Promotional Design for Small Businesses: The Operational Value of Online Deals

Simin Li
Kellogg School of Management, Northwestern University

Kejia Hu
Owen Graduate School of Management, Vanderbilt University

Martin A. Lariviere
Kellogg School of Management, Northwestern University

Among the limited ways for small service providers to balance demand and supply, launching temporary consumer offer may be attractive. However, relatively little work has empirically examined whether and how such offers pay off service providers. In this paper, using a comprehensive dataset from two leading deal platforms in China, we empirically study a new business model: the online deal. Service providers, who face predictable demand swings and capacity constraints, launch online deals for customers to prepay online and redeem later in store. Using a structural model, we show that online deals effectively facilitate demand-supply coordination through two levers, the discount and, more interestingly, the advance sales period. To our knowledge, using the advance sales period as revenue management tool has not been studied in the literature. Tailored to demand fluctuations and the service provider's operating margin, the advance sales period and the discount serve two operational roles to achieve profit maximization – smoothing demand mean and reducing variance-related costs. Furthermore, our model estimates enable us to quantify the operational value of the online deal. Via counterfactual analyses, we show that by using these two levers instead of solely a discount, 82.1% service providers see a mean profit improvement of 23.6%. The additional lever, advance sales period, helping to mitigate the extreme discounts is likely where the profit boosts come from.

Key words: new business model, service operations, supply-demand coordination

1. Introduction

While revenue management tools for large service providers such as airlines have been well-studied, relatively little work has empirically examined how small service firms facing shifts in demand can benefit from designing promotions to balance supply and demand. Achieving such coordination can be challenging for small firms. Given the high level of discretion and training often required from employees in the service industry, capacity is rarely flexible enough to adjust costlessly to extreme demand swings—peaks or valleys—even when the timing of demand fluctuations is known well in advance, for instance, when a shock is tied to a national holiday. Launching a temporary consumer offer to proactively shape demand may be attractive to service providers since solely "chasing demand" is often a sub-optimal option (Sasser 1976).

We focus on a new and controversial form of such consumer offers – the online deal, which is the discounted but prepaid service deal offered via online platforms. The deal opens for purchasing before the holiday but remains active for redemption both before and during the holiday time. Since 2011, online deals have been offered via deal platforms, including *Meituan* in China, and have grown rapidly over that time. In 2017, online deals generated $99.02 billions in revenue with a yearly growth of 18% in China (Analysys 2018). The monthly activate users on the two leading platforms in China reached 310 millions, while the active small service providers has reached 4.4 millions by 2017 (Meituan-Dianping 2017). Despite the growing popularity of online deals among small businesses, questions about the value of online deals to service providers remain unresolved. Experienced service providers believe online promotions can either "propel or smother your business" depending on the deals' design (Goltz 2010). Researchers suspect discounts can be detrimental to service providers under high volume holiday demand (Girotra et al. 2013).

To reconcile the popularity of online deals among small service providers with the doubts raised by researchers and practitioners depends on two understudied questions: Which deal design parameters can be used as demand-shaping levers, and how these demand-shaping levers benefit the firms. Edelman et al. (2016) examine the discounts observed in online vouchers and argue that the two mechanisms by which deep discounts might benefit service providers work only under very narrow conditions. However, are there any other deal parameters that can be used to increase profits and benefit firms, in particular, the timing of the deal's launch? This is a particularly relevant parameter for such deals that are active for redemption during holidays, or more generally across time when demand varies. Additional deal parameters could potentially mitigate the limitations of using only discounts.

In this paper, we empirically investigate these two understudied questions using a comprehensive dataset collected from the two leading online deal platforms in China, *Meituan* and *Dianping*. We study whether online deals pay off for small service providers, facing predictable holiday demand fluctuations. To begin, we find that small service providers use both discount level and launch date of holiday online deals as demand-shaping tools. In particular, discount level is strategically designed for holidays, and this design varies by industry. However more interestingly, we find that the number of days between launch date and holiday, which we term as the **advance sales period**, also varies across service industries. Given service providers want to offer a discount price, choosing an appropriate non-zero advance sales period benefits the service providers. Furthermore, we find that through altering these levers – the discount and advance sales period—two operational roles are implemented: demand mean smoothing and variance-related cost reduction.

We show that depending on demand fluctuation and operating margin, service providers aim to achieve demand smoothing and accordingly variance-related cost reduction by optimally designing discounts and advance sales periods – which are in turn characterized by customers' responsiveness – to optimize profits. Customers' responsiveness to the discount is unsurprising. However customers' responsiveness to the deal launch time is rarely discussed. The closest counterpart is strategic customers' choices over peak or off-peak hours (Dana 1998, Marinesi et al. 2018). In fact, around 49.1% of the service deals in our dataset have customers who are more responsive to the *advance sales period* than to the classic lever, the *discount*. We empirically show that 0.088 out of 0.1 unit of the additional increase in the demand swings magnitude can be smoothed out by re-optimizing both levers. Furthermore, cost-to-price ratio increase with 0.1 unit, adjusting the advance sales period accordingly leads to a further variance-related cost reduction up to 0.15 unit.

This is the first empirical work to quantify the operational value of online deals under demand swings across multiple industries. Moreover, we are the first to document and substantiate the importance of an optimally designed advance sales period. We find two major benefits of having a new lever beyond the discount. First, an optimally designed discount-only deal often ends up with an extreme discount level, i.e. close to the boundaries of $(0,1)$. The advance sales period mitigates such extreme choices. Additionally, discounts only contribute to smooth demand mean fluctuations, while the advance sales period might further help to control demand uncertainty and thus capacity costs. Revenue-maximizing small service providers benefit from both aspects. Counterfactual analyses show that 82.1% of service providers will gain a mean of 23.6% profit improvement if firms set an advance sales period instead of having an "instantaneous" discount.

The rest of the paper is organized in the following way. In Section 2, we summarize the literature that is related to our work. In Section 3 we introduce the dataset we use for analysis. We also present the cleaning and preparation of data used in the later empirical analysis. In Section 4, we build the structural model to characterize how service providers design deals to coordinate demand and supply and thus maximize profits. We estimate the model and present a rule of thumb for optimal deal design, including choosing the advance sales period, in Section 4.2. We discuss the operational value of the online deal in Section 5, especially the operational role of the advance sales period. Lastly, in Section 6, we show by counterfactual analysis that the advance sales period is an effective, profit-boosting revenue management tool that should not be overlooked.

2. Relationship to the Literature

Our work is related to two streams of literature: the first examines promotional strategies to coordinate demand and supply, and the second studies the optimal time to launch promotions.

2.1. Promotional strategies to coordinate demand and supply

Researchers have investigated various forms of promotional strategies to alter customer demand and thereby achieve demand-supply coordination. Although different in design, the promotional strategies take on two similar operational roles: shifting demand mean and controlling demand variance.

Promotional strategies can be designed to boost demand mean balanced with capacity constraints, to smooth demand fluctuations over time. Dating from cents-off paper coupons (Narasimhan 1984) to recent web-based daily deals (Reimers and Xie 2018) and voucher discounts (Edelman et al. 2016), researchers have shown that discounts exploit customers' heterogeneous valuations in service and thus boost demand mean. Moreover, web-based deals facilitate social interactions, which help to expand the market (Lee and Lee 2012, Jing and Xie 2011, Luo et al. 2014, Ye et al. 2012). Further taking capacity constraints into consideration, deals can be designed to smooth demand fluctuations. Sasser (1976) proposes several approaches to balance demand, including a differential pricing scheme (Feldman et al. 2017, Cachon et al. 2017, Lazarev 2013, Leslie 2004) and reservation system to shift demand (Yılmaz et al. 2017). While most promotional strategies are generated toward one objective, we show that the online deal can be designed for both, where holiday deals smooth demand through discount while the advance sales period induces customers to shift demand voluntarily across time.

Apart from demand mean smoothing, online deals also help to reduce demand variance. There are several techniques in the literature that jointly deal with demand mean and demand uncertainty that are closer to our object of study. Advance purchase discounts, adopted by airlines and hotel chains, nudge customers to lock in demand early using lower-than-regular prices (Gale and Holmes 1993, Mesak et al. 2010, Dana 1998). Similarly, service providers might gather demand information during purchasing but postpone delivery according to the prices charged (Nguyen and Wright 2015). One popular online promotion has been threshold discounts. Marinesi et al. (2018) show by a stylized model that having a threshold to activate the discount over a slow period balances the demand from slow to rush hours, which is confirmed via a numerical study using threshold discounting on opera house tickets. The above-mentioned business models smooth demand by offering different prices over a set of discrete time points. Furthermore, such promotions reveal some demand information before demand realization. However, the value of reducing demand uncertainty is never explicitly discussed in the literature.

Our contribution is then three-fold. First, we are the first to study the operational value of a new business model – the online deal, which has risen in popularity since 2011. Like the above business models, online deals move purchase time forward away from the redemption. However, instead of selecting discount levels for some specific time points, online deals optimally choose a continuum of purchasing and redemption date, i.e. the advance sales period. Second, we are the first to empirically recover customers' demand sensitivities to both discounts and advance sales periods. We not only empirically confirm the existence of the demand-smoothing and variance reduction effect, but also quantify how much demand fluctuations are smoothed and how much cost is reduced by controlling demand uncertainty. We prove that advance sales period is the key to achieving both operational roles. Third, we quantify the contribution to the operational values of online deal from the advances sales period separately from the classic lever, discount.

2.2. Optimal time to launch promotions

Unlike the pricing of promotions, the promotion scheduling problem is largely overlooked (Baardman et al. 2018). Operations researchers have started to work on the scheduling of promotions only recently. For example, Cohen et al. (2017) study when and how many rounds of promotions to launch over a finite horizon for fast-moving consumer goods. Baardman et al. (2018) use bipartite matching to schedule promotions to time periods for retail goods in a supermarket setting. In the marketing literature, sales promotion calendars have been studied for packaged goods under different demand models (Silva-Risso et al. 1999, Tellis and Zufryden 1995). However, since the perishable nature of service capacities (Ng et al. 1999) prevents stockpiling effects among customers, the results of promotion planning for consumer products are not applicable to the service industries. To the best of our knowledge, no work in the optimal scheduling of promotions has been done for service industries under time-varying customer demand.

We thus contribute to the literature with a study of the optimal launch time of service industries promotions. It is worth noting that various other aspects of such promotions have been studied in the literature, but not optimal launch time. For instance, Krishna and Zhang (1999) study optimal duration of the promotion; Lee (2014) provides a collaborative decision on pricing and commission rate of promotions; there are also studies on the purchasing and redemption time of promotions (Wu et al. 2014, Song et al. 2016). We fill in the blank by presenting a rule of thumb to design launch time optimally. Furthermore, we substantiate the importance and quantify the operational value of optimal launch time.

3. Data

We use a dataset covering deals offered on two top platforms in China, *Meituan* and *Dianping*[1]. As shown in Table 1, for each deal, we observe information regarding its service category, the original and current service price, time and city of its availability, detailed descriptions of the service, the launch and closing date, the offering duration[2], and the sales accumulated by the end of 1st, 5th and last active day from the launch date. We observed $5,664$ online deals launched on the platforms from January 4, 2012 to March 31, 2012, offered in 13 major cities in China with a total population of 144 million people. The observed deals contribute a total of $66 million sales generated from 6.6 million purchases. The two platforms cover the same set of cities, time windows, holidays, and services. Also, online deals are all non-refundable.

We present two sample observations from the dataset to illustrate what are the data fields we directly observe.

We can then construct the discount, one of the key deal parameters to be used in our study. We denote the classical demand-altering tool – the discount – as α, and it has a natural definition as in Equation (1).

$$\alpha = \frac{\text{Deal Price}}{\text{Original Price}} \tag{1}$$

The timespan of our dataset enables us to study deals offered for holidays versus regular time, since it covers three holidays, including *Spring Festival*, the most celebrated holiday for family

[1] *Meituan* and *Dianping* were ranked top two in market share since 2012 (iiMedia Research 2013).

[2] Duration = Closing date − Launch date

Table 1 Two sample observations from the dataset for illustration.

Platform	Launch Date	Closing Date	Volume (1)	Volume (5)	Volume	City	Service Category
Meituan	2012-01-16	2012-02-01	3	30	80	BJ	Casual Dine-In
Dianping	2012-02-13	2012-02-27	11	69	90	BJ	Outdoor

Deal Description
Only 999 yuan for ten people party at our dine-in *YuGongYuPo* originally at 1702.6 yuan.
Only 99 yuan for two-people ski package at *Yuyang Ski Resort* originally at 280 yuan.

Note: Deal prices, original prices, names of service providers, and detailed services are extracted from the *Deal Description* column via text mining. The column *Service Category* is the segmentation of detailed services suggested by the two deal-offering websites (see Section 3.1). The three *Volume* columns record the total sales volume by the end of the first, fifth and last day from when the deal was launched. The sales by the first, fifth and last day from launch are thus Volume $\times$ Deal Price.

reunion in China from January 22nd to January 28th; *Qing Ming*, the Chinese long weekend memorial holiday from April 2nd to April 4th, and *Valentine's Day*, a highly commercialized festival on February 14th.

We find no evidence of periodical deal offerings in the dataset, suggesting that service providers may be launching deals to manage swings in demand rather than regularly running promotions. It is thus fair to assume that online deals are launched based on service providers' thorough evaluations of their needs at a specific time. In particular, we label the deals that are launched before the holiday but remain active for redemption during the holiday to be **holiday online deals**.

For holiday online deals, in addition to the discount, the timing of the deal launch is another key parameter. That is, service providers need to decide how many days before the first day of the holiday the deal should be released. We name this time window the ***advance sales period***.

Table 2 Summary Statistics of Holiday and Non-holiday deals. Total number of deals $N = 5,664$, of which 67.9% are holiday deals.

	$\mu_{Holiday}$	$\mu_{Non-holiday}$	$(\mu_{nh} - \mu_h)CI_{95\%}$	p_{KS}
Original price	477.73	496.93	(-74.41,112,81)	0.11
Discount	0.414	0.435	(0.009, 0.113)	1.1e-06
Total sales volume	945.66	1535.68	(408.41,771.61)	2.4e-13
Offering duration	30.4	19.0	(-12.1,-10.8)	2.2e-16

Formally, we denote the advance sales period as T, and it can be calculated as:[3]

$$T = \text{First day of the holiday} - \text{Launch date} \tag{2}$$

Note that the advance sales period is only defined for the holiday deals, which has a mean of 12.73 days. For non-holiday deals, T is set to 0.

To ensure that the notion of "holiday online deals" is not tenuous, we confirm that service providers specifically design deal features in response to holidays by two statistical tests. We compare the differences between holiday and non-holiday deals in terms of their deal parameters, including original price, discount, and offering duration, as well as the market responses measured by their sales volume - see Table 2. Apart from the holiday and non-holiday means reported in columns 1 and 2, we also report the 95% confidence interval of the differences between the means obtained from Welch's two sample t-tests and the p-value generated from the two-sided Kolmogorov-Smirnov (KS) tests.

We conclude that both deal discounts and deal sales volumes are statistically significantly different between holiday and non-holiday deals. This is supported by both tests: zero is not included in the t-test confidence interval, and the p-value calculated from the KS test is small. We also see that service providers strategically lengthen the offering duration to cover the holidays. Furthermore,

[3] 0.9% of total deals are launched during the holiday, i.e. $T < 0$. By definition, these are not holiday deals. We exclude them from the analysis.

notice that the null-value is within the confidence interval of the original price and also that the difference in price distribution is not significant, at a 10% significance level according to the KS test. This mitigates the concern that services packages offered from holiday to non-holiday times are different.

3.1. Categorization of service providers

We have confirmed that holiday deals are specially designed, but to further specify the special holiday strategies requires a finer categorization of service providers, since we expect service providers to adopt "similar" deal strategies only if they share "similar" characteristics and offer services at "similar" prices. Experienced service providers identify two industry characteristics as deal design determinants: predictable demand fluctuations during holidays δ_D and service providers' operating margins (Goltz 2010). While the "similarity" is achieved by maximizing the ratio of cross-group variations on these two industry characteristics over the within-group variations, yet still having sufficient deal samples in each group ensures the validity of our further analysis in Section 4.

Formally, the finer categorization is achieved by first adopting the industry categorizations shown on the platforms as a basis[4], then further conducting an algorithmic clustering by service prices. To help customers navigate through the website, the two platforms segment all deals into nine major categories: *Food and restaurants, Movie theater, Hotel, Indoor, Outdoor, Travel, Lifestyle services, Hair salon, and Body care services*[5]. In the merchant guidebook, the two platforms both claim

[4] The only adjustment we made is on the *Lifestyle* service group. The original *Lifestyle* service group covers the most diverse services such as *Photography, Car wash, Health checkups and Tailoring*. However, if we put all these services together, it would violate our intention for categorization since these services likely do not share similar industry specifics, neither holiday demand swings nor operating margins. This can be seen in their largely different demand patterns, from cyclical (*Car wash, Health checkups*) to rare and random occurrences (*Tailoring*). Also, *Photography* will require a much higher labor intensity than that of *Car wash*. To ensure the within-group homogeneity of deals, we separated all service types in group *Lifestyle*. After this decomposition, only the *Photography* group had a sufficient amount of deals for next steps. All the rest of the service groups had less than 30 deals. Thus we only kept *Photography* deals, which contribute 65.2% of all *Lifestyle* deals.

[5] The two platforms share the same deal segmentation levels.

that their categorizations reflect customers' common beliefs about which services are alike – see Dianping (2018). Platform designers' expertise knowledge ensures similarities in demand changes across time and service costs within categories. The nine categories are then further refined based on deals' original prices by a K-means clustering. We select the optimal number of clusters based on the elbow rule, which in turn ensures the within-category homogeneous operating margins. We demonstrate how we use the elbow approach on the SCREE plots in Appendix A.

We present the final 10 service categories[6] and their summary statistics[7] in Table 3, in which deal strategies should be homogeneous within but vary across categories. In total, we use 86.0% of the deals in the original dataset ($5,664$ total deals). We statistically validate the significant differences between the discount and the advance sales period across the rest of 10 service categories using Tukey's Honest Significance Test (see Appendix C).

3.2. The anticipated demand shocks

Within industry categories, service providers design online deals in response to the same anticipated demand swings; therefore, recovering the underlying holiday demand fluctuations service providers are facing is the first step to understand how they optimize deals.

We use the variations in sales among holiday versus non-holiday deals after controlling for price, discounts and other deal characteristics to estimate the baseline change in market demand, i.e. the anticipated demand shock, δ_D. The estimation is presented in details in Appendix B. We present the estimation results in the last column of Table 3. It is worth noting that the five categories that show positive holiday effect, including *Travel, Outdoor, Photography, Indoor* and *Casual Dine-in*, match the top expenditure categories announced in the holiday expenditure report in 2012 (Liu and Sun 2017). There are also industries experiencing holiday demand drops. Cheap hotels, for instance,

[6] We keep only industries with more than 45 deals after price clustering to guarantee sufficient samples within the categories for further analysis.

[7] The price reported in the *Travel* industry is the person and per day price. The mean of package price is $\bar{\mu}_p = 466$, with a standard deviation of 367.

Table 3 Summary statistics by refined categories. In total, we use $4,869$ deals in our analysis, of which $2,619$ are holiday deals, 90.0% of the original number of holiday deals. The italic category names indicate where the refined categories originally stems from. The new names are given based on observing the deals included in the same cluster after the algorithmic price clustering. Low, High in the brackets indicate the price levels.

Service Category (number of deals)	$\bar{\mu}_{\alpha,h}$ (s.d.)	$\bar{\mu}_p$ (s.d.)	$\bar{\mu}_{\text{sales vol}}$ (s.d.)	$\bar{\mu}_T$ (s.d.)	$\widehat{\delta_D}$
Body Care: Beauty (243)	0.26	445.60	563.90	14.26	-0.134
	(0.17)	(351.13)	(951.30)	(12.72)	(0.100)
Body Care: SPA (147)	0.10	1903.70	221.87	12.39	-0.026
	(0.10)	(664.71)	(279.06)	10.32)	(0.109)
Casual Dine: Casual Dine-in (2855)	0.48	94.65	1278.62	12.67	0.010
	(0.18)	(78.85)	(2919.22)	(13.18)	(0.047)
Fine Dine: Sit-down (Low) (419)	0.49	89.70	1279.80	13.17	-0.016
	(0.18)	(53.84)	(2956.56)	(13.30)	(0.063)
Fine Dine: Sit-down (High) (105)	0.32	333.89	453.36	11.82	-0.208
	(0.18)	(113.06)	(749.63)	(11.18)	(0.232)
Hotel: Hotel (273)	0.52	192.31	440.44	10.66	-0.793
	(0.15)	(79.85)	(760.31)	(10.20)	(0.306)
Indoor: Indoor (110)	0.26	177.20	1977.28	12.78	0.039
	(0.21)	(119.24)	(4039.00)	(14.72)	(0.108)
Outdoor: Outdoor (Gate) (431)	0.47	182.82	772.71	9.25	0.100
	(0.18)	(129.65)	(1558.95)	(10.64)	(0.076)
Photography: Photography (120)	0.20	2300.84	170.02	12.92	0.154
	(0.10)	(1512.46)	(242.79)	(11.93)	(0.156)
Travel: Travel (166)	0.39	257.33	422.39	10.05	0.089
	(0.12)	(100.77)	(977.79)	(9.99)	(0.077)

see the largest demand decrease during holidays[8]. This results from the fact that the majority of customers seeking such services temporarily leave the major cities for their hometown during *Spring Festival* and *Qingming*. The cities covered in our dataset are all first and top second-tier cities with significant immigrant populations.

To cope with these wide demand swings, both peaks and valleys, small service providers resort to online deals. In the following section, we construct a model to show how online deals are used as demand-shaping tools and whether these deal strategies pay off.

4. Structural Model

A strategic service provider would leverage a profit-maximizing deal strategy, tailored to its operating margin and the predictable demand swing, to shape its holiday discount demand. The strategy is then carried out by optimally designing the online deal parameters according to an understanding of how customers would respond to the deal design. We formally construct the service provider's decision-making process in a profit-maximizing problem, which demonstrates how customers respond to deal strategies, and how deal strategies in turn drive total profit. Our estimation strategy quantifies not only the optimal deal design but, more interestingly, customers' sensitivities to the deal parameters, which jointly unveil the operational value of the online deal.

4.1. Model

The service provider's profit-maximizing objective function is determined not only by the service provider, its deal strategy, and its intrinsic features including operating margin and demand fluctuation, but also by customers' sensitivities to the deal strategy. In general, a service provider generates revenue from both discount and full-price customers and expends costs on the total volume of customers served. The marginal revenue earned from the discount demand D_1 and the full-price demand D_2 are αp and p respectively. We model both demands as normally distributed.

To start with, the mean of discount demand D_1 and full-price demand D_2 share some common determinants, which consist of their respective baseline population, the non-industry-related fixed

[8] Note $\widehat{\delta_D} > -1$ should always be true.

effects and the original price of the service. We normalize the total baseline population of the potential customers to be 1, of which a proportion of β_6 comes from full-price customers and the rest comes from discount customers, i.e. the deal users in our context. During the holiday, the demand swing results in a δ_D change in the total baseline population, which ends up to be $(1+\delta_D)$. We assume that the holiday demand swing affects discount and full-price demand uniformly. On top of the baseline population, customer demand might vary across cities (M), holidays (H), and also how long the holidays last (L). Lastly, discount and full-price customer demand should both be driven by the original price of the service, p, though with different price sensitivities, which we denote as β_3 and β_7, respectively.

In addition to these common determinants, the mean of discount demand D_1 is further driven by deal-related variables, including the two deal parameters, the competition the deal faces, the deal offering duration, and the platform on which the deal is launched. Given the same set of deal parameters, service providers' discount demands can vary depending on how sensitive their customers are to the deal parameters. We denote discount sensitivity as β_1 and advance sales period sensitivity as β_2. The sign and the magnitude of the sensitivity coefficients reveal the direction and quantify the effectiveness of mean discount demand shifting. Furthermore, each deal faces competitions from "similar" deals. As customers navigate through the online platforms, they browse and choose from a set of competing deals that are offered on the same platform and holiday, in the same city and under the same service category. We summarize the competition level (R) a particular online deal faces as the number of deals that are "similar" to it. We further include offering duration (d) and platform fixed effect (F) to control for any possible unobserved effect on discount demand.

When it comes to demand variance, in addition to the common baseline variation, σ^2, shared by discount and full-price customers, the advance sales period plays a role that is unique to discount demand. Intuitively, on the one hand, discount demand variance might decrease as the advance sales period increases, since early launches can induce more potential customers to reveal their demand for service and therefore reduce uncertainty. On the other hand, demand variance could increase as the advance sales period increases, since more opportunities for pre-holiday redemption granted to customers increase the uncertainty in actual holiday demand. We denote customer sensitivity to the advance sales period in demand variance as β_5. It determines whether the online deal can reduce variance-related costs.

We use a convex cost structure as proxy for small service providers' inflexible capacity. The marginal service cost increases as more customers are served. Additionally, service providers have to pay a constant commission rate to the deal-offering platform. We model a service provider's total payment to the platform as a linear structure in T with a coefficient γ.

Formally, a service provider's profit maximization problem is modeled as follows, the optimal solution of which is the analytic optimal deal strategy a service provider should follow.

$$\max_{\alpha,T} \mathbb{E}V(\alpha,T) = \mathbb{E}[(\alpha p)D_1 + (p)D_2 - (c(D_1+D_2)^2 + \gamma T)] \tag{3}$$

$$\begin{aligned}
\text{s.t.} \qquad & D_1 \sim N((1+\delta_D)(1-\beta_6) - \beta_1\alpha + \beta_2 T - \beta_3 p + \\
& \qquad\qquad \phi^R R + \phi^H H + \phi^M M + \phi^F F + \phi^d d + \phi^L L, \sigma^2 + \beta_5 T) \\
& D_2 \sim N((1+\delta_D)\beta_6 - \beta_7 p + \phi^{H'} H + \phi^{M'} M + \phi^{L'} L, \sigma^2) \\
& \qquad\qquad \text{where} \quad \beta_7 > 0, c > 0, \beta_{i\notin\{7\}}, \phi, \phi' \text{ are free}
\end{aligned}$$

Proposition 1. *Solving the profit optimization problem, we can write down the expressions for the* ***optimal deal strategy*** (α^*, T^*). *Denote* $\kappa = 1 + \delta_D$.

$$\begin{aligned}
\alpha^* =& A_1\frac{1}{p^2} + A_2\frac{1}{p} + A_3 + A^H\frac{1}{p}H + A^M\frac{1}{p}M + A^F\frac{1}{p}F + A^L\frac{1}{p}L \\
T^* =& B_1\frac{1}{p^2} + B_2\frac{1}{p} + B_3 p + B_4 \\
&+ B^R R + (B^H + B^{H'}\frac{1}{p})H + (B^M + B^{M'}\frac{1}{p})M + (B^F + B^{F'}\frac{1}{p})F + B^d d + (B^L + B^{L'}\frac{1}{p})L
\end{aligned} \tag{4}$$

where

$$\begin{pmatrix} A_1 \\ A_2 \\ A_3 \\ B_1 \\ B_2 \\ B_3 \\ B_4 \end{pmatrix} = \begin{pmatrix} \frac{2\beta_1\beta_5}{\beta_2}c^2 + \frac{2\beta_1\gamma}{\beta_2}c \\ 2\beta_6\kappa c + \frac{\beta_5}{\beta_2}c + \frac{\gamma}{\beta_2} \\ -2\beta_7 c \\ \frac{\beta_1}{\beta_2}(\frac{2\beta_1\beta_5}{\beta_2}c^2 + \frac{2\beta_1\gamma}{\beta_2}c) \\ \frac{\beta_1}{\beta_2}2\beta_6\kappa c + \frac{2\beta_1\beta_5}{\beta_2^2}c + \frac{2\beta_1\gamma}{\beta_2^2} \\ \frac{\beta_3}{\beta_2} \\ -\frac{(1-\beta_6)\kappa}{\beta_2} - \frac{2\beta_1\beta_7}{\beta_2}c \end{pmatrix}, \begin{pmatrix} A^M \\ A^F \\ A^L \\ A^H \\ B^R \\ B^d \\ B^H \\ B^M \\ B^F \\ B^L \end{pmatrix} = \begin{pmatrix} 2c\phi^{M'} \\ 2c\phi^{F'} \\ 2c\phi^{L'} \\ 2c\phi^{H'} \\ \phi^R\frac{1}{\beta_2} \\ \phi^d\frac{1}{\beta_2} \\ -\phi^H\frac{1}{\beta_2} \\ -\phi^M\frac{1}{\beta_2} \\ -\phi^F\frac{1}{\beta_2} \\ -\phi^L\frac{1}{\beta_2} \end{pmatrix} \tag{5}$$

Summarized in Proposition 1, the optimal deal strategy is indeed a function of customers' sensitivities and industry characteristics including cost-to-price ratio and predictable holiday demand fluctuation. We can thus empirically identify the optimal deal strategy by linking the observed deal design to the analytic optimal expression in (4).

4.2. Identification strategy

We estimate the key parameters from the structural model in this section in two steps: quantifying the optimal deal design and further identifying customers' sensitivities and other model parameters. Since a particular optimal deal strategy is only expected to be shared among service providers with similar customer sensitivities and industry features, the estimation strategy presented below is conducted on the level of our clustered service categories.

Assuming service providers are rational, we expect the observed discount and advance sales period to align with the analytic optimal solution. Formally, the observed deal designs (α, T) should only minimally deviate from the optimal design (α^*, T^*). Deviations from the optimal solution, which are modeled as the mean zero error terms denoted as $(\varepsilon_\alpha, \varepsilon_T)$. These errors are idiosyncratic and not associated with any of the observed covariates in the demand mean specifications, which include price, holiday, city, platform, and holiday length.

$$\alpha = \alpha^*(A_{i\in\{1,2,3\}}, A^x) + \varepsilon_\alpha, \quad \varepsilon_\alpha \sim N(0, \xi_\alpha^2), \quad \mathbb{E}[\varepsilon_\alpha | \frac{1}{p}, \frac{1}{p^2}, H, M, F, L] = 0$$

$$T = T^*(B_{i\in\{1,2,3,4\}}, B^y) + \varepsilon_T, \quad \varepsilon_T \sim N(0, \xi_T^2), \quad \mathbb{E}[\varepsilon_T | \frac{1}{p}, \frac{1}{p^2}, p, R, d, H, M, F, L] = 0 \tag{6}$$

$$\text{where} \quad x \in \{H, M, F, L\}, y \in \{R, d, H, M, F, L\}$$

To quantify the optimal deal design, we are only left with recovering the auxiliary parameters $(\mathbf{A}, \mathbf{B})$.[9] We find their least squared estimators by minimizing the sum of squared errors of the two regressions as in (7).

$$\min_{\mathbf{A},\mathbf{B}} \quad (\alpha - \alpha^*(A_{i\in\{1,2,3\}}, A^x))^2 + (\beta - \beta^*(B_{i\in\{1,2,3,4\}}, B^y))^2 \tag{7}$$

By construction, the unconstrained problem (7) is a strongly convex programming, which guarantees the existence and uniqueness of the least squared estimator $\mathbf{A}, \mathbf{B}$, and therefore a unique optimal deal design. Given $\mathbf{A}, \mathbf{B}$, we can now invert equations (5). The estimated industry-specific senstivitiy coefficients and cost parameters are shown in Table 4. Detailed proofs of identification are given in Appendix G.

Point identifications of the model parameters are guaranteed by the uniqueness of design coefficients $\mathbf{A}$, $\mathbf{B}$. We estimate these parameters' standard deviations via bootstrapping. Within each service category, for every round of bootstrap, we select 99% of the deals to run the analysis. We repeat the procedure 100 times for each category, and the standard deviation is calculated from the 100 rounds of outputs. The results are robust to a range of 97%, 98%, and 99% of the subsamples[10].

5. Estimation Results and Discussion: The Operational Role of Online Deals

Applying the identification strategy, we first present the estimated results, based on which we analyze the operational role and quantify the operational value of the discount and the advance

[9] $\mathbf{A} = (A_{i\in\{1,2,3\}}, A^{x\in\{H,M,F,L\}})$, $\mathbf{B} = (B_{i\in\{1,2,3,4\}}, B^{y\in\{R,d,H,M,F,L\}})$

[10] We use multi-start from 100 randomly generated initial points of the regression coefficients to ensure that global minimum is achieved.

sales period. We essentially respond to the questions raised at the very beginning: whether and how the online deal pays off service providers.

5.1. Estimation results

Customers' sensitivities to deal parameters We start with quantifying two sets of customers' sensitivities to the deal parameters: the discount and advances sales period sensitivities, β_1 and β_2, which shape demand mean; and the advance sales period sensitivity, $\beta_5 c + \gamma$, that affects uncertainty-related costs.

Non-zero discount sensitivities, β_1, indicate that discount is an effect customer demand-shaping lever (see Table 4). Using discount as a lever to shape demand mean can induce two types of customer reactions, leading to opposing signs of discount sensitivity, β_1. First, customers' preferences over saving money collectively result in a positive discount sensitivity β_1. On the other hand, deep discounts make customers suspicious about the service quality and variety (Grewal et al. 1998, Kardes et al. 2004), which leads to a negative discount sensitivity β_1. Our estimated β_1 appear to align with the theoretical results in literature (see Appendix E) in terms of how the joint effects of the two opposing customer responses would look across industries (see Table 4).

Non-zero advance sales period sensitivities, β_2, indicate advance sales period is also an effective demand-shaping tool. Customers' sensitivity to the advance sales period $\hat{\beta}_2$ mostly comes from online deals' advertising effect, which reflects as the positive $\hat{\beta}_2$ as shown in Table 4. The longer the deals are made available, the more discount customers are attracted to the service. Further elaboration on interpreting β_2 includes in Appendix E.

The significant $\widehat{\beta_5 c + \gamma}$ shown in Table 4 enters the cost structure as the sensitivity coefficient to advance sales period T and controls the effectiveness of T as a variance-related cost-reduction lever. By definition, the sensitivity coefficient $\widehat{\beta_5 c + \gamma}$ is jointly determined by the coefficient β_5, which controls the demand uncertainty, and also by the platform commission rate γ.

Marginal service cost In addition to customer sensitivities, the other determinant of optimal deal strategy to be quantified is the service operating margins, which boils down to estimating the

Table 4 **The number of holiday deals, estimated coefficients of sensitivities to deal parameters, and marginal cost parameter of each service category are reported with their standard deviations.**

Service Category	N	$\hat{\beta}_1$	$\hat{\beta}_2$	$\hat{c}$	$\widehat{\beta_5 c+\gamma}$
Beauty	149	0.027	0.027	2.705	−0.0315
		(0.0090)	(0.0033)	(0.2011)	(0.00797)
SPA	105	−4.803	0.012	0.299	−0.0002
		(0.0342)	(0.0001)	(0.0037)	(0.00000)
Casual Dine-in	1383	−0.364	0.023	0.091	−0.0017
		(0.0130)	(0.0001)	(0.0006)	(0.00007)
Sit-down (Low)	234	−3.338	0.025	0.035	0.0012
		(0.1448)	(0.0003)	(0.0027)	(0.00002)
Sit-down (High)	69	−0.771	0.002	0.669	−0.0001
		(0.0074)	(0.0000)	(0.0086)	(0.00000)
Hotel	146	0.090	0.002	3.113	−0.0019
		(0.0021)	(0.0001)	(0.0386)	(0.00001)
Indoor	58	−3.090	0.055	0.072	−0.0027
		(0.1883)	(0.0015)	(0.0011)	(0.00016)
Outdoor (Gate)	149	−1.939	0.027	0.085	0.0006
		(0.0704)	(0.0003)	(0.0072)	(0.00002)
Photography	86	−23.556	−0.141	0.076	−0.0010
		(1.4584)	(0.0105)	(0.0079)	(0.00011)
Travel	109	2.897	0.006	0.046	−0.0004
		(0.3135)	(0.0009)	(0.0020)	(0.00006)

cost parameter, c, since service price p is observed. The homogeneity of operating margins within service categories ensures the validity of our estimator $\hat{c}$ evaluated at the service category level.

The magnitudes of our estimated marginal service costs align with the results and intuitions summarized in the literature. According to Lovelock (1983), since staffing costs contribute to a major part of service industries' marginal costs, service categories that give little discretion or room for judgment to their employees in terms of altering the service characteristics have lower costs. Such industries include traveling, indoor, outdoor theme parks, and fair price sit-down restaurants and casual dine-ins (see Table 4). On the contrary, serving a customer in a higher end sit-down restaurant with wide food options, hosting a customer in a separate hotel room, and providing beauty services and massages require either highly customized services or high levels of labor intensity and discretion or both. This inevitably leads to higher marginal costs (see Table 4).

The service categories can now be grouped into four quadrants based on the two characteristics that have driven service providers' design of holiday deals, the operating margin (measured by the cost-to-price ratio) and the anticipated demand swings as shown in Figure 1, by which we shall expect that the deal strategies adopted by service categories falling under the same quadrant play "similar" roles. We suspend an elaborate discussion of the operational roles that online deals serve until the next section. It is worth noticing that the service categories with low margins are in general less interested in offering online deals. Particularly, no service industries with low profit margins and stable demand participate in the deal business. This observation aligns with the result of Edelman et al. (2016), in which they conclude that online discount vouchers are more profitable for firms with low marginal costs using a stylized model.

5.2. The operational value of online deals: effective deal parameters and their operational roles

From the structural model estimators, we conclude that both the discount and the advance sales period are effective demand-shaping levers and furthermore that they jointly leverage two operational strategies to benefit service providers.

The online deal's abilities to shift demand mean and control uncertainty translate to two operational roles the online deal plays when it is optimally designed. First is *demand smoothing*, where service providers smooth demand to make it compatible with the limited capacity, thus avoiding

Figure 1 The grouping of service industries based on both operating margin and magnitude of demand fluctuations. We divide the whole space into four quadrants. From the first to the fourth quadrant, the industries are (low operating margin, large demand fluctuations), (large, large), (large, small), (low, small). The solid lines demarcate the full space.

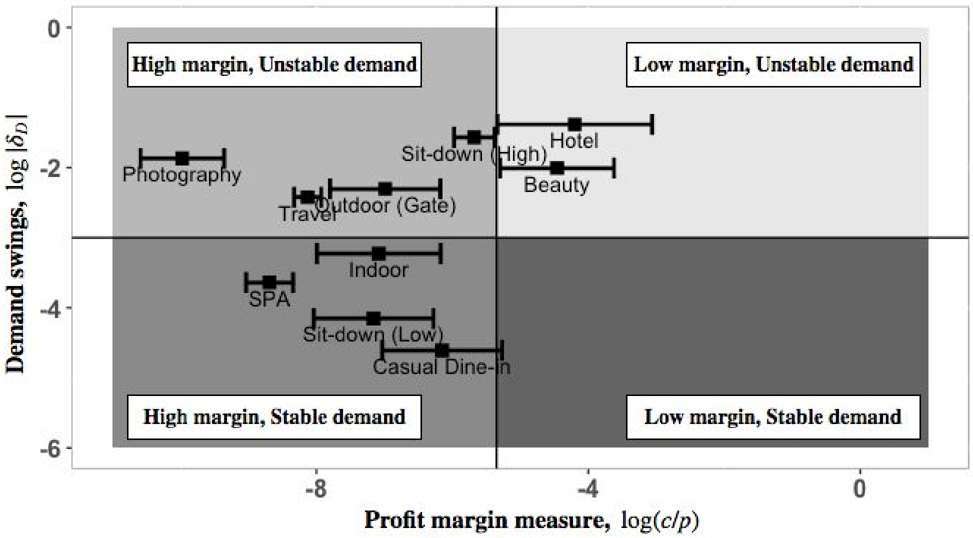

high mismatch costs. Smoothing demand is achieved via the design of both discount α and advance sales period T. Second is *variance-related cost reduction*, where providers achieve direct cost reduction by controlling demand uncertainty via advance sales period T. This joint effect of both roles pays off the service providers.

Before we proceed to quantify these two effects, we argue that the demand-smoothing and variance-related cost reduction effects are independent, which justifies our discussing them separately. Recalling the expected profit expressions in (8), we can break down the service provider's profit maximization problem into two subproblems. First, given capacity cost c, what is the optimal demand level to serve? Second, is there a way to directly reduce capacity costs via offering online deals?

$$\begin{aligned} \mathbb{E}V &= \underbrace{\alpha p\mu_1 + p\mu_2 - c(\mu_1+\mu_2)^2}_{\text{demand-smoothing object}} - 2\sigma^2 c \quad \underbrace{-(\beta_5 c+\gamma)T}_{\text{capacity-cost reduction object}} \\ \mu_1 &= \underbrace{(1-\beta_6)}_{\text{non-holiday discounted demand}} + \delta_D(1-\beta_6) \underbrace{-\beta_1\alpha + \beta_2 T}_{\text{deal-induced demand}} - \beta_3 p + \boldsymbol{\phi}\mathbf{X} \end{aligned} \tag{8}$$

We start by defining a measure for the demand-smoothing effect. We set the status quo as the demand to be served during a regular day, i.e. $\delta_D = 0$. Therefore, if the service provider does nothing different than regular days when a holiday comes, the demand to be served in this case directly changes by the amount of holiday demand swing $\widehat{\delta_D}$. On the contrary, the service provider can take $\widehat{\delta_D}$ into consideration and then launch an optimal holiday online deal. The difference between the final demand to be served when the online deal is available and the demand under the status quo is the post-smoothing-by-online-deal demand swing. The difference in magnitude between the nominal demand swing $\widehat{\delta_D}$ and the post-smoothing demand swing measures the demand-smoothing effect. We formally write the definition as follows.

Definition 1 (Demand-smoothing effect). We define an index $\mathcal{B}$ to measure how strong the demand-smoothing effect is. Suppose the service provider chooses $(\widetilde{\alpha}^*, \widetilde{T}^*)$ and (α^*, T^*) to maximize profit when $\delta_D = 0$ and $\delta_D = \widehat{\delta_D}$ respectively. The total demand to be served is then $\widetilde{\mu} = \widetilde{\mu}_{1;\delta_D=0} + \widetilde{\mu}_{2;\delta_D=0}$ and $\mu^* = \mu_{1;\delta_D=\widehat{\delta_D}} + \mu_{2;\delta_D=\widehat{\delta_D}}$ where $\widetilde{\mu}_{i;\delta_D} \triangleq \mu_i(\widetilde{\alpha}^*, \widetilde{T}^*; \delta_D = 0)$ and $\mu^*_{i;\delta_D} \triangleq \mu_i(\alpha^*, T^*; \delta_D = \widehat{\delta_D})$. $\widetilde{\mu}$ is the demand level to serve at the holiday steady state $\delta_D = 0$, where μ^* is the re-optimized demand-level at $\delta_D = \widehat{\delta_D}$. We define $\mathcal{B}$ as follows:

$$\mathcal{B} = |\widehat{\delta_D}| - |\widetilde{\mu} - \mu^*|$$

Note that an effective demand-smoothing treatment would have $\mathcal{B} > 0$ regardless of whether the demand swing is positive or negative. Furthermore, the larger $\mathcal{B}$, the stronger the demand-smoothing effect.

We see significant demand-smoothing effects across almost all service industries as illustrated in Figure 2. The distance between the dashed nominal demand swing and the solid post-smoothing demand swing represents the demand-smoothing effect, $\mathcal{B}$. Graphically, a successful demand-smoothing treatment results in a point, together with its 90% confidence interval, lying between the two dashed lines. All service categories except *Travel* and *Hotel* experience demand smoothing via online deals. Furthermore, the increasing gaps between the dashed and solid lines in Figure 2 seem to suggest that the larger the demand fluctuation, the larger the demand-smoothing effect. To

Figure 2 **The plot illustrates the demand-smoothing effect observed in each industry. The dashed line plots the magnitude of the predictable holiday demand swings, $\pm|\delta_D|$, while the solid line is the size of the demand swings after smoothing achieved through online deals, i.e. $|\widetilde{\mu} - \mu^*|$. We plot the 90% confidence interval around the post-demand-smoothing line.**

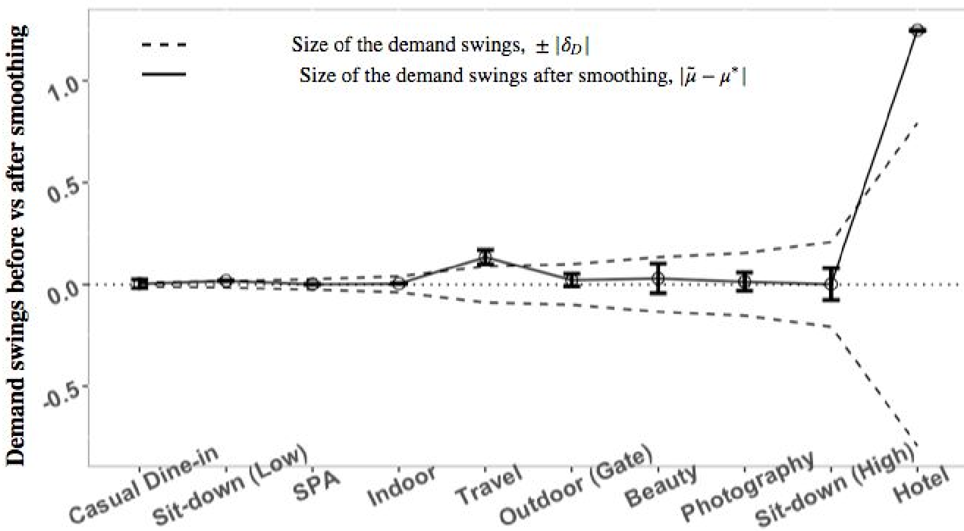

accommodate wide holiday demand fluctuations by their regular capacity[11] with moderate adjustments, service providers have to resort to demand smoothing, which might involve both stimulation and discouragement of demand (Lovelock 1983). We formally state the conjecture as follows.

HYPOTHESIS 1. *The larger the demand fluctuation, the stronger the demand-smoothing effect, through strategically selecting discount α and advance sales period T.*

Across all industries, for demand mean-shaping to happen, we see that the service providers need to be strategic in not only the discount level α^* but also the advance sales period, T^*. Subsequently, we can estimate the benefits contributed specifically by the advance sales period, T, in the mechanism of shaping demand mean. This can be quantified by comparing the value of the demand-smoothing object in the expected profit expression (8) under two scenarios: when the customers only offers a discount price $(\tilde{\alpha}^*, 0)$ and when they launch an optimal online deal, (α^*, T^*). We illustrate the profit gap between the two demand-smoothing objects in Figure 3.

[11] It is reasonable to assume regular capacity is optimally set according to the regular day or non-holiday demand.

Figure 3 **The contribution to one of online deal's operational roles, demand mean-shaping, specifically from selecting an optimal advance sales period, T.**

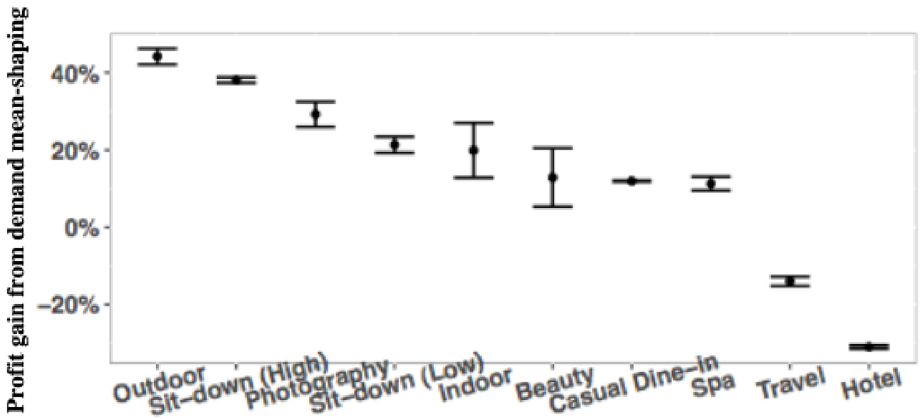

Figure 3 quantifies how much value the advance sales period contributes to the demand-mean shaping operational role. The only two industries that do not see a profit gain, *Travel* and *Hotel* are also the ones that do not see a significant demand-smoothing effect.

Next, we define a measure for the variance-related cost reduction effect.

Definition 2 (Variance-related Cost Reduction Effect). We define an index $\mathcal{R}$ to measure how strong the capacity-cost reduction effect is. That is,

$$\mathcal{R} = -(\widehat{\beta_5 c + \gamma})T$$

A negative sensitivity coefficient $\widehat{\beta_5 + \gamma}$ along with a non-zero advance sales period T ensure an effective cost reduction. Note that, mathematically, $\mathcal{R} \in (-\infty, \infty)$. We say there is a strong cost reduction effect when $\mathcal{R}$ is large.

In essence, the variance-related cost reduction effect offers a direct way – leveraging an optimally designed advance sales period – for the service provider to reduce its marginal service cost. The new marginal cost parameter is $\tilde{c} = \hat{c} - \mathcal{R}/(\mu_1 + \mu_2)^2$. We overlay the percentage cost reduced by the cost reduction effect, $\frac{\mathcal{R}/(\mu_1+\mu_2)^2}{\hat{c}}$, with the original marginal cost parameter $\hat{c}$ in Figure 4. One can argue that the source of cost reduction comes from customers' early purchases, which reduce demand uncertainty and therefore variance-induced costs – see (8). Therefore, the total service

Figure 4 **The plot illustrates the cost reduction effect. The y-axis on the left and the bar chart represent the percentage cost reduced, $\frac{\mathcal{R}/(\mu_1+\mu_2)^2}{\hat{c}}$, with the error bar representing the 90% confidence interval. That is, all service categories except for the Sit-down (Low) and Outdoor (Gate) see statistically significant cost reductions. The y-axis on the right and the curve plot the original marginal cost parameter, $\hat{c}$.**

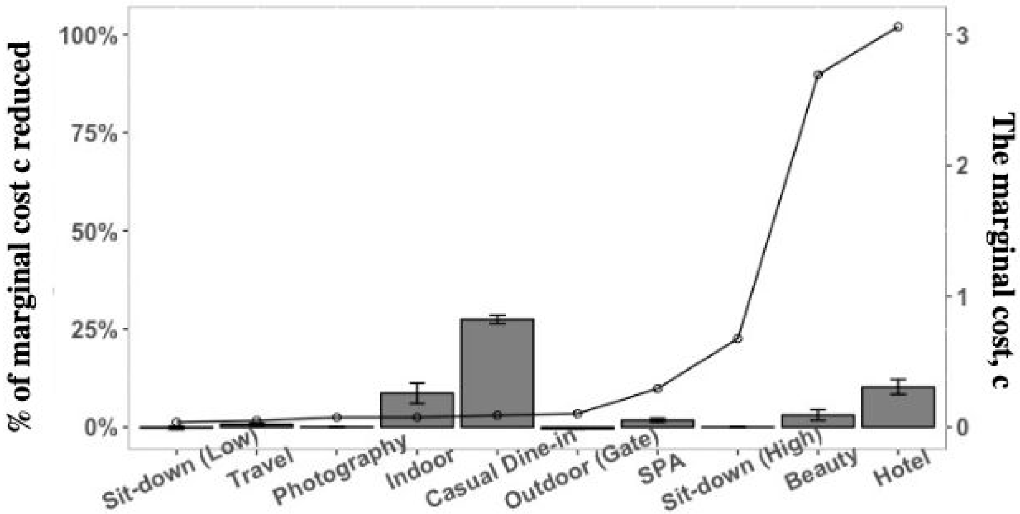

costs can be cut down without changing the expected demand to be served. This is especially attractive to service providers with small operating margins. We have the following conjecture.

Hypothesis 2. *The lower the operating margin, the stronger the variance-related cost reduction effect, which is achieved by re-optimizing the advance sales period T.*

We rigorously test the two hypotheses by pooling all deals together. Additionally, we control all other fixed effects contained in our structural model.

$$\mathcal{B} \sim \beta^{DSmooth}|\boldsymbol{\delta}_{\mathbf{D}}| + \beta_1^{CITY}M + \beta_1^{PLATFORM}F + \beta_1^{COMPETITION}R + \beta_1^{HOLLEN}L + \beta_1^{DUR}d$$

$$\mathcal{R} \sim \beta^{CostRed}\mathbf{c}/\mathbf{p} + \beta_2^{CITY}M + \beta_2^{PLATFORM}F + \beta_2^{COMPETITION}R + \beta_2^{HOLLEN}L + \beta_2^{DUR}d$$

The regression results shown in Table 5 not only support the directional relationships stated in our hypotheses, but further quantify the relations. In particular, with 0.1 unit increase in the magnitude of the demand swing, an average of 0.088 of the increased demand swing can be smoothed out by re-optimizing the deal parameters. Furthermore, by re-optimizing the advance sales period, a 0.1 increase in the magnitude of cost-to-price ratio leads to a 0.15 unit reduction in the variance-related cost.

Table 5 **Regression tests of the two hypotheses. We omit the other covariates in the table for clarify of display.**

	Operational Role of Deals:	
	Demand-smoothing (1) $\mathcal{B}$	Capacity-cost red. (2) $\mathcal{R}$
DEMAND SWING, $\lvert\delta_D\rvert$	0.878^{***}	
	(0.003)	
MARGIN, c/p		1.505^{***}
		(0.097)
Note:		*p<0.1; **p<0.05; ***p<0.01

We believe that an "instantaneous" customer offer, with no advance sales period, would not be as effective in maximizing profits compared to the online deal, since the effects and benefits of an optimal advance sales period T cannot be fully performed by the discount α alone. Advance sales period T has a unique feature: it alters variance and thus reduces capacity costs, which is particularly valuable under low operating margins. In addition, the advance sales period serves as an alternative tool to achieve demand smoothing, which is especially meaningful when an extreme discount is unwanted. The two benefits of a non-zero advance sales period lead to the profit gap between the two business models: the online deal versus the "instantaneous" discount.

6. Counterfactual Analysis: The Value of Online Deals' Optimal Launch Time

We proceed to perform several counterfactual analyses to substantiate the importance and quantify the operational value of the advance sales period.

To quantify the value of the advance sales period, we compare online deals to "instantaneous" consumer offers, which are also widely used. "Instantaneous" consumer offers are made available to customers when they visit the store. Service providers put up a barcode linking to the deal at the check-out counter for the customers who are aware of the deal-offering platforms to use. A customer's purchase and redemption activities happen at the same time, i.e. $T = 0$. Therefore, service providers would neither have the buffer time to smooth out demand swings nor gain any demand information in advance. The operational value of having an advance sales period is lost. A service provider can only shape demand via pure *pricing* strategy.

PROPOSITION 2. *The optimal price discount solves the model as follows in (9). The optimal deal strategy is to set an optimal discount level:*

$$\tilde{\alpha}^* = \arg\max \mathbb{E}[\tilde{\alpha} p \tilde{D}_1 + pD_2 - c(\tilde{D}_1 + D_2)^2]$$
$$\tilde{D}_1 \sim N((1+\delta D)(1-\beta_6) - \beta_1\alpha - \beta_3 p + \phi^R R + \phi^H H + \phi^M M + \phi^L L, \sigma^2), D_2 \sim N(\beta_6 - \beta_7 p, \sigma^2) \tag{9}$$

Solving from the model above,

$$\tilde{\alpha}^* = (p(1-\beta_6)\kappa - \beta_3 p^2 + 2\beta_1(1-\beta_6)\kappa c - 2\beta_1\beta_3 pc - 2\beta_1\beta_7 pc + p\phi^R R + 2\beta_1\phi^R cR + \sum_{X\in\{H,M,L\}} (p\phi^X X + 4\beta_1 c\phi^X X))/(2\beta_1 p + 2\beta_1^2 c)$$

A natural measure of the advance sales period's operational value is the percentage expected profit increase between using online deals and "instantaneous" consumer offers, as formulated in (10).

$$\frac{\mathbb{E}V(\alpha^*, T^*) - \mathbb{E}V(\tilde{\alpha}^*)}{|\mathbb{E}V(\tilde{\alpha}^*)|} \tag{10}$$

We find that by having a non-zero advance sales period, almost all industries see profit improvements on the industry level. Moreover, the majority of individual service providers gain profit improvements - see columns $2-3$ of Table 6. Pooling all service providers, among the 82.1% of service providers who see profit improvements, the mean profit gain is 23.6% with a standard deviation of 16.9%. The service industries that fail to see profit improvements, i.e. *Hotel* and *Travel*, are exactly the industries that have one of the smallest sensitivity coefficients to the advance sales period but at the same time face one of the larger demand fluctuations (see Table 3), which also fail demand smoothing as shown in Figure 2. The insensitivity to launch time and thus the failure of demand smoothing treatment may explain why the advance sales period loses its edge in boosting profit for the *Hotel* and *Travel* industries.

Table 6 The % of profit or revenue improvements per industry and the % of service providers that see profit improvements when switching from using the "instantaneous" discount to online deals. In the last three columns, we present the normalized profit sensitivity with respect to the discount and time sensitivities β_1, β_2. The last column reports the ratio of the two profit sensitivities over customers' responsiveness.

Service Category	median % improvements in profit	% service providers see profit improvement
Beauty	35.4%	59.7%
SPA	13.5%	67.0%
Casual Dine-in	12.1%	98.6%
Sit-down (Low)	22.1%	98.1%
Sit-down (High)	37.7%	100%
Hotel	-30.2%	7%
Indoor	11.6%	53.8%
Outdoor (Gate)	43.9%	98.6%
Photography	26.2%	64.2%
Travel	-14.9%	0.0%

Furthermore, we find that the non-zero advance sales period successfully helps service providers to mitigate the problems of being forced to launch extreme discounts, which helps us to reconcile the popularity of deals with the criticism on the deep discounts and explain the profit boosts. This can be seen from Figure 5, that unlike the optimal discount level of the online deal which nicely spreads out in $(0,1)$, the distribution of the optimal "instantaneous" discount has modes that are close to the boundaries of $(0,1)$. Service providers are either being over-conservative by bringing in discount customers or being over-aggressive in expanding the market via deep discounts. However, extreme discounts are often not desirable. In the short term, deep discounts of online vouchers are profitable only under narrow conditions (Edelman et al. (2016)), while in the long run, they might hurt a service provider's brand image (Kumar and Rajan (2012)). On the other hand, not

Figure 5 Histogram of the optimal discount under two business models, online deals and "instantaneous" price discounts. Online deal discounts have a nicely spread-out shape over $(0,1)$, while in contrast, price discounts see extreme values at their optimal points.

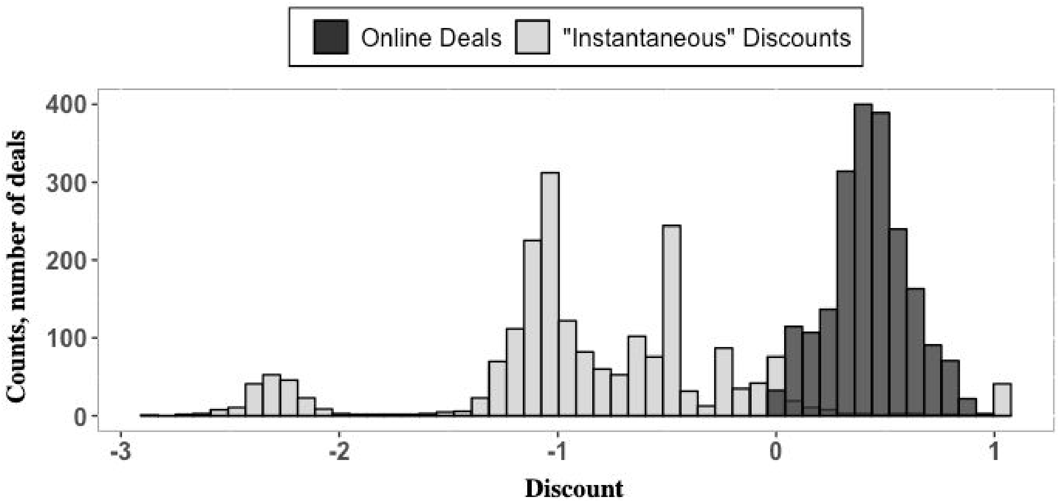

launching any promotional strategy on such platforms will reduce a service provider's visibility to a potentially larger market, which is not desirable either.

7. Conclusion

In this paper, using a comprehensive deal-level dataset from the two leading deal platforms in China, we empirically investigate whether and how online deals are designed to maximize profits under capacity constraints and predictable demand fluctuations. Although popular in practice, whether service providers can benefit from using online deals is unclear. Therefore, adding to the large stream of literature on electronic customer offers, we quantify the operational value of this new business model – the online deal.

We compare and complement the literature by confirming that the advance sales period is another effective demand management tool apart from the well-studied classic lever, the discount. We contribute the first study on the optimal launch time for promotions in the service industry from an operations perspective. We show that the launch time is relevant and important to manage demand. Relevant, because from estimating the structural model, we find that, the advance sales period is also an effective demand-shaping lever. Important, because from the counterfactual analysis, we see that the advance sales period mitigates unwanted extreme discounts, which occur when promotional strategies are designed to be “instantaneous.”

The importance of the advance sales period is further substantiated by revealing the two profit-boosting operational roles that it implements jointly with the discount. In particular, tailoring to their operating margins and demand swings, service providers alter the two levers to achieve both demand mean smoothing and variance-related cost reduction. Specifically, 0.088 out of 0.1 unit of increase in the demand swing can be smoothed out by re-optimizing the discount and the advance sales period, while a 0.1 increase in the cost-to-price ratio can be compensated by re-optimizing the advance sales period and thus decreasing variance-related costs by 0.15 unit. Furthermore, we see that an average of 23.6% profit gain can be achieved when service providers use both the discount and the advance sales period as levers to manage demand instead of using solely discounts.

Our work presents the two operational roles online deal play that can benefit the small service providers and particularly sheds light on the importance of strategically selecting the timing to launch deals. For the small service providers studied in our project, it is better to launch promotions that leverage both the discount and the advance sales period.

References

Analysys (2018) Yearly report: Chinese o2o market (2017). URL `http://www.199it.com/archives/683935.html`.

Baardman L, Cohen MC, Panchamgam K, Perakis G, Segev D (2018) Scheduling promotion vehicles to boost profits. Management Science 65(1):50–70.

Cachon GP, Daniels KM, Lobel R (2017) The role of surge pricing on a service platform with self-scheduling capacity. Manufacturing & Service Operations Management 19(3):368–384.

Chandon P, Wansink B, Laurent G (2000) A benefit congruency framework of sales promotion effectiveness. Journal of marketing 64(4):65–81.

Cohen MC, Leung NHZ, Panchamgam K, Perakis G, Smith A (2017) The impact of linear optimization on promotion planning. Operations Research 65(2):446–468.

Dana JD Jr (1998) Advance-purchase discounts and price discrimination in competitive markets. Journal of Political Economy 106(2):395–422.

Dianping (2018) A guidebook to self-service deal launching. URL http://evt.dianping.com/merchant/%E5%95%86%E6%88%B7%E8%87%AA%E5%8A%A9%E4%B8%8A%E5%8D%95%E6%95%99%E7%A8%8B.pdf.

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

Feldman P, Li J, Tsai HT (2017) Welfare implications of congestion pricing: Evidence from sfpark .

Gale IL, Holmes TJ (1993) Advance-purchase discounts and monopoly allocation of capacity. The American Economic Review 135–146.

Girotra K, Marinesi S, Netessine S (2013) Can groupon save its business model? Harvard Business Review URL https://hbr.org/2013/12/can-groupon-save-its-business-model.

Goltz J (2010) Doing the math on a groupon deal. URL https://boss.blogs.nytimes.com/2010/11/23/doing-the-math-on-a-groupon-deal/.

Grewal D, Krishnan R, Baker J, Borin N (1998) The effect of store name, brand name and price discounts on consumers' evaluations and purchase intentions. Journal of Retailing 74(3):331 – 352, ISSN 0022-4359, URL http://dx.doi.org/https://doi.org/10.1016/S0022-4359(99)80099-2, research Perspective on Retail Pricing.

iiMedia Research (2013) Yearly report: O2o market in china (year 2012). URL http://www.iimedia.cn/36436.html.

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

Kardes FR, Posavac SS, Cronley ML (2004) Consumer inference: A review of processes, bases, and judgment contexts. Journal of Consumer Psychology 14(3):230 – 256, ISSN 1057-7408, URL http://dx.doi.org/https://doi.org/10.1207/s15327663jcp1403_6.

Krishna A, Zhang ZJ (1999) Short- or long-duration coupons: The effect of the expiration date on the profitability of coupon promotions. Management Science 45(8):1041–1056, ISSN 00251909, 15265501, URL http://www.jstor.org/stable/2634806.

Kumar V, Rajan B (2012) Social coupons as a marketing strategy: a multifaceted perspective. Journal of the Academy of Marketing Science 40(1):120–136, ISSN 1552-7824, URL http://dx.doi.org/10.1007/s11747-011-0283-0.

Larsson R, Bowen DE (1989) Organization and customer: managing design and coordination of services. Academy of Management Review 14(2):213–233.

Lazarev J (2013) The welfare effects of intertemporal price discrimination: an empirical analysis of airline pricing in us monopoly markets. New York University .

Lee FSJ (2013) Hospitality products and the consumer price–perceived quality heuristic: An empirical perspective. Services Marketing Quarterly 34(3):205–214, URL http://dx.doi.org/10.1080/15332969.2013.798194.

Lee I (2014) A collaborative management of social commerce deals: Decisions on pricing and commission rate with a capacity consideration. Journal of Revenue and Pricing Management 13(3):233–246, ISSN 1477-657X, URL http://dx.doi.org/10.1057/rpm.2013.30.

Lee I, Lee K (2012) Social shopping promotions from a social merchant's perspective. Business Horizons 55(5):441–451.

Leslie P (2004) Price discrimination in broadway theater. RAND Journal of Economics 520–541.

Liu L, Sun S (2017) Expenditure during spring festival. URL http://news.xinhuanet.com/fortune/2017-02/05/c_1120411585.htm.

Lovelock CH (1983) Classifying services to gain strategic marketing insights. The Journal of Marketing 9–20.

Luo X, Andrews M, Song Y, Aspara J (2014) Group-buying deal popularity. Journal of Marketing 78(2):20–33.

Marinesi S, Girotra K, Netessine S (2018) The operational advantages of threshold discounting offers. Management Science 64(6):2690–2708, URL http://dx.doi.org/10.1287/mnsc.2017.2740.

Meituan-Dianping (2017) Meituan-dianping prospectus. URL https://www1.hkexnews.hk/listedco/listconews/sehk/2018/0907/ltn20180907011.pdf.

Mesak HI, Zhang H, Pullis JM (2010) On optimal service capacity allocation policy in an advance selling environment in continuous time. European Journal of Operational Research 203(2):505–512.

Narasimhan C (1984) A price discrimination theory of coupons. Marketing Science 3(2):128–147, ISSN 07322399, 1526548X, URL http://www.jstor.org/stable/183747.

Ng IC, Wirtz J, Sheang Lee K (1999) The strategic role of unused service capacity. International Journal of Service Industry Management 10(2):211–244.

Nguyen TH, Wright M (2015) Capacity and lead-time management when demand for service is seasonal and lead-time sensitive. European Journal of Operational Research 247(2):588–595.

Reimers I, Xie C (2018) Do coupons expand or cannibalize revenue? evidence from an e-market. Management Science .

Sasser WE (1976) Match supply and demand in service industries. Harvard Business Review 54(6):133–140.

Silva-Risso JM, Bucklin RE, Morrison DG (1999) A decision support system for planning manufacturers' sales promotion calendars. Marketing Science 18(3):274–300.

Song M, Park E, Yoo B, Jeon S (2016) Is the daily deal social shopping?: An empirical analysis of customer panel data. Journal of Interactive Marketing 33:57 – 76, ISSN 1094-9968, URL http://dx.doi.org/https://doi.org/10.1016/j.intmar.2015.12.001.

Stafford MR, Kilburn AJ, Allen CA (2006) The use of pricing information in service advertisements: A comparison of retail versus professional services. Services Marketing Quarterly 27(1):1–22.

Tellis GJ, Zufryden FS (1995) Tackling the retailer decision maze: Which brands to discount, how much, when and why? Marketing Science 14(3):271–299.

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

Ye M, Sandholm T, Wang C, Aperjis C, Huberman BA (2012) Collective attention and the dynamics of group deals. Proceedings of the 21st International Conference on World Wide Web, 1205–1212 (ACM).

Yılmaz Ö, Pekgün P, Ferguson M (2017) Would you like to upgrade to a premium room? evaluating the benefit of offering standby upgrades. Manufacturing & Service Operations Management 19(1):1–18.

Figure 6 **Clustering analysis by service prices on the industry level. The x-axis represents the number of clusters; the y-axis represents the within-cluster sum of squares after clustering. The dotted vertical line on each panel indicates the optimal number of clusters chosen for the specific industry.**

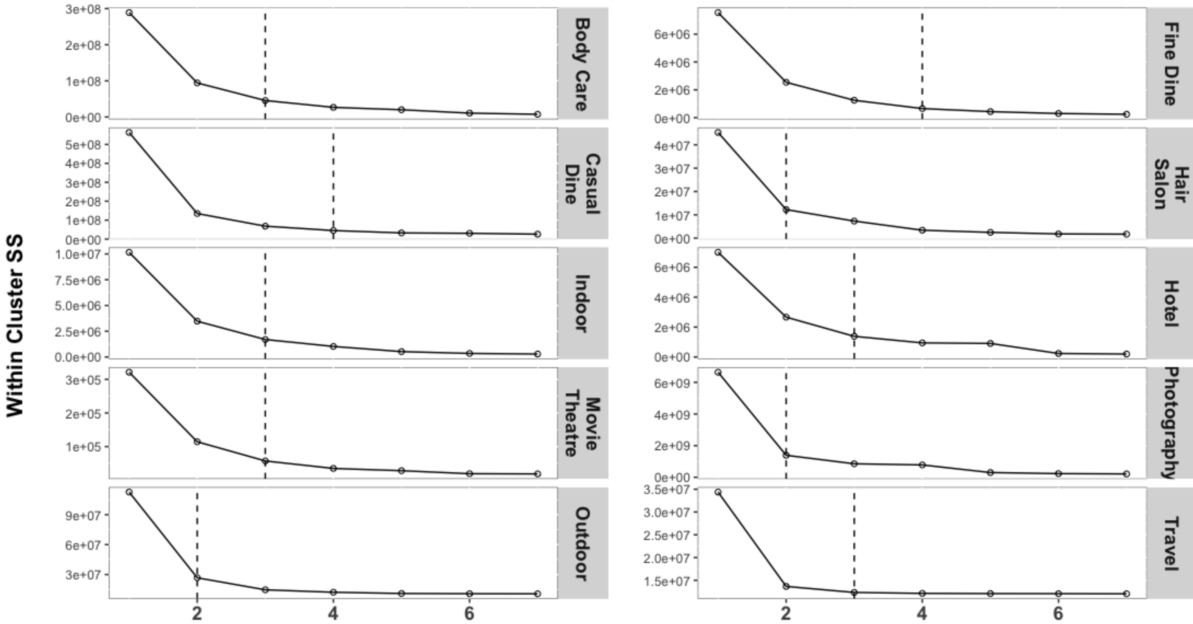

Appendix A: Clustering analysis by services' original prices

The Food and Restaurants are the largest group. We first adopt a sub-categorization suggested by the platform and split the group into *Fine Dine* and *Casual Dine* for further analysis.

We cluster the service deals by their original prices using k-means clustering to achieve homogeneity in the within-cluster operating margins of deals, $p - c$. The optimal number of clusters is determined by the elbow method though SCREE plots as shown in Figure 6.

The elbow method suggests that the turning point that makes an angle on the SCREE plot is a good choice of the number of optimal clusters. We take the variance explained by the between-cluster sum of squares into consideration as well - see Table 7. The larger the percentage of data variance shown in the between-cluster sum of squares, the more homogeneous the price within the clusters. For the categories with a larger sample size, say *Casual Dine-in* and *Hotel*, we break the category into more clusters by price to get more homogeneous operating margins within the clusters, which ensures a better structural model estimation. The final grouping of categorizations shown in Table 3 in Appendix D is obtained by first clustering by original prices and then eliminating the clusters with less than 45 deals. It turned out that, even with a less-than-ideal number of clusters, *Movie Theater* and *Hair Salon* did not have a large enough sample size to conduct further analysis and were thus removed entirely from any further estimation. All small clusters that have less than 45 holiday deals are all dropped.

Table 7 The percentage of variance explained by between-cluster variations on the industry level. Greater variance means greater homogeneity in service prices within clusters.

Service Category	Body Care	Casual Dine	Indoor	Movie Theatre	Outdoor
Variance explained by between-cluster SS	84.2%	92.0%	82.8%	75.4%	76.2%
Service Category	Fine Dine	Hair Salon	Hotel	Photography	Travel
Variance explained by between-cluster SS	91.3%	73.0%	80.0%	79.0%	91.0%

Appendix B: Anticipated demand shocks: a key determinant of strategic deal design

We first model the customer demand function and then tease out those demand changes contributed solely by the holiday effect.

B.1. The demand function

We use a second order spline model with a linear knot to estimate the customer demand over the life of both holiday and non-holiday deals. For a deal i in category j, the sales volume of deal i in category j up to day t is denoted as $S_{i,j,t}$ - see (11). That is, we model the accumulative demand $S_{\cdot,\cdot,t}$ as a baseline demand accumulated over time t with a holiday demand shift[12], if any, and further adjusted by the price effects and other fixed effects of the online deal. We now proceed to explain each term in the demand function.

$$\begin{aligned} S_{i,j,t} = &\underbrace{\theta_{j,1} + \theta_{j,2} t_{i,j} + \theta_{j,3} t_{i,j}^2}_{\text{baseline demand}} + \underbrace{\theta_{j,h} t_{i,j} \mathbf{I}\{t_{i.j} \leq T_{i,j}\} + \theta_{j,h} T_{i,j} \mathbf{I}\{t_{i,j} > T_{i,j}\}}_{\text{holiday demand shift}} + \\ &\underbrace{\theta_{j,p,h} p_{i,j} * h_{i,j} + \theta_{j,p,nh} p_{i,j} * h_{i,j} + \theta_{j,\alpha,h} \alpha_{i,j} * h_{i,j} + \theta_{j,\alpha,nh} \alpha_{i,j} * h_{i,j}}_{\text{price effects on demand}} + \\ &\underbrace{\theta_{j,R} R_{i,j,t} + \theta_{j,M} M_{i,j} + \theta_{j,F} F_{i,j}}_{\text{other fixed effects}} + \varepsilon_s \end{aligned} \tag{11}$$

[12] Holiday deals are indexed by the subscription h, while non-holiday deals are indexed by nh.

The *baseline demand* is quadratic in time, which is increasing and concave in the days elapsed since the deal is launched, denoted as $t_{i,j}$. The holiday effect imposes a *constant shift*, either jump or drop,[13] on top of the baseline demand. The holiday demand shift is not assumed to be identical across categories, which is captured by the holiday effect coefficient $\theta_{j,h}$. The holiday effect vanishes once the holiday starts, which is $T_{i,j}$ days after the deal launch.[14]

Apart from the time effect, customer demand is also affected by price effects such as competition as well as other non-deal-related fixed effects. We do not assume that customers respond to price or discount in the same way across holiday and non-holiday periods. This is captured by two price sensitivities, $\theta_{j,p,h}, \theta_{j,p,nh}$, and two discount sensitivities, $\theta_{j,\alpha,h}, \theta_{j,\alpha,nh}$. $h_{i,j}$ is the binary variable indicating whether the deal is a holiday deal. Additionally, each online deal faces competition from other similar deals that are active during the same time, which is captured in the coefficient $R_{i,j,t}$. Lastly, we account for the intrinsic difference in customer behavior across cities, $M_{i,j}$, and platforms, $F_{i,j}$. Note that we use holiday-driven sales changes as a proxy for the natural changes in holiday-redeemed demand since the redemption information is unavailable.

B.2. Identifying the holiday-induced demand changes

In the construction of the demand function, holiday demand shift coefficients $\theta_{\cdot,h}$ capture the industry-specific holiday demand swings. We normalize the shift by the corresponding industry quadratic baseline demand. The normalized demand change is free from scales. We formally state the estimation strategy of the predictable demand change in Proposition 3.

PROPOSITION 3 **(Nominal Demand Change)**. *For industry j, define the nominal demand jump $\delta_{D,j}$ as*

$$\delta_{D,j} = \frac{\theta_{j,h} T}{\theta_{j,1} + \theta_{j,2} T + \theta_{j,3} T^2} \tag{12}$$

The standard deviation of the nominal demand change is

$$\sigma(\widehat{\delta_{D,j}}) = \sqrt{Var[\frac{\Delta_{D,j}}{D_{b,j}}]} = \sqrt{\mathbb{E}[Var(\frac{\Delta_{D,j}}{D_{b,j}}|T)] + Var[\mathbb{E}(\frac{\Delta_{D,j}}{D_{b,j}}|T)]} \tag{13}$$

where the baseline demand is $D_{b,j} \triangleq \theta_{j,1} + \theta_{j,2} T + \theta_{j,3} T^2$, $\Delta_{D,j} \triangleq \theta_{j,h} T$.

[13] Examples of demand drop include manicure services, which are not popular during a family-oriented holiday like *Spring Festival*.

[14] A graphical illustration of the baseline demand and its holiday demand shift is given in Appendix F.

Figure 7 **These two figures show the results from the Tukey-Kramer tests on all pairwise mean comparisons between industries. On the x-axis are the mean differences between paired industries and their** 90% **confidence interval. The y-axis indicates which pair of industries is being considered. is The red bars are the significant differences; the rest are non-significant.** 75% **of the discount mean comparisons are significant, while** 45% **of the advance sales period mean comparisons are significant.**

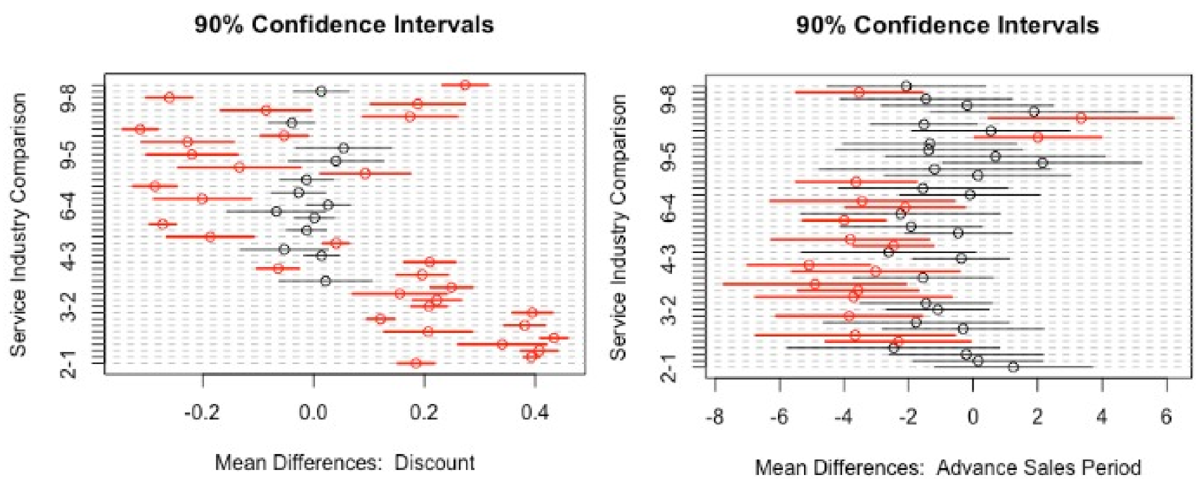

The identification of industry-specific nominal demand change $\delta_{D,j}$ *is achieved by identifying* $\theta_{j,h}$*,* $\theta_{j,1}$*,* $\theta_{j,2}$*, and* $\theta_{j,3}$ *from the demand function (11) on the industry level.*

We estimate the demand function (11) for each service category using deals' sales volumes on the 1st, 5th, and last day of their durations, $(1, S_1)$, $(5, S_5)$, (d, S_d) and then present the estimated holiday demand changes and their standard deviations in column 6 of Table 3.

Appendix C: Test of significantly different deal designs across service categories: the Tukey-Kramer method

We conclude that there are statistically significant differences in the mean of the discount and the advance sales period across industries, categorized in Table 3. We use the Tukey-Kramer method to conduct all pairwise mean comparisons (see Figure 7). This method is preferred over a series of two sample t-tests since it is more robust. Furthermore, since we have unequal sample sizes across industries, Tukey's method returns a conservative result. That is, the actual significance level controlled for the test is even smaller than what we initially set for multiple testing, $\alpha = 10\%$.

The significant mean comparisons show that deals are indeed uniquely designed across industries characterized by their operating margins and demand fluctuations. Note that the analysis is conducted using raw data, i.e. without controlling for cities, platforms, holidays, or other fixed effects that are not industry-related, which partially explains the variations within service industries (length of the error bar). The non-significance among some pairs suggests that they might have similar operating margins and demand fluctuations. These two industry characteristics determine deal strategies, not a general industry fixed-effect.

Appendix D: Summary statistics of the refined categories after price clustering

We present the summary statistics of the refined categories (after price clustering) in this section. The grouping of deals are shown in Table 3.

Appendix E: Interpreting β_1 and β_2

In this section, we show that the values of our estimated β_1 and β_2 align with the theoretical results of how customers respond to discounts and launch time across different service industries in the past literature.

Using discount as a lever to shape demand mean can induce two types of customer reactions, leading to opposing signs of discount sensitivity, β_1. First, customers' preferences over discount prices collectively result in a positive discount sensitivity β_1. Such preferences come from the utilitarian benefits (money saving) and hedonic benefits (self-perception of being a strategic shopper) that customers would gain from purchasing at a discount price (Chandon et al. 2000). Moreover, the higher the original price, the stronger the preference, and the larger the discount sensitivity in magnitude. On the other hand, deep discounts make customers suspicious about the service quality and variety (Grewal et al. 1998, Kardes et al. 2004), which leads to a negative discount sensitivity β_1. Such suspicion is observed in hospitality services, including hotels, restaurants, theaters, and outdoor trips (Lee 2013), as well as services whose quality are hard to evaluate, such as photography (Stafford et al. 2006, Larsson and Bowen 1989). Additionally, our data suggest that restaurant deals with deeper discounts often come with more limited variety. Our estimated β_1 appear to align with the theoretical results in literature (see Appendix E) in terms of how the joint effects of the two opposing customer responses would look across industries (see Table 4).

Customers' sensitivity to the advance sales period $\hat{\beta}_2$ mostly comes from online deals' advertising effect, which reflects as the positive $\hat{\beta}_2$ as shown in Table 4. The longer the deals are made available, the more discount customers are attracted to the service. The only exception is the photography service. One can argue that such exception happens because customers who seek for photography deals are avoiding a highly

Figure 8 The timeline of a typical online deal. We align the holiday and non-holiday baseline demand over time. The increments are identified as the holiday effect. We use the estimated holiday demand swings $\theta_h = 0.09$ to illustrate.

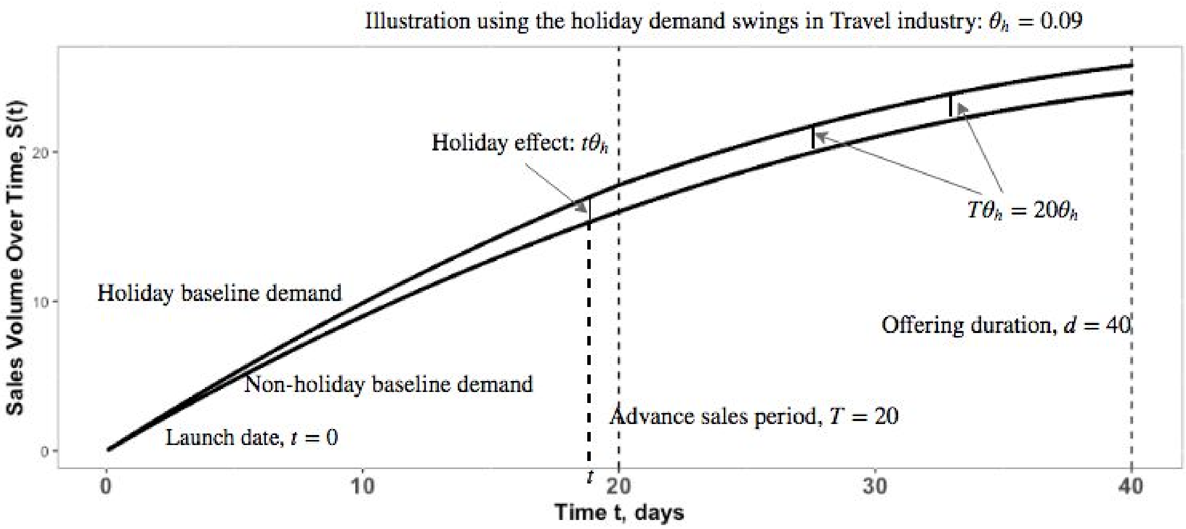

probable crowd during holidays by pre-holiday redemption. Customers form expectations of crowdedness based on large holiday demand fluctuation (see Table 3) and commonly-acknowledged high labor intensity (Lovelock 1983). Such strategic customer behavior to avoid crowdedness is not new to researchers (Marinesi et al. 2018, Song et al. 2016).

Appendix F: Graphical illustration of holiday demand swing identification

We showcase both holiday and non-holiday baseline demand as in Figure 8, where the distance between the two lines accumulated over time represents the accumulated holiday effect on sales volume, i.e. $t\theta_h$. We start from t=0, which is the launch date of the deal. The baseline demand is an increasing and concave quadratic function. Before the actual holiday, holiday effect θ_h adds up to the baseline demand per day. That is, on the tth day before the holiday, we see an accumulated holiday effect of $t\theta_h$. The daily boost in sales owing to the holiday effect vanishes after the holiday ends, $t = T = 20$. The total holiday effect on the sales volume for the specific deal sums up to and stays at $T\theta_h$. Essentially, the percentage increase in sales compared to the baseline demand over the advance sales period is defined as the holiday demand swings, that is $\frac{T\theta_h}{\text{baseline demand}}$, which is Expression (12).

Appendix G: Identification of the parameters

In this section we give a detailed proof of identification of the sensitivity coefficients and the marginal service cost parameter.

Proof Recall that we normalize the total population to be 1. From the optimal expressions (α^*, T^*), we can construct the following expressions. Recall that $\kappa = 1 + \delta_D$.

$$\frac{\beta_1}{\beta_2} = \frac{B_1}{A_1} \tag{14}$$

$$\frac{\beta_1}{\beta_2} A_2 + \frac{1}{\beta_2}\frac{\beta_1}{\beta_2}(\beta_5 c + \gamma) = B_2 \tag{15}$$

$$2\frac{\beta_1}{\beta_2}(\beta_5 c + \gamma) c = A_1 \tag{16}$$

$$2\beta_6 \kappa c + \frac{\beta_5 c + \gamma}{\beta_2} = A_2 \tag{17}$$

$$\frac{-(1-\beta_6)\kappa}{\beta_2} = B_4 - \frac{B_1}{A_1} A_3 \tag{18}$$

From (15) and (17) we can find

$$\frac{\beta_5 c + \gamma}{\beta_2} = \frac{A_1 B_2 - B_1 A_2}{B_1} = A_2 - 2\beta_6 \kappa c \Rightarrow \beta_6 = \frac{2A_2 B_1 - A_1 B_2}{2\kappa c B_1} \tag{19}$$

Also, we have $\beta_5 c + \gamma = \frac{A_1 B_2 - B_1 A_2}{B_1}\beta_2$; substituting into (16), we then have,

$$c = \frac{A_1^2}{2}\frac{1}{(A_1 B_2 - B_1 A_2)\beta_2} \overset{(18)}{=} \frac{A_1^2}{2}\frac{B_4 A_1 - B_1 A_3}{(A_1 B_2 - B_1 A_2)(-A_1 \kappa(1-\beta_6)}$$

substitute the above expression to equation (19).

$$c = \frac{2B_1 A_2 - A_1 B_2}{2B_1 \kappa} - \frac{B_4 A_1^2 - A_1 B_1 A_3}{2\kappa(A_1 B_2 - B_1 A_2)}$$

Subsequently, we can work out β_6 by substituting c back to (19). Similarly, we have $\beta_2 = \frac{-A_1\kappa(1-\beta_6)}{B_4 A_1 - B_1 A_3}$, $\beta_5 c + \gamma = \frac{A_1 B_2 - B_1 A_2}{B_1}\beta_2$, $\beta_1 = \frac{B_1}{A_1}\beta_2$, $\beta_3 = \beta_2 B_3$, $\beta_7 = \frac{A_3}{-2c}$.

Appendix H: The Irrelevant σ

Note from Appendix G that we do not point identify σ, β_5, γ. In this section, we will show that the percentage improvements in profit given in Appendix **??** are lower bounds of the profit improvements regardless of the choice of σ, β_5, γ. That is, our results holds regardless of whether these three parameters can be point identified or not.

The percentage increase in profit follows the expression below.

$$\frac{\mathbb{E}[V(\alpha^*, T^*)] - \mathbb{E}[V(\tilde{\alpha}^*)]}{\mathbb{E}[V(\tilde{\alpha}^*)]} = \frac{\alpha^* p\mu_1 - \tilde{\alpha}^* p\tilde{\mu_1}^* + p\mu_2^* - p\mu_2 - c(\mu_1^* + \mu_2^*)^2 + c(\tilde{\mu_1}^* + \mu_2)^2 - (\beta_5 c + \gamma)T^*}{\tilde{\alpha}^* p\tilde{\mu_1}^* + p\mu_2 - c(\tilde{\mu_1} + \mu_2)^2 - 2c\sigma^2} \tag{Sec. 6}$$

Note that we only need to identify $(\beta_5 c + \gamma)$ as a whole. Also, we set $\sigma = 0$ when calculating the percentage improvements. It is easy to see that the percentage increase calculated when $\sigma = 0$ is the strict lower bound of the exact percentage improvement.