's past consumption information would be important to the curated box providers. Finally, there likely exists a non-linear relationship between the box delivery duration and the corresponding sales conversion rate. Finding the right time interval to deliver consecutive curated boxes would be another interesting problem to investigate.

Bibliography

Allcott, H. and Rogers, T. (2014). The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, 104(10):3003–37.

Allon, G., Federgruen, A., and Pierson, M. (2011). How much is a reduction of your customers' wait worth? an empirical study of the fast-food drive-thru industry based on structural estimation methods. *Manufacturing & Service Operations Management*, 13(4):489–507.

Andonova, Y., Anaza, N. A., and Bennett, D. H. (2021). Riding the subscription box wave: Understanding the landscape, challenges, and critical success factors of the subscription box industry. *Business Horizons*, (0007-6813). `https://doi.org/10.1016/j.bushor.2021.02.024`.

Aouad, A. and Segev, D. (2020). Display optimization for vertically differentiated locations under multinomial logit preferences. *Management Science*. `https://doi.org/10.1287/mnsc.2020.3664`.

Auto Laundry News (2016). *Results from the Auto Laundry News: Exterior Conveyor Survey 2016*. `http://www.carwashmag.com/fileadmin/pdfs/aln_Conveyor_Survey_July_2016.pdf`.

Avci, B., Girotra, K., and Netessine, S. (2014). Electric vehicles with a battery switching station: Adoption and environmental impact. *Management Science*, 61(4):772–794.

Avery, J., Steenburgh, T. J., Deighton, J., and Caravella, M. (2012). Adding bricks to clicks: Predicting the patterns of cross-channel elasticities over time. *Journal of Marketing*, 76(3):96–111.

Baddeley, A. (2007). *Working memory, thought, and action*, volume 45. Oxford University Press.

Bai, J., So, K. C., Tang, C. S., Chen, X., and Wang, H. (2019). Coordinating supply and demand on an on-demand service platform with impatient customers. *Manufacturing & Service Operations Management*, 21(3):556–570.

Bakos, Y. and Brynjolfsson, E. (1999). Bundling information goods: Pricing, profits, and efficiency. *Management science*, 45(12):1613–1630.

Bala, R. and Carr, S. (2009). Pricing software upgrades: The role of product improvement and user costs. *Production and Operations Management*, 18(5):560–580.

Barua, A., Kriebel, C. H., and Mukhopadhyay, T. (1991). An economic analysis of strategic information technology investments. *MIS quarterly*, pages 313–331.

Bawa, K. and Shoemaker, R. (2004). The effects of free sample promotions on incremental brand sales. *Marketing Science*, 23(3):345–363.

BBC (2017). *Samsung confirms battery faults as cause of Note 7 fires*. `https://www.bbc.com/news/business-38714461`.

Belavina, E., Girotra, K., and Kabra, A. (2017). Online grocery retail: Revenue models and environmental impact. *Management Science*, 63(6):1781–1799.

Bell, D., Gallino, S., Moreno, A., Yoder, J., and Ueda, D. (2018a). *The store is dead-long live the store*. MIT Sloan Management Review.

Bell, D. R., Gallino, S., and Moreno, A. (2018b). Offline showrooms in omnichannel retail: Demand and operational benefits. *Management Science*, 64(4):1629–1651.

Bell, D. R., Gallino, S., and Moreno, A. (2020). Customer supercharging in experience-centric channels. *Management Science*, 66(9):4096–4107.

Bernstein, F. and Martínez-de Albéniz, V. (2017). Dynamic product rotation in the presence of strategic customers. *Management Science*, 63(7):2092–2107.

Bertsimas, D., Gupta, V., and Kallus, N. (2018). Data-driven robust optimization. *Mathematical Programming*, 167(2):235–292.

Bhargava, H. K. and Choudhary, V. (2001). Information goods and vertical differentiation. *Journal of Management Information Systems*, 18(2):89–106.

Bhargava, H. K. and Choudhary, V. (2008). Research note – when is versioning optimal for information goods? *Management Science*, 54(5):1029–1035.

Bhaskaran, S. R. and Krishnan, V. (2009). Effort, revenue, and cost sharing mechanisms for collaborative new product development. *Management Science*, 55(7):1152–1169.

Boada-Collado, P. and Martínez-de Albéniz, V. (2020). Estimating and optimizing the impact of inventory on consumer choices in a fashion retail setting. *Manufacturing & Service Operations Management*, 22(3):582–597.

Boksem, M. A. and Tops, M. (2008). Mental fatigue: costs and benefits. *Brain research reviews*, 59(1):125–139.

Boston Consulting Group (2010). *Batteries for electric cars: challenges, opportunities and* *outlook to 2020*. https://www.bcg.com/documents/file36615.pdf.

Brynjolfsson, E., Hu, Y., and Rahman, M. S. (2009). Battle of the retail channels: How product selection and geography drive cross-channel competition. *Management Science*, 55(11):1755–1765.

Brynjolfsson, E., Hu, Y., and Smith, M. D. (2010). Research commentary—long tails vs. superstars: The effect of information technology on product variety and sales concentration patterns. *Information Systems Research*, 21(4):736–747.

Brynjolfsson, E., Hu, Y. J., and Rahman, M. S. (2013). *Competing in the age of omnichannel retailing*. MIT Cambridge, MA. pp. 1-7.

Cachon, G. P. and Feldman, P. (2011). Pricing services subject to congestion: Charge per-use fees or sell subscriptions? *Manufacturing & Service Operations Management*, 13(2):244–260.

Cachon, G. P., Gallino, S., and Olivares, M. (2019). Does adding inventory increase sales? evidence of a scarcity effect in us automobile dealerships. *Management Science*, 65(4):1469–1485.

Calvo, E., Cui, R., and Wagner, L. (2020). Disclosing product availability in online retail. *Manufacturing & Service Operations Management, Forthcoming*. https://doi.org/10.1287/msom.2020.0882.

Calzolari, G. and Nardotto, M. (2017). Effective reminders. *Management Science*, 63(9):2915–2932.

Caro, F. and Gallien, J. (2010). Inventory management of a fast-fashion retail network. *Operations Research*, 58(2):257–273.

Caro, F. and Gallien, J. (2012). Clearance pricing optimization for a fast-fashion retailer. *Operations research*, 60(6):1404–1422.

Caro, F., Kök, A. G., and Martínez-de Albéniz, V. (2020). The future of retail operations. *Manufacturing & Service Operations Management*, 22(1):47–58.

Charness, G. and Gneezy, U. (2009). Incentives to exercise. *Econometrica*, 77(3):909–931.

Chellappa, R. K. and Shivendu, S. (2005). Managing piracy: Pricing and sampling strategies for digital experience goods in vertically segmented markets. *Information Systems Research*, 16(4):400–417.

Chen, C. (2001). Design for the environment: A quality-based model for green product development. *Management Science*, 47(2):250–263.

Chen, Q., Jasin, S., and Duenyas, I. (2016). Real-time dynamic pricing with minimal and flexible price adjustment. *Management Science*, 62(8):2437–2455.

Chen, T., Fenyo, K., Yang, S., and Zhang, J. (2018). Thinking inside the subscription box: New research on e-commerce consumers. *McKinsey & Company*, pages 1–9. https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/thinking-inside-the-subscription-box-new-research-on-ecommerce-consumers.

Chen, Y.-J., Tomlin, B., and Wang, Y. (2013). Coproduct technologies: Product line design and process innovation. *Management Science*, 59(12):2772—-2789.

Cheng, H. K. and Liu, Y. (2012). Optimal software free trial strategy: The impact of network externalities and consumer uncertainty. *Information Systems Research*, 23(2):488–504.

Cohen, M. A., Cui, S., and Gao, F. (2017). Performance, reliability or time-to-market? innovative product development and the impact of government regulation. *Innovative Product Development and the Impact of Government Regulation (August 1, 2017)*.

Crunchbase (2021). *Y-Closet: Crunchbase Company Profile & Funding.* https://www.crunchbase.com/organization/ycloset.

Cui, R., Li, J., and Zhang, D. J. (2020a). Reducing discrimination with reviews in the sharing economy: Evidence from field experiments on airbnb. *Management Science*, 66(3):1071–1094.

Cui, R., Li, M., and Li, Q. (2020b). Value of high-quality logistics: Evidence from a clash between sf express and alibaba. *Management Science*, 66(9):3879–3902.

Cui, R., Lu, Z., Sun, T., and Golden, J. (2020c). Sooner or later? promising delivery speed in online retail. *Available at SSRN*. http://dx.doi.org/10.2139/ssrn.3563404.

Cui, R., Zhang, D. J., and Bassamboo, A. (2019). Learning from inventory availability information: Evidence from field experiments on amazon. *Management Science*, 65(3):1216–1235.

Danaher, P. J. (2002). Optimal pricing of new subscription services: Analysis of a market experiment. *Marketing Science*, 21(2):119–138.

Data & Marketing Association (2015). *National Client Email Report of U.K.* https://dma.org.uk/uploads/ckeditor/National-client-email-2015.pdf.

Desai, P. (2001). Quality segmentation in spatial markets: When does cannibalization affect product line design? *Marketing Science*, 20(3):265—283.

Dhebar, A. (1994). Durable-goods monopolists, rational consumers, and improving products. *Marketing Science*, 13(1):100–120.

Du, N., Li, L., Lu, T., and Lu, X. (2020). Prosocial compliance in p2p lending: A natural field experiment. *Management Science*, 66(1):315–333.

Dubé, J.-P., Fang, Z., Fong, N., and Luo, X. (2017). Competitive price targeting with smartphone coupons. *Marketing Science*, 36(6):944–975.

Duflo, E. (2001). Schooling and labor market consequences of school construction in indonesia: Evidence from an unusual policy experiment. *American Economic Review*, 91(4):795–813.

Fader, P. S. and Hardie, B. G. (2007). How to project customer retention. *Journal of Interactive Marketing*, 21(1):76–90.

Feldman, J., Zhang, D., Liu, X., and Zhang, N. (2018). Customer choice models versus machine learning: Finding optimal product displays on alibaba. *Available at SSRN*. `http://dx.doi.org/10.2139/ssrn.3232059`.

Feldman, P., Papanastasiou, Y., and Segev, E. (2019). Social learning and the design of new experience goods. *Management Science*, 65(4):1502–1519.

Ferreira, K. J., Lee, B. H. A., and Simchi-Levi, D. (2016). Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management*, 18(1):69–88.

Fisher, M., Gallino, S., and Li, J. (2018). Competition-based dynamic pricing in online retailing: A methodology validated with field experiments. *Management science*, 64(6):2496–2514.

Forbes (2020). *Retail Subscriptions Thrive During COVID-19, Kaleigh Moore, 2020/07/15*. `https://www.forbes.com/sites/kaleighmoore/2020/07/15/retail-subscriptions-thrive-during-covid-19/`.

Gal-Or, E. (1983). Quality and quantity competition. *The Bell Journal of Economics*, 14(2):590–600.

Gallino, S. and Moreno, A. (2014). Integration of online and offline channels in retail: The impact of sharing reliable inventory availability information. *Management Science*, 60(6):1434–1451.

Gallino, S. and Moreno, A. (2018). The value of fit information in online retail: Evidence from a randomized field experiment. *Manufacturing & Service Operations Management*, 20(4):767–787.

Gao, F. and Su, X. (2017a). Omnichannel retail operations with buy-online-and-pick-up-in-store. *Management Science*, 63(8):2478–2492.

Gao, F. and Su, X. (2017b). Online and offline information for omnichannel retailing. *Manufacturing & Service Operations Management*, 19(1):84–98.

Ghose, A., Li, B., and Liu, S. (2019). Mobile targeting using customer trajectory patterns. *Management Science*, 65(11):5027–5049.

Han, B. R., Chu, L. Y., Sun, T., and Wu, L. (2020). Commercializing the package flow:

Cross-sampling physical products through e-commerce warehouses. *Available at SSRN*. http://dx.doi.org/10.2139/ssrn.3566756.

Han, B. R., Sun, T., Chu, L. Y., and Wu, L. (2019). Connecting customers and merchants offline: Experimental evidence from the commercialization of last-mile stations at alibaba. *Available at SSRN*. http://dx.doi.org/10.2139/ssrn.3452769.

Heiman, A., McWilliams, B., Shen, Z., and Zilberman, D. (2001). Learning and forgetting: Modeling optimal product sampling over time. *Management Science*, 47(4):532–546.

Ho, S. Y. and Bodoff, D. (2014). The effects of web personalization on user attitude and behavior. *MIS quarterly*, 38(2):497–A10.

Hu, K., Acimovic, J., Erize, F., Thomas, D. J., and Van Mieghem, J. A. (2019). Forecasting new product life cycle curves: Practical approach and empirical analysis: Finalist–2017 m&som practice-based research competition. *Manufacturing & Service Operations Management*, 21(1):66–85.

Hu, Y.-S. (2016). Batteries: getting solid. *Nature Energy*, 1(4):1–2.

Huang, T. and Van Mieghem, J. A. (2014). Clickstream data and inventory management: Model and empirical analysis. *Production and Operations Management*, 23(3):333–347.

Hui, S. K., Inman, J. J., Huang, Y., and Suher, J. (2013). The effect of in-store travel distance on unplanned spending: Applications to mobile promotion strategies. *Journal of Marketing*, 77(2):1–16.

Kalyanam, K., Borle, S., and Boatwright, P. (2007). Deconstructing each item's category contribution. *Marketing Science*, 26(3):327–341.

Karlan, D., McConnell, M., Mullainathan, S., and Zinman, J. (2016). Getting to the top of mind: How reminders increase saving. *Management Science*, 62(12):3393–3411.

Kornish, L. J. (2001). Pricing for a durable-goods monopolist under rapid sequential innovation. *Management Science*, 47(11):1552–1561.

Krishnan, V. and Bhattacharya, S. (2002). Technology selection and commitment in new product development: The role of uncertainty and design flexibility. *Management Science*, 48(3):313–327.

Krishnan, V. and Ulrich, K. T. (2001). Product development decisions: A review of the literature. *Management science*, 47(1):1–21.

Krishnan, V. and Zhu, W. (2006). Designing a family of development-intensive products. *Management science*, 52(6):813–825.

Levinthal, D. A. and Purohit, D. (1989). Durable goods and product obsolescence. *Marketing Science*, 8(1):35–56.

Li, S., Lu, L. X., Lu, S. F., and Huang, S. (2020). Estimating the demand spillover effect of inventory stockouts in a fashion footwear retail setting. *Available at SSRN*. http://dx.doi.org/10.2139/ssrn.3696682.

Liang, C., Çakanyıldırım, M., and Sethi, S. P. (2014). Analysis of product rollover strategies in the presence of strategic customers. *Management Science*, 60(4):1033–1056.

Lim, M. K., Mak, H.-Y., and Rong, Y. (2014). Toward mass adoption of electric vehicles: impact of the range and resale anxieties. *Manufacturing & Service Operations Management*, 17(1):101–119.

Lim, W. S. and Tang, C. S. (2006). Optimal product rollover strategies. *European Journal of Operational Research*, 174(2):905–922.

Lin, Z., Zhang, Y., and Tan, Y. (2019). An empirical study of free product sampling and rating bias. *Information Systems Research*, 30(1):260–275.

Mak, H.-Y., Rong, Y., and Shen, Z.-J. M. (2013). Infrastructure planning for electric vehicles with battery swapping. *Management Science*, 59(7):1557–1575.

Marc S. Reisch (2017). *Solid-state batteries inch their way toward commercialization*. https://cen.acs.org/articles/95/i46/Solid-state-batteries-inch-way.html.

Marks, L. J. and Kamins, M. A. (1988). The use of product sampling and advertising: Effects of sequence of exposure and degree of advertising claim exaggeration on consumers' belief strength, belief confidence, and attitudes. *Journal of Marketing Research*, 25(3):266–281.

Martínez-de Albéniz, V., Planas, A., and Nasini, S. (2020). Using clickstream data to improve flash sales effectiveness. *Production and Operations Management*, 29(11):2508–2531.

McKinsey (2019). *Recharging economies: The EV-battery manufacturing outlook for Europe*. https://www.mckinsey.com/industries/oil-and-gas/our-insights/recharging-economies-the-ev-battery-manufacturing-outlook-for-europe.

Moorthy, K. S. and Png, I. P. (1992). Market segmentation, cannibalization, and the timing of product introductions. *Management Science*, 38(3):345–359.

Mussa, M. and Rosen, S. (1978). Monopoly and product quality. *Journal of Economic theory*, 18(2):301–317.

Niculescu, M. F. and Wu, D. J. (2014). Economics of free under perpetual licensing: Implications for the software industry. *Information Systems Research*, 25(1):173–199.

Pauwels, K. and Neslin, S. A. (2015). Building with bricks and mortar: The revenue impact of opening physical stores in a multichannel environment. *Journal of Retailing*, 91(2):182–197.

Ramachandran, K. and Krishnan, V. (2008). Design architecture and introduction timing for rapidly improving industrial products. *Manufacturing & Service Operations Management*, 10(1):149–171.

Randhawa, R. S. and Kumar, S. (2008). Usage restriction and subscription services: Operational benefits with rational users. *Manufacturing & Service Operations Management*, 10(3):429–447.

Retana, G. F., Forman, C., and Wu, D. J. (2016). Proactive customer education, customer retention, and demand for technology support: Evidence from a field experiment. *Manufacturing & Service Operations Management*, 18(1):34–50.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688.

Rubin, D. C. and Wenzel, A. E. (1996). One hundred years of forgetting: A quantitative description of retention. *Psychological review*, 103(4):734.

Singer, J. D., Willett, J. B., Willett, J. B., et al. (2003). *Applied longitudinal data analysis: Modeling change and event occurrence*. Oxford University Press.

Song, M. and Montoya-Weiss, M. M. (2001). The effect of perceived technological uncertainty on japanese new product development. *Academy of Management journal*, 44(1):61–80.

Staats, B. R., Dai, H., Hofmann, D., and Milkman, K. L. (2017). Motivating process compliance through individual electronic monitoring: An empirical examination of hand hygiene in healthcare. *Management Science*, 63(5):1563–1585.

SUBTA (2019). *2019 State of the Subscription Commerce Economy Annual Report*. `https://subta.com/wp-content/uploads/2019/10/2019_SUBTA_Annual_Report.pdf`.

Sun, T., Gao, G., and Jin, G. Z. (2019a). Mobile messaging for offline group formation in prosocial activities: A large field experiment. *Management Science*, 65(6):2717–2736.

Sun, T., Shi, L., Viswanathan, S., and Zheleva, E. (2019b). Motivating effective mobile app adoptions: Evidence from a large-scale randomized field experiment. *Information Systems Research*, 30(2):523–539.

Terwiesch, C. (2019). Om forum—empirical research in operations management: From field studies to analyzing digital exhaust. *Manufacturing & Service Operations Management*, 21(4):713–722.

Thatcher, M. E. and Oliver, J. R. (2001). The impact of technology investments on a firm's production efficiency, product quality, and productivity. *Journal of Management Information Systems*, 18(2):17–45.

Thatcher, M. E. and Pingry, D. E. (2004). An economic model of product quality and it value. *Information Systems Research*, 15(3):268–286.

The Economist (2010). *Monster power Lithium-ion batteries start to take on the big stuff*. `https://www.economist.com/technology-quarterly/2010/10/26/monster-power`.

Thomke, S. H., Krishnan, V. V., and Nimgade, A. (1998). *Product development at dell computer corporation*.

UPI news (2017). *Battery technology keeping electric car adoption in the slow lane*. https://www.upi.com/Science_News/2017/10/06/.

U.S. Department of Commerce (2017). *New Insights on Retail E-Commerce (ESA Issue Brief No. 04-17)*. https://www.esa.gov/reports/new-insights-retail-e-commerce.

Wang, K. and Goldfarb, A. (2017). Can offline stores drive online sales? *Journal of Marketing Research*, 54(5):706–719.

Wang, W., Ferguson, M. E., Hu, S., and Souza, G. C. (2013). Dynamic capacity investment with two competing technologies. *Manufacturing & Service Operations Management*, 15(4):616–629.

Wang, Y., Lu, L. X., and Shi, P. (2020). Does continuous email engagement improve profitability? evidence from a field experiment of a subscription-based service provider. *Available at SSRN*. http://dx.doi.org/10.2139/ssrn.3636595.

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.

Xin, M. and Choudhary, V. (2018). It investment under competition: The role of implementation failure. *Management Science*, 65(4):1909–1925.

Xu, X. (2009). Optimal price and product quality decisions in a distribution channel. *Management Science*, 55(8):1347–1352.

Yin, S., Ray, S., Gurnani, H., and Animesh, A. (2010). Durable products with multiple used goods markets: Product upgrade and retail pricing implications. *Marketing Science*, 29(3):540–560.

Zhang, D. J., Allon, G., and Van Mieghem, J. A. (2017). Does social interaction improve learning outcomes? evidence from field experiments on massive open online courses. *Manufacturing & Service Operations Management*, 19(3):347–367.

Zhang, D. J., Dai, H., Dong, L., Wu, Q., Guo, L., and Liu, X. (2019). The value of pop-up stores on retailing platforms. *Management Science*, 65(11):5142–5151.