



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

The Operational Advantages of Threshold Discounting Offers

Simone Marinesi, Karan Girotra, Serguei Netessine

To cite this article:

Simone Marinesi, Karan Girotra, Serguei Netessine (2017) The Operational Advantages of Threshold Discounting Offers. Management Science

Published online in Articles in Advance 05 Jun 2017

<https://doi.org/10.1287/mnsc.2017.2740>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2017, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

The Operational Advantages of Threshold Discounting Offers

Simone Marinesi,^a Karan Girotra,^b Serguei Netessine^c

^a The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104; ^b Technology and Operations Management, INSEAD, 77305 Fontainebleau, Cedex, France; ^c Technology and Operations Management, INSEAD, Singapore 138676

Contact: marinesi@wharton.upenn.edu (SM); karan@girotra.com (KG); serguei.netessine@insead.edu (SN)

Received: July 11, 2013

Accepted: September 22, 2016

Published Online in Articles in Advance:
June 5, 2017

<https://doi.org/10.1287/mnsc.2017.2740>

Copyright: © 2017 INFORMS

Abstract. We study threshold discounting, or the practice of offering a discounted-price service if at least a prespecified number of customers signal interest in it, as pioneered by Groupon. We model a capacity-constrained firm, a random-sized population of strategic customers, a desirable hot period, and a less desirable slow period. Compared to a more traditional approach (slow period discounting or closure), threshold discounting has two operational advantages. First, the contingent discount temporally balances demand when the market for the service is large, and reduces supply of the service (preserving higher margins) when the market is small, allowing the firm to respond to the service's unobserved market potential. Second, activation of the threshold discount signals the market state and the consequent service availability to strategic customers, inducing them into self-selecting the consumption period to one that improves the firm's capacity utilization. Yet, threshold discounting can be harmful in situations with chronically low demand. In contrast with past work on strategic customers, their presence is advantageous to firms in our context. A calibrated numerical study shows that threshold discounting improves firm profits over a traditional approach by as much as 33% (7% on average).

History: Accepted by Yossi Aviv, operations management.

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/mnsc.2017.2740>.

Keywords: Groupon • strategic customers • group buying • revenue management • business model innovation

1. Introduction

Many firms operate in environments with highly variable demand while capacity is fixed in the short term. This setting has particularly acute consequences for service firms, as spare capacity in low-demand periods typically cannot be used to serve customers in high-demand periods. Prominent industries that struggle with this problem include the movie theater industry (\$10.2 billion of revenues in 2016), the restaurant industry (\$783 billion of revenues in 2016),¹ and a wide variety of retail services, such as beauty salons, bowling clubs, museums, opera houses, etc.

Over the last decade, the rise of online customer engagement technologies gave service providers in these industries new tools to interact with customers, most notably online discounted deals. The most famous online deal website is Groupon: founded in late 2008 and the first of an innovative breed of online deal firms, it went public in 2011, raising \$700 million to become the largest IPO by a U.S. Internet company after Google (see Barr and Baldwin 2011, Ovide 2011).

Groupon's value proposition since its foundation has been to help service providers attract customers during off hours and better use their capacity. To lure customers into off hours, Groupon has relied on deep discounts coupled with the use of an innovative discount structure, in which the discounted deals were

valid only if a certain number of customers showed interest in the offering. The benefits of such deals, henceforth referred to as *threshold discounting* offers, have been celebrated in the business press as a way to leverage "network effects" and economies of scale (see, e.g., Mourdoukoutas 2011), and they have more recently received attention from the academic community as well.

The picture that emerges around threshold discounting offers is incomplete and controversial. Incomplete, because such offers have so far been studied only from a marketing perspective—ignoring, for example, capacity constraints—and no analysis on their operational implications has been attempted. Controversial, because it is difficult to reconcile the celebrated advantages of threshold discounting offers with their progressive discontinuation by many merchants. Hence, two important questions remain open. When do threshold discounting offers provide value and when can they be harmful? What are the phenomena that drive the value (or lack thereof) of such offers?

Our work provides potential answers to both questions. We consider a capacity-constrained firm that offers its services to a random-sized population of strategically acting customers who prefer to be served on a desirable "hot" time period over a less desirable "slow" time period (e.g., a movie theater on Saturday

versus Monday evening); demand is thus variable but substitutable between the two periods. The traditional approach used by firms in this context is to either close shop on the slow period, or open on the slow period at a discounted price. Our analysis isolates the operational advantages of threshold discounting from the demand expansion effects of discounts. We find that threshold discounting can often outperform the traditional approach on account of two key operational phenomena, both of which arise from information communication.

First, threshold discounting allows communication *from the customers to the firm*: setting an activation threshold on the number of subscriptions required to activate the deal endows the firm with a built-in, demand-responsive mechanism (which we refer to as *responsive duality*) that matches different market states with appropriate discounting/opening decisions, resulting in higher profits thanks to its ability to better match supply and demand, while at the same time maintaining efficiency in terms of margins and fixed costs. Second, the announcement of an active deal also communicates information *from the firm to the customers*, who use it to their own advantage. In particular, this induces a *strategic scarcity effect* that increases customers' responsiveness to slow-period discounts by exploiting their strategic behavior, inducing them into self-selecting their consumption period to one that (typically) better serves the firm's interests of managing capacity and margins. Notably, threshold discounting often leads to a superior firm performance compared to traditional approaches even in the absence of additional beneficial effects including word-of-mouth, market segmentation, or economies of scale: the only benefits of threshold discounting that have been studied so far.

Our analysis on the operational advantages of threshold discounting also provides a plausible explanation for why threshold discounting offers may not work in every setting. We show that firms with chronically low demand, which arguably constitute a big proportion of the service providers featured in daily deals websites, may see their profit decrease when employing threshold discounting offers. This finding suggests that threshold discounting offers may have been discontinued because of a lack of fit between the (operational) advantages that they provide and the needs of those (low-demand) service firms that were attracted by daily deal websites.²

We conclude with a numerical study, calibrated in part on data from a service provider of the nature considered in our analysis. In the over 800 parameter scenarios examined, we find that threshold discounting increases firm profits by as much as 33% (average 7%) over the profit earned with the traditional approach. We also run a specific analysis on the communication

role played by the activation threshold and its demand-balancing property: we refer the interested reader to Section 6 for all results of our numerical study.

Our work makes several contributions. This is the first study to examine the operational performance of threshold discounting. We find that threshold discounts do provide value from an operational standpoint on account of two novel effects—the contingent nature of the discount that reduces supply demand mismatches, and the ability of these discounts to use strategic customer behavior to their advantage. We provide a detailed analysis of when and why such effects create value, and we provide prescriptions on when threshold discounting destroys value, which allows us to provide a plausible explanation on why threshold discounts have been discontinued.

In contrast with most strategic customer behavior literature in operations, we show that in our setting strategic customer behavior is often beneficial. Our numerical analysis quantifies the profit increase from offering threshold discounts (up to 33%) and finds that these benefits are highest in situations characterized by high market uncertainty, strategically acting customers, and intermediate levels of fixed costs and demand seasonality. Overall, our paper argues that threshold discounting is an overlooked method for managing capacity that, if used correctly, can substantially improve a firm's performance.

2. Literature Review

Our work is related to three streams of literature: group buying and quantity discounts, strategic customers, and demand manipulation via pricing.

2.1. Group Buying and Quantity Discounts

In the early 2000s, several group-buying websites (e.g., Mercata.com, LetsBuyIt.com) were founded with the objective of aggregating the buying power of customers to obtain quantity discounts (see Anand and Aron 2003 and the references therein). While these group-buying websites bear some resemblance to daily deal websites like Groupon, the latter's growth has been fueled by a very different (and much more popular) business model, based on offering discounts on retail services rather than on inventories of physical products (the difference is of substance, as we discuss in Section 5) and on employing threshold discounts within specific time windows, which our paper aims to study.

Closer to our work are recent papers that look specifically at threshold discounts. Jing and Xie (2011) show that in a threshold discounting offer, informed players act as sales representatives with their friends in an attempt to reach the threshold and earn the discount, to the firm's benefit. Chen and Zhang (2015) show analytically that, under some conditions, threshold discounts

are the optimal mechanism to price discriminate a population of customers. Wu et al. (2014) and Li and Wu (2014) empirically study the dynamic evolution of customers' subscriptions over time, with the former studying threshold-driven effects, and the latter looking at herding and word-of-mouth effects. Hu et al. (2013) models the impact of pledge timing, sequential or simultaneous, on customers' pledge behavior and firm profit. The success of Groupon has also spurred works that have focused on issues other than threshold discounts (see Edelman et al. 2016 and references therein). While similarly motivated, none of these papers take an *operational* perspective on threshold discounting schemes, consider intertemporal demand substitution, or provide explanations for why major players have discontinued threshold discounting offers.

2.2. Strategic Customers

In recent years, many authors have studied the implications of strategic customer behavior in operations management (OM) settings. For instance, Su (2007) analytically studies pricing policies when customers have different degrees of patience; Aviv and Pazgal (2008) evaluate the benefit of preannounced and contingent pricing strategies; Su and Zhang (2008) study the effectiveness of quantity and price commitment strategies; and Cachon and Swinney (2009) evaluate the benefit of quick response (see Netessine and Tang 2009 for more references). More recent topics include conspicuous consumption (Tereyağoğlu and Veeraraghavan 2012), product variety (Parlaktürk 2012), online click tracking (Huang and Van Mieghem 2013), social comparisons (Roels and Su 2013), social learning (Papanastasiou and Savva 2017), and the informative power of queue length in Veeraraghavan and Debo (2011).

Like many of the above papers, our customers time their purchases accounting for the strategic behavior of the other players. Unlike the above papers, however, we explore the consequences of such strategic behavior in a novel setting, namely, a firm that employs threshold discounting offers in a context characterized by seasonal demand over substitute consumption periods. The implications of strategic behavior in our context are unexpected and in contrast with the main findings from this large literature, which generally finds that strategic customer behavior is harmful to the firm.

2.3. Demand Manipulation via Pricing

A large body of literature has studied how a capacity- or inventory-constrained firm can use different pricing strategies to better match supply with demand. All the literature on revenue management, for instance, focuses on this topic (see Talluri and Van Ryzin 2005 for a survey), including all papers on peak load pricing (see Crew et al. 1995 for a survey). Similar in spirit

to our work are perhaps Lus and Muriel (2009), who study pricing (and technology choices) when dealing with substitutable products, Boyacı and Özer (2010), who investigate the benefit of using a capacity management technique (advanced selling) together with pricing, and Kong (2016), who studies the impact of time-restricted discounts on service congestion, demand cannibalization, and firm profit in a two-period queueing model. Our paper departs from the existing literature in that we study a way to reduce the supply demand mismatch through a novel pricing approach: namely, we study the use of optimal threshold discounting offers in the presence of strategic customers.

3. The Model

3.1. Preliminaries

We consider a capacity-constrained service provider that offers a service over two periods—a hot period h and a slow period s —to a market comprised of infinitesimal customers. The market size, X , is a random variable with probability density function g , cumulative distribution function G , and support $[\underline{x}, +\infty)$, $\underline{x} \geq 0$, where a strictly positive \underline{x} corresponds to situations in which there is surely a minimum interest in the firm's service. Customers are heterogeneous in their valuation for the two periods; specifically, customer i 's valuations, v_h^i and v_s^i , are drawn from a continuous bivariate distribution with probability density function f over the support $[\underline{v}_h, \bar{v}_h] \times [\underline{v}_s, \bar{v}_s]$, $\underline{v}_h \geq \underline{v}_s$ and $\bar{v}_h \geq \bar{v}_s$. The hot period is more popular than the slow period—if the two periods were priced the same, more customers would prefer to be serviced in the hot period than in the slow period.³ Customers derive no value from consuming the service more than once.

The service is offered in the hot period at an exogenously set price r_h , with $r_h < \bar{v}_h$, and in the less preferred slow period at a chosen discounted price, $(1 - \theta)r_h$, where $\theta \geq 0$ is the discount offered.⁴ The service provider can choose whether to stay open in each service period. If open for service, the provider incurs fixed costs $c_F \geq 0$ in wages, utilities, etc. and can serve at most k customers. We assume that fixed costs are not so high as to preclude profit in the best-case scenario ($kr_h > c_F$) and that capacity is not always binding ($\underline{x} < k$). We normalize the marginal cost of serving a customer to zero. When demand outstrips capacity, the provider rations capacity randomly among customers. We place no restrictions on which period precedes the other.

The population of customers consists of a fraction γ of strategic customers, and a fraction $1 - \gamma$ of nonstrategic customers, $\gamma \in [0, 1]$. Strategic customers take into account the strategies of other customers while making their decisions—in particular, this allows them to form rational expectations on the probability of being

served in each service period. Nonstrategic customers on the other hand simply choose the service period in which their valuation exceeds price the most, effectively ignoring availability differences driven by other customers' choices.

The setup described above is a stylized representation of the consumption of many services such as movie theaters, opera houses, museums, etc. These services share the key characteristics of our setup—demand uncertainty, desirable and less desirable service periods, single consumption, seasonal opening/pricing, and a per-period capacity that is fixed in the short run. We first examine the traditional approach typically employed by firms in similar circumstances—a choice between closing down or discounting in the slow period. Next, we compare this traditional approach with threshold discounting offers as popularized by online deal sites such as Groupon, LetsGroop, BigDeal, etc.

3.2. The Traditional Approach: Seasonal Closure or Regular Discounting

Traditionally, service providers faced with seasonal demand patterns either shut down in slow periods or remain open but offer a discounted price. For example, in many cities of mainland Europe where fixed costs of operation are high, restaurants and museums are typically closed on Mondays.⁵ On the other hand, in London, service providers often stay open on Mondays, but offer discounts and promotions to attract customers.⁶

Our benchmark model, *the traditional approach*, captures these two (mutually exclusive) schemes. The decisions and sequence of events are provided in Figure 1. First, nature draws the market size X , whose realization x is not observed by the customers or by the firm, and each customer in the market observes her private hot and slow period valuations, and whether she is strategic or not. Then, the service provider decides whether to offer the service in the slow period and, provided the service is offered, what discount θ to offer. Finally, customers respond with their choice of visit timing and are served based on available capacity and demand. The formal equilibrium solution is provided in Section A.1 in the online appendix, which contains the proofs of all results in the paper. The service

provider's decision of whether to open in the slow period is driven by a comparison of the profits from closing in the slow period (seasonal closure) and the profits from opening and offering a discount (regular discounting).

3.2.1. Seasonal Closure. Under seasonal closure, strategic and nonstrategic customers with a hot period valuation higher than r_h visit during the hot period, and the service provider serves them up to capacity (Figure 2, panel (a)). The profit earned by the firm when the market size realization is x is given by $\pi_c(x) = r_h \min(k, \alpha_h^c x) - c_F$, where $\alpha_h^c = \int_{r_h}^{\bar{v}} \int_0^{\bar{v}} f(v_h, v_s) dv_s dv_h$ is the fraction of the market with valuation for the hot period higher than r_h , and $\alpha_h^c x$ is the hot period demand under closure. The expected profit is then

$$\Pi_c = \int_x^{+\infty} \pi_c(x) dG(x). \quad (1)$$

3.2.2. Regular Discounting. Alternatively, the service provider may decide to also open in the slow period, albeit at a lower price $r_h(1 - \theta)$, and incur additional fixed costs c_F . All customers, strategic or not, whose valuations are below the prices in the corresponding period, i.e., with $v_h^i < r_h$ and $v_s^i < r_h(1 - \theta)$, do not visit the firm. Similarly, all customers who can make a positive surplus only in one period, i.e., either $v_h^i > r_h$ and $v_s^i \leq r_h(1 - \theta)$, or vice versa, visit in that period. The remaining customers visit in either the hot or the slow period, whichever maximizes their surplus. More specifically, nonstrategic customers visit in the slow period iff $v_s^i + \theta r_h > v_h^i$, while strategic customers visit in the slow period iff $v_s^i > \hat{v}_d(v_h^i | \theta)$. Here $\hat{v}_d(v_h^i | \theta)$ is the slow period valuation that makes a strategic customer with hot period valuation v_h^i indifferent between visiting in either period, once her rational expectation on the probability of being served in each period is taken into account (see Section A.1 in the online appendix for its formal characterization).

Overall, for a given discount θ , the market can be divided accordingly into a fraction of slow period visitors $\alpha_s^d(\theta)$, a fraction of hot period visitors $\alpha_h^d(\theta)$, and a fraction of nonvisitors $\alpha_\emptyset^d(\theta)$, each obtained as the weighted average of their counterparts for the strategic ($\sigma_m^d(\theta)$) and nonstrategic ($\mu_m^d(\theta)$) portions of the

Figure 1. Timeline for the Traditional Approach

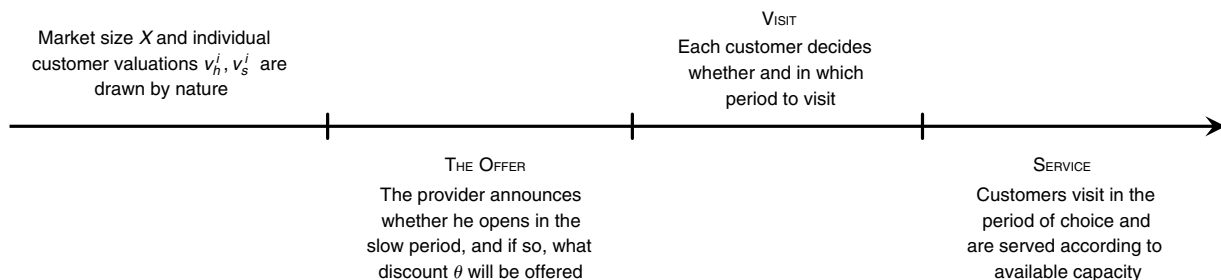
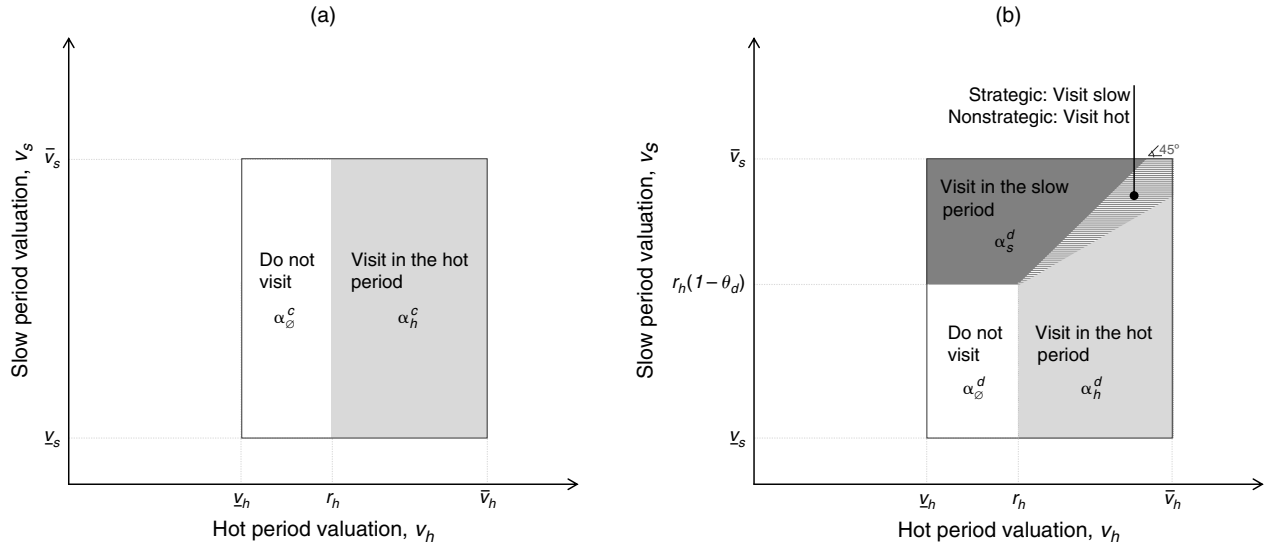


Figure 2. Customer Equilibrium Visit Strategies as a Function of Their Valuation Vector Under (a) Seasonal Closure and (b) Regular Discounting



population, i.e., $\alpha_m^d(\theta) = \gamma\sigma_m^d(\theta) + (1 - \gamma)\mu_m^d(\theta)$, $m \in \{s, h, \emptyset\}$. Panel (b) in Figure 2 depicts these customer segments. Analogously, $\alpha_h^d(\theta)x$ and $\alpha_s^d(\theta)x$ represent the hot and slow period demand under regular discounting when the discount is θ .

The profit earned by the firm when the discount is θ and the market size realization is x is given by

$$\pi_d(x | \theta) = r_h \min(k, \alpha_h^d(\theta)x) + r_h(1 - \theta) \min(k, \alpha_s^d(\theta)x) - 2c_F.$$

The equilibrium discount θ_d is the one that solves the firm's profit maximization problem:

$$\Pi_d = \max_{\theta} \int_x^{+\infty} \pi_d(x | \theta) dG(x), \quad \text{s.t. } \theta \geq 0. \quad (2)$$

Throughout the paper, we omit the firm's action(s) from the argument of a function when referring to the equilibrium outcome; for example, we simply use α_m^d and $\pi_d(x)$, $m \in \{s, h, \emptyset\}$ to mean $\alpha_m^d(\theta_d)$ and $\pi_d(x | \theta_d)$.

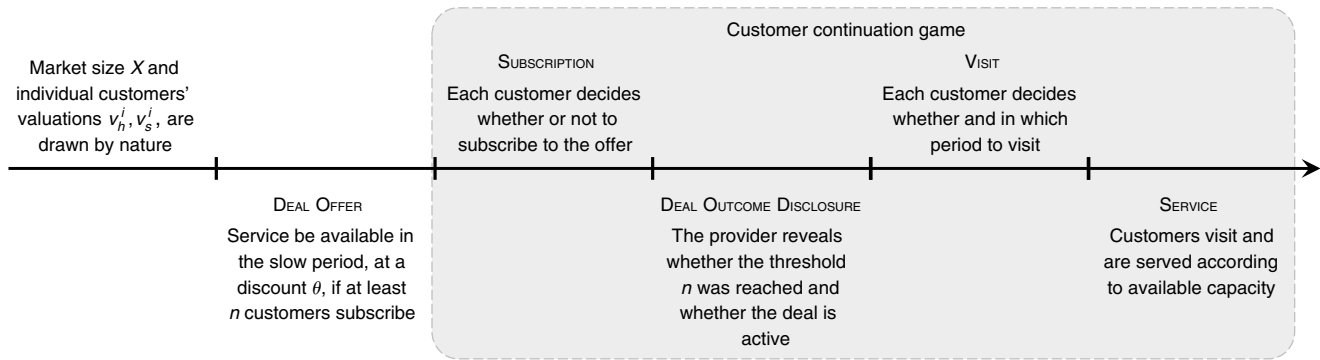
The optimal discount under regular discounting is driven by a trade-off among three effects. First, a higher discount reduces the margins earned in the slow period. Second, a higher discount expands demand in the slow period as more customers can now afford the service. Third, a higher discount shifts demand across periods: as long as the discount is not excessive, this typically rebalances demand, increasing capacity utilization.⁷ Mathematically, these three effects correspond to the respective three terms in the marginal profit Equation (3):

$$\frac{d}{d\theta} \Pi_d(\theta) = - \underbrace{\left(\int_x^{+\infty} r_h \min(k, \alpha_s^d(\theta)x) dG(x) \right)}_{\Pi'_{d-s}(\theta)}$$

$$+ \underbrace{\left(\int_x^{k \cdot \alpha_h^d(\theta)^{-1}} r_h(1 - \theta) \left(-\frac{d\alpha_h^d(\theta)}{d\theta} \right) x dG(x) \right)}_{\Pi'_{d-h}(\theta)} + \underbrace{\left(\int_x^{+\infty} \frac{d\alpha_h^d(\theta)}{d\theta} r_h(1_{\{x < k(\alpha_h^d(\theta))^{-1}\}} - 1_{\{x < k(\alpha_s^d(\theta))^{-1}\}}(1 - \theta))x dG(x) \right)}_{\Pi'_{d-sh}(\theta)}. \quad (3)$$

We hereafter make two assumptions. First, the discount always shifts demand from the hot to the slow period, i.e., $d\alpha_h^d(\theta)/d\theta < 0$. This assumption rules out unrealistic irregular valuation distributions in which, following a small increase in the discount, a large fraction of the market can suddenly afford the slow period, thus reducing availability and driving more strategic customers to the hot period. Second, in equilibrium, the firm does not charge so high a discount to make the slow period busier than the hot period, i.e., $\alpha_h^d \geq \alpha_s^d$. This assumption is merely expositional—our results do not make use of this assumption. This assumption ensures that seasonality does not get overturned once the discount decision is made, and indeed the hot period is still the hot period. In our extensive numerical study (Section 6) based on a wide range of plausible parameters, we find that the first assumption holds in all scenarios studied, for every discount level $\theta \in [0, 1]$, and the second assumption also holds in all scenarios examined.

The formal comparison between regular discounting and closure is provided in Lemma 1 in the online appendix. Essentially, regular discounting is a better choice than closure when fixed costs are lower, slow period valuation of customers is higher, and a large

Figure 3. Timeline of Threshold Discounting

market size is more likely, since the cost/benefit ratio of opening at a discount in the slow period becomes more favorable. A firm resorting to the *traditional approach* chooses between regular discounting or closure to maximize its profit, $\Pi_a = \max(\Pi_c, \Pi_d)$.

3.3. Threshold Discounting

Threshold discounting allows customers to avail themselves of the service in the slow period at a discounted price if and only if enough other customers show interest in doing the same. Figure 3 illustrates the sequence of events for a threshold discounting scheme. As with the traditional approach, nature draws the market size X , whose realization x is not observed by either customers or by the firm, and customer i in the market observes her private valuation vector (v_h^i, v_s^i) . Next, the service provider announces the threshold discounting deal: the service will be offered in the slow period at a discount $\theta > 0$ to all customers who subscribe to the offer, but only if at least n of them sign up for it; if less than n customers sign up, the service provider will instead close during the slow period. Each customer then decides whether to subscribe to the deal or not. We assume that customers incur a positive but arbitrarily small cost to subscribe, i.e., they subscribe if and only if their expected gain is strictly positive.⁸ The firm then communicates whether the deal's activation threshold was reached (the deal is active) or not, and closes/discounts as per the preannounced threshold discounting deal. Customers then choose a period to visit and consume the service according to the available capacity.


We conservatively assume that the use of an activation threshold does not lead to beneficial marketing effects, such as fostering word of mouth or increasing customer valuations (see Jing and Xie 2011 for an analysis of such cases): we do so in order to evaluate threshold discounting offers based on their *operational* properties. We now proceed to examine the equilibrium outcome.

*Customer strategy for a given deal (θ, n) .*⁹ Once the firm announces the discount θ and the threshold n ,

customers respond depending on which of four segments their valuations belong to (Figure 4, segments separated by dashed lines): (i) Customers in the first segment, with $v_h^i \leq r_h$ and $v_s^i \leq r_h(1 - \theta)$, are priced out of the market, so they do not subscribe and do not visit. (ii) Customers in the second segment, with $v_h^i \leq r_h$ and $v_s^i > r_h(1 - \theta)$, make a positive surplus only in the discounted slow period, hence they subscribe to the deal, visit during the slow period if the deal is active, and they do not visit otherwise. (iii) Customers in the third segment, with $v_h^i > r_h$ and $v_s^i \leq r_h(1 - \theta)$, make a positive surplus only in the hot period, so they do not subscribe and visit in the hot period regardless of the outcome of the deal; strategic and nonstrategic customers in the preceding three clusters do not differ in their strategies. (iv) Customers in the remaining segment can make a positive surplus both in the hot and in the discounted slow period, so they choose the period that yields the highest surplus. Specifically, those among them who are nonstrategic, subscribe and visit in the slow period if and only if $v_s^i + \theta r_h > v_h^i$ and the deal is active, and they visit in the hot period otherwise (see panel (a)). More interestingly, strategic customers account for their beliefs on service availability in each period, hence they subscribe and visit in the slow period if and only if their slow period valuation is higher than $\hat{v}_t(v_h^i | \theta, n)$ and the deal is active, visiting in the hot period otherwise (see panel (b)). The threshold $\hat{v}_t(v_h^i | \theta, n)$ represents the slow period valuation that makes a strategic customer with hot period valuation v_h^i indifferent between visiting in the two periods (see Section A.2 in the online appendix).

The market can consequently be divided into a fraction of subscribers $\alpha_s^t(\theta, n)$, nonsubscribers $\alpha_h^t(\theta, n)$, and priced out customers $\alpha_\emptyset^t(\theta, n)$, by combining the corresponding fraction of strategic ($\sigma_m^t(\theta, n)$) and nonstrategic ($\mu_m^t(\theta)$) customers in the population, i.e., $\alpha_m^t(\theta, n) = \gamma \sigma_m^t(\theta, n) + (1 - \gamma) \mu_m^t(\theta)$, $m \in \{s, h, \emptyset\}$ (respectively, the dark gray, pale gray, and white areas in Figure 4).

Figure 4. Customers' Equilibrium Strategies Under Threshold Discounting as a Function of Their Hot and Slow Period Valuations, for (a) Nonstrategic and (b) Strategic Customers



Firm strategy. Given the deal (θ, n) , the profit of the firm conditional on the market size being x is given by

$$\pi_t(x | \theta, n) = \begin{cases} r_h \min(k, \alpha_h^c x) - c_F & \text{if } x < n / \alpha_s^t(\theta, n), \\ r_h \min(k, \alpha_h^t(\theta, n)x) \\ \quad + r_h(1 - \theta) \min(k, \alpha_s^t(\theta, n)x) - 2c_F & \text{if } x \geq n / \alpha_s^t(\theta, n), \end{cases}$$

while the equilibrium deal (θ_t, n_t) is the solution to the firm's profit maximization problem

$$\Pi_t = \max_{\theta, n} \int_x^{+\infty} \pi_t(x | \theta, n) dG(x), \quad (4)$$

s.t. $\theta > 0$ and $G(n / \alpha_s^t(\theta, n)) > 0$.

The constraints in the firm's maximization problem ensure a strictly positive probability that the deal may or may not be active. In other words, these constraints preserve uncertainty in the deal outcome, without which threshold discounting would be equivalent to the traditional approach. This allows us to discern whether using an activation threshold delivers value to the firm, which is at the core of our research question.

4. Threshold Discounting v/s Traditional Approach

Theorem 1. *Threshold discounting outperforms the traditional approach, i.e., $\Pi_t > \Pi_a$, if conditions (5) and (6) hold:*

$$r_h(1 - \theta_d) \alpha_s^d x < r_h(\alpha_h^c - \alpha_h^d) x + c_F, \quad (5)$$

$$\Pi'_{d-sh} > 0, \quad (6)$$

where α_m^j is the fraction of the demand that visits in period $m \in \{h, s\}$ under scheme $j \in \{c, d\}$, Π'_{d-sh} is the demand balancing effect of discounting in equilibrium, defined in Equation (3), and θ_d is the equilibrium discount under regular discounting and is given by Equation (2).

The advantages of threshold discounting over the traditional approach arise from its most characteristic feature, the activation threshold, and in particular how this threshold allows the firm to best use its capacity to serve a market that is both uncertain in its size and heterogeneous in its preferences for the slow and hot period. To understand the role of multiple effects of the threshold, we build the comparison at the heart of Theorem 1 by considering special cases of our setup along two dimensions: customer strategicity, and the demand expansion effect of discounts. Specifically, we start with the case that has no strategic customers and where the price is such that discounting only shifts, but does not expand, demand for the service, progressively building up to the general case that includes both these features.

4.1. All Customers Are Nonstrategic, $\gamma = 0$

4.1.1. Without Demand Expansion via Discounting.

Consider the case in which the hot period service yields a positive surplus to all customers, $v_h \geq r_h$, so that all customers derive value from visiting the firm. In this setting, the discount offered by the firm does not affect what fraction of the market the firm captures and demand is simply the market size. Note, however, that the size of the market is still uncertain and not all demand may be served because of capacity constraints. The traditional purpose of setting

Downloaded from informs.org by [128.91.109.54] on 21 February 2018, at 09:12. For personal use only, all rights reserved.

discounts—increasing demand by reducing margins—is not relevant in this case and the sole advantage of any discounting, regular or threshold, is to increase sales by best serving demand with the existing capacity. In other words, the purpose of discounting is now “purely operational,” that is, shifting demand across periods to increase capacity utilization. We will next demonstrate how threshold discounting is better than regular discounting in achieving this goal.

It should be noted that, if the hot period price were endogenized and in the absence of other direct or indirect pricing constraints (e.g., regulation, competition, ...) the firm would never be better off pricing lower than the lowest customer valuation. The discussion of such a case, while potentially unrealistic, is valuable because it is instrumental in decoupling the consequences of threshold discounting on the *demand-shift* effect of discounts from those on the *demand-expansion* effect of discount. We henceforth use the notation $\pi_j^m(x|\cdot)$ and $\Pi_j^m(\cdot)$ to indicate, respectively, the conditional (on market size) and expected profit under scheme $m \in \{c, d, a, t\}$ when customers are nonstrategic.

Corollary 1. *When all customers are nonstrategic ($\gamma = 0$) and discounting does not expand demand ($v_h \geq r_h$), threshold discounting always outperforms the traditional approach, i.e., $\Pi_t^m > \Pi_a^m$.*


In contrast with the general case of Theorem 1, this special case that focuses on the operational effects of discounting shows that threshold discounting *always* dominates the traditional approach.

If the firm decides to close in the slow period, the firm can only serve up to k customers, but it earns a full margin r_h on each of them, and incurs limited fixed costs. This leads to a higher profit than discounting when the market size is low, because it allows the firm to be efficient (fixed costs incurred only in one period) and get the most from each customer served (no discount). As the market size increases, however, the firm cannot serve all customers that visit in the hot period because of the limited capacity available. By opening

in the slow period and discounting it, the firm can instead attract customers in both periods, albeit serving some at a lower margin and incurring additional fixed costs. When the market is large enough, the extra sales accrued with regular discounting make it more profitable than closure. Formally, there exists a critical market size x^+ such that the conditional profit is higher for closure if the market is smaller than this size $x < x^+$, and the profit is higher for discounting if the market is bigger $x > x^+$ (Figure 5, panel (a)).

Firms employing the traditional approach must choose between closure or discounting ex ante without knowing the market size, and pick the scheme that maximizes profit in expectation over all possible market realizations. This choice may turn out to be wrong in retrospect, once the market size is revealed. Firms employing threshold discounting, on the other hand, can overcome this limitation by always making the “right” choice: in threshold discounting, the firm discounts the slow period only when a certain threshold number of customers subscribe, and closes it down otherwise. Our analysis reveals that once the terms of the deal (θ, n) are announced, the number of customers subscribing is an increasing function of the market size x . Moreover, a firm can intelligently choose the threshold n such that the deal is active when the market size is higher than a level of choice.¹⁰ Effectively, pairing the discount θ_d with an appropriate threshold n^+ ensures that threshold discounting becomes regular discounting when the market size is adequately high to warrant the same, and becomes closure when the market size is not high enough—an infallible traditional approach, one that always makes the right choice. Figure 5, panel (b) illustrates the conditional profit that captures this ability: this is simply the upper envelope of the conditional profits from closure and discounting. Note that when the firm announces the threshold discounting deal, the information available on the market size is the prior density function g , so the firm has *no informational advantage* compared to employing a traditional approach. Note also that, by announcing the

Figure 5. Profit of the Firm as a Function of the Realized Market Size x Under (a) Regular Discounting and Closure, in Equilibrium; (b) Threshold Discounting, When Employing the Best Between Regular Discounting and Closure in Every Market State; and (c) Threshold Discounting, in Equilibrium



deal, the firm gives up all decision power, in the sense that all future actions are merely an application of the terms of the deal (open at a discount if the threshold is met, close otherwise). Therefore, the advantage of threshold discounting is not about making better decisions once more information (about the market size) is available. Nonetheless, the choice of an appropriate activation threshold results in the same profit as if the firm could choose whether to discount or close once the market size is observed.

The ability of threshold discounting to deploy the best between closure and regular discounting makes it more profitable than the traditional approach. Threshold discounting can, however, go even further. In general, the optimal *level* of the discount to offer differs depending on which market states the discount is activated in. So, while regular discounting must choose the discount that optimally trades off margins and sales across all market states, the threshold discount level can be set such that it makes said trade-off only for the subset of cases where the market is high enough for the discount to be active. The latter is a more targeted discount and earns the firm higher profits. Essentially, given the fact that in threshold discounting the discount is active only when needed, the discount level itself (and consequently the threshold) can be tailored to best exploit this knowledge and target the offer to the market regimes (higher and lower market states) in which the firm, respectively, discounts and closes—leading to an even higher profit. This discount adjustment effect is illustrated in Figure 5, panel (c), which shows the equilibrium conditional profit in threshold discounting that leads to an even greater expected profit than the upper envelope of closure and regular discounting.

Together, we refer to the ability of threshold discounting to mimic discounting or closure as needed and further to target the discount as *responsive duality*. Both effects universally favor threshold discounting, boosting its profit compared to the traditional approach when discounting does not expand demand and in the absence of strategic customers. Under such a setting, responsive duality summarizes the distinctive advantages of threshold discounting over the traditional approach.

4.1.2. With Demand Expansion via Discounting. Next, we consider the case in which some customers are priced out of the hot period, $v_h < r_h$. In this setting, discounting—regular or threshold—can expand the demand for the product by capturing some of these customers who would never visit at the full price. This additional effect attenuates the above described advantages of threshold discounting.

Corollary 2. *When all customers are nonstrategic ($\gamma = 0$) threshold discounting outperforms the traditional approach, $\Pi_i^t > \Pi_a^t$, if condition (5) holds.*

When discounting expands slow period demand, the fact that closure is preferred when the market size is small and discounting is preferred otherwise may no longer hold, compromising the responsive duality advantage of threshold discounting. Specifically, it is now possible that for some (limited) parameter values, regular discounting, on account of its ability to expand demand, earns a higher profit than closure, even for the worst market scenario, becoming preferred to closure in *all* market states. In the absence of any trade-off between closure and discounting, threshold discounting is no longer guaranteed to outperform the traditional approach.

Condition (5) simply excludes those parameter values where the demand expansion effect of discounting is so strong that even for the lowest possible market size it is more profitable to remain open on both days and discount heavily to expand the demand, rather than preserve margins and reduce fixed costs. Mathematically, the condition states that in the lowest possible market scenario, the gain from demand expansion is not large enough to offset the combined effect of the cannibalization of demand from the hot period to the slow period—on account of discounting—plus the increased fixed costs. This is a very mild condition—as long as there is a possibility that the market can be small enough, there are cases where one would like to close and threshold discounting will dominate the traditional approach.

Before we proceed, note that the analysis of the nonstrategic case contains insights that make it valuable beyond its role of benchmark for the case of a market with strategic customers. First, note that the trade-off between discounting and closure arises out of lower margins *and* higher fixed costs in the former. As such, the trade-off can exist even in the absence of fixed costs, so that *economies of scale are not required* for threshold discounting to deliver value to the firm.

Second, it can be shown that when demand expansion is not excessive, i.e., $\alpha_s^d(\theta)(1 - \theta) < \alpha_h^c - \alpha_h^d(\theta) \forall \theta$, the activation threshold condenses all the required information for choosing optimally between opening at a discount and closing. This implies that no other activation rule can improve over a threshold-activation rule, that is, a firm that could decide to activate the deal if the market size were to belong to a set of choice $\mathbb{A} \subset [x, +\infty)$ cannot earn a higher profit than under a threshold discounting offer. This is because of the fact that (i) the optimal choice between opening at a discount and closing is going to be a threshold choice even once the market state is observed, and (ii) the critical market level at which discounting becomes a better choice than closing is known in advance by the firm (it is only a function of the discount offered). The interested reader can find the formal result in Proposition 1 in the online appendix. This observation highlights how a threshold discounting offer is a simple

yet very effective way for a capacity-constrained firm to serve a market that is both uncertain in its size and heterogeneous in its preferences for the slow and hot periods.

4.2. Some Customers Are Strategic, $\gamma > 0$

In making their subscription and visit decisions, strategic customers take into account how other customers make these decisions by assessing service availability in the two periods. This is particularly relevant in the context of threshold discounting offers—the fact that the “deal is on” signals the market size, not only to the firm (which leads to the advantage discussed in the previous subsection), but also to the customers. Customers can now use this information to build their assessment of availability in the two periods, and this influences their choice of period in which to visit. In addition, since customers faced with threshold discounting offers anticipate that such information might be available in the future, they take it into account also when deciding whether to subscribe to the deal.

Taken together, these observations suggest that on account of such strategic customer behavior, threshold discounting is no longer equivalent to regular discounting when the deal is active. Our next results indicate (i) how customers’ response to discounts differs between threshold discounting and regular discounting, and (ii) how this difference impacts the firm profits. Understanding the two will help us highlight a second potential source of advantage for threshold discounting, and in the meantime understand the reason for condition (6) in Theorem 1.

Theorem 2 (Strategic Scarcity). *In the presence of strategic customers, $\gamma \in (0, 1]$,*

(1) *a given discount θ leads to the same total sales in regular discounting and threshold discounting, $\alpha_h^t(\theta, n) + \alpha_s^t(\theta, n) = \alpha_h^d(\theta) + \alpha_s^d(\theta) \forall n$;*

(2) *a given discount θ leads to more temporally balanced demand in threshold discounting, $\alpha_h^t(\theta, n) - \alpha_s^t(\theta, n) < \alpha_h^d(\theta) - \alpha_s^d(\theta) \forall n$;*

(3) *the difference in demand between the hot and the slow periods in threshold discounting decreases in the activation threshold n ; specifically, for any θ , $(\partial/\partial n)(\alpha_h^t(\theta, n) - \alpha_s^t(\theta, n))$ is strictly negative for $n < k$ and equals zero afterward.*

The theorem compares threshold discounting and regular discounting when the same discount is offered under both schemes. First, note that total demand is the same in the two schemes (item 1), because the portion of the market that visits the firm comprises of customers who are not priced out in both periods, hence it depends only on the discount θ .

The second item of the theorem states that while total demand is the same, demand is more balanced (or the difference between the hot and the slow period’s

demand is lower) in threshold discounting than in regular discounting. This finding builds on the fact that a threshold discounting offer, by announcing that the deal is active, signals to customers that the market is high enough to trigger the deal, or in other words, it informs customers that the market is higher than what their prior information suggested.¹¹ As a direct consequence of customers’ updated beliefs on the market size, the ratio of service availability between the hot and the slow periods—henceforth *relative service availability*, defined as hot/slow—decreases (this is explained below) making the slow period relatively more attractive to customers, in essence giving the slow period an *availability* advantage over the hot period. As a result, more customers visit in the slow period instead of visiting in the hot period, and item 2 follows.

To see why signaling a higher market to customers reduces the relative service availability, consider how this changes under different market sizes. When the market size is low and capacity is not binding in either period, service availability is the same (and equal to 1) in both periods—the relative availability is also 1. As the market size increases, hot period capacity becomes binding (availability in hot < 1) while there is still spare capacity in the low period (availability in slow = 1)—the relative availability now decreases. Once the market size is so high that capacity is tight in both periods, the relative availability is at its lowest, and is equal to the ratio of slow and hot period demand. It now follows that the relative service availability decreases when customers place more weight on higher market realizations, as they do upon learning that the deal is active.¹²

The third item of the theorem states that the higher relative availability effect just discussed is more pronounced when the activation threshold is higher because a higher market size is required to trigger the deal, and strategic customers revise their prior belief on the market size distribution upward even further. Conversely, there is nearly no difference in balancing between regular and threshold discounting when the activation threshold is set so low that the deal is nearly always active.

Finally, note that an active deal not only makes customers more willing to visit in the slow period, but also makes them more willing to subscribe to the offer, because subscriptions grant discounts in the slow period only when the deal is active—that is, when this is more attractive on account of the above property. We refer to strategic customers’ response to a threshold discounting offer and its consequences on the demand pattern of the firm described in Theorem 2 as the *strategic scarcity effect*.

Having established that threshold discounting is more effective than regular discounting at balancing because of strategic customer behavior, we now discuss

its impact on profit. At first, this balancing of demand may appear to be a good thing, since, after all, one of the purposes of discounting is to balance demand by shifting it from the hot to the slow period. Further, this balancing is essentially “free,” i.e., threshold discounting achieves better balancing not by offering higher discount but simply by communicating to customers that the deal is active. “Too balanced” a demand, however, may not be in the interests of the firm. More specifically, it can be shown that every unit of demand shifted from the hot period to the slow period has diminishing returns on profit that eventually become negative when demand is perfectly balanced (see Lemma 5 in the online appendix).

To see why, first note that demand balancing improves profit in certain contingencies, but reduces profit in others. Specifically, it improves profit only when there is *excess demand* in the hot period and *excess capacity* in the slow period, simply because it feeds the slow period with excess demand from the hot period, thereby increasing revenues. On the other hand, demand balancing reduces profit when there is excess capacity in both periods, because it feeds the slow period with demand that could be met in the hot period, thereby reducing revenues; and it has no effect when both periods are fully utilized. Note also that the first case is more likely to arise, and the second is less likely to arise, when demand is more *unevenly* distributed across periods. Hence, as demand becomes more and more even, any further shift of demand to the slow period is less likely to benefit the firm and more likely to harm it. In the limit, when demand is almost perfectly balanced, the odds of beneficial contingencies approach zero and demand balancing invariably hurts profit.

To summarize, we have established two important aspects of how strategic customer behavior influences the efficacy of threshold discounting. First, it makes threshold discounting more effective than regular discounting by shifting demand from the hot period to the slow period (on account of strategic scarcity). Second, the benefits of this shift diminish—and eventually hurt the firm—as demand becomes more evenly balanced. With these facts in mind, we now consider how threshold discounting compares to the traditional approach with and without demand expansion, starting from the latter.

4.2.1. No Demand Expansion via Discounting. As discussed before, when $v_h \geq r_h$, the discount does not allow the firm to capture a larger share of the market. The discount is now solely a way for the firm to give away margins to achieve more balanced demand, which leads to higher sales. On account of higher discount effectiveness, a threshold discounting firm can achieve the same level of demand balancing as a firm

using regular discounting, while offering a lower discount. This is always beneficial for the firm as it leads to same sales and higher margins, hence to a higher profit. The corollary now follows:

Corollary 3. *When discounting does not expand demand ($v_h \geq r_h$), threshold discounting always outperforms the traditional approach, i.e., $\Pi_t > \Pi_a$.*

Note that this “higher effectiveness of discounts” advantage of threshold discounting is over and above the responsive duality advantage discussed in the non-strategic case. Threshold discounting then allows the firm to put in place a smart combination of closure in low market states and a demand balancing-enhanced discounting in high market states, which increases profit even further.

4.2.2. With Demand Expansion via Discounting. Consider now the case in which discounting, in addition to shifting demand, also expands demand on account of capturing a larger share of the market, i.e., $v_h < r_h$. In this case, the discount is a way for the firm to give away margins not just to get more balanced demand (as before) but also to expand demand. While the discount in threshold discounting is still better at achieving the first objective, it may now complicate achieving the second.

A firm that utilizes threshold discounting can, as before, use the higher effectiveness of the discount to obtain as much demand balancing while offering a lower discount (as compared to regular discounting), but this higher effectiveness may hurt it. Now the firm will realize the same demand balancing as regular discounting and lower discounts (both good), but the firm will also realize lower demand expansion than it would have with regular discounting. This lower demand expansion may overcome all the other advantages of threshold discounting and lead to lower profits as compared to regular discounting.

Alternatively, a threshold discounting firm can use the higher effectiveness of the discount to obtain better demand balancing for the same discount (as compared to regular discounting), but this may also hurt it. If the higher discount effectiveness balances demand too much, the firm’s profit may again get reduced, because of demand shift having diminishing returns on the profit of the firm, as discussed above. Taken together, these two scenarios point out that the higher effectiveness of the discount may hurt the firm.

A potential remedy for the higher effectiveness lies in the second design instrument, which can be used to mitigate some of the potential downsides of higher discount effectiveness—the activation threshold. Recall that the firm can lower the activation threshold to reduce the effectiveness of threshold discounts as much as needed (item 3 of Theorem 2), thus reducing

the potential downsides associated with it. Yet, the activation threshold already serves an important purpose, controlling when threshold discounting is akin to regular discounting or to closure, and setting it too low may compromise the firm's ability to exploit the responsive duality advantage, i.e., it may trigger the deal for too low a market state, when closing would have been preferred instead. When a regular discounter would benefit from more demand balancing on the margin (condition (6) holds), threshold discounting can always set the same discount and an appropriately low threshold that simultaneously ensures (i) that the additional demand balancing (relative to regular discounting) benefits the firm, and (ii) that the responsive duality advantage is not excessively compromised.

To summarize, when there is the possibility of demand expansion, threshold discounting dominates regular discounting if demand expansion via discounting is not excessive (condition (5)), and regular discounting would benefit from further demand balancing (condition (6)). Together, these conditions say that threshold discounting delivers value when demand balancing is a primary driver both for the firm's discounting decision and for its profitability. In some situations, however, threshold discounting may actually hurt the firm: the next section identifies situations in which this will be the case.

4.3. Chronically Low Demand

The above analysis has considered a general market size distribution G with support $[x, +\infty)$. However, in some cases, a firm may face market prospects that are consistently low, and capacity is always in excess, i.e., $\bar{x} \leq k$ with $\bar{x} \triangleq \sup(\text{supp}(X))$. A newly opened business, for example, in the presence of rigid capacity (e.g., square footage) typically chooses its capacity to a level that is optimal in the long run, but may be too high in the short run, when the business has not yet built a reputation. A mature business, on the other hand, may face a decrease in demand because of an increase in competition, or because of a change in macroeconomic conditions (economic recession, change in consumers' taste) that may last for some time. In such cases, as we are going to show, threshold discounting may actually reduce the profit of the firm relative to the traditional approach.

Theorem 3. *For a firm whose capacity always exceeds demand, threshold discounting strictly reduces profit compared to the traditional approach if and only if, between closure and regular discounting, one dominates the other. Formally, when $\bar{x} \leq k$, then $\Pi_t < \Pi_a$ if and only if*

$$\nexists x^+ \in (x, \bar{x}): r_h \alpha_h^c x^+ = r_h (\alpha_h^d + (1 - \theta_d) \alpha_s^d) x^+ - c_F. \quad (7)$$

When capacity always exceeds demand, all advantages of threshold discounting are moot, and inefficiencies arise. To see why, first note that when capacity

always exceeds demand, availability is assured in both periods, hence strategic customers behave like non-strategic customers. The only potential advantage of threshold discounting lies in responsive duality (4.1), that is (i) exploiting the ex post trade-off between regular discounting and closure, and (ii) tailoring the discount and the threshold to further exploit the two market regimes (higher and lower market states) in which the firm, respectively, discounts and closes.

Consider the second advantage: when capacity is never binding, the firm cannot possibly gain anything by pricing differently than regular discounting. In fact, the reason why it may be optimal for the firm to charge different discounts in relation to different market states is due to the existence of capacity constraints that may or may not be binding depending on the size of the market. When capacity is never binding, the same discount is optimal in every market state.¹³ If, in addition, either discounting or closure are optimal in all possible market states (as would be the case when condition 7 holds) any approach that mixes the two—as threshold discounting does—is necessarily going to reduce the profit of the firm.

As noted above, situations with a chronically low demand—as the ones captured in Theorem 3—are likely to hold for young businesses and struggling businesses alike. These categories probably represent a large portion of the firms featured on daily deals websites. Hence, a plausible reason for why threshold discounting offers have been discontinued by many players in the industry could have been the lack of fit between the (operational) benefit associated with threshold discounting offers and the needs of those firms (with chronically low demand) that seek to be featured on daily deal websites. Nonetheless, whether discontinuing threshold discounting was a savvy long-term decision remains questionable, as long-term value is less likely to come from demand-starved businesses, and is more likely to come from healthier businesses with conspicuous seasonal demand—the type of businesses that would benefit the most from the operational advantages of threshold discounting offers.¹⁴

5. Effect of Strategic Customers on the Profit of the Firm

In this section, we want to investigate the impact of customers strategic behavior on the profit of the firm. Most of the existing literature on strategic customers (Su and Zhang 2008; Liu and van Ryzin 2008; Cachon and Swinney 2009, 2011) has either proven that strategic customers reduce firm profits, or has taken it as a given and developed countermeasures to reduce their negative effect.¹⁵ The typical setting often invoked is the one in which an apparel retailer sells a finite inventory over a finite season, and may resort to price markdowns at the end of the season to dispose of leftover inventory. By anticipating price markdowns, strategic

customers can decide to postpone their purchases until the end of the season, thus reducing profits for the firm. Our model shares many characteristics with this typical setting. In Cachon and Swinney (2009), for example, strategic customers can decide to purchase in two different periods—during the season, when their valuation for the product is higher, or at the end of the season, when their valuation is lower—which map exactly into the hot and slow periods in our framework. As in our paper, in Cachon and Swinney (2009) the firm offers a reduced price in the period that customers value the least. Finally, as in our paper, strategic customers take into account the actions of other customers and act to maximize their expected surplus. Despite these similarities, the effects of strategic customers in our setting are different from those in the classic settings studied in the literature.

Theorem 4. *Impact of strategic customers on profit:*

- When $r_h < v_h$ (no demand expansion via discounts), the profits of a threshold discounting firm always increase with more strategic customers, i.e., $(d/d\gamma)\Pi_t > 0$.
- When $r_h \geq v_h$ (with demand expansion via discounts), if the antiderivative of $\Pi'_{d-mg}(\theta) + \Pi'_{d-ex}(\theta)$ is unimodal in θ and condition (6) holds, then $(d/d\gamma)\Pi_t^\mu > 0$.

Note that an increase in γ (proportion of strategic customers) increases the amount of the population that accounts for capacity availability, thus making the strategic scarcity effect more relevant. Recall that this effect is unambiguously good for the firm in the absence of demand expansion via discounts, or when condition (6) holds. In this sense, the finding that strategic customers improve the profit of a firm employing threshold discounting is in line with our discussion above around increased effectiveness of threshold discounting enabled by the strategic customers. More

interesting is to reflect on what is different in our setting relative to the traditional OM setting that leads to this drastic departure from the well studied, harmful consequences of strategic customers in the existing literature.¹⁶

To begin with, observe that in our case the firm’s initial commitment to a price reduction has nothing to do with the use of price commitment strategies as a countermeasure to strategic customers, as, for example, studied in Su and Zhang (2008). In their setting, the firm commits to high enough prices at the end of the season to induce strategic customers to purchase in season, i.e., in the “hot” period. In our setting, the firm announces price reductions to achieve the opposite effect, i.e., redirect customers from the hot period into the slow period. This is because the setting considered in the existing literature modeled durable products, specifically physical units of inventory, while our setting captures perishable units of service capacity. For a firm selling inventory of a physical product, a customer who purchases a unit during the low season rather than during the high season is always harmful, because it reduces the profit earned on that same unit of inventory. Capacity, on the other hand, is not transferable across time—excess capacity in one period cannot be carried over to the next period—so it is possible that the firm runs out of capacity in the hot period while still having capacity in the slow period, in which case a customer visiting in the slow period instead of the hot period increases sales and profit for the firm. Hence, the time-specific nature of capacity transforms strategic customers’ intertemporal purchasing decision from a certain threat to margins to a potential opportunity to increase capacity utilization and profit! Figure 6 provides a summary of the main differences across the two settings.

Figure 6. Inventory and Capacity Settings: Differences and Implications on the Impact of Strategic Customer Effect on Firm Profits

Inventory setting	Capacity setting
<ul style="list-style-type: none"> • Inventory survives over time: One fewer unit sold in the high price period is one more unit available in the low price period • A customer who visits in the low price period instead of visiting in the high price period always reduces firm’s profits on account of lower margins earned on the same inventory unit • The firm does not advertise discounts to avoid a surge in visits in the low price period • Strategic customers can anticipate price rebates and availability; nonstrategic customers do not anticipate either • Strategic customers hurt the firm because they are more inclined to visit in the low price period • The low price period follows the high price period 	<ul style="list-style-type: none"> • Capacity perishes over time: One fewer unit sold in one period has no impact on the amount of units available in the other • A customer who visits in the low price period instead of visiting in the high price period may improve firm’s profits by increasing sales if the high price period is full but the low price period is not • The firm advertises discounts to induce visits in the low price period • Strategic and nonstrategic customers observe prices on both periods; strategic customers can anticipate availability • Strategic customers may benefit the firm because they are more inclined to visit in the low price period • The high price period can follow the low price period

Downloaded from informs.org by [128.91.109.54] on 21 February 2018, at 09:12. For personal use only, all rights reserved.

Table 1. Parameter Values Employed in the Numerical Study

Parameter	Value(s) considered	Source
Capacity, k	1,500	The capacity of Teatro Regio is 1,582 seats, of which 1,530 are proper seats (the rest being stools). We rounded down to 1,500.
Market size density function, g	Uniform $U[a, b]$ Mean μ : 3,000; 4,000; 5,000; $b - a$: 2,000; 4,000; 6,000	Given the popularity of the Teatro Regio in the last years, we consider the average potential market size to be between 2 and 3.33 times the single-period capacity. We have no information on demand variability, thus we use different mean-preserving spreads to study the impact of market uncertainty on performance metrics of interest. Bimodal parameters are chosen to replicate the mean and standard deviation of the uniform distribution above. ^a
Fixed cost, c_f	€30,000; €50,000; €70,000; €90,000; €110,000	The fixed costs that could be saved by closing down on a given night at Teatro Regio are estimated to be about 50K€–70K€, which comprises the per-show payroll for external performers and the cost of utilities. We also consider other values in order to capture a broad range of situations.
Full price, r_h	€130	The price charged for prime-time performances during the season.
Upper bound valuation, \bar{v}	€180	An educated guess based on ticket prices; same for both periods.
Lower bound valuation, \underline{v}	€50	Slightly above the lowest price charged at Teatro Regio for off-peak periods; same for both periods.
Customer valuation density function, f	$f(v_h, v_s)$ equal to $(\bar{v} - \underline{v})^2(1 - \eta)^2/2$ if $v_h \geq v_s + (\bar{v} - \underline{v})\eta$, 0 otherwise	We consider a uniform density over a triangular support, which captures customers' preference for the hot period; η is the seasonality coefficient.
Seasonality coefficient, η	0; 0.15; 0.30	0 (weaker) corresponds to valuations being below the diagonal, 0.30 (stronger) corresponds to a customer valuation area that is about half of it, 0.15 is an intermediate value.
Fraction of strategic customers, γ	0; 0.5; 1	We consider the two extreme cases of a strategic population and of a nonstrategic population, and the intermediate case.

^aObtained as an equally weighted sum of normal pdfs $\mathcal{N}(\psi, \sigma)$ and $\mathcal{N}(\psi + \Delta, \sigma)$, truncated at zero, where $\sigma = \Delta/4$, and ψ and Δ are chosen to replicate the mean and std. dev. of each of the uniform distributions listed above. Formally, $g(x) = \hat{g}(x)(1 - \hat{G}(0))^{-1}$, where $\hat{g}(x) = \Delta^{-1}\sqrt{2/\pi}[\exp(-8(x - \psi)^2\Delta^{-2}) + \exp(-8(x - \psi - \Delta)^2\Delta^{-2})]$.

6. Numerical Analysis

In this section, we present the results of a numerical study that helps us illustrate the advantages of threshold discounting. We consider the usage of threshold discounting at a potential service provider, the opera house Teatro Regio located in Torino, Italy. We extrapolate cost data from their 2014 balance sheet, and we use their pricing data to guide our choice of customers' intertemporal preference parameters. Table 1 illustrates the values chosen for each parameter and the sources employed.


In the absence of complete data on customer preferences, we capture the heterogeneity of customers' valuations with a uniform distribution over the support $\{(v_h, v_s): v_h \geq v_s + \eta(\bar{v} - \underline{v})\}$, where \bar{v} and \underline{v} are the highest and lowest customer valuations (same for both periods), $\eta = 0$ refers to the case when the support coincides with the area below the diagonal $v_h = v_s$, and higher values of η "shrink" the support toward the point (\bar{v}, \underline{v}) , thus making the hot period increasingly more preferred by customers (due to increasing $v_h - v_s$), and making the demand more seasonal. For the market size density function, g , we employ both a uniform distribution and a "twin hills" bimodal distribution obtained as the normalized sum of two normal distributions with the same standard deviation and

shifted means, censored at zero (see Table 1 for the full description of the parameters employed).

Overall, we consider 18 different market size distributions, three different levels of seasonality, three levels of strategic customers in the population, and five different cost structures—including the actual cost structure of Teatro Regio—and we simulate all possible combinations of these parameters, for a total of 810 scenarios examined. Before we proceed with our results, it is important to acknowledge some limitations of our analysis. Like all models, ours is meant to provide a simplified representation of reality that is amenable to study and, as such, it necessarily leaves out some realistic phenomena, e.g., there can be repeat customers, or customers who are unable to access the firm in one period may spill over to later periods. The use of same data from a real enterprise to calibrate our model is not meant to deliver a detailed prescriptive analysis, rather it is meant to improve the accuracy of both the direction and the magnitude of the effects found in the numerical analysis.

Figure 7, panel (a) shows the profit gains of threshold discounting over the traditional approach (i.e., the best between closure and regular discounting) for each of the 810 scenarios simulated, distinguishing between the case of a uniform (gray line) and a bimodal (black line) market size density function. Profit gains range

Figure 7. Profit Gains of Threshold Discounting over the Traditional Approach



Notes. Profit gains are measured as the increase in expected profit when employing threshold discounting compared to the best between regular discounting and closure. Market uncertainty refers to the standard deviation of the market size density function, g . Seasonality refers to the seasonality coefficient, η .

from -0.3% to a sizable 33% , and are similar for uniform and bimodal market distributions, though typically slightly higher in the latter case. In 14% of the scenarios, threshold discounting performs worse than the traditional approach, but the difference in performance is small, with an average profit reduction of 0.1% (Figure 7, panel (b)). In the remaining 86% scenarios, threshold discounting performs better than the traditional approaches and leads to profit gains that are often substantial, on average about 8.3% higher, which results in a total average profit gain of 7% across all scenarios.

Going from a nonstrategic to a strategic population of customers has a remarkably consistent positive impact on profit gains, increasing them in 82% of the scenarios with low market uncertainty, and in 100% of the other scenarios (Figure 7, panel (c), line, right axis). These gains contribute to 12% of the advantage of threshold discounting on average, going from 8% when the market uncertainty is lower, up to over 16% when market uncertainty is higher (Figure 7, panel (c), bars, left axis). These observations show that the strategic scarcity effect described in Section 4.2 is, in many cases, beneficial for the firm, especially when market uncertainty is high, possibly because of the increased effectiveness of discounts in these cases (this will be

investigated in Section 6.1), and overall reinforces the importance of modeling strategic customer behavior in our context.

Figure 7, panel (d) shows the advantage of threshold discounting for different levels of fixed costs and average market size. Fixed costs have an inverted-U shape impact on the advantage of threshold discounting, as does the average market size—see, for example, how the darkest bar first increases, and then decreases, as the average market size increases. Moreover, the fixed costs associated with the highest profit gain is higher for larger markets because when the average market size is higher, the preferred traditional approach tends to be regular discounting, which suffers more from fixed costs than threshold discounting; conversely, when the average market is lower, closure is often better than discounting, and is affected the least by fixed costs.

Interestingly, the impact on profit gains of a joint increase in fixed costs and market size by a given factor turns out to be equivalent to a proportional decrease in capacity by the same factor (see Lemma 9 in the online appendix). The reason is that these three primitives of the model determine the “scale” of the system, with capacity and market size distribution regulating

the variable component of profit and fixed costs affecting the fixed component of profit. If all three change by the same factor, profit gains are not affected; similarly, a change of two of these primitives in one direction is equivalent to a proportional change in the third primitive in the opposite direction. Indeed, our analysis shows that capacity also has an inverse-U-shape impact on profit gains.

Figure 7, panel (e) shows that the advantage of threshold discounting decreases when customer preference for the hot period is stronger (higher η), and increases when market uncertainty is higher. The former finding may appear surprising in light of Theorem 1, which basically establishes the superiority of threshold discounting when demand balancing is beneficial for the firm—or when there is strong demand seasonality. While asymmetric customer preferences and the ensuing seasonal demand are required for demand balancing to be beneficial at all, too strong customer preferences make it difficult and very costly for the firm to balance demand: in these cases, a simple approach that does not rely on demand balancing, such as closure, can deliver most of the profit by serving customers in the period that they strongly prefer. The second observation, i.e., that higher market uncertainty increases the value of threshold discounting, can be explained by the fact that making the wrong choice between discounting and closure ex ante is typically more costly in more extreme market states: as these contingencies become more likely, the value of threshold discounting increases.

Interestingly, Figure 7, panel (f) shows that higher market uncertainty can improve profit not just relative to the traditional approach, but also in absolute terms. The advantage that threshold discounting derives from market uncertainty is akin to similar effects observed in the context of option value. As discussed in Subsubsection 4.1.1, a firm employing threshold discounting, despite making decisions in the face of an uncertainty market size, can be thought of as if choosing between (demand balancing-enhanced) discounting and closure once market size uncertainty is realized. There are, however, some differences, for example, (i) higher market uncertainty in our context also reduces the expected gain from both options, i.e., from discounting and closure, and (ii) the firm has to choose exactly one approach, so the value of using one option is linked to the value of not using the other (choosing both and choosing neither are not viable choices).

To summarize, we find that threshold discounting can provide substantial profit gains, that strategic customers are mostly beneficial for a firm employing threshold discounting, and that the advantage of threshold discounting is highest when market uncertainty is high, fixed costs are neither very high nor very low, and seasonality is present but not excessive.

6.1. The Demand Balancing Effect of the Activation Threshold

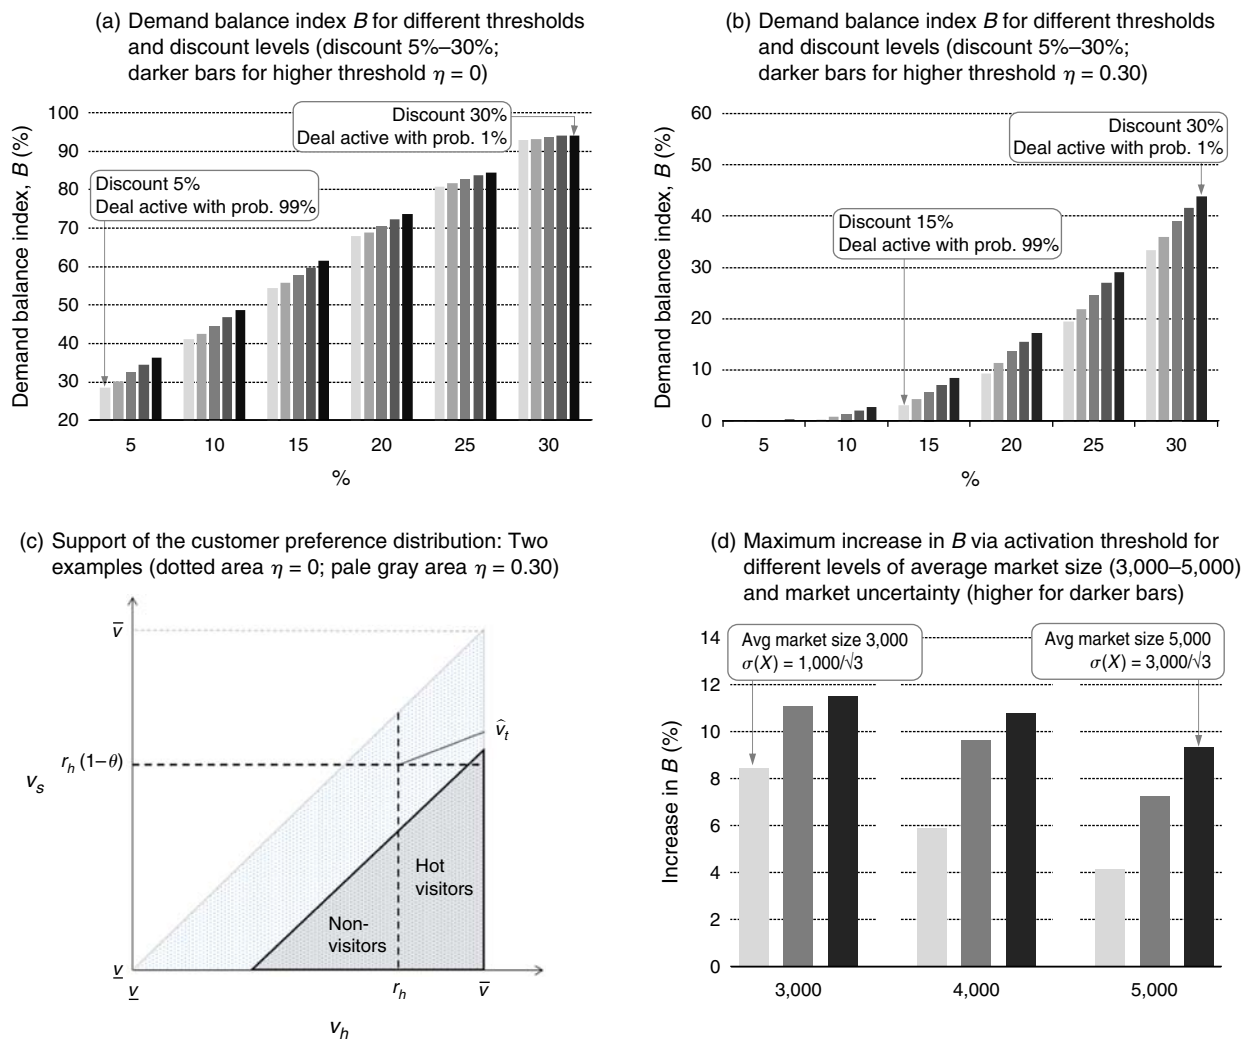
Finally, to further investigate the novel, surprising role of strategic customers in our model, we numerically investigate the effect at the center of it all—more efficient demand balancing under threshold discounting. This analysis highlights how learning of the deal activation changes strategic customers' visit decisions and the consequent demand patterns. To do so, we define a measure of the degree of balance of a given demand pattern across the two periods: the demand balance index B is equal to two times the fraction of demand that visits in the slow period divided by total demand, i.e., $B \triangleq (2\alpha_s^t(\theta, n))/(\alpha_s^t(\theta, n) + \alpha_h^t(\theta, n))$, $B \in [0, 1]$. When $B = 0$, all customers visit in the hot period, while when $B = 1$ an equal amount of customers visit in each period, so demand is perfectly balanced across periods.

Using a subset of 135 scenarios from the initial set (only strategic customers, uniform market size density function g), we consider a firm employing threshold discounting, and we change the hot and slow period prices and the activation threshold (expressed as probability of the deal being active to allow for a comparison across different scenarios) to study how they affect the demand balance index B .

Panels (a) and (b) of Figure 8 consider six possible discount levels, ranging from 5% to 30%, and for each of them shows how the demand balance index changes when the activation threshold is increased from a very low level (deal active with 99% probability, light gray) to a very high level (deal active with 1% probability, dark gray) using three intermediate levels (deal active with 25.5%, 50%, and 74.5% probability) for a total of five thresholds. As expected, in both panels (a) and (b), for any given discount, a higher threshold leads to a more evenly distributed demand (darker bars correspond to higher B) on account of strategic scarcity. However, the two panels display different patterns, with the threshold being more effective when discounts are low in panel (a), and being more effective when discounts are high in panel (b).

As shown in Theorem 2, an active deal signals to customers a higher-than-expected market size, which leads more customers to visit in the slow period on account of their updated beliefs that imply a decrease in relative availability. This balancing effect ceases to exist under two conditions. The first condition is when demand is already balanced, for example, when a high discount is offered, since in this case learning that the market size is higher does not affect the relative availability in the two periods (they are equally available in every market state). As we can observe in panel (a), the activation threshold is increasingly less effective at shifting demand as the discount grows, that

Figure 8. Demand-Balancing Effect of the Activation Threshold



is, as demand becomes more balanced. This explanation, however, does not fit the pattern observable in panel (b), in which the opposite seems to happen—the activation threshold is effective only once the discount is high enough. This happens because panel (b) considers a scenario with strong customer preferences for the hot period, hence a minimum discount is required before customers consider the possibility of visiting in the slow period, and higher discounts increase the fraction of the market that is (potentially) willing to shift to the slow period in exchange for a discount. Figure 8, panel (c) depicts the case of strongest customer preferences for the hot period ($\eta = 0.30$): note how in this example, when the discount θ is offered, no customer wants to subscribe or visit in the slow period. In general, the activation threshold is most effective at shifting demand when the discount is high enough to make the slow period interesting, but not so high that would already make demand even under regular discounting.

Finally, Figure 8, panel (d) reports the increase in B (demand more balanced) obtainable by moving the activation threshold from a 1% deal-on probability to a 99% deal-on probability, as a function of market uncertainty and average market size. We observe that the demand balancing effect of the activation threshold is more effective with a higher market uncertainty, because in these cases the informational update due to the deal activation is stronger (strategic customers' posterior and prior beliefs differ to a larger extent) and so is customers' response to it—their shift to the slow period. Demand balancing is also stronger when the average market size is lower. The reason is that customer response to the deal activation is stronger when the information that it carries is most different from their initial expectation. Since an active deal signals a high market, this leads to a stronger customer response (and more demand balancing) when their initial belief was of a lower market.

To summarize, our numerical analysis reinforces the existence of the strategic scarcity effect. Overall, we find

that this effect increases when the threshold is higher, market uncertainty is higher, and expected market size is lower. The impact of discounts is surprisingly complex, and appears to be mediated by demand seasonality, i.e., by how strong customer preference is for the hot period relative to the slow period (coefficient η). In particular, we find that the demand-balancing effect is the strongest when the discount is high enough to stir customers' interest in the slow period, but not so high to render demand balanced across periods "by itself," i.e., if customers were nonstrategic.

7. Discussion

In this paper, we have complemented previous work on threshold discounting by investigating the advantages of these offers from a—so far unexplored—operational standpoint. To this end, we have chosen a setting characterized by an uncertain market size, fixed capacity, and demand seasonality, in order to study demand-supply mismatches while accounting for demand substitution effects across periods. The choice of modeling the demand function at a micro level by directly characterizing raw customer preferences on both service periods has allowed us to discover important implications of strategic customer behavior in such settings. Overall, we find that threshold discounting offers can deliver value by allowing communication both from customers to the firm (responsive duality advantage) and from the firm back to customers (strategic scarcity advantage), and need not rely on networking effects or economies of scale, the only advantages studied so far for threshold discounts.

With respect to assumptions, we have tried to be as general as possible within the boundaries chosen for our stylized model. In particular, we have assumed no specific functional form for either the market size density function g or the customer valuation density function f , and the conditions imposed on the latter were instrumental to capturing demand seasonality, a defining element in our setting. The hot period price has been assumed exogenous in our model to keep exposition clearer—nearly all theorems in the paper can be extended analytically to incorporate an endogenous hot period price with minor, cosmetic changes in our results.¹⁷ This study focused on when and why threshold discounting delivers value (or not); a fruitful direction for future research could be the provision of practical guidelines on how to set the price and the activation threshold in order to get the most benefit from such offers.

Acknowledgments

The authors are grateful to Steve Chick, Jürgen Mihm, and Tim Van-Zandt; seminar participants at Georgetown University, Harvard University, Indiana University, New York University, Rutgers University, University College of London,

University of Pennsylvania, and University of Utah; and three anonymous referees and an associate editor for many helpful comments and suggestions that helped to improve the paper.

Endnotes

¹ Sources: <https://goo.gl/emx1jR> and <https://goo.gl/7DvBqD> (last accessed December 28, 2016). Expected revenues in the United States.

² Another interesting explanation for the discontinuation of threshold discounting offers relates to an incentive misalignment (and asymmetric bargaining power) between the firm and the intermediary through which the offer is channeled; while this analysis is not featured in the paper, it follows as an easy extension of the analysis developed. The interested reader can find this extension in the online appendix.

³ Formally, $\int_r^{+\infty} (\int_{v_h}^{+\infty} f(v_h, v_s) dv_h) dv_s > \int_r^{+\infty} (\int_{v_s}^{+\infty} f(v_h, v_s) dv_s) dv_h \forall r \leq r_h$.

⁴ The assumption of exogenous r_h is made for expositional clarity. Our results extend to the case with endogenous r_h ; see Section 7.

⁵ For example, see "Practical Paris: What's closed on Sundays, Mondays, and Tuesdays?" by Meg Zimbeck, October 3, 2012, <http://goo.gl/M52do>.

⁶ For example, The Lexi Cinema and Cavendish Conference Venues run "Monday madness" promotions, reducing their prices on Mondays, when they expect fewer customers.

⁷ See Lemma 2 in the online appendix for the precise conditions that delineate when this happens.

⁸ Section C.2 in the online appendix extends the model to the case where this cost is more substantial; as expected, we find that when customers incur substantial transaction costs, the profit of a threshold discounting offer is lower compared to the case in which transaction costs are negligible.

⁹ We discard equilibria in the continuation game in which no customer subscribes to the deal because no one else does and the deal is never active: any of these equilibria is Pareto dominated by the equilibrium in which a nonzero fraction of the customers subscribe and the deal is active with a positive probability, which is what we consider next (see Section A.2 in the online appendix).

¹⁰ Formally, we show that for any desired market level $\hat{x} > 0$ and any discount $\theta \leq \bar{\theta}$, there exists a unique threshold $\hat{n}(\theta, \hat{x})$ such that, if the firm announces the deal $(\theta, \hat{n}(\theta, \hat{x}))$, then the discount θ is offered in the slow period (the deal is active) if and only if the market size is higher than \hat{x} , or equivalently $\hat{n}(\theta, \hat{x})\alpha_s^t \cdot (\theta, \hat{n}(\theta, \hat{x}))^{-1} = \hat{x}$, where θ is defined as the discount level that makes demand even across the two service periods (see Lemma 3 in the online appendix).

¹¹ More precisely, customers' posterior belief on the market size after learning that the deal is active has first-order stochastic dominance over their prior belief.

¹² Let $A_m(x | \theta, n) = \min(1, k/(\alpha_m^t(\theta, n)x))$ be the service availability in period m when the discount is θ , the threshold is n , and the market size is x . Then, the monotonicity of $A_h(x | \theta, n)/A_s(x | \theta, n)$ with respect to x implies the monotonicity of $\mathbb{E}_x[A_h(x | \theta, n) | x \geq \hat{x}]/\mathbb{E}_x[A_s(x | \theta, n) | x \geq \hat{x}]$ with respect to \hat{x} , for any $G(x)$. See Lemma 6 in the online appendix for the formal proof.

¹³ In such a case, the profit of the firm when open on both periods is proportional to x for every θ , i.e., it is equal to $x \cdot r_h(\alpha_h^t(\theta) + (1 - \theta)\alpha_s^d(\theta))$, which has a unique optimal discount θ for every market size x .

¹⁴ Another potential explanation for why threshold discounting offers may have been discontinued is the disparity in bargaining power (and the incentives misalignment) that exists between the large intermediaries that dominate the industry, like Groupon, and the plethora of small service businesses that wanted to be featured in

their websites. We find that, compared to the firm, an intermediary prefers deals with a lower threshold—and a higher discount—which is also consistent with casual observations of many deals observed in practice. The interested reader can find a formal analysis for mediated threshold discounting offers in Section C.2.1 in the online appendix.

¹⁵There are two exceptions. One is the empirical work by Li et al. (2014), which argues that if, on the one hand, strategic customers reduce margins, on the other hand they increase demand, either by forcing the firm to reduce prices, which in itself raises demand, or by making customers postpone purchases and thus have a second purchasing opportunity. Hence, the effect on profit may go either way. The second exception is the working paper Chun and Ovchinnikov (2015), who show that airlines can use loyalty programs to take advantage of customers' strategic behavior, inducing them to fly more than needed, thereby increasing demand and profit. As it will soon be clear, our result arises because of very different underlying dynamics.

¹⁶One may wonder how strong is the regularity condition that the antiderivative of $\Pi'_{d-mg}(\theta) + \Pi'_{d-ex}(\theta)$ be unimodal in θ . Actually, not too strong. In fact, the demand expansion effect of discount on profit (primitive of $\Pi'_{d-ex}(\theta)$) tends to have an inverse-U-shape because increasing sales has progressively less value as the discount increases, and the margin effect of discount on profit (primitive of $\Pi'_{d-mg}(\theta)$) is always concave, since for higher discounts the margin is lost on a higher amount of sales.

¹⁷We could extend all results except the second point of Theorem 4, which has been confirmed numerically; see Section C.2.3 in the online appendix for the extended results and their proofs.

References

Anand KS, Aron R (2003) Group buying on the web: A comparison of price-discovery mechanisms. *Management Sci.* 49(11):1546–1562.

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.

Barr A, Baldwin C (2011) Groupon's IPO biggest by U.S. Web company since Google. *Reuters* (November 4), <http://goo.gl/h8VFj>.

Boyacı T, Özer Ö (2010) Information acquisition for capacity planning via pricing and advance selling: When to stop and act? *Oper. Res.* 58(5):1328–1349.

Cachon GP, Swinney R (2009) Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Sci.* 55(3):497–511.

Cachon GP, Swinney R (2011) The value of fast fashion: Quick response, enhanced design, and strategic consumer behavior. *Management Sci.* 57(4):778–795.

Chen Y, Zhang T (2015) Interpersonal bundling. *Management Sci.* 61(6):1456–1471.

Chun SY, Ovchinnikov A (2015) Strategic consumers, revenue management, and the design of loyalty programs. Working paper, McDonough School of Business, Washington, DC.

Crew MA, Fernando CS, Kleindorfer PR (1995) The theory of peak-load pricing: A survey. *J. Regulatory Econom.* 8(3):215–248.

Edelman B, Jaffe S, Kominers S (2016) To Groupon or not to Groupon: The profitability of deep discounts. *Marketing Lett.* 27(1):39–53.

Hu M, Shi M, Wu J (2013) Simultaneous vs. sequential group-buying mechanisms. *Management Sci.* 59(12):2805–2822.

Huang T, Van Mieghem JA (2013) The promise of strategic customer behavior: On the value of click tracking. *Production Oper. Management* 22(3):489–502.

Jing X, Xie J (2011) Group buying: A new mechanism for selling through social interactions. *Management Sci.* 57(8):1354–1372.

Kong G (2016) Promotion strategy for a service firm with delay sensitive customers. Working paper, University of Minnesota, Minneapolis.

Li J, Granados N, Netessine S (2014) Are consumers strategic? Structural estimation from the air-travel industry. *Management Sci.* 60(9):2114–2137.

Li X, Wu L (2014) Herding and social media word-of-mouth: Evidence from Groupon. Working paper, HEC Paris, Paris.

Liu Q, van Ryzin GJ (2008) Strategic capacity rationing to induce early purchases. *Management Sci.* 54(6):1115–1131.

Lus B, Muriel A (2009) Measuring the impact of increased product substitution on pricing and capacity decisions under linear demand models. *Production Oper. Management* 18(1):95–113.

Mourdoukoutas P (2011) Groupon's problem. *Forbes* (August 14), <http://goo.gl/AjG1x>.

Netessine S, Tang CS (2009) *Consumer-Driven Demand and Operations Management Models: A Systematic Study of Information-Technology-Enabled Sales Mechanisms*, Vol. 131 (Springer, Berlin).

Ovide S (2011) Groupon IPO: Growth rate is 2,241%. *Wall Street Journal* (June 2), <http://goo.gl/UFFwK>.

Papanastasiou Y, Savva N (2017) Dynamic pricing in the presence of social learning and strategic consumers. *Management Sci.* 63(4):919–939.

Parlaktürk AK (2012) The value of product variety when selling to strategic consumers. *Manufacturing Service Oper. Management* 14(3):371–385.

Roels G, Su X (2013) Optimal design of social comparison effects: Setting reference groups and reference points. *Management Sci.* 60(3):606–627.

Su X (2007) Intertemporal pricing with strategic customer behavior. *Management Sci.* 53(5):726–741.

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

Talluri KT, Van Ryzin GJ (2005) *The Theory and Practice of Revenue Management* Vol. 68 (Springer-Verlag, Berlin).

Tereyağoğlu N, Veeraraghavan S (2012) Selling to conspicuous consumers: Pricing, production, and sourcing decisions. *Management Sci.* 58(12):2168–2189.

Veeraraghavan SK, Debo LG (2011) Herding in queues with waiting costs: Rationality and regret. *Manufacturing Service Oper. Management* 13(3):329–346.

Wu J, Shi M, Hu M (2014) Threshold effects in online group buying. *Management Sci.* 61(9):2025–2040.