MANAGEMENT SCIENCE
Vol. 62, No. 6, June 2016, pp. 1778–1799
ISSN 0025-1909 (print) | ISSN 1526-5501 (online)

http://dx.doi.org/10.1287/mnsc.2015.2189
© 2016 INFORMS

Optimizing Product Launches in the Presence of Strategic Consumers

Ilan Lobel
Leonard N. Stern School of Business, New York University, New York, New York 10012, ilobel@stern.nyu.edu

Jigar Patel
School of Business, Montclair State University, Montclair, New Jersey 07043, patelji@mail.montclair.edu

Gustavo Vulcano
Leonard N. Stern School of Business, New York University, New York, New York 10012, gvulcano@stern.nyu.edu

Jiawei Zhang
Leonard N. Stern School of Business, New York University, New York, New York 10012; and
New York University Shanghai, Pudong, Shanghai, China 200122, jzhang@stern.nyu.edu

A technology firm launches newer generations of a given product over time. At any moment, the firm decides whether to release a new version of the product that captures the current technology level at the expense of a fixed launch cost. Consumers are forward-looking and purchase newer models only when it maximizes their own future discounted surpluses. We start by assuming that consumers have a common valuation for the product and consider two product launch settings. In the first setting, the firm does not announce future release technologies and the equilibrium of the game is to release new versions cyclically with a constant level of technology improvement that is optimal for the firm. In the second setting, the firm is able to precommit to a schedule of technology releases and the optimal policy generally consists of alternating minor and major technology launch cycles. We verify that the difference in profits between the commitment and no-commitment scenarios can be significant, varying from 4% to 12%. Finally, we generalize our model to allow for multiple customer classes with different valuations for the product, demonstrating how to compute equilibria in this case and numerically deriving insights for different market compositions.

Keywords: new product introduction; new product development; strategic consumer behavior; technology products; noncooperative game theory
History: Received February 25, 2013; accepted November 24, 2014, by Yossi Aviv, operations management. Published online in *Articles in Advance* October 15, 2015.

1. Introduction

Firms continuously strive to improve the products they sell to their consumers, either by enhancing the quality of their products or by incorporating new features into them. Companies take these improvements to market by releasing newer and better generations of their products over time. The cycle of firms releasing ever better products on the market is a particularly visible phenomenon in the technology industry, where companies upgrade the hardware or software they sell on a regular basis, but it is also prevalent in many other sectors of the economy where firms sell technologically enabled products to their customers, ranging from medical device manufacturers to the auto industry.

A product launch is an expensive endeavor, involving complex manufacturing, logistics, and marketing efforts, and a mistimed product release could have significant consequences on a firm's profit stream (see Hendricks and Singhal 1997 and Lilien and Yoon 1990). Furthermore, firms cannot release new generations of a product in rapid succession and expect consumers to willingly pay to upgrade each time a new version hits the market. A consumer will purchase the latest version on the market only if it is sufficiently more technologically advanced than the product he already owns. In this paper, we focus on how consumers' forward-looking behavior affects a firm's launch policy optimization problem. Consumers value newer and better technologies more than older ones and are often strategic in anticipating the introduction of future generations of a product when considering purchasing the current version on the market. A properly optimized launch policy should take this behavior into account when deciding the appropriate time to launch new products and whether to give consumers information about upcoming launches.

An illustrative case of consumer forward-looking behavior with respect to upcoming product launches

RIGHTSLINK

is the story of Apple's iPhone. Apple Inc., the world's largest corporation in 2014, currently earns over 50% of its revenue from iPhone sales. Since the release of the original iPhone in 2007, Apple has launched new generations of the smartphone every 12 to 16 months. Each new generation brings a more innovative design and/or better technology than its predecessor.

Figure 1 shows sales data and release dates of different generations of the iPhone. The numbers show a significant boost in sales at the moment of launching a new generation and relatively low sales in the quarter previous to a new launch. For instance, in the third quarter of 2011, Apple sold 3 million fewer iPhone 4's than analysts had expected. At the time, Bloomberg News (Satariano 2011) argued that the upcoming launch of the newer iPhone 4S caused consumers to withhold their purchases and that this postponement behavior was the cause behind Apple's first missed sales estimate in over six years. On October 7, 2011, the iPhone 4S preorders started, and according to Apple, more than 4 million orders were received within three days, setting a record in the history of mobile phone sales. This gives credence to the claim that consumers were withholding their purchases until the new generation of the iPhone came on the market. The same phenomenon occurred again in the third quarter of 2012, when earnings below analysts' expectations were blamed by Apple's chief executive officer (CEO) Tim Cook (Wells 2012) on the "incredible anticipation out there for a future product," a reference to the upcoming iPhone 5. The consumers, who had been holding back their smartphone purchases, rushed to buy the iPhone 5 as soon as it came out, and Apple reported that more than 5 million units were sold over the weekend of the product launch. This anecdotal evidence supports the observation that technology-savvy consumers look forward and internalize the value of delaying or skipping purchases with the goal of maximizing the total surplus they obtain from utilizing the firm's products. It is also clear, as acknowledged by Apple's CEO, that this strategic behavior has a profound impact on the company cash flow. How should a firm account for the consumer strategic behavior in its launch policy to mitigate its potential negative impact? What mechanism can be used to deter such behavior? How much money left on the table can be recovered when implementing such a mechanism? In this paper, we propose a stylized model to address these questions.

1.1. Overview of Main Results

We consider a monopolistic firm that launches successive generations of a given product, where generations that are introduced later have superior technology and are thus more valuable to consumers. The technology process evolves according to a Brownian motion with a positive drift. At any point in time, there is only one product generation available in the market, and the firm has the option to replace it with a new product generation that captures the increased technology level, at the expense of incurring a fixed launch cost. The demand is represented by a mass of infinitesimal consumers who decide which generations of the product to purchase. Consumers are forward-looking and may decide to hold back their purchases until a newer generation of the model is released. In this regard, the demand for the current generation model may be postponed and realized as demand for future product generations.

Figure 1 (Color online) Apple iPhone Sales and Release Dates in the United States

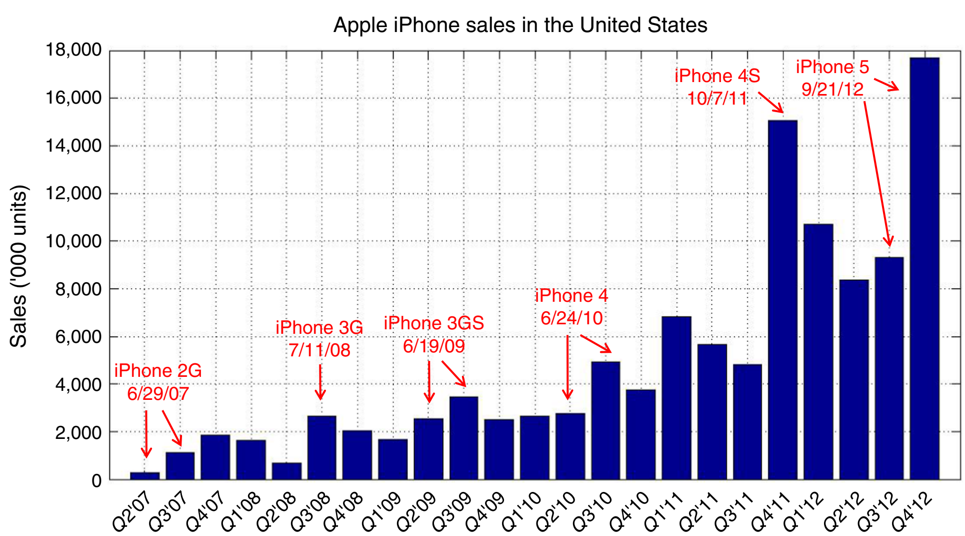

Notes. Data obtained from http://statista.com and http://apple.com. Note that when the release date is close to the end of a quarter, the boost in sales is also reflected in the next quarter.

RIGHTSLINK

The focus of our paper is on the characterization of the firm's optimal launch policy when facing such a postponement behavior on the consumers' side, which adds an important new dimension to the models available in the literature to study the problem of optimizing new product introductions. The firm decides the technology levels of its product releases and does so with the objective of maximizing the net present value of its cash flow. Consumers similarly optimize their own total discounted utilities.

We analyze a few variations of the general setting. We start by assuming that consumers are homogeneous in that they all have a common valuation for the product and that the product price is exogenous. In the first scenario, the firm makes product launch decisions "on the go," as time passes and technology improves. Both the firm and the consumers make decisions based on a Markovian state that represents the gap between the technology the firm has developed in the lab and the technology in the product currently available in the marketplace, as well as how outdated the technology the consumers currently own is. We find that in the unique equilibrium path of all Markov perfect equilibria, the firm releases products whenever the developed technology is better than the one available in the market by a given margin, a type of policy we describe as the *1-cycle*. In equilibrium, the firm utilizes the 1-cycle policy that maximizes its own utility.

The second scenario is one where the firm has the ability to commit to future products' technology levels. The firm first preannounces technology levels, and the consumers follow by optimizing their purchasing decisions to maximize their own utilities. We characterize the firm's optimal launch policy and the consumers' best response. Depending on system parameters, the optimal launch policy is either a single introduction over the entire horizon or multiple launches with two alternating cycles where a small technology increment is followed by a larger one, a policy we call the *2-cycle*. The latter is the case for more realistic system parameters. The rationale that supports this strategy is that, by exercising its commitment power, the firm can deter consumers' postponement behavior by promising a longer cycle (i.e., a major technology improvement) after two consecutive, relatively close introductions. This release policy is repeated over time and, interestingly, to a certain extent resembles Apple's sequence of introducing an iPhone with major improvements followed by an iPhone with minor changes. The iPhones 3G, 4, 5, and 6 represented major redesigns of the iPhone product line, whereas iPhones 3GS, 4S, 5S, and 6S correspond to fairly minor improvements. Depending on the product and the particular launch, Apple announces the product either several months (the original iPhone, the Apple Watch, the Mac Pro) or a few weeks (most iPhone generations) in advance of the launch. Therefore, in practice, Apple utilizes an announcing policy that is somewhere in between the two extreme cases that we analyze here (on the go or precommitment), but that leads to a sequence of releases with a similar structure to our 2-cycle, preannounced policy.

Next, we endogenize the firm's price and characterize the joint optimal launch and pricing policy for both scenarios. We verify that the difference in profits between the commitment and no-commitment scenarios can be significant, varying from 4% to 12% depending on the problem parameters. These findings imply that the firm's financial performance can be improved if the firm is able to commit to a launch policy in advance. To achieve these gains, the firm does not need to commit to an entire launch path but only to not releasing a new product too soon after two consecutive, rather close introductions.

We then generalize the special case of the model to the multiple consumer classes that have different valuations for the product. Using a duality argument, we show that the optimal launch policy under equilibrium constraints can be formulated as a single mixed-integer program and solve the problem for different market compositions. We also provide a recursive scheme for computing equilibria for the case of launches on the go. We find that even when launch costs are insignificant or zero, the consumer classes are coupled, and the firm earns less revenue than it would if it could sell to each consumer class separately. We show that the firm's profit decreases with an increase in consumer heterogeneity. Despite the additional complexity of the model with multiple customer classes, the output of the mixed-integer program still reflects the benefits the firm obtains from precommitting to the technology levels of the product launches in order to deter consumers' speculation about the future launch events.

1.2. Organization

The remainder of the paper is organized as follows: we begin with a review of the related literature in the next section. In §3, we introduce the model with a single consumer class (i.e., consumers are homogeneous with respect to valuations). The firm's launch policy optimization problem when the firm launches products on the go is analyzed in §4, followed by the case where the firm commits to a launch policy in advance, studied in §5. We solve the joint launch and pricing optimization problem in §6. Section 7 extends our model to incorporate multiple consumer classes. Our concluding remarks are reported in §8. All the proofs can be found in the online appendix (available as supplemental material at http://dx.doi.org/10.1287/mnsc.2015.2189).

RIGHTSLINK

2. Literature Review

Product launch policies have been explored in the operations management literature but usually in contexts where demand is exogenous and treated in aggregated terms or where consumers are myopic in the sense that they do not account for the benefit of waiting into their utility function.

Perhaps the closest paper to ours is the work by Krankel et al. (2006). As in our paper, they consider a firm that introduces successive generations of a product over an infinite time horizon with fixed introduction costs. In their model, the firm's technology evolution is exogenous and stochastic, and the demand is given by a Bass-type diffusion process. Their model leads to a trade-off similar to ours: delaying introduction to a later date may lead to the capture of further technology improvements, possibly at the cost of slowing sales for the existing product (and a decline in market potential for the product to be introduced, given their focus on durable products). They prove the optimality of a state-dependent threshold policy that is defined based on the technology level of the incumbent generation, cumulative sales of the current generation, and the ongoing technology level of the firm's research and development (R&D).

Cohen et al. (1996) consider a finite horizon model of product launch. They analyze the firm's trade-off between the reduction of new product introduction cycle time and improvements in product performance. As in our case, their model assumes that the newer generation replaces the previous generation. A product is introduced at the beginning of the time horizon, and the firm needs to determine when to introduce the new generation and what the target performance level should be for the new product. Their model concludes that faster is not necessarily better if the new product market potential is large and if the existing product (to be replaced) has a high margin. In addition, they show that it is better to take time to develop a superior product when the firm is faced with an intermediate level of competition.

The work by Paulson Gjerde et al. (2002) considers a firm's decision-making process regarding the level of innovation to incorporate in successive product generations and discusses the framework under which the firm should innovate to the technology frontier compared with adopting incremental improvements. Klastorin and Tsai (2004) develop a game-theoretic model with two profit-maximizing firms that enter a new market with competing products that have finite, known life cycles. The first entrant sets a price for its product and enjoys a monopoly situation until the second firm enters the market. When the second firm enters the market, both firms simultaneously set (or reset) their product prices knowing the design of both products at that time. They argue that a subgame perfect equilibrium occurs under certain conditions defined by the parameters of the model. Their model shows that product differentiation always arises at equilibrium because of the joint effects of resource utilization, price competition, and product life cycle. An important implication of their paper is that a profit-maximizing firm would be unwise to arbitrarily shorten its product life cycle for the sake of competition. Inspired by the interactions between Intel and Microsoft, Casadesus-Masanell and Yoffie (2007) analyze a dynamic duopoly model between producers of complementary products. They study the timing of the release of two consecutive PC generations. They find that the original investments in R&D are similar, but the timing of new product releases is misaligned: Intel wants to release a new generation microprocessor early, whereas Microsoft prefers to delay the new release so as to build and profit from the installed base of the first generation.

Kumar and Swaminathan (2003) consider a product innovation problem under capacity constraints. Their model of demand, modified from the original model of Bass (1969), captures the effect of unmet past demand on future demand. The firm's finite production capacity plays a major role. Under this constraint, they propose a "build-up" heuristic that is a good approximation to the optimal policy: the firm does not sell for a period of time and builds up enough inventory to never lose sales once it begins selling.

The effect of environmental regulation on product introduction policies was studied by Plambeck and Wang (2009). They consider a manufacturer who chooses the expenditure level and the development time for the next product generation, which together determine its quality. Consumers purchase the new product and dispose of the previous generation product, which becomes e-waste. The price of a new product strictly increases with its quality, and consumers form rational expectations about the timing of the next launch. In a monopolistic case, to maximize its profit, the firm introduces new products too quickly and spends too little on R&D for each product. As the authors point out, if the firm could publicly commit to increasing the development time for the next new product, customers would anticipate using the current new product for longer and would therefore be willing to pay more for it. A similar result is observed under a duopoly. The duopolists introduce new products too quickly in the sense that if they could jointly commit to longer development times, both firms would earn greater discounted profits. Plambeck and Wang (2009) find that the firm releases products at a rate that is too fast for its own profits but too slow for the consumers. They do not consider, however, the class of asymmetric launch policies that the firm can utilize when it has commitment

RIGHTSLINK

power. With commitment and arbitrary launch policies, we find that the firm prefers to alternate between minor and major launches, rather than extend the time between launches. Even without commitment, we find a result different from Plambeck and Wang (2009), with the firm releasing products on the go at its optimal rate, which is faster than the consumer optimal rate. We believe this difference is due to the distinct assumptions on the technology process. We assume the technology is given by an exogenous Brownian motion with a positive drift, whereas Plambeck and Wang assume technology is generated according to a Cobb–Douglas production function, where a shorter development time can be compensated by an increased research expenditure.

The problem of optimal introduction of new products has also been studied by the marketing community. Most marketing papers build on the seminal work by Bass (1969) on product diffusion, by incorporating multiple product generations into their models. Bayus (1992) considers two product generations with overlapping diffusion of the generations. The paper analyzes the prices for the two generations that maximize the discounted profit after the second product is launched. Norton and Bass (1987, 1992) study the substitution effect of having multiple product generations in the market simultaneously. Pae and Lehmann (2003) empirically analyze the impact of intergeneration time on product diffusion in two industries, random access memory chips and steel. Stremersch et al. (2010) empirically addresses the question of whether introducing new product generations accelerates demand growth, finding that the passage of time, not product launches, is the main driver of demand growth. Wilson and Norton (1989) show that product line extensions should be either introduced early in a product's life cycle or not introduced at all, depending on the degree of substitutability between the original product and its extension. Mahajan and Muller (1996) propose an extension of the Bass diffusion model that captures substitution effects between different generations of a product. In contrast to Wilson and Norton (1989), they find that launching a new generation when the earlier product becomes mature is sometimes the optimal strategy.

There is also related work in the information systems literature about software release planning. Ruhe and Saliu (2005) consider the problem of assigning a finite set of features to different software releases. Their objective is to maximize the average satisfaction of the stakeholders subject to resource and dependency constraints on development. The problem is modeled as an integer programming problem, and the solution offered is based on a linear programming relaxation heuristic. Greer and Ruhe (2004) consider the release planning problem in a dynamic environment, where the number of releases is not fixed. After every release, they reoptimize the problem to decide on the next release features. Another related problem studied in information systems is the optimal bundling of information goods (see Bakos and Brynjolfsson 1999), which is a static counterpart to the problem of optimal bundling of technologies developed over time.

Our work is also closely related to the economics literature on adoption of new technologies. Balcer and Lippman (1984) consider the problem of the adoption of new technologies by a firm facing an exogenous technology improvement process. Similar to our optimal policy under no commitment, they show that the firm will adopt the current best practice if the difference between the available technology and the current technology exceeds a certain threshold. They further show that as time passes without new technological advances, it may become profitable to incorporate a technology that has been available even though it was not profitable to do so in the past. Farzin et al. (1998) consider an infinite horizon dynamic programming framework to investigate the optimal timing of technology adoption. In their model, technology follows a stochastic jump process. They explicitly address the option value of delaying adoption and compare the results to those using traditional net present value methods, in which this delaying option is ignored and technology adoption takes place as long as the resulting net cash flows are positive. They analyze optimal switching times to a newer technology for the case of a single and finite known number of switches allowed, and they observe that the firm's optimal timing is greatly influenced by technological uncertainties.

Our paper contributes to the literature on forward-looking consumers in the context of operational settings. Since the mid-2000s, there has been growing interest within the revenue management community in modeling the strategic behavior of consumers and developing ways to mitigate the adverse impact of this forward-looking behavior on the firm's revenues (e.g., see Aviv et al. 2009). Several of the proposed mechanisms rely on commitment devices related to inventory availability and preannounced prices (e.g., Elmaghraby et al. 2008, Aviv and Pazgal 2008, Su and Zhang 2008, Correa et al. 2011, Borgs et al. 2014, Besbes and Lobel 2015), using internal price-matching policies (as in Lai et al. 2010), or using time-binding reservations (e.g., Osadchiy and Vulcano 2010). In a similar spirit to the aforementioned papers, the preannouncing of launch technology levels of the successive products that we propose serves as a commitment device that enables firms to achieve significantly more surplus compared with the on-the-go case.

RIGHTSLINK

3. Model

3.1. Description

A firm continually develops technology over time. The technology developed by the firm follows a stochastic process, which we assume to be a Brownian motion $Z(t) = \mu t + \sigma B(t)$, with positive drift μ and variance σ^2, and where $B(t)$ denotes the standard Brownian motion. We represent a sample path of the Brownian motion from time 0 up until time t by $\omega_t = \{Z(s): 0 \le s \le t\} \in \Omega_t$ and an entire sample path of the Brownian motion by $\omega = \{Z(s): s \ge 0\} \in \Omega$.

At any time $t \in \mathbb{R}^+$, the firm can launch a new product in the market with technology level $Z(t)$. The firm is a monopolist and chooses its launch policy to maximize the net present value of its cash flow. Whenever the firm decides to launch a new product, it incurs a fixed launch cost $K > 0$. We represent the set of times the firm launches new products in the market by $\boldsymbol{\tau}$ and the set of technology levels of these products by $\mathbf{z}$. Whenever the set of launch dates forms a sequence, we represent the time of the ith product release by τ_i and its technology level by $z_i = Z(\tau_i)$. Policies where the firm introduces only finitely many generations over time can also be captured by this representation by letting $z_j = \infty$ for some j, with the understanding that only $j-1$ generations of the product were available before its final disappearance from the marketplace.

We assume that only one product generation is available in the market at any point in time, so the technology level introduced at time τ_j remains active during the time $\tau_j \le t < \tau_{j+1}$, and it expires at time τ_{j+1}. We denote $w(t)$ as the technology available for the consumers at time t. Formally stated,

$$w(t) = \sup_{\{t' \in \boldsymbol{\tau}\} \cap \{t' \le t\}} Z(t'),$$

with the convention that $w(t) = 0$ for any t before the first product launch. We further assume that the firm has the unlimited capacity to produce and deliver its product to its customers.

We assume consumers are infinitesimal and normalize their mass to 1, representing each consumer by a location θ in the $[0,1]$ line segment. We assume all consumers have homogeneous preferences, with a common valuation v, an assumption we relax in §7. Consumers are assumed to be strategic, in that they take into account the value of delaying their purchases and optimize their discounted total surplus over time. For any given consumer $\theta \in [0,1]$, we represent his set of purchase times by $\boldsymbol{\kappa}^\theta$, with the ith purchase being represented by κ_i^θ, whenever the set of purchases forms a sequence. The set of technology levels of consumer θ's purchases is represented by $\mathbf{q}^\theta$, with the ith purchase having technology level $q_i^\theta = w(\kappa_i^\theta)$. We represent the technology level owned by consumer θ at time t by $C^\theta(t)$, which is equal to the technology available on the market at the time of the latest purchase; i.e.,

$$C^\theta(t) = \sup_{\{t' \in \boldsymbol{\kappa}^\theta\} \cap \{t' \le t\}} w(t'),$$

with $C^\theta(t) = 0$ at any time t before the first purchase. Whenever the consumer makes a purchase, he pays a price $p > 0$ to the firm, and the firm incurs a cost $c \ge 0$ to manufacture and deliver the good to the consumer. The consumer accrues utility as he uses the product over time. He does so at a higher rate when owning a technologically more advanced product. The consumer's instantaneous consumption value at time t is $vC^\theta(t)$; i.e., the rate is proportional to the technology he owns. We assume that both the consumers and the firm discount the future at rate δ (we relax this assumption in §7).

We represent a strategy of the firm by $s^f \in S^f$ and a strategy of a consumer θ by $s^\theta \in S^c$. We postpone the formal definition of the strategy spaces S^f and S^c until §§4 and 5, where we study two distinct kinds of launch strategies (on the go and preannoucing, respectively) the firm might adopt. For a given strategy of the firm s^f, the strategies of all consumers $\{s^{\theta'}\}_{\theta'}$, and a sample path of the Brownian motion ω, a consumer θ's total discounted utility U^θ will be equal to his discounted value obtained from using the firm's products minus the payments he makes to the firm:

$$U^\theta(s^f, \{s^{\theta'}\}_{\theta'}, \omega) = \int_{t=0}^{\infty} vC^\theta(t)e^{-\delta t}\,dt - p\sum_{i=1}^{\infty} e^{-\delta\kappa_i^\theta}.$$

We can solve the integration in the equation above by noting that the technology owned by a consumer $C^\theta(t)$ is constant in between purchase times $\boldsymbol{\kappa}^\theta$:

$$\begin{aligned} U^\theta(s^f, \{s^{\theta'}\}_{\theta'}, \omega) &= \sum_{i=1}^{\infty}\left(\int_{\kappa_i^\theta}^{\kappa_{i+1}^\theta} vq_i^\theta e^{-\delta t}dt - pe^{-\delta\kappa_i^\theta}\right) \\ &= \sum_{i=1}^{\infty}\left(q_i^\theta\left(e^{-\delta\kappa_i^\theta} - e^{-\delta\kappa_{i+1}^\theta}\right)\frac{v}{\delta} - pe^{-\delta\kappa_i^\theta}\right). \quad (1) \end{aligned}$$

The firm's utility U^f is given by the total sales profit minus the launch costs; i.e.,

$$U^f(s^f, \{s^{\theta'}\}_{\theta'}, \omega) = (p-c)\int_0^1 \sum_{i=1}^{\infty} e^{-\delta\kappa_i^\theta}\,d\theta - K\sum_{i=1}^{\infty} e^{-\delta\tau_i}. \quad (2)$$

We assume $(p-c) > K$ since otherwise the firm's optimal policy would be to never launch any product. As is generally done when representing random variables, we will often suppress the dependence on ω from U^θ and U^f. To ensure that all limits are well defined without excessive mathematical formalism,

RIGHTSLINK

Figure 2 (Color online) Product Launches and the Consumer Utility

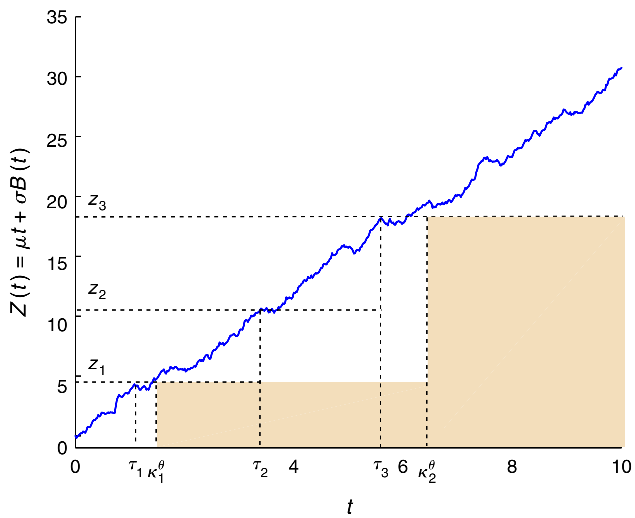

we focus on equilibria where the consumer population adopts at most finitely many different strategies. Given the homogeneity of the consumer market, we will focus on characterizing symmetric equilibria, where all consumers adopt the same strategy. In such situations, we will further simplify the notation and represent the strategy adopted by all consumers by $s^c \in S^c$. Thus, in symmetric equilibria, the firm's and the consumers' utilities will be represented, respectively, by the random variables $U^f(s^f, s^c)$ and $U^c(s^f, s^c)$.

Figure 2 schematically shows the model dynamics, showing a possible realization of the technology development $Z(t)$, the launch dates $\boldsymbol{\tau}$, the launch technology levels $\mathbf{z}$, and the purchases of a given consumer θ. The firm launches the first product with technology level z_1 at time τ_1. This product will be in the market until time τ_2, which is the launch time of the second product generation. A consumer θ purchases this product at time κ_1^θ; he thereafter earns utility at a rate proportional to z_1. At time τ_2, the firm launches the second product generation with technology level z_2, which replaces the first product. In the realization represented in Figure 2, the consumer θ never purchases the second product released. The firm eventually launches the third product generation with technology level z_3 at time τ_3, which the consumer purchases at time κ_2^θ.

3.2. Discussion of the Model Assumptions

When developing our model, we deliberately simplified several aspects of the problem. For example, we assume that the firm operates as a monopoly, focusing on the game-theoretic analysis of the interaction between the firm and its consumers. This is a reasonable assumption for settings where the firm has a loyal customer base that prefers to buy the firm's product to a competitor's. This is particularly prevalent in markets such as the ones for smartphones, tablets, and laptop computers, where customers often buy into entire product ecosystems. These ecosystems enable seamless integration across products and encourage customers to invest in apps that can easily be ported across a firm's product line but not to a competitor's product, making the true cost of switching firms potentially quite high.

In our model, the firm has infinite capacity. We make this assumption since most technology firms have sufficient initial inventory of the new product at the launch times, or the initial high demand is satisfied by the firm in a relatively short time frame. For example, despite the extremely high demand for new iPhones on launch dates, Apple has been able to satisfy the excess demand within a few weeks. Considering that new product launches occur roughly once a year, a few weeks is a sufficiently short response time, making it reasonable to assume that the initial high demand is fulfilled. We make these assumptions both for tractability purposes and to emphasize that even in the absence of such confounding factors, not having a publicly announced launch policy when customers are forward-looking is potentially costly for the firm.

We also assume that all product generations are sold at the same price. Changing prices could be a risky proposition for technology firms that seek a long-term relationship with their customer base, as the consumer outrage that followed the price drop of the original iPhone less than three months after its introduction in July 2007 demonstrated. After this price imbroglio forced Apple to issue rebates to early iPhone buyers, Apple has kept the price with a service contract of the lowest storage version of its flagship phone at a constant US$199 in the United States, from the release of the iPhone 3G in 2008 at least through the launch of the iPhone 6 in 2014. Adner and Levinthal (2001) provide good evidence that the prices of different technology product generations tend to become constant over time. Furthermore, this assumption of fixed marginal revenue is a commonly adopted one in the operations literature (see, for example, Krankel et al. 2006, Kumar and Swaminathan 2003).

The modeling of the technology process as a Brownian motion with positive drift captures the uncertainty in the R&D process. The drift represents the progressive increment that occurs as a result of the technological investments made by the firm and its suppliers. The risk and inherent uncertainty associated with the R&D process is captured by the standard deviation. Note that though the firm faces setbacks in its technology development in our model, with the Brownian motion moving downward as well as upward, the technology encountered by the consumer in the marketplace is always improving over time, as the

RIGHTSLINK

firm will never choose to launch a product that is inferior to the one it currently offers to consumers. The model also assumes the technological development occurs publicly. In reality, firms sometimes hide their R&D projects from consumers and competitors. The assumption of modeling the technology level as a scalar is common in the literature (see, for example, Kouvelis and Mukhopadhyay 1999, Cohen et al. 1996, Moorthy 1988, Liu and Özer 2009).

Finally, in our model we assume a fixed market size, which is normalized to have unit size. In reality, successful products tend to attract growing customer bases. We did not attempt to model the impact of a launch policy on attracting new customers, but it is plausible that having new customers arrive over time as a function of the technology available on the market would accentuate the firm's incentive to introduce new products at a faster rate.

4. Launching Products On the Go

In this section, we analyze the scenario where the firm does not make any announcements about upcoming product launches and decides when to release a new product *on the go*. In practice, a firm might choose not to announce future product launches in order to maintain operational flexibility or as an attempt to forestall consumers from waiting until a new product comes on the market. Without a launch announcement, the firm and the consumers play a Markov perfect equilibrium (MPE), where both sides make their decisions based on their beliefs about each others' future behavior. An MPE is a special case of a subgame perfect equilibrium where the players make their decisions using Markovian strategies. Even though the players have the option of deviating to non-Markovian strategies, they do not choose to do so.

We consider a continuous-time formulation of the game where the firm dynamically decides whether to launch a new product generation or not, and consumers respond by deciding whether to purchase the current product on the market or not. We define the state of the game $M(t)$ at time t to be the pair $M(t) = (Z(t) - w(t), O(t))$, where $Z(t) - w(t)$ refers to the accumulated technology that has not yet been deployed in a product available on the market by time t, and $O(t)$ captures the aggregate outdatedness profile of the consumers' technology at time t. A single consumer θ's outdatedness $O^\theta(t)$ at time t is given by $O^\theta(t) = w(t) - C^\theta(t)$, which is the market technology minus the technology the consumer owns. The outdatedness profile $O(t)$ captures the fractions of the market with each level of outdatedness but does not distinguish individual customers. For example, if $w(t) = 5$ and $C^\theta(t) = 5$ for all $\theta \in [0, 0.9]$, and $C^\theta(t) = 3$ for all $\theta \in (0.9, 1]$, then $O(t) = ((0, 0.9), (2, 0.1))$, as 90% of the consumers have the most current technology on the market—an outdatedness level of zero—and 10% of the consumers have a technology that is two units of progress behind the current one available for purchase. The set of feasible values of $M(t)$ is denoted by $\mathcal{M}$. We consider MPEs with respect to the state space $\mathcal{M}$ to ensure that the firm only responds to the accumulated technology development and the aggregate outdatedness of the consumers, not individual consumers' purchasing histories.

At any time $t \in \mathbb{R}^+$, the firm decides whether to launch a new generation or not. Mathematically speaking, the firm's strategy is a mapping $s^f\colon \mathcal{M} \to \{0, 1\}$ such that $t \in \boldsymbol{\tau}$ whenever $s^f(M(t^-)) = 1$, where $M(t^-) = \lim_{t' \uparrow t} M(t')$. That is, the firm launches a new product at time t depending on the state of the market at the time immediately before it, which is represented by the limit t^- in our continuous-time game. At time t, a consumer θ observes the same state as the firm, $M(t^-)$, and his individual outdatedness $O^\theta(t^-) = \lim_{t' \uparrow t} O^\theta(t')$. Formally, a consumer θ's strategy is $s^\theta\colon \mathcal{M} \times \mathbb{R}^+ \to \{0, 1\}$, where $t \in \boldsymbol{\kappa}^\theta$ whenever $s^\theta(M(t^-), O^\theta(t^-)) = 1$. As is standard in game theory, we assume the players know each other's strategies. Therefore, a consumer knows when making his purchase decision whether a launch occurs at time t, since he knows the state $M(t^-)$ and, consequently, $s^f(M(t^-))$.

The firm and the consumers play a Markov perfect equilibrium $(s^f, \{s^{\theta'}\}_{\theta'})$ in the following manner. Every time a consumer makes a purchasing decision, he compares the utility from the potential purchase with the expected utility derived from continuing to use the current product he owns. When making this choice, he takes into account the firm's future launch strategy and evaluates how long he will continue using either the current or the newly purchased product. The firm, meanwhile, also takes into account consumers' purchasing strategies when deciding whether to launch a new product. As a result, the firm will not introduce a new product during the time in which the consumers' net utility from continuing the usage of the old product is higher than the incremental utility from the new purchase. In an MPE $(s^f, \{s^{\theta'}\}_{\theta'})$, at any state $M \in \mathcal{M}$, the firm's best response is to play according to s^f, assuming the consumers will play according to $\{s^{\theta'}\}_{\theta'}$ in the future, and at any state $M \in \mathcal{M}$ and outdatedness $O^\theta \in \mathbb{R}^+$, a consumer θ's best response is to play according to s^θ, assuming the firm will play according to s^f in the future and the other consumers will act according to $\{s^{\theta'}\}_{\theta' \neq \theta}$ now and in the future.

To analyze the MPE of this game, we first focus on the consumer best response. Suppose the firm utilizes a simple threshold strategy where it launches a new product whenever $Z(t) - w(t) \geq z$ for some threshold value $z > 0$. That is, if the technology available in the

RIGHTSLINK

lab is sufficiently better than the technology available in the market, the firm releases a new product. Naturally, if z is large, the consumers will respond by buying every time there is a new release, but consumers will prefer to skip generations of the product if z is close to zero. The following lemma characterizes the lowest threshold z at which consumers buy all products released by the firm.

Lemma 1. *Assume the firm releases a new product whenever $Z(t) - w(t) \geq z$ for some threshold $z \geq 0$. Then, consumers will purchase every release if and only if $z \geq z^*$, where z^* is the unique positive solution of the equation*

$$z(1 - e^{-\Delta z}) = \frac{p\delta}{v} \quad (3)$$

and Δ is the technology-dependent discount rate

$$\Delta = \frac{2\delta}{\sqrt{\mu^2 + 2\delta\sigma^2} + \mu}. \quad (4)$$

Equation (3) determines the threshold policy that leads consumers to become indifferent to any single purchase. To be clear, consumers still earn positive utility from their overall interactions with the firm whenever the products released exhibit an improvement of z^*; only the marginal purchase yields zero expected utility for the consumer under this release policy, in the sense that if the consumer were to skip one product launch, he would not lose any utility, but if he skipped two consecutive launches, he would earn a strictly lower utility. In Equation (3), the fact that consumers anticipate future upgrades when considering purchasing today appears in the formula as the term $e^{-\Delta z}$.

The parameter Δ, which we call the technology-dependent discount rate, is defined based on both the discount rate δ and the parameters of the Brownian motion $Z(t)$. Naturally, the term Δ is increasing in δ and decreasing in μ, as both increasing the discount rate or decreasing the positive drift of the technological progress makes the future releases less valuable today. Surprisingly, the term Δ is decreasing in the variance σ^2 of the technology Brownian motion. This implies that a more predictable technological process makes future technology less valuable in expectation. This is a consequence of the fact that launching a new product is a real option for the firm, as it will only release new products after technological progress is achieved, never after a technological setback. The interesting implication is that the firm does not face a trade-off between the drift μ and the variance σ^2 and, in fact, should be willing to invest in order to increase both parameters.

We now turn to our main result in this section, where we show the unique equilibrium path of any MPE of the game being played by the firm and the consumers.

Theorem 1. *In the equilibrium path of any MPE, product launches occur exclusively when $Z(t) - w(t) = z^*$, where z^* is defined in Equation (3). Furthermore, all consumers (except for a set of measure zero) buy at every product launch.*

The theorem above states that the launch cycle with threshold z^* is implementable in an MPE and is, in fact, the only possible Markovian equilibrium outcome. This is a positive result for the firm since z^* is the minimum threshold for the firm, as consumers would balk at purchasing new products if they were introduced any more often. This is driven by the fact that the firm has market power: all consumers are infinitesimal and no one individual per se affects the firm's profit. Along the equilibrium path, the consumers choose to purchase at every time when there is a launch, so the outdatedness profile is always identically equal to $O(t) = ((0, 1))$. Consumers, being infinitesimal, cannot individually change the outdatedness profile of the market. Therefore, verifying whether a deviation from the equilibrium path is profitable does not involve determining what consumers would do in other states of the game with different outdatedness profiles. We highlight that even though we do not determine what happens in many of the off-the-equilibrium-path states, our theorem applies in a setting with subgame perfection since the set of MPEs is a subset of the class of subgame perfect equilibria. Even though the equilibrium path is unique, there is multiplicity of equilibria. There exist asymmetric MPEs where a mass of consumers of measure zero does not buy at every launch. There is also the possibility of multiplicity due to different actions being taken at off-the-equilibrium-path states that cannot be reached via unilateral deviations.

In the next section, we explore whether the firm can do better than release a new product every time the technology improvement reaches the threshold z^*. In particular, we explore the optimal launch strategy of a firm that is able to commit to technology levels of future launches.

5. Optimizing Launch Cycles

We now consider the case where the firm decides on a launch policy in advance and announces it to the consumer market. The firm decides the sequence of technology levels of its products in advance, and the consumers respond by selecting the products that optimize their own utilities. We assume that the firm has commitment power, so that it is able to make credible promises about future generations' technology levels. In practice, firms do not announce entire sequences of product releases in advance, but they often do make announcements about upcoming products. The formulation studied in this section, with the firm being able to commit to an entire path of future

RIGHTSLINK

launches, models a best-case scenario for the firm and thus allows us to consider what types of announcements are valuable for the firm. We show below that the firm does not in fact need to be able to commit to entire technology paths in advance in order to implement the optimal policy.

In our model, the firm first announces all technology levels of future product generations. The firm commits to a launch policy, which is represented by its set of technology levels $\mathbf{z}$. In our framework, the firm does not make commitments with respect to the future timing of product launches, only technology levels. The set of policies available to the firm S^f is thus equal to the set of subsets of the positive real numbers $2^{\mathbb{R}^+}$. The launch date τ_i associated with a given launch technology level z_i is thus the hitting time of the Brownian motion that represents the technology process. The consumers then respond by making purchases as a function of the launch policy and the stochastic technology process. Therefore, a consumer's strategy set at time t is represented by $S_t^c\colon 2^{\mathbb{R}^+} \times \Omega_t \to \{0,1\}$, with $t \in \boldsymbol{\kappa}^\theta$ whenever $s_t^\theta(\mathbf{z}, \omega_t) = 1$. The consumer strategy space is the collection of strategy sets available to him at all times; i.e., $S^c = \{S_t^c\}_t$.

The firm would like to maximize its expected utility subject to the consumers' optimal response. The firm's product launch optimization problem can thus be formulated as follows:

$$\begin{aligned} \underset{\substack{\mathbf{z}\in 2^{\mathbb{R}^+} \\ s^\theta \in S^c,\, \theta\in[0,1]}}{\text{maximize}} \quad & \mathbb{E}[U^f(\mathbf{z}, \{s^\theta\}_\theta)] \\ \text{s.t.} \quad & \mathbb{E}[U^c(\mathbf{z}, (s^{\theta'}, \{s^\theta\}_{\theta\neq\theta'})) \mid \omega_t] \\ & \geq \mathbb{E}[U^c(\mathbf{z}, (\bar{s}^{\theta'}, \{s^\theta\}_{\theta\neq\theta'})) \mid \omega_t] \\ & \text{for all } t,\ \omega_t \in \Omega_t,\ \theta' \in [0,1],\ \bar{s}^{\theta'} \in S^c. \end{aligned} \tag{OPT-1}$$

In the formulation above, the firm selects its launch policy $\mathbf{z}$ and its consumers' purchasing behavior $\{s^\theta\}_\theta$, but it is restricted by a consumer incentive compatibility constraint. This constraint ensures that no focal consumer θ' wants to deviate from his strategy $s^{\theta'}$ to a different strategy $\bar{s}^{\theta'}$ at any possible sample path ω_t of the technology evolution.

The firm and the consumers clearly have conflicting interests regarding a desirable outcome. The firm would like to release products and have consumers buy them as frequently as possible. Because of discounting, the firm would also like purchases to occur earlier rather than later. Meanwhile, consumers need to accumulate value from each purchase and therefore prefer to wait until a product release contains enough technological improvement to justify a new expenditure. Thus, the firm's challenge is to find a launch policy that leads consumers to buy as often and as early as possible.

The first step in our analysis is to simplify the consumer's incentive compatibility constraint. Our first lemma shows that there exists a symmetric equilibrium where consumers buy all products the firm releases. We introduce a subset of the consumer strategy space S^c, to which the consumer can restrict himself without loss of optimality. These are strategies where the consumer buys products if, and only if, they have certain technology levels, and where consumers buy these products immediately after they are released. With a slight abuse of notation, we say the consumer is playing strategy $\mathbf{q} \in 2^{\mathbb{R}^+}$ if he buys exactly the set of products released with technology level in the set $\mathbf{q}$ and does so immediately upon each product release.

Lemma 2. *There exists an optimal solution of OPT-1 where all consumers purchase at every time when there is a new product launch; i.e., $\mathbf{q} = \mathbf{z}$. In this solution, the consumers' policy $\mathbf{q} = \mathbf{z}$ constitutes a best response if, and only if, $\mathbb{E}[U^c(\mathbf{z}, \mathbf{q})] \geq \mathbb{E}[U^c(\mathbf{z}, \bar{\mathbf{q}})]$ for all $\bar{\mathbf{q}} \subseteq \mathbf{z}$. In this solution, $\mathbf{z}$ can be represented by a sequence of launch technology levels $\{z_i\}_{i\in\mathbb{N}}$, where $z_i < z_{i+1}$ for all $i \in \mathbb{N}$.*

The consumers have no incentive to buy products at any point other than launch dates as they can anticipate purchases to their respective release dates and thus use more advanced products earlier and for longer periods of time. Meanwhile, if a firm introduces a product that the consumers do not buy, the firm can improve on its policy by removing this particular product launch. The lemma allows us to simplify OPT-1 to a search over a single sequence $\mathbf{z}$:

$$\begin{aligned} \max_{\mathbf{z}\in 2^{\mathbb{R}^+}} \quad & \mathbb{E}[U^f(\mathbf{z}, \mathbf{z})] \\ \text{s.t.} \quad & \mathbb{E}[U^c(\mathbf{z}, \mathbf{z})] \geq \mathbb{E}[U^c(\mathbf{z}, \bar{\mathbf{q}})] \\ & \text{for all } \bar{\mathbf{q}} \subseteq \mathbf{z}. \end{aligned} \tag{OPT-2}$$

The constraints in the formulation above are still unwieldy since they require the consumers to compare their set of purchase technologies $\mathbf{z}$ with any possible alternative set of purchase technologies $\bar{\mathbf{q}} \subseteq \mathbf{z}$. The next lemma and the subsequent proposition allow us to significantly simplify the consumer's problem.

Lemma 3. *For any given $\mathbf{z} = \{z_i\}_{i\in\mathbb{N}}$, the function $\mathbb{E}[U^c(\mathbf{z}, \mathbf{q})]$ is submodular in $\mathbf{q}$ when restricted to the domain of subsets of $\mathbf{z}$.*

The lemma above shows that the expected utility of the consumers is submodular in their purchases. That is, for two given consumer policies $\mathbf{q}' \subset \mathbf{q}'' \subset \mathbf{z}$, adding an extra purchase of technology $\hat{q} \in \mathbf{z}\backslash\mathbf{q}''$ adds more value to policy $\mathbf{q}'$ than to policy $\mathbf{q}''$. In other words, from a consumer's perspective, multiple generations of the firm's product constitute a set of imperfectly substitutable goods: early generations of the product create value by being available early, whereas late generations create value by containing advanced technology. Thus, the more products a consumer purchases,

RIGHTSLINK

the less value he obtains from adding a new purchase. Note that, though the function $\mathbb{E}[U^c(\mathbf{z},\mathbf{q})]$ is submodular in $\mathbf{q}$ for any given $\mathbf{z}$, it does not follow that $\mathbb{E}[U^c(\mathbf{z},\mathbf{q})]$ is a monotone function of $\mathbf{q}$. In fact, adding additional purchases to a set $\mathbf{q}$ can certainly have a negative impact on the consumer's overall utility. The following proposition leverages this submodularity result to show that the consumers can consider one product at a time when deciding whether to make purchases.

Proposition 1. *For a firm's launch policy $\mathbf{z}$, the consumer's purchasing policy $\mathbf{q}=\mathbf{z}$ is a best response if, and only if, $\mathbb{E}[U^c(\mathbf{z},\mathbf{z})] \geq \mathbb{E}[U^c(\mathbf{z},\mathbf{z}\backslash\{z_i\})]$ for all $i\in\mathbb{N}$.*

This proposition is one of the key simplifying ideas that makes the launch policy optimization problem tractable. Instead of considering whether a large group of purchases are collectively worth their price, a consumer need only think if each one of them is individually worth it. Its proof relies on the submodularity established in Lemma 3: an additional purchase z_i adds less value to the set of purchases $\mathbf{z}\backslash\{z_i\}$ than to any other smaller set of purchases $\bar{q}\subset\mathbf{z}\backslash\{z_i\}$. This proposition is particularly useful because the net utility from excluding purchase z_i from the consumer policy $\mathbf{z}$ depends only on the technologies z_{i-1}, z_i, and z_{i+1}, as given by

$$\mathbb{E}[U^c(\mathbf{z},\mathbf{z})]-\mathbb{E}[U^c(\mathbf{z},\mathbf{z}\backslash\{z_i\})] = \mathbb{E}\left[e^{-\delta\tau_i}\left[\frac{(z_i-z_{i-1})v}{\delta}(1-e^{-\delta(\tau_{i+1}-\tau_i)})-p\right]\right]. \quad (5)$$

The only relevant pieces of information from the set of launch technologies $\mathbf{z}$ for deciding whether purchasing the product with technology z_i is worth it for the consumer are z_{i-1}, z_i, and z_{i+1}. Proposition 1, combined with Equation (5), has the following important implication.

Observation 1. The firm can implement any launch policy $\mathbf{z}=\{z_i\}_{i\in\mathbb{N}}$ by committing to the technology level z_{i+1} whenever it releases a new product of technology level z_i.

That is, there is no need for the firm to commit to its entire launch sequence at time $t=0$. Instead, the firm can implement its launch policy by simply announcing the technology level of the following launch every time it introduces a product. In fact, we argue below that the firm does not even need to commit to a technology level after every launch in order to implement the optimal policy.

The formula in the right-hand side (RHS) of Equation (5) can be simplified. To determine whether it is nonnegative, we can ignore the always positive term $e^{-\delta\tau_i}$. Furthermore, $\mathbb{E}[e^{-\delta(\tau_{i+1}-\tau_i)}]$ is the moment-generating function of the hitting time of $z_{i+1}-z_i$ by a Brownian motion with drift μ and variance σ^2. In the proof of Lemma 1, we use a result from Karatzas and Shreve (1991) to argue that $\mathbb{E}[e^{-\delta(\tau_{i+1}-\tau_i)}]=e^{-\Delta(z_{i+1}-z_i)}$, where Δ is defined in Equation (4). We can further simplify the formula by writing it in terms of the difference in the technology levels, rather than the levels themselves. With a change of variables to $r_1=z_1$ and $r_i=z_i-z_{i-1}$ for all $i\geq 2$, Equation (5) is nonnegative if, and only if,

$$\frac{r_i v}{\delta}(1-e^{-\Delta r_{i+1}})-p\geq 0.$$

We can also simplify the firm's utility function under a firm policy $\mathbf{z}$ and a consumer policy $\mathbf{z}$ to

$$U^f(\mathbf{z},\mathbf{z})=(p-c-K)\sum_{i=1}^{\infty}\mathbb{E}[e^{-\delta\tau_i}]=(p-c-K)\sum_{i=1}^{\infty}e^{-\delta\sum_{j=1}^{i}r_j},$$

where the second equality combines the fact that $z_i=\sum_{j=1}^{i}r_j$ with the same moment-generating function argument from the proof of Lemma 1. The term $p-c-K$ is a positive constant and can thus be ignored. With the two equations above, we can recast the firm's launch optimization problem as the following deterministic problem:

$$\max_{r_i\in[0,\infty],\, i\in\mathbb{N}} \sum_{i=1}^{\infty}e^{-\Delta\sum_{j=1}^{i}r_j} \quad \text{s.t.}\quad r_i(1-e^{-\Delta r_{i+1}})\geq\frac{\delta p}{v}\quad \text{for all } i\in\mathbb{N}. \quad \text{(OPT-3)}$$

In the formulation above, we explicitly allow the decision variables to take value infinity, which will occur if the optimal policy contains only finitely many product launches.

The next theorem establishes the optimal solution of the firm's launch policy optimization problem and is one of our main results. It involves a launch policy that we denote as the 2-*cycle*, where the firm introduces new generations of its product according to interlaunch technology levels $(r_1,r_2,r_1,r_2,\ldots,)$ for some values r_1 and r_2.

Theorem 2. *The unique optimal launch policy is the following*:

- *If $v/(p\Delta\delta)\leq 1/\lambda$, then the firm should launch a single generation of the product with technology level $(\delta p)/v$ and subsequently exit the market.*
- *If $v/(p\Delta\delta)>1/\lambda$, then the firm should launch new product generations in accordance with a 2-cycle policy with cycles $r_1<r_2$, where r_1 is the unique solution of the equation*

$$p\delta(e^{-\Delta t}+1+\Delta t)=2v\Delta t^2, \quad (6)$$

and $r_2=(1/\Delta)\log(vr_1/(vr_1-p\Delta))$,
where λ is the unique solution of the equation $(\lambda-1)e^{\lambda}=1$.[1]

[1] We can write λ in terms of the Lambert W function as $\lambda=W(1/e)+1$, which yields $\lambda\approx 1.27846$. The Lambert W function is defined implicitly by the equality $z=W(z)e^{W(z)}$ for all z.

RIGHTSLINK

Figure 3 (Color online) Values of r_1 and r_2 as Determined by Theorem 2

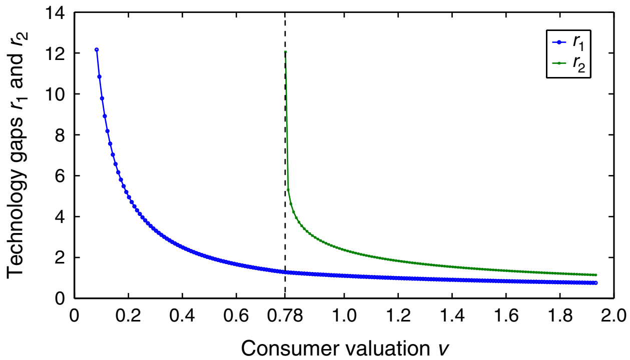

Notes. The parameters p, δ, and Δ are kept constant at 1. The value $v = 0.782$ is the threshold between the single launch and the 2-cycle regimes.

To prove this theorem, we relax OPT-3 by removing all the constraints associated with even values of i. This relaxed problem can be solved explicitly because all remaining constraints ($i = 1, 3, 5, \ldots$) must be binding as the objective function is decreasing in all r_i's, the constraints are increasing in all r_i's, and each variable r_i appears in exactly one constraint. Note that there are two ways in which the constraint

$$r_1(1 - e^{-\Delta r_2}) = \frac{\delta p}{v}$$

can be true: either $r_1 = \delta p / v$ and $r_2 = \infty$ or $r_1 > \delta p / v$ and $r_2 < \infty$. This gives rise to the two possible optimal regimes. The solution of the relaxed problem is completed by a decoupling argument, which shows by symmetry that all odd r_i's take the same value (thus, all even r_i's also take the same value). The final step of the proof is to verify that the solution of the relaxed problem is feasible in OPT-3 and, therefore, is an optimal solution of this problem as well.

The technology gaps of the optimal launch policy are depicted in Figure 3. The theorem shows that the key ratio for determining the optimal launch policy is $v/(p\Delta\delta)$, which can be interpreted as the product's *long-term technological value*, with the understanding that a low value means that the consumer expenditure is amortized slowly over time. For instance, this occurs when v/p is low, and the product gains little value over time compared with its relatively high price. Low drift μ and low variance σ^2 also lead to low long-term technological value through high Δ. It might seem counterintuitive that low drift has a similar effect to low variance, but both lead to new products being introduced less often. This phenomenon is similar to what occurs with financial options, where both higher expected returns and higher variance lead to higher option prices. This occurs because releasing a new product is an option for the firm, one that it will not choose to exercise unless the technology level goes up from the previous release. Beyond that, if the future is heavily discounted (δ is high), then the product is of low long-term technological value since the next time when consumers would purchase an upgrade of the product will occur at a time that is so far into the future that the net present value of this eventual upgrade is small. In this scenario, the optimal policy for the firm is to introduce the product once and stop its development immediately after. That is, for a product of low long-term technological value, the firm is better off discontinuing the product than producing successive generations of it. Note that this occurs even though our model assumes the firm's technology evolves exogenously, and therefore, the firm does not save any sort of research expenditure by discontinuing a product. This occurs solely because the firm can bring forward in time the consumers' (first and only) purchase by committing to not releasing upgrades. If no upgrades will ever be available, the consumers will expect a higher surplus from the purchase and thus be willing to buy it earlier, thus benefiting the firm, though overall its total profits will be on the low side.

By contrast, the firm's optimal policy is to introduce a series of upgrades over time if the product has a high long-term technology value $v/(p\Delta\delta)$. In this case, the firm's optimal launch policy is a 2-cycle one where a major technology improvement follows a minor one. Recall from §4 that when the firm does not announce a launch policy, it ends up in an equilibrium where it introduces products according to a 1-cycle policy. By committing to future launch dates, the firm can improve its profits by shortening the length of half of its launches (i.e., by launching a product with a minor technology improvement with respect to the previous release). In fact, as a practical matter, the firm only needs to precommit to the major cycle, whereas the minor cycle could be executed on the go. Furthermore, it only needs to commit not to launch a new product too early. That is, at the moment of releasing a new product on the go (i.e., ending a minor cycle), the firm announces that the next launch will not occur until a given date or until a certain technological level is achieved. This announcement will trigger early purchases from the consumers that otherwise would choose to wait. Once the major cycle is completed, the firm does not need to make any announcements but can instead release the next product as an on-the-go launch. Note that our analysis does not recommend in any way that the firm withhold technology from a launch or that it launch substandard products. Under the optimal policy, the firm always releases the best technology it currently has available and the firm only releases products the consumers will find desirable, because all products launched satisfy consumer incentive compatibility constraints.

RIGHTSLINK

The magnitude of $v/(p\Delta\delta)$ defines the shape of the two cycles. When the long-term technological value $v/(p\Delta\delta)$ is high, the firm will release products often. In this case, attempting to anticipate the first customer purchase by delaying the second release date is not a particularly valuable tactic for the firm, and r_1 will be only slightly smaller than r_2. In the limit as $v/(p\Delta\delta)$ goes to infinity, r_1 is equal to r_2, and it turns out that preannouncing product launches converges to the on-the-go launch policy z^* characterized in Theorem 1, so that $r_1 = r_2 = z^*$ (see the proof of Theorem 2). On the other hand, if $v/(p\Delta\delta)$ is only slightly above $1/\lambda$, the two cycles will be very asymmetric. As the long-term technological value approaches $1/\lambda$ from above, the second cycle r_2 becomes increasingly large, eventually converging to the single launch solution.

6. Price Optimization

Our analysis has assumed so far that the price in both scenarios (launching products on the go and at preannounced times) is exogenously given. The optimal policy in each case is defined by cycles whose definition depends on a given price p. In this section, we extend the analysis to a situation where the price is endogenously determined. Namely, we characterize the optimal joint pricing and product launch policy for both scenarios. In both cases, we assume that a single price p is chosen by the firm at $t = 0$ and that the firm uses this price throughout the game.

6.1. Optimal Pricing for Launching Products On the Go

In §4, we considered the firm's launch policy in the Markov perfect equilibrium that occurs when the firm launches products on the go and showed, in Theorem 1, that the equilibrium path is given by a single cycle z^*, where $\mathbf{z} = \boldsymbol{\kappa}^\theta = \{z^*, 2z^*, 3z^*, \ldots\}$ for all θ. The firm's expected utility function (see Equation (2)) is thus given by

$$\mathbb{E}[U^f(\mathbf{z},\mathbf{z})] = \sum_{i=1}^{\infty}((p-c-K)e^{-\Delta i z^*}) = (p-c-K)\frac{e^{-\Delta z^*}}{1-e^{-\Delta z^*}}. \quad (7)$$

Recall from Lemma 1 that, for a given price p, the optimal cycle length z^* is the unique solution to Equation (3). Hence, the associated optimal price function p must verify

$$p = \frac{vz^*(1-e^{-\Delta z^*})}{\delta}. \quad (8)$$

Since v, δ, and z^* are all nonnegative, the RHS of the equation above is increasing in z^*, and thus, there is a one-to-one correspondence between cycle length z^* and associated optimal price p. Substituting Equation (8) into Equation (7), we obtain the firm's joint price and cycle length optimization problem just as a function of z^*. We denote the firm's expected utility by $\mathbb{E}[U^f(z,z)]$; i.e.,

$$\mathbb{E}[U^f(z,z)] = \left(\frac{vz^*(1-e^{-\Delta z^*})}{\delta} - c - K\right)\frac{e^{-\Delta z^*}}{1-e^{-\Delta z^*}}. \quad (9)$$

This function is unimodal in the cycle length z^*, and given the optimality condition of Equation (8), we get the following.

Proposition 2. *For the equilibrium path obtained in Theorem 1, the firm's expected utility function is unimodal in price.*

The optimal price does not have a closed-form expression, but it can be computed by a simple one-dimensional line search.

6.2. Optimal Pricing for Preannounced Introductions

In §5, we considered the optimal launch policy when the firm commits to a preannounced schedule of launch technology levels. For a given price p, Theorem 2 characterizes the two possible outcomes: either a single introduction or a 2-cycle sequence of minor and major interlaunch technology levels. When there is a single introduction, the technology level of this introduction r_1 is equal to $p\delta/v$, which is strictly increasing in p. Likewise, when there are multiple introductions, the unique solution r_1 to Equation (6) is strictly increasing in p. In either case, there is a one-to-one correspondence between r_1 and p. Moreover, in the 2-cycle policy, the long cycle r_2 is also uniquely determined from r_1 (and, therefore, from p).

The following theorem characterizes the optimal price for each of the two possible optimal launch policies.

Theorem 3. *Let λ be defined as in Theorem 2. The firm's profit function is unimodal in price, conditionally on the firm utilizing the optimal launch policy for any given price. In particular, the optimal price is given by the following*:

- *If $v/((c+K)\delta\Delta) \le 1/(\lambda-1)$, then the single launch policy is optimal, with price*

$$p^* = c + K + \frac{v}{\delta\Delta}. \quad (10)$$

- *If $v/((c+K)\delta\Delta) > 1/(\lambda-1)$, then the 2-cycle launch policy is optimal, with price*

$$p^* = \frac{2v\Delta(r_1^*)^2 e^{\delta r_1^*}}{\delta(1+e^{\delta r_1^*}(1+\delta r_1^*))}, \quad (11)$$

where r_1^* is the unique positive solution to the equation

$$\Delta^3(c+K) + \frac{2\Delta^2(c+K)(e^{-\Delta r_1}+1)}{r_1} + \frac{\Delta(c+K)(e^{-\Delta r_1}+1)^2}{(r_1)^2} + \frac{2v\Delta(1-e^{-2\Delta r_1})}{\delta r_1} - \frac{2\Delta^3 v r_1}{\delta} = 0. \quad (12)$$

RIGHTSLINK

Figure 4 (Color online) Joint Launching and Pricing Policy for the Preannounced Introduction Case

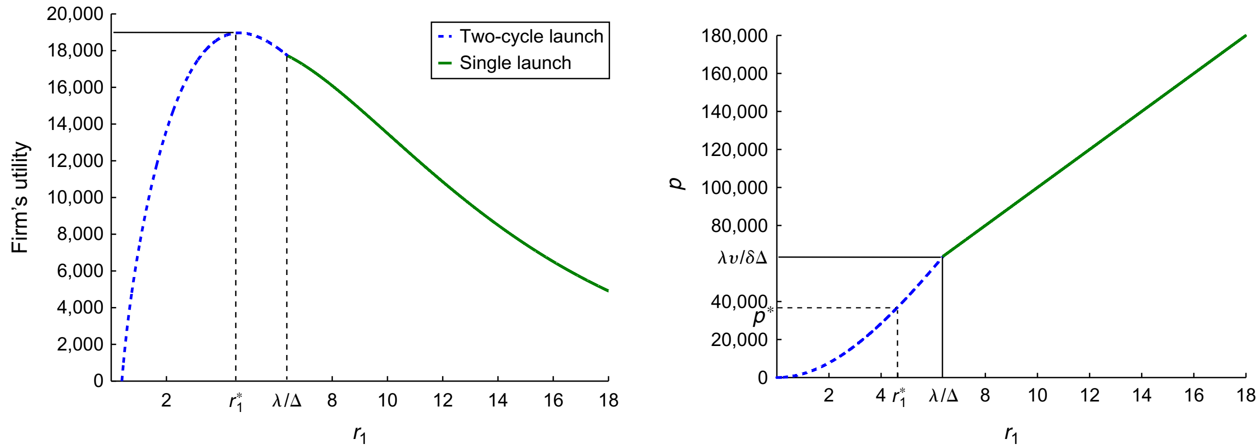

Notes. Parameter values are $v = 2{,}000$, $c = 0$, $\delta = 0.2$, $\mu = 1$, $\sigma = 0$, and $K = 300$. Left: firm's expected utility as a function of r_1, assuming that the price charged is the unique optimal $p(r_1)$. In this case, the optimal minor launch is given by $r_1^* = 4.64$. Right: bijection between r_1 and $p(r_1)$. In this case, the optimal minor launch is given by r_1^* and maps onto a unique optimal price $p^* = p(r_1^*) = 37{,}065$. The associated major launch cycle is given by $r_2^* = 8.02$.

The proof establishes a bijection between r_1 and $p(r_1)$, and then it rests on the unimodality of the firm's utility function in r_1. Specifically, for any given $p > c + K$, there is a unique corresponding $r_1(p)$ that solves the firm's profit optimization problem (OPT-1) or after taking the inverse, there is a well-defined function $p(r_1)$.

Figure 4 illustrates the firm's utility function for the joint launching and pricing problem. The graph in the left plots the firm's utility as a function of the minor cycle phase r_1 when the price charged is the unique $p(r_1)$. The graph in the right shows the bijection between r_1 and $p(r_1)$, so that the unique optimal r_1^* maps onto a unique optimal price $p^* = p(r_1^*)$. This is given by the linear relation $p^* = vr_1^*/\delta$ in case r_1^* turns out to be big enough (and the firm implements a single introduction policy). Otherwise, the optimal p^* is given by Equation (11).[2]

Observe that for the single-cycle regime, the optimal price p^* in Equation (10) is increasing in the fixed launch cost K, in the consumers' valuation v, and in the variable cost c, and it is decreasing in the discount factor δ. The same holds for p^* in the two-cycle regime, which can be verified from the following two facts: the RHS in Equation (11) is increasing in r_1^*—and hence p^* is increasing in r_1^*—and the left-hand side in Equation (12) is strictly decreasing in r_1.

In Theorem 3, the expression $v/((c+K)\delta\Delta)$ plays the role of the product's long-term technological value. Note that when the price is endogenously determined, the aggregate cost per launch $c + K$ replaces p compared with the expression in Theorem 2. The aggregate cost per launch, which contains the production cost c as well as the fixed cost K, does not affect the long-term technological value of the product when the price is exogenous, since it does not affect the consumer base. When the price is set by the firm, however, the aggregate cost per unit is an important element in determining the retail price of the product and, thus, impacts the long-term technological value of the product. Note also that the cost-based version of the long-term technological value $v/((c+K)\delta\Delta)$ that determines the optimal launch policy in Theorem 3 requires a higher bar of $1/(\lambda - 1)$ for a two-cycle launch policy to be optimal compared with the bar of $1/\lambda$ for the price-based long-term technological value from Theorem 2.

6.3. The Value of Commitment

We now compare the profit the firm obtains from utilizing a preannounced policy as opposed to launching products on the go. Recall that the optimal preannounced policy only requires the firm to announce the launch of technology levels for the upcoming product and only to do so at the end of a minor launch cycle, not to announce the entire path of future launches. In both scenarios, we assume the firm utilizes the optimal joint pricing and launch policy.

Table 1 shows for a given aggregate cost per unit $c + K$, consumer value v, and discount factor δ the optimal pricing policy and the associated profit for each scenario. We can see that the utility difference in favor of the commitment case can be significant, with values reaching 12% in the cases listed there. Thus, we observe that when the firm can commit to an introduction policy, it can recover a significant portion of the revenues that were left on the table in case

[2] For the ease of exposition, in our plots in Figure 4, we have extended the domain for r_1 so that $p(r_1)$ goes below the threshold $c + K$. The strict monotonicity of $p(r_1)$ is still preserved in the extended domain, though it could lead to a negative utility as indicated in the left panel of Figure 4 for r_1 sufficiently close to zero.

RIGHTSLINK

Table 1 Profit Comparison Between On-the-Go and Preannounced Policies Under Optimal Prices

			On the go (OTG)		Preannounced (Prean)		
$c+K$	v	δ	p^*_{OTG}	Profit (p^*_{OTG})	p^*_{Prean}	Profit (p^*_{Prean})	% Gap
30	400	0.20	6,395.74	3,661.44	7,333.63	3,809.09	4.03
		0.60	773.15	392.15	882.69	409.89	4.52
		1.00	318.22	131.76	361.53	138.96	5.46
	700	0.20	11,136.88	6,420.49	12,773.77	6,677.67	4.01
		0.60	1,301.60	698.37	1,488.77	728.33	4.29
		1.00	511.25	241.34	582.50	253.02	4.84
200	400	0.20	6,799.80	3,566.70	7,771.56	3,723.03	4.38
		0.60	1,089.59	318.92	1,232.58	342.62	7.43
		1.00	572.31	79.59	600	89.25	12.15
	700	0.20	11,549.52	6,324.02	13,220.05	6,590.13	4.21
		0.60	1,649.14	616.61	1,870.41	653.55	5.99
		1.00	799.53	176.89	900	193.52	9.40

of on the go. We also observe that the discount factor δ appears to be the most important factor in determining the gap in profits. This is consistent with the cost-based long-term technological value being proportional to $1/(\delta\Delta)$ while being linear in $v/(c+K)$.

7. Heterogeneous Markets

In this section, we extend our analysis to a setting with multiple customer classes, where different classes are characterized by different valuation levels for the product. We look first at the case of preannounced launches and then at the on-the-go case. We provide both algorithms for computing strategies for these two settings, as well as numerical insights about the role of effects of different market compositions and discount rates. For simplicity, we assume in this section that the technology evolves deterministically ($\sigma=0$). Without loss of generality, we assume that the technology drift μ is equal to 1, so the technology evolution reduces to $Z(t)=t$. Throughout this section, we denote as Θ the number of consumer classes and θ the class of a consumer, not the individual identity of a consumer. We let N_θ be the mass of consumers in class θ and let v_θ denote the valuation of class θ.

7.1. Preannounced Launches with Multiple Customer Classes

We start our analysis by focusing on the case of preannounced launches. We demonstrated in Theorem 2 that the long-term technological value of the firm's product is the key determinant of the optimal launch policy. The long-term technological value depends, however, on the customer valuation parameter v. When the market is composed of multiple customer classes, with some fraction having a high value for the firm's product and another fraction having a low value, should the firm follow a 2-cycle launch policy (as is optimal when selling to high value customers), a single introduction (as is in the optimal policy for selling to low value customers), or something altogether different? In a market where multiple consumer classes coexist, it turns out that the firm launches products on a regular basis and cannot lure the low type consumer to purchase early by committing to a single launch. Since products launched for one consumer class are available for all classes, the product launch problem invariably gets coupled and needs to be solved accounting for the full market composition. Because of this coupling, the heterogeneous market problem is hard to address in continuous time. As an approximation, we present here a finite horizon, discrete time version of the same problem.

7.1.1. Integer Programming Formulation. All consumers are still assumed to value the product in proportion to its technology level, but we now allow for different customer classes to have different parameters v_θ scaling the value they obtain from the product. The time horizon is finite and discrete, with time slots $t=0,1,\ldots,T$, where T is some large positive integer. We let the consumers and the firm have different discount factors: δ_c and δ_f, respectively. We introduce the terms $\gamma_c=e^{-\delta_c}$ and $\gamma_f=e^{-\delta_f}$, which are more natural representations of the discount factors in a discretized setting, where a smaller value of γ corresponds to a more impatient player.

We now argue that the equilibrium of the game played between the firm and the consumers can be computed via a single mixed-integer linear program (MIP). We first consider the second stage of the game. That is, we analyze the purchasing problem of a consumer of type θ for a given sequence of product launches. Let $x_\theta(t)$ be the indicator of the decision to purchase the product at time t. The first problem we formulate establishes a consumer's optimal purchasing policy $\{x_\theta(t)\colon t=0,1,\ldots,T\}$ for a given launch policy $\{y(t)\colon t=0,1,\ldots,T\}$, where $y(t)\in\{0,1\}$ represents whether a launch occurs at time t:

$$
\begin{aligned}
&\underset{x_\theta(\cdot),\,w_\theta(\cdot,\cdot)}{\text{maximize}}\ \sum_{t=0}^{T}\gamma_c^t\left(-px_\theta(t)+v_\theta\sum_{t'=1}^{t}t'w_\theta(t,t')\right)\\
&\text{subject to}\quad w_\theta(t,t')\le x_\theta(t')\le y(t')\\
&\qquad\qquad\qquad\text{for } t=0,\ldots,T \text{ and } t'=0,\ldots,t,\\
(\mathscr{P}_1)\quad &\qquad w_\theta(t,t'),\ x_\theta(t)\in\{0,1\}\\
&\qquad\qquad\qquad\text{for } t=0,\ldots,T, \text{ and } t'=0,\ldots,t,\\
&\qquad \sum_{t'=0}^{t}w_\theta(t,t')=1\quad\text{for } t=0,\ldots,T.
\end{aligned}
$$

In the formulation above, $w_\theta(t,t')$, with $t'\le t$, is an indicator of whether a consumer of type θ uses technology t' at time t. The first inequality in the first constraint ensures that a consumer of type θ who uses

RIGHTSLINK

technology t' at time t indeed purchased that technology at time t'. The second inequality in that constraint prohibits the consumer from purchasing in period t' if a launch did not occur in that period. Though a consumer could potentially purchase a product at a period other than its launch date, that is always a suboptimal consumer policy. For any time t' where the firm does not launch a product (i.e., $y(t')=0$), no consumer will own the respective technology level; that is, $x_\theta(t')=0$ and $w_\theta(t,t')=0$ for all $t\ge t'$. The last equality ensures that the consumer owns a technology product at any time t. To meet this constraint before the first purchase actually occurs, we impose the ad hoc condition $x_\theta(0)=y(0)=1$. This anchoring is without loss of optimality since it only shifts the consumers' and firm's utility functions by constants ($-p$ for the consumers and $p-c-K$ for the firm), which are added back when reporting results in the next subsection.

Problem ($\mathcal{P}_1$) can be seen as a single-dimensional, uncapacitated facility location with unimodal profits. To cast this as a facility location problem, assume there is a customer in each location from 1 up to T and assume that all periods where $y(t)=1$ correspond to the possibility of opening a facility at location t. The values of $x_\theta(t)$ then represent whether a facility is built on location t, and the values of $w_\theta(t,t')$ represent whether location t' serves a customer in location t. This problem has unimodal profit since the benefit obtained from utilizing facility t' to serve customer t is increasing in $w_\theta(t,t')$ for a fixed t (in this problem, a facility can only serve a customer if $t'\le t$). This problem is known to have no integrality gap (see Goemans and Skutella 2004 or Kolen and Tamir 1990), and thus we can instead consider the following relaxation without loss of optimality:

$$
(\mathcal{P}_2)\quad
\begin{aligned}
\underset{x_\theta(\cdot),\,w_\theta(\cdot,\cdot)}{\text{maximize}}\quad & \sum_{t=0}^{T}\gamma_c^t\Big(-px_\theta(t)+v_\theta\sum_{t'=1}^{t}t'w_\theta(t,t')\Big)\\
\text{subject to}\quad & w_\theta(t,t')\le x_\theta(t')\le y(t')\\
&\qquad \text{for } t=0,\dots,T \text{ and } t'=0,\dots,t,\\
& 0\le w_\theta(t,t'),\ x_\theta(t)\le 1\\
&\qquad \text{for } t=0,\dots,T \text{ and } t'=0,\dots,t,\\
& \sum_{t'=0}^{t}w_\theta(t,t')=1 \quad \text{for } t=0,\dots,T.
\end{aligned}
$$

We can simplify ($\mathcal{P}_2$) by removing unnecessary constraints. The constraints $0\le x_\theta(t')\le 1$ are redundant since $0\le w_\theta(t,t')\le x_\theta(t')\le y(t')\le 1$. The constraint $\sum_{t'=1}^{t}w_\theta(t,t')=1$ together with $w_\theta(t,t')\ge 0$ also implies that $w_\theta(t,t')\le 1$, making this constraint redundant as well. We now utilize the big-M method to move the constraint $x_\theta(t')\le y(t)$ to the objective function, obtaining

$$
(\mathcal{P}_3)\quad
\begin{aligned}
\underset{x_\theta(\cdot),\,w_\theta(\cdot,\cdot)}{\text{maximize}}\quad & \sum_{t=1}^{T}\gamma_c^t\Big((My(t)-p-M)x_\theta(t)\\
&\qquad +v_\theta\sum_{t'=1}^{t}t'w_\theta(t,t')\Big)\\
\text{subject to}\quad & w_\theta(t,t')-x_\theta(t')\le 0\\
&\qquad \text{for } t=0,\dots,T \text{ and } t'=0,\dots,t,\\
& w_\theta(t,t')\ge 0\\
&\qquad \text{for } t=0,\dots,T \text{ and } t'=0,\dots,t,\\
& \sum_{t'=1}^{t}w_\theta(t,t')=1 \quad \text{for } t=0,\dots,T.
\end{aligned}
$$

We now consider the dual of problem ($\mathcal{P}_3$). Let $d_\theta(t,t')$ and $c_\theta(t)$ be the dual variables associated with the inequality and equality constraints, respectively. The dual is given by

$$
(\mathcal{D})\quad
\begin{aligned}
\underset{c_\theta(\cdot),\,d_\theta(\cdot,\cdot)}{\text{minimize}}\quad & \sum_{t=1}^{T}c_\theta(t)\\
\text{subject to}\quad & c_\theta(t)+d_\theta(t,t')\ge\gamma_c^t v_\theta t'\\
&\qquad \text{for } t=0,\dots,T \text{ and } t'=0,\dots,t,\\
& d_\theta(t,t')\ge 0 \quad \text{for } t=0,\dots,T,\\
& \sum_{r=t}^{T}d_\theta(r,t)=\gamma_c^t(M+p-My(t))\\
&\qquad \text{for } t=0,\dots,T.
\end{aligned}
$$

We can now formulate the firm's profit maximization problem, accounting for the consumers' best responses, as a single MIP. The firm chooses a launch policy $\{y(t)\colon 0\le t\le T\}$, subject to a constraint that all consumer types respond optimally. We can enforce the optimality of the consumer response by including both the primal and dual variables and constraints into the firm's problem and by adding a new constraint that ensures that a given type θ's objective is identical in the primal and dual problems. The MIP that finds the optimal firm's policy is given by

$$
\begin{aligned}
\underset{y,\,x,\,w,\,c,\,d}{\text{maximize}}\quad & \sum_{t=0}^{T}\gamma_f^t\Big(-Ky(t)+(p-c)\sum_{\theta=1}^{\Theta}N_\theta x_\theta(t)\Big)\\
\text{subject to}\quad & y(t)\in\{0,1\} \quad \text{for } t=0,\dots,T,\\
& w_\theta(t,t')\ge 0 \quad \text{for } t=0,\dots,T \text{ and}\\
&\qquad t'=0,\dots,t,\\
& x_\theta(t)\le y(t) \quad \text{for } \theta=1,\dots,\Theta \text{ and}\\
&\qquad t=0,\dots,T,
\end{aligned}
$$

RIGHTSLINK

$$w_\theta(t,t') \le x_\theta(t') \quad \text{for } \theta = 1,\ldots,\Theta,\quad t=0,\ldots,T,\ t'=0,\ldots,t,$$

$$(\mathcal{F})\quad \sum_{t'=1}^{t} w_\theta(t,t') = 1 \quad \text{for } \theta=1,\ldots,\Theta \text{ and } t=0,\ldots,T,$$

$$\sum_{t=1}^{T} \gamma_c^t\left(-px_\theta(t) + v_\theta \sum_{t'=1}^{t} t' w_\theta(t,t')\right) = \sum_{t=1}^{T} c_\theta(t) \quad \text{for } \theta=1,\ldots,\Theta,$$

$$c_\theta(t) + d_\theta(t,t') \ge \gamma_c^t v_\theta t' \quad \text{for } \theta=1,\ldots,\Theta,\quad t=0,\ldots,T,\ t'=0,\ldots,t,$$

$$\sum_{r=t}^{T} d_\theta(r,t) = \gamma_c^t[M(1-y(t))+p] \quad \text{for } \theta=1,\ldots,\Theta \text{ and } t=0,\ldots,T,$$

$$d_\theta(t,t') \ge 0 \quad \text{for } \theta=1,\ldots,\Theta,\quad t=0,\ldots,T,\ t'=0,\ldots,t.$$

In summary, by using a duality argument, we can show that the seemingly difficult problem of designing an optimal launch policy under equilibrium constraints can be formulated as a single MIP. Once we solve this problem, to undo the anchoring assumption that $x_\theta(0)=y(0)=1$, we add $K-(p-c)\sum_{\theta=1}^{\Theta} N_\theta$ to the optimal objective function value in $(\mathcal{F})$.

7.1.2. Numerical Experiments. In this subsection, we utilize the MIP formulation above and solve the firm's utility maximization problem using IBM ILOG CPLEX Optimization Studio V12.4. We test several market effects on the profit performance of the firm and report our findings below.

Effect of multiple consumer classes: Our first experiment is aimed at understanding the effect that a multiple-class market has on the firm's optimal launch policy. Table 2 compares single-class markets, analogous to the ones studied in §5, to two-class markets (i.e., where $\Theta=2$). We assume a horizon of length $T=35$, which is long enough to avoid end effects given the high discount rate considered ($\gamma_f=\gamma_c=0.8$). The product variable cost is set at $c=0$.

Note that for the single-class markets, the optimal launching policy of this discrete time, finite horizon approximation mimics the characterization of the optimal policy for the continuous time, infinite horizon case. For instance, in the first row in Table 2, consumers in a homogeneous, low-value market alternate between cycles of lengths 2 and 4 (i.e., a minor launch is followed by a major launch). Under this policy, the firm can induce low value consumers to purchase as early as in period 2 by pushing the next launch to period 6. Similarly, in a homogeneous high-value market, consumers alternate cycles of lengths 1 and 3 (see the second row in Table 2), though now, since consumers' valuation is higher (and therefore, the long-term technology value $v/(p\Delta\delta_c)$ is higher), the seller is able to introduce products more frequently.

When both consumer classes coexist in the market, they become coupled, completely altering the product introduction and purchasing patterns, even if there is no fixed cost for launching a product. The third row in Table 2 reports the purchasing times of both classes when coexisting in the market. The effective product introduction times are the union of these two sets of purchase times, which is different from the union of the launch times of the two separate markets. Recall from §5 that the rationale behind the 2-cycle policy is to trigger earlier purchases from strategic consumers who would otherwise wait. However, when the firm faces such distinct segments, it has to introduce products less often compared with the union of both purchasing patterns in isolation (e.g., it does not introduce at $t=4$). Since the firm cannot target just one of the segments, it has to compromise the willingness to buy of both classes. For a group of customers with a given valuation, the addition of a second cohort with a different valuation can be beneficial, because it restricts the firm's ability to use an optimal launch policy directed at extracting more revenue from that particular customer class. In Table 2, the low value customers are thwarted from buying

Table 2 Comparison of the Firm's Launch Policy and Consumers' Purchasing Times in Homogeneous and Heterogeneous Markets

Firm's parameter	Consumers' parameters		Introduction and purchasing times		
K	Θ	v_θ	Type 1	Type 2	Firm's profit
0	1	2,000	2, 6, 8, 12, …, 32	—	12,229
	1	5,000	—	1, 4, 5, 8, …, 32, 33	20,477
	2	(2,000; 5,000)	Every three periods	2, 3, 6, 8, …, 32, 33	29,641
9,000	1	2,000	2, 6, 8, 12, …, 32	—	1,222
	1	5,000	—	1, 4, 5, 8, …, 32, 33	2,047
	2	(2,000; 5,000)	2, 6, 8, 12, …, 32	2, 6, 8, 12, …, 32	13,441

Notes. Discount rates are $\gamma_c=\gamma_f=0.8$, price $p=10{,}000$, unit cost $c=0$, $T=35$, and $N_\theta=1$. Type 1 consumers have valuation $v_1=2{,}000$, and type 2 have valuation $v_2=5{,}000$.

RIGHTSLINK

Table 3 Effect of Different Degrees of Consumers' Heterogeneity

Consumers' parameters			Introduction and purchasing times			
Θ	N_θ	$\mathbf{v} \triangleq (v_\theta)$	Type 1	Type 2	Type 3	Firm's profit
1	3	12,000	3, 7, ..., 31	Same as type 1	Same as type 1	245,451
3	1	(10,000; 12,000; 14,000)	3, 8, 14, ..., 31	3, 7, ..., 31	Same as type-2	231,246
3	1	(8,000; 12,000; 16,000)	6, 10, ..., 32	3, 7, ..., 21, 26, 32	3, 7, ..., 20, 23, 26, 32	201,699
3	1	(6,000; 12,000; 18,000)	Every six periods	3, 7, 11, 15, 18, 22, 25, 29, 34	Every three periods	203,552
3	1	(4,000; 12,000; 20,000)	7, 15, 23, 29	5, 10, 15, 18, 22, 26, 29	2, 5, ..., 26, 29, 33	194,672
3	1	(2,000; 12,000; 22,000)	13, 26	Every four periods	2, 5, ..., 32, 34	188,487

Note. Parameters are $T = 35$, $c = 0$, $K = 1{,}000$, $p = 90{,}000$, and $\gamma_f = \gamma_c = 0.8$.

every three periods when they are the only customers around (first row), but they are able to do so when they coexist with high value customers (third row). Overall, the total profit the firm obtains when both classes of customers are present together is 9% less than the sum of the profits it would obtain when facing two separate single-class markets.

The product launch cost also plays an important role in the firm's launch policy. For instance, when comparing the scenario without launch costs vis-à-vis the one where $K = 9{,}000$, we see that in the former, there are time epochs when just one consumer class makes a purchase (e.g., at $t = 2$ and $t = 8$, only the high value consumers buy when $K = 0$; see the third row in Table 2). Since the product launch cost is smaller than the revenue from any single consumer class, the firm may choose to introduce at time epochs in which only one consumer class will purchase the product. However, when $K = 9{,}000$ and $p = 10{,}000$, the profit of $1{,}000$ from selling to a single class of customers is far below the profit of $2 \times 10{,}000 - 9{,}000 = 11{,}000$ from selling to both classes, thus leading the firm to only introduce at times when both segments of the market buy together, as indicated by the purchase times of both classes in the last row in Table 2. In such situations, the firm launches products according to the optimal policy for selling to low value customers since the high value class who prefers more frequent introductions will always buy. With high fixed costs, the firm benefits from the presence of multiple consumer classes since the firm can pool classes to reduce the impact of the fixed launch costs.

Effect of consumer's heterogeneity: Table 3 illustrates the effect of different degrees of market heterogeneity on the consumers' purchasing behavior and on the optimal introduction policy. The consumer valuation $v = 12{,}000$ is set as the reference point (first row), and from there onward, different ranges of dispersion of the valuations are evaluated. The market size is the same for all cases, but its composition varies. When there are three different types present in the market, each type accounts for a third of the total market size.

Because the market composition of the consumers has wider spread in terms of valuations, it is difficult for the firm to target any single consumer class to extract surplus. Since the consumers with low valuation would like to purchase new versions of the product less often compared with the consumers with high valuation, it is very likely that, when there are high valuation consumers present in the market, the optimal launch policy for the firm is characterized by frequent launches. Low valuation consumers would free ride over high valuation ones, taking advantage of these frequent launch times to pick their purchasing times and skip some of the launches so as to retain much of their surplus. This is verified, for instance, when comparing the frequent purchases of low valuation type 1 consumers who purchase with cycles of lengths 3 and 4 in the first row in Table 3 but space their purchases as the market becomes more heterogeneous, with a cycle of length 13 in the last row. Hence, the firm is incapable of extracting as much surplus from the low valuation consumers, as in the case when it tailors the product launch dates targeting only such consumers. We can see this effect in the decreasing profits of the firm from Table 3 as the valuations of the different consumer classes get farther apart.[3]

Effect of discount factors: For a single-class market, we study the effect of varying the consumers' discount rate. The results are reported in Table 4, and they verify what was sketched in Figure 3. Here, the number of periods in the horizon is set at $T = 45$. As γ_c decreases (i.e., as δ_c increases, consumers become more impatient), we observe a shift in the launch policy regime from multiple launches for $\gamma_c = 0.8$ and $\gamma_c = 0.7$, to a single launch for $\gamma_c \leq 0.6$.[4] Since the

[3] Note that the values in the third and fourth rows are inverted with respect to a decreasing order of profits, which is due to end effects of the finite horizon formulation of the problem. Nevertheless, the difference between both values is only of the order of 1%.

[4] The consecutive introduction times toward the end of the horizon for $\gamma_c \leq 0.6$ are due to the deep discount rates that translate into extremely low long-term technology values for the consumers, contributing to the solution within the margin of error of the CPLEX solver.

RIGHTSLINK

Table 4 **Consumers' Patience and Launch Policy**

Consumers' discount factor γ_c	Introduction and purchasing times	Firm's profit
0.80	6, 15, 21, 30, 36	6,162
0.70	8, 16, 24, 32, 42	4,030
0.60	10, 41–45	2,154
0.56	11, 37–45	1,740
0.52	12, 37–45	1,396
0.48	13, 34–45	1,146

Note. Parameters are $T = 45$, $K = 30{,}000$, $p = 50{,}000$, $\gamma_f = 0.8$, $c = 0$, $\Theta = 1$, $N_\theta = 1$, and $v = 2{,}000$.

Table 5 **Effect of Firm's and Consumers' Patience**

Firm's discount factor γ_f	Consumers' discount factor γ_c	Introduction and purchasing times		Firm's profit
		Type 1	Type 2	
0.5	0.8	Every 4 periods	$1, 4, 5, 8, \ldots, 32, 34$	6,606
0.7	0.8	Every 3 periods	$2, 3, 6, 8, 9, 12, \ldots, 32, 33$	15,886
0.8	0.8	Every 3 periods	$2, 3, 6, 8, 9, 12, \ldots, 32, 33$	29,449
0.8	0.7	Every 3 periods	$2, 3, 6, 8, 9, 12, \ldots, 32, 33$	29,449
0.8	0.5	Every 4 periods	Every 2 periods, 33, 35	24,549

Note. Parameters are $K = 100$, $p = 10{,}000$, $\Theta = 2$, $N_1 = N_2 = 1$, $c = 0$, and $v = (2{,}000; 5{,}000)$.

consumers would buy less often to overcome the purchase cost and gain utility through use of the product, the firm would prefer to induce a single early purchase from the consumers compared with very sparse purchases later in the horizon.

In Table 5, in a heterogeneous market with two consumer classes, we vary the values of $\gamma_f = \exp(-\delta_f)$ and $\gamma_c = \exp(-\delta_c)$ (all else being equal) to illustrate the effect of the firm's and consumers' discount factors on the launching and purchasing policies. We assume a horizon of length $T = 35$.

For a given degree of consumers' patience γ_c, when the firm is more impatient than the consumers (i.e., $\gamma_f < \gamma_c$), its profits grow as the discount factor approaches γ_c. On the other hand, in the more realistic setting where consumers are more impatient than the firm, and moreover, when they become increasingly impatient (i.e., γ_c decreases), their utility from a given purchase decreases. This affects both consumer types in different ways: the low type 1 must use the product for a longer time to recover enough utility from it—this is observable in the longer interpurchasing cycles in the last row of Table 5—whereas the high type 2 will buy more regularly. Overall, the firm's profits go down as consumers become more impatient.

Value of market information: Table 6 shows the value of precise market information for the firm in terms of consumer valuations. Suppose that the firm believes that the consumers are homogeneous with a certain valuation, and it launches products assuming so, whereas in reality the market composition is different. In this case, the consumers' purchasing time epochs may be different from what the firm would expect. In Table 6, the total population mass is set at 2. Throughout, the firm assumes that the market is homogeneous with valuation $v = 12{,}000$. The "Consumers' parameters" column shows the true valuation mix defining the market composition, with $N_\theta = 1$ for the second to sixth data row.

We observe that misinformation regarding the consumer valuations can severely hurt the firm. To illustrate, from the last row, we observe that a broad spread between the firm's belief and the actual valuations lowers the firm's profits by 58%. That is, the firm would have gained an additional profit of 49,202 if it knew that the true valuation mix of the market is (2,000; 22,000) rather than misestimating consumers as being homogeneous at $v = 12{,}000$.

7.2. Launches On the Go with Multiple Customer Classes

We now turn our attention to on-the-go launch policies when there are consumers with different valuations in the market. Like we did for preannounced launches, we first propose an algorithm for finding equilibria and then proceed to provide some numerical insights.

Table 6 **Value of Market Information**

Consumers' parameters			Actual purchasing times		Firm's outcome		
Θ	N_θ	v_θ	Type 1	Type 2	Launches	Realized profit	Profit under complete info
1	2	12,000	$3, 7, \ldots, 31$	$3, 7, \ldots, 31$	$3, 7, \ldots, 31$	163,330	163,330
2	1	(10,000; 14,000)	7, 14, 21, 28, 31	$3, 7, \ldots, 31$	$3, 7, \ldots, 31$	105,126	152,960
2	1	(8,000; 16,000)	7, 14, 21, 28	$3, 7, \ldots, 31$	$3, 7, \ldots, 31$	105,037	140,676
2	1	(6,000; 18,000)	7, 14, 21, 28	$3, 7, \ldots, 31$	$3, 7, \ldots, 31$	105,037	141,132
2	1	(4,000; 20,000)	7, 14, 21, 28	$3, 7, \ldots, 31$	$3, 7, \ldots, 31$	105,037	150,827
2	1	(2,000; 22,000)	14, 28	$3, 7, \ldots, 31$	$3, 7, \ldots, 31$	85,333	134,535

Note. Parameters are $T = 35$, $c = 0$, $K = 1{,}000$, $p = 90{,}000$, and $\gamma_f = \gamma_c = 0.8$.

RIGHTSLINK

7.2.1. A Backward Induction Approach. In this subsection, we propose a method for computing equilibria for the on-the-go launch case. As we did for the case of preannounced launches, we replace our continuous time, infinite horizon model with a discrete time, finite horizon version.

Let Θ be the number of consumer classes, with N_θ being the mass of consumers with valuation v_θ for any $\theta \in \{1, 2, \ldots, \Theta\}$. The methodology we offer applies to any finite number of consumer classes, but the computational burden does grow exponentially in the number of classes. Our method involves solving for the subgame perfect equilibrium of a slightly different game than the one we studied before, where we replace each consumer class by a single representative consumer. Each consumer class θ is replaced by a representative consumer who makes purchases at a price pN_θ, at a cost to the firm of cN_θ. The equilibrium of this game between representative consumers is not necessarily the same as the equilibrium of the original game with infinitesimal consumers. However, it does provide a candidate equilibrium path. We can test whether any infinitesimal consumer would want to deviate from the equilibrium of this game. Note that the deviation is different here, because an infinitesimal consumer deviating does not change anyone else's utility except his own. Fortunately, this methodology worked in our numerical experiments, by generating an equilibrium for the game with infinitesimal consumers in over 90% of instances tried. We do not know whether symmetric equilibria—where all consumers belonging to the same class make identical decisions—exist in the instances where this method fails.

The main advantage of replacing the consumer classes by representative consumers is that we reduce the state space of the problem to a finite set. For every period $t \in \{0, 1, \ldots, T\}$, the set of possible states are represented by the technology available in the market $w(t)$ and the technologies owned by the different consumers $C^\theta(t)$. We solve this problem by the standard technique for finite horizon dynamic games: backward induction. We expand the set of periods to twice their original horizon, to account for the fact that in each period the firm first chooses whether to release a new product, with consumers subsequently deciding whether or not to buy. Call the periods where the firm acts t^f and the periods where the consumers act t^c. We start from the last period, T^c, and consider for every possible state whether each particular representative consumer buys the product at that state. Once the values of those states are calculated, we take a step back to period T^f. For each possible state at time T^f, we calculate whether the firm launches given the continuation value of that state. We then naturally proceed to state $(T-1)^c$, where the consumers play a bimatrix game to decide whether to buy the product at each possible state. We repeat these steps inductively all the way back to time 0^f, where the firm will obviously not launch a product with zero technology. By following the states chosen by the firm and the consumers starting from time 0^f, we find the equilibrium path of the game.

7.2.2. Numerical Experiments. We now present a few numerical results for the multiple-consumer class market when the firm announces products on the go. We note the effects found here mostly replicate the case of preannouncing.

Effect of multiple consumer classes: As with preannouncing, having multiple consumer classes can lead to coupling of purchasing behavior, since a product release for one class is also available for the other. However, our numerical experiments show that coupling is less likely here than in the preannounced case. See that with $K = 0$, coupling does not occur in Table 7. In the on-the-go case, coupling typically occurs at relatively higher values of launch cost.

Effect of consumer heterogeneity: When the market is composed of only one type of consumer, the firm can easily optimize its launch cycle to that consumer group. As Table 8 shows, diversity in valuations makes this kind of consumer targeting difficult and, thus, reduces the firm's profit.

Table 7 Comparison of the Firm's Launch Policy and Consumers' Purchasing Times in Homogeneous and Heterogeneous Markets

Firm's parameter	Consumers' parameters			Introduction and purchasing times		
K	Θ	N_θ	v_θ	Type 1	Type 2	Firm's profit
0	1	1	2,000	3, 6, 9, ..., 27, 32, 35	—	10,483
	1	1	5,000	—	Every two periods, 35	17,776
	2	1	(2,000; 5,000)	Every three periods, 35	Every two periods, 35	28,259
9,000	1	1	2,000	3, 6, 9, ..., 27, 32, 35	—	1,048
	1	1	5,000	—	Every two periods, 35	1,777
	2	1	(2,000; 5,000)	3, 6, 9, ..., 27, 31, 35	3, 6, 9, ..., 27, 29, 31, 33, 35	11,534

Notes. Discount rates are $\gamma_c = \gamma_f = 0.8$, price $p = 10{,}000$, unit cost $c = 0$, and $T = 35$. Type 1 consumers have valuation $v_1 = 2{,}000$, and type 2 have valuation $v_2 = 5{,}000$.

RIGHTSLINK

Table 8 Effect of Different Degrees of Consumers' Heterogeneity

Consumers' parameters			Introduction and purchasing times		Firm's profit
Θ	N_θ	$\mathbf{v} \triangleq (v_\theta)$	Type 1	Type 2	
1	2	12,000	4, 8, 12, …, 35	Same as type 1	108,240
2	1	(10,000; 14,000)	4, 8, 12, …, 35	Same as type 1	108,240
2	1	(8,000; 16,000)	Every five periods	3, 6, 9, …, 31, 35	116,850
2	1	(6,000; 18,000)	6, 12, 18, …, 35	3, 6, 9, …, 35	108,170
2	1	(4,000; 20,000)	8, 16, 24, 34	3, 6, 9, …, 35	95,064
2	1	(2,000; 22,000)	12, 24, 35	3, 6, 9, …, 35	85,650

Notes. Parameters are $T = 35$, $c = 0$, $K = 4{,}000$, $p = 80{,}000$, and $\gamma_f = \gamma_c = 0.8$. The market size is $N_1 = N_2 = 1$ for all cases, with different compositions.

Table 9 Consumers' Patience and Launch Policy

Consumers' discount factor γ_c	Introduction and purchasing times	Firm's profit
0.70	Every 6 periods	7,104.6
0.60	Every 7 periods	5,306.9
0.50	Every 8 periods	4,031.3
0.40	Every 9 periods	3,100.6
0.30	Every 10 periods	2,405.5

Note. Parameters are $T = 45$, $K = 30{,}000$, $p = 50{,}000$, $\gamma_f = 0.8$, $c = 0$, $\Theta = 1$, $N_1 = 1$, and $v = 4{,}000$.

Effect of consumer patience: For a single-class market, reduced consumer discount factor reduces the frequency at which they purchase and, therefore, lowers the firm's profits, as Table 9 shows.

8. Conclusion

In this paper, we study how a technology firm should introduce successive generations of a product over time to maximize the net present value of its cash flow given that its customers are also making purchasing decisions to maximize their own discounted surpluses. We show that when the firm releases products on the go, the technology increments are constant in equilibrium. We also show that the firm can increase its profits between 4% and 12% by committing in advance to the next product generation's technology level after two consecutive, relatively close introductions. The optimal preannounced launch strategy involves alternating minor and major technology improvement cycles when the consumers' benefit of owning the product (measured by its long-term technology value) is high. This strategy enables the firm to anticipate some consumer purchases by promising those customers that no new products will be released with a small technology increment, thus reducing their opportunity value of waiting for the next product release.

Both the preannounced and on-the-go launch policies can be seen as extreme cases of the practices executed by technology firms in the real world. Microsoft, Intel, and Apple sometimes announce expected technology levels for their major product releases far in advance, but they may not do so for the successive generations. Concurrently, and resuming the Apple story described in §1, consumers build expectations based on the past history of releases and are not surprised by new iPhones coming up on the market in a sequence of major followed by minor improvements (3G to 3GS, 4 to 4S, 5 to 5S, 6 to 6S), which is not captured by the Markovian foundation of our on-the-go setting either. Indeed, one could argue that the time lag between the announcement and the major release in our preannounced setting can be seen as a proxy for the implicit commitment of not releasing soon after a minor launch. All in all, preannouncement as one possible commitment device is important after releases with relatively little technology improvement in order to convince consumers to purchase regardless. It is particularly valuable when either the expected rate or the variance of technology progress are high, since consumers may expect product launches to occur relatively often in these two cases.

The framework introduced in this paper opens several new avenues for potential future work. One interesting direction is empirical: estimating long-term technology values of real products and matching these values to their launch policies would allow us to understand whether firms are currently successful in managing strategic customers. Extending our model to a competitive setting or to a setting with finite production capacity would also be important topics for future research.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/mnsc.2015.2189.

Acknowledgments

The authors are grateful to Yossi Aviv, the associate editor, and three anonymous referees for their constructive feedback during the review process. I. Lobel thanks the National Science Foundation for its support [Grant CCF-1216004], and J. Zhang thanks the National Natural Science Foundation of China for its support [Grant 71331004].

References

Adner R, Levinthal D (2001) Demand heterogeneity and technology evolution: Implications for product and process innovation. *Management Sci.* 47(5):611–628.

Aviv Y, Pazgal A (2008) Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing Service Oper. Management* 10(3):339–359.

Aviv Y, Levin Y, Nediak M (2009) Counteracting strategic consumer behavior in dynamic pricing systems. Netessine S, Tang C, eds. *Consumer-Driven Demand and Operations Management Models*, International Series in Operations Research and Management Science, Vol. 131 (Springer, New York), 323–352.

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

RIGHTSLINK

Balcer Y, Lippman SA (1984) Technological expectations and adoption of improved technology. *J. Econom. Theory* 34(2):292–318.

Bass FM (1969) A new product growth model for consumer durables. *Management Sci.* 15(5):215–227.

Bayus BL (1992) The dynamic pricing of next generation consumer durables. *Marketing Sci.* 11(3):251–265.

Besbes O, Lobel I (2015) Intertemporal price discrimination: Structure and computation of optimal policies. *Management Sci.* 61(1):92–110.

Borgs C, Candogan O, Jennifer C, Lobel I, Nazerzadeh H (2014) Optimal multiperiod pricing with service guarantees. *Management Sci.* 60(7):1792–1811.

Casadesus-Masanell R, Yoffie DB (2007) Wintel: Cooperation and conflict. *Management Sci.* 53(4):584–598.

Cohen MA, Eliashberg J, Ho T-H (1996) New product development: The performance and time-to-market tradeoff. *Management Sci.* 42(2):173–186.

Correa J, Montoya R, Thraves C (2011) Contingent preannounced pricing policies with strategic consumers. Working paper, University of Chile, Santiago.

Elmaghraby W, Gülcü A, Keskinocak P (2008) Designing optimal preannounced markdowns in the presence of rational customers with multiunit demands. *Manufacturing Service Oper. Management* 10(1):126–148.

Farzin Y, Huisman K, Kort P (1998) Optimal timing of technology adoption. *J. Econom. Dynam. Control* 22(5):779–799.

Goemans MX, Skutella M (2004) Cooperative facility location games. *J. Algorithms* 50(2):194–214.

Greer D, Ruhe G (2004) Software release planning: An evolutionary and iterative approach. *Inform. Software Tech.* 46(4):243–253.

Hendricks KB, Singhal VR (1997) Delays in new product introductions and the market value of the firm: The consequences of being late to the market. *Management Sci.* 43(4):422–436.

Karatzas I, Shreve S (1991) *Brownian Motion and Stochastic Calculus*, 2nd ed. (Springer, New York).

Klastorin T, Tsai W (2004) New product introduction: Timing, design, and pricing. *Manufacturing Service Oper. Management* 6(4):302–320.

Kolen A, Tamir A (1990) *Discrete Location Theory* (John Wiley & Sons, New York).

Kouvelis P, Mukhopadhyay S (1999) Modeling the design quality competition for durable products. *IIE Transactions* 31(9): 865–880.

Krankel RM, Duenyas I, Kapuscinski R (2006) Timing successive product introductions with demand diffusion and stochastic technology improvement. *Manufacturing Service Oper. Management* 8(2):119–135.

Kumar S, Swaminathan JM (2003) Diffusion of innovations under supply constraints. *Oper. Res.* 51(6):866–879.

Lai G, Debo LG, Sycara K (2010) Buy now and match later: Impact of posterior price matching on profit with strategic consumers. *Manufacturing Service Oper. Management* 12(1):33–55.

Lilien GL, Yoon E (1990) The timing of competitive market entry: An exploratory study of new industrial products. *Management Sci.* 36(5):568–585.

Liu H, Özer Ö (2009) Managing a product family under stochastic technological changes. *Internat. J. Production Econom.* 122(2): 567–580.

Mahajan V, Muller E (1996) Timing, diffusion, and substitution of successive generations of technological innovations: The IBM mainframe case. *Tech. Forecasting Soc. Change* 51(2): 109–132.

Moorthy KS (1988) Product and price competition in a duopoly. *Marketing Sci.* 7(2):141–168.

Norton J, Bass F (1987) A diffusion theory model of adoption and substitution for successive generations of high-technology products. *Management Sci.* 33(9):1069–1086.

Norton J, Bass F (1992) Evolution of technological generations: The law of capture. *Sloan Management Rev.* 33(2):66–77.

Osadchiy N, Vulcano G (2010) Selling with binding reservations in the presence of strategic consumers. *Management Sci.* 56(12):2173–2190.

Pae J, Lehmann D (2003) Multigeneration innovation diffusion: The impact of intergeneration time. *J. Acad. Marketing Sci.* 31(1): 36–45.

Paulson Gjerde KA, Slotnick SA, Sobel MJ (2002) New product innovation with multiple features and technology constraints. *Management Sci.* 48(10):1268–1284.

Plambeck E, Wang Q (2009) Effects of e-waste regulation on new product introduction. *Management Sci.* 55(3):333–347.

Ruhe G, Saliu MO (2005) The art and science of software release planning. *IEEE Software* 22(6):47–53.

Satariano A (2011) Apple falls after wait for iPhone causes missed estimates. *Bloomberg News* (October 19), http://www.bloomberg.com/news/articles/2011-10-18/apple-s-results-miss-estimates-as-iphone-sales-fall-short-shares-decline.

Stremersch S, Muller E, Peres R (2010) Does new product growth accelerate across technology generations? *Marketing Lett.* 21(2):103–120.

Su X, Zhang F (2008) Strategic customer behavior, commitment, and supply chain performance. *Management Sci.* 54(10): 1759–1773.

Wells C (2012) iPhone 5 release date hints dropped by Apple CEO Tim Cook. *New York Daily News* (July 26), http://www.nydailynews.com/news/money/iphone-5-release-date-hints-dropped-apple-ceo-tim-cook-article-1.1122429.

Wilson LO, Norton JA (1989) Optimal entry timing for a product line extension. *Marketing Sci.* 8(1):1–17.

RIGHTSLINK