

Setting Prices on Priceline

Chris K. Anderson

School of Hotel Administration, Cornell University, Ithaca, New York 14853,
canderson@cornell.edu

Priceline is best known for its *name-your-own-price* format, in which consumers bid for services but not for service providers. Because Priceline serves as an opaque selling mechanism, it attracts price-conscious consumers. Sellers also benefit because they can price into multiple market segments without worrying that they are diluting revenue they might receive from customers who are willing to use conventional selling channels and pay more. A firm that releases its inventory to Priceline must manage the trade-off of pricing its inventory too low (and forgoing revenue) versus pricing it too high and forgoing a sale. In this paper, we outline the mechanism that Priceline uses to determine if customer bids are successful and, given this mechanism, establishes optimal prices and inventory allocations for Kimpton Hotels.

Key words: dynamic programming; marketing; pricing, buyer behavior.

History: This paper was refereed. Published online in *Articles in Advance* June 10, 2009.

Kimpton develops, owns, or manages 43 independent boutique lifestyle hotels in the United States and Canada; its website, <http://www.kimptonhotels.com>, includes a complete listing of its hotels and details on each property. Kimpton takes a unique approach to style and four-star service. Each of its hotels has a distinctive energy and personality, and offers distinguished restaurants and lounges, complimentary wireless Internet access, hosted evening wine hours, and great rates.

Many are located in urban business districts, and most Kimpton guests are business travelers. Unlike leisure travelers, business travelers tend to make travel plans close to their planned arrival date—generally fewer than 14 days before arrival. This short reservations window means that Kimpton has little time to recover from weak demand periods or periods in which reservations fall substantively below forecasted demand. However, because many of its guests are business travelers, Kimpton cannot offer the large discounts that stimulate demand; doing so would decrease rate yields from these business travelers who would have booked a reservation without the discount and are willing to use traditional booking channels. Priceline's opacity (i.e., consumers do not know the service provider, such as the airline or hotel, until they have paid for the service) provides

a natural method of segmenting business and leisure travelers. Thus, it allows Kimpton to price into the leisure market while not diluting its business market.

Kimpton releases rooms to Priceline to help increase occupancy on days when forecasted occupancy is low (i.e., <90–95 percent). However, it frequently restricts the number of rooms released to maintain customer perceptions within the marketplace; i.e., it puts a cap on the number of rooms released, regardless of forecasted occupancy or revenue implications. Currently, Kimpton uses prices recommended by its vendor-supplied revenue-management system to determine how to set prices on Priceline. In the following sections, we outline our efforts to improve Kimpton's Priceline price-setting process. We discuss how Priceline's name-your-own-price (NYOP) mechanism works, develop a model for setting optimal Priceline prices, and discuss an implementation at a Washington, DC-based Kimpton hotel. Our model develops a price-dependent arrival process and extends the framework of Lee and Hersh (1993) to include both setting optimal prices and allocating inventory.

Online Pricing and Priceline.com

Pricing services, such as hotel rooms, rental cars, and airline seats, online has dramatically changed how

service firms reach customers; online travel sales were expected to exceed offline (i.e., traditional) sales channels by 2008 (2008 data are not yet available). Service providers initially reacted positively to pricing online; they believed that it would provide them with new channels to reach customers and thus enable increased segmentation opportunities. However, over time these service providers have tried to move customers back to firm-managed distribution channels (e.g., company websites and call centers) to control sales costs and commissions, and maintain direct contact with their customers to facilitate loyalty programs and other marketing efforts.

In parallel with this migration to firm-managed channels, service providers have tried to streamline prices and create price parity (i.e., equivalent prices for all booking methods) across all distribution channels. They believe that price parity instills trust and comfort in customers because they do not need to shop around for the best prices by a given service provider. From the consumer's viewpoint, online pricing and online travel agents, such as Orbitz, Expedia, and Travelocity, and metasites, such as Kayak, have greatly simplified shopping because they allow easy price comparisons. The opaque nature of Priceline's NYOP mechanism allows a service provider to effectively post different prices simultaneously while maintaining price parity on full-information-posted price channels. Service firms look to use Priceline's NYOP process to stimulate price-sensitive demand while not cannibalizing their existing sales. Cannibalization is limited because, although prices on Priceline are much lower than regular posted prices, the customer does not select the service provider, and the purchased product includes more restrictions (e.g., it might not be refundable or transferable) than products purchased from other distribution channels.

Priceline presents two mechanisms to allow customers to acquire services; the first is similar to other online travel agents in which a customer selects from posted prices by multiple service providers; the second allows customers to submit bids for services without knowing the service provider's identity (or certain details of the service). In this paper, we concentrate on the second Priceline mechanism.

Consider a customer who submits a bid for an overnight stay at a three-star hotel in New York's

midtown Manhattan, with an arrival date of November 1st. Priceline determines if it can meet the customer's offer. If it rejects the offer, the customer can alter the bid by changing an attribute, such as changing three-star to four-star, midtown to Times Square, or the arrival date. Alternatively, the customer could rebid on the three-star hotel in midtown after 24 hours have elapsed. Priceline's method of determining if a bid is acceptable is unique and greatly favors the service provider. Following a customer bid p_c , Priceline creates a list of all n qualifying properties, e.g., all three-star or better hotels in midtown Manhattan. It then randomly selects one of these qualifying properties (each has an equally likely probability of being selected). Priceline then checks the prices that the hotel has loaded; each property is allowed to load three rate classes (all prices are loaded into a global distribution system). If the selected property has a price p that is lower than the consumer's bid price p_c , then a transaction occurs. The hotel receives p , the customer pays p_c , and Priceline keeps $p_c - p$ as its share. If the property has more than one room rate that is less than the customer's bid, the property receives the highest price minus p_c . For example, if a customer bids \$100 and the selected property has loaded rates of \$100, \$90, and \$70, the customer pays \$100, the property receives \$90, and Priceline receives \$10.

If the property selected in this "first round" does not have a price such that $p < p_c$, then Priceline goes back to the $n - 1$ remaining properties; again, it randomly selects a property and checks prices. In this second round, each property's probability of being selected depends upon its success rate in the first round, i.e., the percentage of time when, based on random selection, it has a price that is less than the customer's bid. A property that, when selected in the first round, has a rate in which $p < p_c$ 50 percent of the time would have twice as likely a chance of being randomly selected in the second round as one that has a rate in which $p < p_c$ 25 percent of the time. Priceline calculates this success rate (or batting average) based only in the first round, not on subsequent rounds. It repeats rounds until it finds a property that meets the customer's bid or exhausts all properties, resulting in no sale. This mechanism favors the property because it selects the highest price that yields Priceline

a profit. In addition, the random nature of property selection does not require the properties to compete with each other on price; they only compete with the customer because a firm's price relative to that of another firm does not impact its probability of being selected (at least in the first round).

Opaque selling and the NYOP mechanism have begun to generate interest in the academic literature, most of which has focused on its structure. Fay (2004) develops a stylized model of a monopolist firm that uses an NYOP channel and investigates whether repeat bidding should be allowed. Although Priceline does not allow repeat bidding within a 24-hour period, there are numerous methods to circumvent this limitation (<http://www.BiddingforTravel.com> shows examples). Fay (2004) indicates that partial repeat bidding, i.e., repeat bidding by knowledgeable customers, might be less profitable than complete repeat bidding. Hann and Terwiesch (2003) use data from a European NYOP retailer to investigate consumer transaction costs (i.e., the cost of resubmitting bids) of using a repeat-bidding NYOP channel. In a related paper, Spann et al. (2004) use data from a German NYOP seller of flights between Germany and Spain to investigate consumer frictional or transaction costs and consumer willingness to pay.

Wang et al. (2005) develop a game-theoretic model of a supplier that uses both regular posted-price full-information channels and NYOP channels to reach heterogeneous customers. They develop a two-stage game in which suppliers set posted prices in period 1, observe demand in period 1, and set minimally acceptable prices at the NYOP channel in period 2. Posted prices are rigid in period 2. Consumers observe posted prices in period 1 and then decide to buy or bid in period 2. Because of the rigidity of posted prices and the demand uncertainty, the NYOP channel generates increased revenues for the service provider.

Wilson and Zhang (2008) look at a retailer that sets prices on an NYOP channel. They develop ϵ optimal policies for the retailer to encourage customers to bid their maximum reservation price. In this paper, we look at a less-stylized situation in which the retailer posts multiple prices and must allocate inventory across these price classes when customers are arriving randomly.

Related research looks at opaque selling in which prices are posted; however, some aspect of the service or service provider is hidden from customers. Jiang (2007) develops a Hotelling model (Hotelling 1929) to illustrate how a firm should price on regular full-information channels versus opaque channels. Jiang (2007) indicates that opaque selling can be Pareto-improving for both customers and suppliers when customers are differentiated based on their willingness to pay. Jerath et al. (2007) also develop a Hotelling model of opaque pricing. In their model, two firms of equal capacity offer a differentiated service using three channels: regular posted price, posted last-minute sales, and last-minute sales through an opaque intermediary. Their objective is to investigate the market conditions under which a firm should directly offer last-minute discounts versus offering those discounts through an intermediary. Jerath et al. (2007) relax the posted-price rigidity of Wang et al. (2005) by introducing direct last-minute discounts. They conclude that direct last-minute selling is preferable to using the opaque intermediary if consumer valuations are high or the service offerings are relatively homogeneous.

Priceline gives daily activity reports to the service providers who provide it with their prices and inventory. In the following sections, we outline these activity reports and use the information in the reports to develop optimal prices for firms to release to Priceline. We use data from hotels; however, we could generalize our method to all services that release inventory via Priceline's NYOP mechanism.

Modeling Approach

Understanding Bids on Priceline

In this section, we describe the development of a model based on the customer bid data that Priceline sends to firms each day. We use the bid data to create demand distributions, which are a function of price. In subsequent sections, we then structure a dynamic program to determine the optimal prices to post on Priceline, and the inventory allocation across these prices. We illustrate our model development using data from a Kimpton hotel in Washington, DC.

A customer bids on a room at Priceline by selecting a city and a subarea within that city (e.g., New York

and its subarea, Times Square), and a property quality (e.g., two-star), and places a bid. Priceline determines if it has a willing service provider at the requested price. If the bid is successful, Priceline notifies the bidder of the specific property and charges the customer's credit card. If the bid is unsuccessful, the customer may rebid on the same item (city, subarea, star quality, and arrival and departure dates) after 24 hours; the customer can also alter the bid by changing any of the above features and rebid instantly. Each property receives a daily bid report for the specific market (e.g., city, subarea, and star quality) in which the property is located (Figure 1).

The bid report shows all bids made the prior day for all future arrival days (including the prior day). The

report has a line for each bid posted by each customer. It indicates how many days the customer requested to stay (Length of stay), the number of rooms requested, the hotel's lowest rate (Your Priceline rate), the customer's bid (Offer price), and whether a competitor accepted the bid (Booked by others). For example, line 1 shows that a customer wanted to check in on 9/3 and stay for two nights; the hotel rejected the customer's \$52 bid (its price was \$63), but a competitor accepted the bid. Firms can then combine these reports from multiple days and use them to create a targeted Priceline pricing strategy.

Arrival Distribution

Priceline can sort bids by arrival day and then by the number of days prior to arrival (DBA) that the bid was

Attn: Revenue manager
Director of sales
General manager

DAILY DEMAND REPORT ALL CURRENCIES IN \$USD

	Date	Offered RN	Booked RN	Offered revenue	Booked revenue	Conversion	Avg. offer price	Avg. booked price
OPAQUE	09/03/07	108		\$6,878			\$64	

MISSED OPPORTUNITIES ON 9/3/07

Room Nights = 108

Revenue = \$6,828

Check in date	Check in day of week	Length of stay	Number of rooms	Your Priceline rate* (\$)	Offer price (\$)	Declined revenue (\$)	Reason**	Booked by others
09/03/07	Mon	2	1	63	52	104	HI	Y
09/05/07	Wed	3	1	195	45	135	HI	
09/06/07	Thu	1	1	179	44	44	HI	
09/07/07	Fri	2	1	65	48	96	HI	Y
09/08/07	Sat	4	1	130	88	352	HI	
09/09/07	Sun	4	1	162	88	352	HI	
09/09/07	Sun	4	1	162	66	264	HI	
09/13/07	Thu	7	1	253	46	322	HI	
09/14/07	Fri	4	1	217	84	336	HI	
09/15/07	Sat	4	1	300	66	264	HI	
09/16/07	Sun	2	1	332	70	140	HI	
09/16/07	Sun	3	1		70	210	NR	
09/16/07	Sun	3	1		70	210	NR	
09/16/07	Sun	1	1	195	61	61	HI	

Figure 1: The table shows a sample daily report of customer bid data from Priceline.com.

*If your Priceline rate changes during the requested length of stay, we will display the average of the available Priceline rates.

**HI, your Priceline rate was higher than the customer's offer price; NR, your hotel appeared sold out and did not display any rates. (Reasons not shown: NP, your hotel did not display a Priceline rate though other non-Priceline rates were available; RR, a nonqualified rate equal to or less than the Priceline rate was available at your property.)




Figure 2: The graph shows the distribution of the average number of bids placed at Priceline by days prior to arrival.

placed, thus creating arrival distributions. Figure 2 displays a sample arrival distribution; it shows the average number of bids placed for Monday arrival dates for up to 14 days prior to arrival. Demand for services is seasonal; each day of the week (DOW) and DBA have potentially different demand levels.

Given the discrete nature of the number of bids placed, we use the Poisson distribution to model the frequency of bid activity. Modeling bids by DBA as Poisson distributions, we have

$$f(x | \lambda) = \frac{\lambda^x e^{-\lambda}}{x!}, \quad (1)$$

where x is the number of bids placed and λ is the average number of bids placed for the arrival date in question. Each DBA will have a different bid distribution because the number of bids placed usually increases as the arrival date approaches. Similarly, we model weekends (Friday and Saturday arrivals) as different from weekdays because more business travelers book weekday stays and more leisure travelers book weekend stays. This segmentation in bid frequency will be very service specific, with some hotels or services allowing more or less aggregation of data, i.e., potentially treating each weekday as its own segment. The number of bids placed does not necessarily reflect the number of customers because allowing customers to rebid 24 hours after an unsuccessful bid could cause a demand overestimation.

	DBA							
	0	1	2	3	4	5	6	7
Weekend	12.3	13.3	8.4	6.3	5.5	3.9	2.4	1.7
Weekday	8.5	8.5	4.4	3.1	2.3	2.2	1.8	1.7

Table 1: The table data show the average number of bids placed per day by days prior to arrival, and separated by weekends and weekdays.

Table 1 displays mean daily arrival rates for weekday and weekend arrival dates as a function of the number of days prior to arrival that the bids are placed for this same property.

Bid Distribution

Similar to how they track the number of bids made, firms can track the prices of the actual bids. Figure 3 displays a histogram of bids posted by customers for same-day (DBA = 0) weekend arrival; Figure 4 displays a similar histogram for a weekday arrival. Potentially, each DOW will have a different bid distribution by DBA, thus reflecting customers who will potentially bid lower the further out they are from their planned arrival date (because they will have opportunities to rebid if unsuccessful). Kimpton’s pricing strategy might induce changes in distributions of bids because consumers might use secondary sites (e.g., <http://www.BiddingforTravel.com>) to communicate bid successes and failures; therefore, the hotel might need to update these distributions in the future as it collects new data that reflect potential changes in bid behavior.




Figure 3: The graph shows the distribution of prices for customer bids on Priceline for a weekend arrival at a Kimpton hotel.




Figure 4: The graph shows the distribution of prices for customer bids on Priceline for a weekday arrival at a Kimpton hotel.

Figures 3 and 4, in addition to showing the frequency of bids by price, also display gamma probability density functions that are suitable to the data and use maximum-likelihood estimators.

Demand as a Function of Price

In the following section, we develop demand distributions using the data as summarized in the *Understanding Bids on Priceline* section. Let the size of the customer base interested in placing a bid for a room for a particular arrival date and particular DBA be represented by the random variable N following a Poisson distribution. Let F_B , following a gamma distribution, represent the distribution function of customer bids for a randomly selected customer posting a bid (for the same arrival date and DBA). If the service provider posts a price p , then anyone with a bid greater than p would secure the room. The probability that a randomly selected customer is willing to bid more than p is $1 - F_B(p)$. The bid distribution $F_B(\cdot)$, similar to the arrival distribution, might have separate distributions by DOW and DBA because customers might bid less farther in advance, and bid higher for weekdays (assuming business demand) than for weekends (leisure demand). In our example, we use separate bid distributions for weekends and weekdays, but not by DBA, because we are currently looking only at setting prices for seven days prior to arrival.

To simplify the dynamic programming formulation in the *Determining Optimal Room Rates* section, similar to Gerchak et al. (1985) and Lee and Hersh (1993),

we subdivide each DBA into T subperiods such that the probability (following Poisson arrivals) that more than one bid is placed within each subperiod is smaller than ϵ . For an arrival date, the average number of bids placed at j days in advance of arrival is λ_j , then choose T_j such that $1 - f_j(0) - f_j(1) \leq \epsilon$, where

$$f_j(x) = \frac{(\lambda_j/T_j)^x e^{-\lambda_j/T_j}}{x!}. \quad (2)$$

For a firm posting m prices to Priceline, let P_j^i be the probability that a successful bid is made j DBA for price i , $i = 1, 2, \dots, m$. Without loss of generality, assume $p_j^m > p_j^{m-1} > \dots > p_j^1$. Then

$$P_j^m = [1 - F_B(p_j^m)] \cdot f_j(1), \quad (3)$$

where $f_j(1)$, as Equation (2) shows, is the probability of a customer placing a bid, j DBA. This formulation is similar to the single-product dynamic-pricing model of Zhao and Zheng (2000), in which $F_t(\cdot)$ denotes the cumulative probability distribution of customer reservation prices. Then, $1 - F_t(p)$ is the probability that a customer arriving at time t would purchase an item priced at p . The result is $\lambda(t)[1 - F_t(p)]$ as the demand process for Poisson arrivals $\lambda(t)$.

Under our multiple price framework, for prices $m-1$ to 1, the probability of a sale in period n becomes

$$P_j^i = [F_B(p_j^{i+1}) - F_B(p_j^i)] \cdot f_j(1), \quad i = m-1, \dots, 1, \quad (4)$$

$P_j^0 = 1 - \sum_{i=1}^m P_j^i$ is the probability of no sale.

Determining Optimal Room Rates

A firm that releases a preset inventory to Priceline has two sets of decisions to make: (1) the prices to post and (2) the inventory allocation or booking limits for each price. For a given arrival or check-in date (departure date for an airline), the firm must set prices and allocations for each day prior to arrival. It sets prices and allocations on each DBA, knowing that it can revise these prices and allocations at each subsequent DBA prior to arrival, and accomplishes setting these prices and allocations via backward induction and dynamic programming. Let $V_r^T(j, \mathbf{p}_j)$ represent the expected value for a firm releasing r rooms to Priceline at period j days before arrival with DBA j subdivided into T decision periods, and with \mathbf{p}_j the vector of m prices for DBA j . A request for rate

class i in subperiod t , where $t = 1, 2, \dots, T$, will only be accepted if $p_j^i + V_{r-1}^{t-1}(j, \mathbf{p}_j) \geq V_r^{t-1}(j, \mathbf{p}_j)$. $V_r^{t-1}(j, \mathbf{p}_j) - V_{r-1}^{t-1}(j, \mathbf{p}_j)$ represents the expected opportunity cost of the r th room in decision period t of DBA j . For fixed prices, \mathbf{p}_j , Lee and Hersh (1993) show that $V_r^{t-1}(j, \mathbf{p}_j) - V_{r-1}^{t-1}(j, \mathbf{p}_j)$ is nonincreasing in r for a given t and nondecreasing in t given an r . This monotonicity results in the determination of critical values for r and n for which inventory is restricted from rate class i . Because of the monotonicity if a price class is closed, all lower-priced classes are closed; i.e., it allows the determination of booking limits or allocations across the m rate classes with prices \mathbf{p}_j .

The expected revenue $V_r^t(j, \mathbf{p}_j)$ can be expressed as

$$V_r^t(j, \mathbf{p}_j) = P_j^0 V_r^{t-1}(j, \mathbf{p}_j) + P_j^m [p_j^m + V_{r-1}^{t-1}(j, \mathbf{p}_j)] + \sum_{i=1}^{m-1} P_j^i \max[p_j^i + V_{r-1}^{t-1}(j, \mathbf{p}_j), V_r^{t-1}(j, \mathbf{p}_j)]$$

for $r > 0, t > 0$, (5)

with $V_r^0(0, \mathbf{p}_0) = 0$ and $V_r^0(j, \mathbf{p}_j) = V_r^T(j-1, \mathbf{p}_{j-1})$ for all r . The first term on the right side of Equation (5), $P_j^0 V_r^{t-1}(j, \mathbf{p}_j)$, is the expected revenue from no bids being placed in period t , i.e., the revenue with the same inventory and one fewer opportunity to sell it— $V_r^{t-1}(j, \mathbf{p}_j)$ multiplied by the probability of no sale. The second term, $P_j^m [p_j^m + V_{r-1}^{t-1}(j, \mathbf{p}_j)]$, is the expected revenue from selling a room at the firm's highest posted price, p_j^m , in period t with $r-1$ rooms to sell in period $t-1$, and P_j^m the probability that a bid is placed at a price $\geq p_j^m$. The third term, $\sum_{i=1}^{m-1} P_j^i \max[p_j^i + V_{r-1}^{t-1}(j, \mathbf{p}_j), V_r^{t-1}(j, \mathbf{p}_j)]$, is the expected value across all remaining $m-1$ prices. The expected value is the product of the probability of a request at each price P_j^i , and the max of accepting the request at price p_j^i plus the value of having one fewer room in the remaining periods $V_{r-1}^{t-1}(j, \mathbf{p}_j)$ and not accepting the request, and having all r rooms in the next $t-1$ periods $V_r^{t-1}(j, \mathbf{p}_j)$.

The above formulation gives the optimal expected value for a given set of prices, with $V_r^{t-1}(j) - V_{r-1}^{t-1}(j, \mathbf{p}_j)$ used to calculate optimal booking limits. For firms that load a single set of rates to Priceline daily (as is common practice), optimal prices can be determined by searching over the vector of prices \mathbf{p}_j for $*V_r^T(j, \mathbf{p}_j)$, where $*V_r^T(j, \mathbf{p}_j) = \max_{\mathbf{p}_j} V_r^T(j, \mathbf{p}_j)$. This search is quite feasible because Kimpton only posts three rate classes,

and its prices are usually only whole numbers. Our experience shows that performing this search and the resulting optimal price and booking limit determination take only a few minutes on a standard desktop computer for a model implemented in Excel and coded in Visual Basic for Applications.

Opportunity Cost of a Room

Because Priceline is only one of several distribution channels that firms use to sell inventory, it cannot set its prices in isolation. Firms can integrate a Priceline pricing strategy into their overall inventory or revenue-management strategy by considering the opportunity cost of releasing rooms on Priceline. The deterministic linear program is a common form of inventory control in many revenue-management approaches (Liu and van Ryzin 2008) and is especially relevant for hotel revenue management. The deterministic LP is

$$V(x) = \max p^T y$$

s.t. $Ay \leq x$, (6)

$$0 \leq y \leq \mu,$$

where the decision variables, $y = (y_1, y_2, \dots, y_n)$, represent allocations of capacity across n product classes (for a property, these product classes are different rate classes and lengths of stay) at rates p with demand μ and available capacity x . As Talluri and van Ryzin (2004) discuss, firms often use the dual variables π associated with capacity constraints x . The dual variables represent the marginal value of capacity and provide an estimate of the opportunity cost of releasing inventory on Priceline (and making it unavailable to other distribution channels). Forecasted demand μ is updated regularly as new reservations are received; Weatherford and Kimes (2002) provide details on hotel room forecasting.

The number of rooms to release to Priceline can be evaluated by comparing the expected marginal value of rooms released to Priceline to the shadow price π of the capacity constraint for the arrival day in question from the deterministic LP. The firm can continue to release rooms to Priceline if $*V_r^T(j, \mathbf{p}_j) - *V_{r-1}^T(j, \mathbf{p}_j) > \pi$, where $*V_r^T(j, \mathbf{p}_j)$ is the optimal expected value at DBA j if r rooms are released to Priceline, and $*V_{r-1}^T(j, \mathbf{p}_j)$ is the optimal expected value if only $r-1$ rooms are released.

		Decision period															
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Capacity	1	108	110	112	114	115	117	118	119	121	122	123	124	125	126	127	127
	2	84	86	88	90	92	93	94	96	97	99	100	101	103	104	105	107
	3	64	67	70	73	75	78	80	82	83	85	86	88	89	90	91	92
	4	49	51	54	57	59	61	64	66	68	71	73	75	77	78	80	82
	5	33	36	40	44	47	49	52	54	56	57	59	61	63	65	67	69

<div style="border: 1px solid black; width: 15px; height: 10px; display: inline-block; margin-right: 5px;"></div> Accept only class 3 (\$142)	<div style="border: 1px solid black; width: 15px; height: 10px; background-color: #808080; display: inline-block; margin-right: 5px;"></div> Accept class 2 and 3 (\$92 and \$142)
<div style="border: 1px solid black; width: 15px; height: 10px; background-color: #d3d3d3; display: inline-block; margin-right: 5px;"></div> Accept all classes (\$55, \$92, and \$142)	

Figure 5: This table illustrates expected marginal room values and corresponding optimal inventory policies for a Kimpton hotel.

Illustrative Example

Figures 3 and 4 and Table 1 are used to determine the probability of a request P_j^i , given a price p_i , j days prior to arrival. The following section illustrates these calculations for a weekday arrival for prices posted to Priceline on the arrival day.

For Poisson arrivals with a mean of 8.5 (from Table 1, DBA = 0 for weekdays) per day and $P[\text{number requests} \geq 2 < 0.1]$ from Equation (2)], the DBA must be subdivided into 16 decision periods. Then, for prices $\mathbf{p}_0 = \{p_3, p_2, p_1\} = \{\$142, \$92, \$55\}$ with bids following a gamma distribution with parameters 3.37, 29 results in probability of requests $P_1^0 = 0.0568$ using Equation (3), $P_2^0 = 0.0900$ $P_3^0 = 0.0965$ using Equation (4), where the probability of a request f_1 in Equations (3) and (4) is calculated assuming Poisson arrivals with mean arrivals of 8.5/16 per decision period.

Figure 5 displays the expected opportunity costs, $V_r^{t-1}(0, \mathbf{p}_0) - V_{r-1}^{t-1}(0, \mathbf{p}_0)$, for inventory levels up to five rooms across the 16 decision periods. These opportunity costs can then be used to determine optimal inventory policies for the five rooms potentially being released, i.e., how many rooms should be available at each of the three prices. The policies are critical capacity-time (decision-period) pairs indicating when a price class is open. The prices, $\mathbf{p}_0 = \{\$142, \$92, \$55\}$, are determined by searching $V_r^T(0, \mathbf{p}_0)$ over all possible prices p_0 .

Impact

Perfect experiments to test the effectiveness of a new pricing or inventory strategy do not exist as they do

in science laboratories. To estimate model effectiveness and because of seasonality in most service settings, firms commonly use year-to-year comparisons; i.e., they compare a given date to the same date of the prior year (e.g., the third Monday in September 2008 and the third Monday in September 2007). To evaluate model performance, we used the model to compare Priceline prices for a Kimpton hotel in Washington, DC to a similar Kimpton property also in Washington, DC, but priced as usual to serve as an additional form of comparison. For the second property, we used its current revenue-management system (all properties use the same vendor-developed revenue-management system) to provide Priceline rates.

Since the model’s introduction, rooms sold via Priceline increased 11 percent between 2007 and 2008; the average rate for these rooms increased 3.7 percent. This compares to a price reduction of 8 percent on other electronic distribution channels to maintain regular sales volumes. Room rates in the base case property and sold on Priceline declined by 40 percent; Priceline rates also decreased by 3.7 percent, and regular electronic distribution channel prices declined by 9.5 percent. Although we cannot report on the statistical significance of the differences given the ad hoc nature of the comparison, the results are encouraging because the model has increased Priceline volume and prices simultaneously while not adversely affecting other channel prices.

Conclusions and Limitations

We have outlined some details of the allocation of inventory on the NYOP mechanism that Priceline

uses. The unique nature of its process for selecting the service provider and the price paid to that service provider requires novel approaches to the inventory-allocation decision. Using real data from Kimpton Hotels, we have illustrated how Priceline can use these data to set optimal prices and inventory policies. Our example illustrates how a firm used Priceline to improve its volume and prices, even in a market in which regular prices are declining.

The current formulation is monopolistic in nature because it does not include any analysis of the probability that Priceline selects a specific firm. Including this factor would require the firm to have some insight into the number of competitive firms that also use the NYOP mechanism because Priceline's first random-selection process uses NYOP. In addition, the mechanism is also relevant to the pricing decision because it is necessary to determine each firm's success rate in the first round; Priceline uses this success rate in subsequent rounds. Excluding the selection probability from the formulation, given that it is partially price dependent, might result in overestimating expected revenues (and the resulting prices). We leave this interesting aspect of the problem for future investigation.

Last, the bid reports that firms receive do not allow them to distinguish between bidders or identify repeat bidding. The use of travel websites, such as <http://www.BiddingforTravel.com>, could exacerbate repeat bidding because they provide forums that allow customers to discuss bid success. Excluding repeat bidding overestimates demand and might be an issue for firms that receive few bids; these are often lower-valued service providers, such as one-star hotels. This is less of an issue for three- to five-star hotels, such as Kimpton, because they receive considerably more bids than the rooms they allocate to Priceline.

References

Fay, S. 2004. Partial-repeat-bidding in the name-your-own-price channel. *Marketing Sci.* 23(3) 207–218.

- Gerchak, Y., M. Parlar, T. K. M. Yee. 1985. Optimal rationing policies and production quantities for products with several demand classes. *Canadian J. Admin. Sci.* 2(1) 161–176.
- Hann, I., C. Terwiesch. 2003. Measuring the frictional costs of online transactions: The case of a name-your-own-price channel. *Management Sci.* 49(11) 1563–1579.
- Hotelling, K. 1929. Stability in competition. *Econom. J.* 39(153) 41–57.
- Jerath, K., S. Netessine, S. K. Veeraraghavan. 2007. Revenue management with strategic customers: Last-minute selling and opaque selling. Working paper, The Wharton School, University of Pennsylvania, Philadelphia.
- Jiang, Y. 2007. Price discrimination with opaque products. *J. Revenue Pricing Management* 6(2) 118–134.
- Lee, T. C., M. Hersh. 1993. A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Sci.* 27(3) 252–265.
- Liu, Q., G. J. van Ryzin. 2008. Strategic capacity rationing to induce early purchase. *Management Sci.* 54(6) 1115–1131.
- Spann, M., B. Skiera, B. Schafers. 2004. Measuring individual frictional costs and willingness to pay via name-your-own-price mechanisms. *J. Interactive Marketing* 18(4) 22–36.
- Talluri, K. T., G. J. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, New York.
- Wang, T., E. Gal-Or, R. Chatterjee. 2005. Why a name-your-own-price channel makes sense for service providers (or: Who needs Priceline, anyway?). Working paper, Katz Graduate School of Business, University of Pittsburgh, Pittsburgh.
- Weatherford, L. R., S. E. Kimes. 2002. A comparison of forecasting methods for hotel revenue management. *Internat. J. Forecasting* 19(3) 401–415.
- Wilson, J. G., G. Zhang. 2008. Optimal design of a name-your-own-price channel. *J. Revenue Pricing Management* 7(3) 281–290.
- Zhao, W., Y. Zheng. 2000. Optimal dynamic pricing for perishable assets with nonhomogeneous demand. *Management Sci.* 46(3) 375–388.

Sarita A. Moore, Area Director of Revenue Management, Kimpton Hotels & Restaurants, 1315 16th Street, NW, Washington, DC 20036, writes: "I am writing this letter in support of the research that Professor Anderson has provided Kimpton Hotels & Restaurants. We have been using Professor Anderson's model to set rates and room availability of our hotels for several months. The model has greatly improved our understanding of Priceline's name-your-own-price distribution channel as well as improved our revenue and sales through Priceline."