Choice-Based Dynamic Pricing for Vacation Rentals

In this paper we propose a new dynamic pricing approach for the vacation rental revenue management problem. The proposed approach is based on a conditional logistic regression that predicts the purchasing probability for rental units as a function of various factors, such as lead time, availability, property features and market selling prices. In order to estimate the price sensitivity throughout the booking horizon, a rolling window technique is provided to smooth the impact over time and build a consistent estimation. We apply a non-linear optimization algorithm to determine optimal prices to maximize the revenue, considering current demand, availability from both the rental company and its competitors, and the price sensitivity of the rental guest. A booking curve heuristic is used to align the booking pace with business targets and feed the adjustments back into the optimization routine. We illustrate the proposed approach by successfully applying it to the revenue management problem of Wyndham Destinations vacation rentals. Model performance is evaluated by pricing two regions within the Wyndham network for part of the 2018 vacation season indicating revenue per unit growth of 3.5% and 5.2% (for the two regions) through model use.

Key words: Revenue management, vacation rental, dynamic pricing, price sensitivity, conditional logit regression, non-linear optimization

---

Introduction

Revenue management is both the science and art of utilizing price and other marketing levers to smooth out supply and demand imbalances for fixed capacity firms in an effort to maximize revenue or profit. It was first adopted in the airline industry, and then widely applied to the hospitality industry and other service firms. The primary objective of revenue management is to sell the right product to the right consumer at the right moment for the right price with the right package. The key philosophy of this discipline is to understand how different customers value products differently and then to design a strategy to align their expectation with the product placement, availability and prices. Revenue management is a multidisciplinary field, combining the core techniques of analytics, machine learning and operations research into the decision processes of business operations such as pricing and marketing.

In the context of the hospitality industry, vacation rentals is a relatively new form of product offering, and has seen rapid growth over the past decade largely as a function of new travel platforms like Airbnb and HomeAway. It brings a unique experience and innovative lifestyle that blurs the lines between "home" and "vacation", where customers not only travel to their destinations but also live a life there. Generally speaking, there are two types of vacation rental businesses:

- homeowners listing and self-managing their own properties on distribution channels, such as Airbnb and HomeAway, and
- homeowners contracting with professional property management companies to manage/rent their properties.

Wyndham Destinations (WYND) provides professional property management services and vacation exchange and rental services on behalf of independent owners. Its 25,000 employees help travelers find vacations across 4300 destinations in 110 countries, making Wyndham Destinations the world's largest vacation ownership, exchange and rental company. With a network of brands acting as an intermediary between owners and consumers, its diverse collection of private lodging around the world offers something for every traveler, whether they are looking for an adventure, quiet family getaway, urban cityscape or anything in between.

Unlike airlines and hotels, which have relatively established revenue management systems, the vacation rental business has not formed a systematic way to build pricing and inventory management solutions. In this paper, we will propose a choice-model based dynamic pricing system to help the rental company set the optimal price in order to maximize their revenue. The pricing model incorporates both a conditional logit choice model to estimate the purchasing probability and price sensitivity, and a non-linear optimization together with booking-curve heuristic to effectively optimize revenue and manage inventory.

Owing to the fragmented nature of most vacation rentals there does not exist a vast literature discussing the nuances of pricing for vacation rentals, as unlike hotels or airlines where rooms in a hotel or seats on a given flight are similar, each vacation rental unit is unique. At some level vacation rental pricing is similar to hotel pricing making the extant revenue management (RM) literature relevant. Anderson and Xie (2016) provide a review of RM within hospitality while Talluri and van Ryzin (2004) provide a well-balanced introduction to many aspects of RM, with Elmaghraby and Keskinocak (2003) providing insight into dynamic pricing in general. So, while there exists an established literature around RM, the vacation rental RM problem is quite distinct. An airline starts selling seats on a flight, and then may use transactional data about those seat sales to dynamically price the remaining seats on that same flight, whereas a vacation rental only has one unit to sell for a particular stay period; therefore its dynamic pricing problem is different as it may use sales of other (similar) rental units to help price unsold units or it may use sales (or lack thereof) for other stay periods to help price the particular unit/stay period in question.

Much of the pricing related literature around vacation rentals is based upon scraped data from platforms like Airbnb. Examples of such research are Gibbs et al. (2018) and Magno et al. (2018) where they take posted price information of rental units and they fit OLS models to posted prices in an effort to measure the marginal impact of unit attributes upon posted prices. Under this stream of literature, the intent is to understand the impacts of unit characteristics upon prices not the impact of price upon demand – which is critical in a typical RM setting. Kwok and Xie (2018) empirically investigate the impact of more active pricing strategies upon hotel and vacation rental unit performance on Airbnb, finding that multi-unit versus single unit hosts engage more actively in dynamic pricing and that this dynamic pricing leads to improved performance. As discussed in Kwok and Xie (2018), Airbnb provides pricing tools to rental unit hosts through its recommended smart pricing. The intuitive, data intensive approach of Ye et al. (2018) uses a binary logit model to predict unit purchase probability and then a regression model which minimizes a loss function for rental units. The approach compares system recommended prices (P_{sug}) to actual booked prices P to assess loss to hosts. Unlike Ye et al. (2018) we propose a traditional revenue management approach whereby we try to maximize profit whereas they propose a pricing solution which suggests prices based upon a purchase probability model and losses estimated by comparing system prices to historic list prices. The framework in Ye at al. (2018) works well for AirBnB as not all hosts utilize their suggested prices enabling ongoing estimation of losses as they compare system prices to (host) posted prices, whereas in our setting we are setting prices for all units and as a result deploy typical approaches like use of booking curves as a feedback mechanism (versus using losses) to adjust prices over time.

Unlike traditional hotels (or airlines), the vacation rental industry is one with very heterogeneous products as each unit to be rented may be unique, with different amenities, furnishings, locations and owners. The challenge for pricing the vacation rental business is to find an efficient way to price the various products for the rental companies or home-owners in order to effectively compete in a highly transparent world without diluting the uniqueness and core values of those products. The unique nature of individual vacation rental units lends to the application of discrete choice models when trying to estimate the impact of unit attributes (and price) upon demand.

Choice models are naturally well designed for the vacation rental problem as they help quantify or measure the value of product attributes when consumers face a purchase decision across a set of differentiated options. Discrete choice models have been employed to study consumer choice behavior and understand consumer's brand preferences, market structure as well product attributes. Coretjens and Gautschi (1983) provide a general survey of discrete choice models in marketing and a systematic introduction to discrete choice modelling theory. In a discrete choice model, an individual decision maker (consumer) makes a choice among a feasible set of alternatives. The decision is often made by maximizing random utilities of the alternatives. The multinomial logit model (MNL) is the most commonly used random utility model. While choice models have been used extensively in the marketing literature their application in dynamic pricing settings is very limited. The RM literature has only more recently started to investigate the potential of choice models. Garrow and Koppleman (2004) apply MNL models to airline data to estimate the factors impacting passenger no-shows and cancellations. Perhaps the most comprehensive approach to the choice-based network RM problem is that of Gallego, Ratliff and Shebalov (2015) where they generalize the choice base formation to more accurately account for sales lost to lack of availability and more naturally model the recapture of demand by other (available) product classes.

In this paper, we will focus on a specific rental company within Wyndham Destinations portfolio with two types of inventory – Type I rentals where Wyndham takes full inventory risk as owners are prepaid regardless of utilization, and Type II rentals being contracted inventory where owners are paid contracted rates for rented stay weeks. Both of these products have their advantages and disadvantages: Type I products with prepaid costs ensure a stable supply for the incoming demand with a relatively low cost (to Wyndham) but it will incur financial risk if the inventory is unsold, whereas Type II products have reduced financial risk but incur inventory risk as units may be sold elsewhere or removed from the market by owners. Most of the rental company's inventory is rented out for weekly durations with the fixed check-in day options. Under this type of the mixed product structure, the rental company may have different pricing strategies for Type I products (over Type II) during the selling horizon and across different stay dates as demand is seasonal in nature and selling prices fluctuate over the selling horizon. The objective for this pricing approach is to dynamically set prices for the inventory to maximize the total profit from these two products by fully utilizing the weekly updated information from the inventory availability, demand and competitor booking information.

Modeling Framework

The model flow for this dynamic pricing system is composed of multiple stages with five sequential modules scheduled to run once per week to update the recommended price for the rental company by an automatic job scheduling program. The five modules are:

- Competitive scraping: this module creates the weekly updated information for the inventory availability, booking and pricing from the major competitors to feed into choice modeling in order to evaluate the customer's buying vs. non-buying choice for the rental assortments.
- Market price forecasting: this module projects future market selling prices grouped by similar products over a period using a weighted moving average across the recently sold inventories on the market from both the rental company and its competitors. The market selling price will be used as a market reference price to calculate the relative market price index for the similar product offerings in the choice model and set the boundaries for the pricing optimization as constraints.
- Booking curve heuristics: this module builds a target booking pace from the historical transactions which are used as feedback into the pricing optimization.
- Discrete choice model: this module utilizes a conditional logit regression (CLR) to model the customer purchasing behaviors under different product offerings. In this choice model, we estimate the booking probability for a product with a specific check-in week and booking lead time in days.
- Price optimization: this module develops a myopic (static) optimal price for the coming week using the (price dependent) purchase probability from choice model. At the same time the booking curve heuristics in module 3) will provide a feedback signal from the target booking curve to indicate how aggressive price adjustments should be.

Figure 1 Model flow overview.

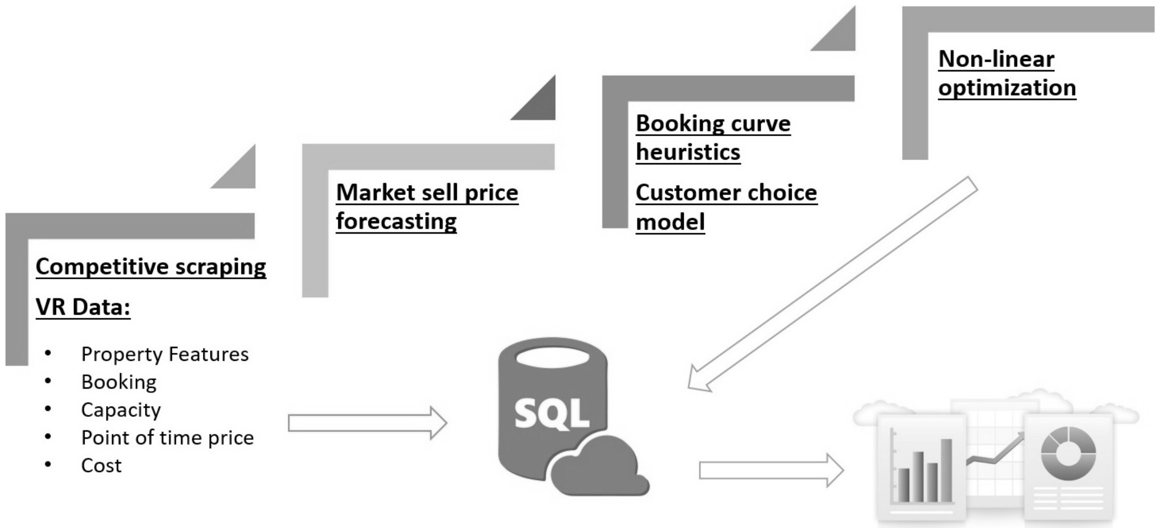

The entire model flow is shown in the Figure 1. After these five modules are processed sequentially, the final recommended price will be written into the SQL server tables, creating a user report to feed into the pricing/distribution system. The major components of all those modules will be discussed in detail in the following several sessions.

Booking Curve Heuristics

Booking or pace curves are common visuals used in many revenue management settings with the curves depicting how demand (sales) for a given product builds over the selling horizon. Orkin (1988) provides an early view of RM and the use of booking curves in managing hotel prices and availability. To benchmark the booking pace, sometimes a target booking curve will be set in advance to make a real-time comparison of current demand. Usually, the target booking curve can be determined by several ways using historical booking data, e.g.:

- a market booking curve by calculating the inventory utilization across all the players in the same market from a certain region,
- rental company's booking curve based on their own historical reservations and trend,
- or a weighted booking curve from both the rental company and other market players.

Comparing the current booking pace (to the benchmark) can provide useful signals about the potential pricing strategies, for example "too fast" or "too slow". A "too fast" booking pace may indicate that the price is relatively too low compared with the customers' expectation so that the inventories are sold out in a much faster than expected pace, while a "too slow" booking curve may on the other side indicate that the price is too high or the inventories are less attractive compared with the product offerings on the market from other players so that more-than-expected inventories will face an unsold risk.

After taking multiple factors into consideration, such as revenue targets, risk tolerance (owing to Type I prepaid inventory), minimization of negative competitor reactions and overall acceptance of system prices, and thoroughly discussing the potential pros and cons with our key business stakeholders, we decided to choose the rental company's historical booking curve as our target booking pace. The booking curve is constructed separately using time series analysis for different product offering, with products grouped by bed types and quality rating with a certain fixed check-in week. In order to make a reasonable adjustment on the booking pace and avoid potential over-correction on price adjustment, both point estimation and confidence intervals with different percentiles are projected here to retain the pace within a certain target range, providing in-time signals back to the optimization as a function of the current pace following within one of these three regions:

- Normal range: this range indicates that the booking pace is within a target range, neither "too slow" nor "too fast". If the current inventory utilization is within this range, we will not feed any additional constraints into the optimization part from this module.
- Warning range: this range indicates the booking pace is either slightly faster or slower than the expected target utilization. If the current inventory utilization lies in the upper band of the warning region, a signal of "no price reduction" will feedback into the optimization part; while on the other hand, if the utilization lies in the lower warning band, a signal of "no price lift" will feedback to the optimization.
- Alarming range: the range indicates that the booking pace is noticeably deviated from the target and needs special attention. Therefore, an active price adjustment based on the price sensitivity from choice modeling will be fed into the optimization to bounce the booking pace back into the targeted normal range.

The booking curve band is designed to be narrower with the selling window moving closer to the check-in date in an effort to minimize the unsold inventory risk. A visualization of the booking curve with three levels of controlling bands across the leading time is shown in Figure 2.

Use of these booking curves, while heuristic in nature and bound by historic actions, provides useful guidance to how demand is materializing and comfort with pricing actions.

Figure 2 Booking curve heuristics with control bands.

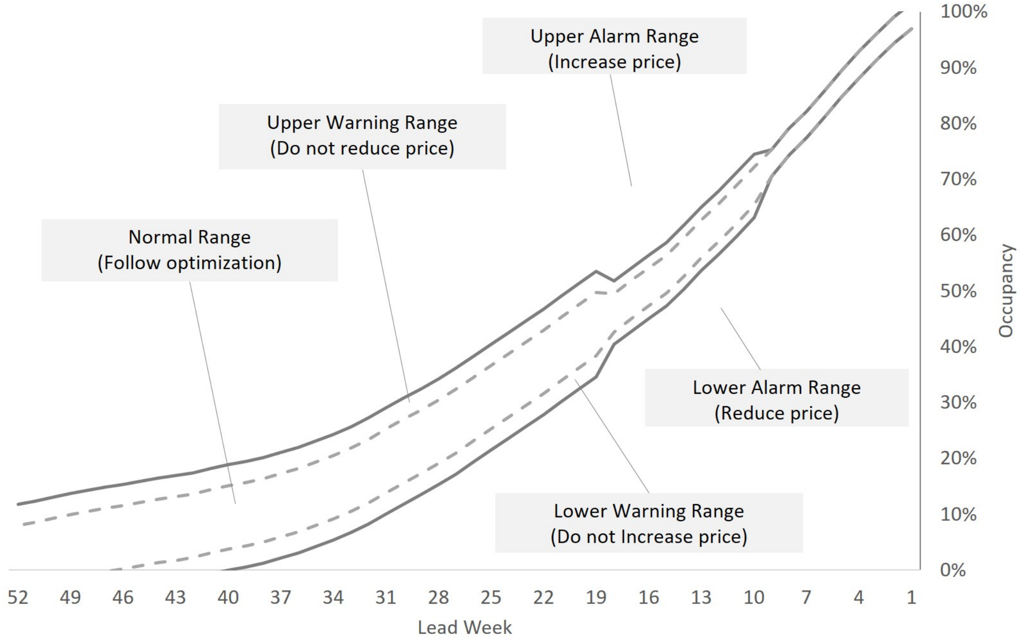

Given their heuristic nature, they may not finally result in a global optimal pricing strategy, but instead provide an avenue for better pricing decisions and convergence over time.

Discrete Choice Model

As discussed earlier MNL is a widely-used discrete choice model in marketing research. We utilize a specific variant of choice modeling, called conditional logit regression (CLR) proposed by McFadden (1973). CLR models the expected utility η_{ij} in terms of the characteristics of the alternatives rather than the attributes of the customers. In this discrete choice model, we utilize a conditional logistic regression to predict the consumers' buying probability π_{ij}, where the probability of choosing the *jth* individual product depends on a vector of variables x_j related to the characteristics of the *jth* alternatives. We estimate the choice model using reservation or individual customer level data - with reservations tracked at Wyndham directly and competitors indirectly using scraped data. We anticipate seasonal differences so different models are estimated for each arrival month during the vacation season.

Generally, in the multinomial dataset, one of the response categories will be nominated as the baseline and the odds-ratio noted as $\frac{\pi_{i1}}{\pi_{iJ}}$, representing the purchase probability between a product and the baseline product for a specific customer. The log-odds of each response in this model follow a linear form:

$$\eta_{ij} = log\frac{\pi_{ij}}{\pi_{iJ}} = \alpha + x_j\beta \tag{1}$$

where α is an intercept and β is a vector of the parameter estimates, for $i = 1$, ..., I and $j = 1$,..., J-1. Conditional logistic regression is simply a form analogous to a binary logistic regression - a special case of the conditional logistic regression with $J = 2$. However, under a binary logistic regression the probability distribution of the response follows a binomial distribution with only one equation that contrasts one product against the other, while a conditional logistic regression follows a multinomial distribution with J-1 equations, with each of the product categories, from 1 to J-1, contrasted to the baseline category J. The conditional logit model can be written in terms of the odds π_{ij} instead of log-odds as follows:

$$\pi_{ij} = \frac{e^{\eta_{ij}}}{\sum_{k=1}^{J} e^{\eta_{ik}}} \tag{2}$$

For each customer, the summation of the probabilities for all product categories will add to one. Like binary logistic regression, conditional logit regression utilizes the MLE (maximum likelihood estimation) to estimate the parameters of the model given observations by finding the parameter sets to maximize the likelihood of making such observations. In this paper, we use the historical booking data from Wyndham and the implied booking (unit changing from available to unavailable implying a reservation) data from competitors to fit our choice model and obtain parameter estimates for a series of attributes. Various factors are taken into consideration in the model, such as the property's features, the availability of inventories from both the rental company and its competitors, and the relative price index calculated as a price ratio between the property and other comparable products on the market. The different types of the independent variables included in the model are listed in Table 1. As Table 1 indicates, the choice model includes property level features as well as price and seasonal factors. Seasonality, being a key feature for vacation rentals is modeled in two manners - a fixed effect through season factor S and through price effects where price effects are both a function of arrival month and when (during the booking cycle) a reservation is made. The fixed effects help model seasonal differences as a function of the time of year when consumers are shopping/researching their vacation, we then add

Table 1 Attribute categories in conditional logit regression.

#	Variable Categories	Description
1	Competitor Indicator	A binary variable to indicate whether the property is from the rental company or its competitors
2	Property Features	Customer reviews, quality rating, location convenience rating, bedroom size, swimming pool and property types (house vs. apartment)
3	Relative Price Ratio	The relative price ratio is a price index, defined as the properties current list price divided by current market selling price, which is calculated by a weighted moving average of the recently sold inventory
4	Seasonal Factors	Given differences in consumer behavior across the time of year when consumers are making vacation reservations, the booking period was divided into three windows: pre-shopping season, shopping season, and post-shopping season. The shopping season is from January to February with an after-holiday booking peak. The pre-shopping season is the lead time before January, while the post-shopping season is the lead time after February and before check-in date.

seasonal price effects owing to supply and demand effects by stay period and changes in price sensitivity as the stay period approaches.

Booking choices arise from the following individual utility functions that are determined from the various characteristics of the property and an idiosyncratic unobserved component specific to the customer that is captured by the error term $\varepsilon_{i,j,m}$. In that case, individual utility is specified as:

$$U_{i,j,m} = \alpha_m + \beta_{1,m} C_j + \beta_{2,m} F_j + \beta_{3,m} S_m + \beta_{4,m} PR_{j,m} + \varepsilon_{i,j,m} \tag{3}$$

Equivalently,

$$U_{i,j,m} = \eta_{j,m} + \varepsilon_{i,j,m} \tag{4}$$

Under individual utility maximization we get the following probabilities for the estimated individual choice probabilities in the special case of the binary logit model:

$$\hat{p}_{i,j,m} = \frac{e^{\hat{\eta}_{j,m}}}{\sum_{j=1}^{J} e^{\hat{\eta}_{j,m}}} = \frac{e^{\hat{\eta}_{j,m}}}{1 + e^{\hat{\eta}_{j,m}}}$$

Where $\hat{p}_{i,j,m}$ is the estimated probability that customer i chooses property j for month m. This derivation also assumes that those who did not book with Wyndham booked with a competitor, and that base utility is set to 0.

The price sensitivity of the rental properties in a certain arrival month, calculated from the coefficient $\beta_{4,m}$, may vary across different check-in months due to the uneven balance

Figure 3 Price sensitivity by arrival month.

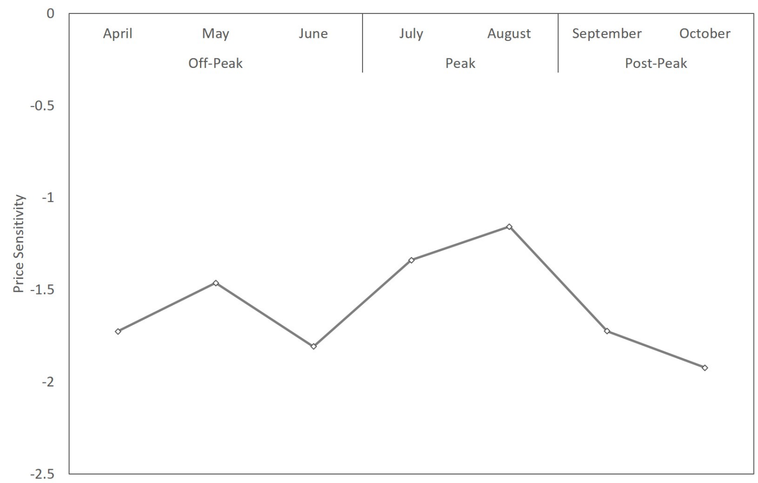

between supply and demand. For example, for months like August, because demand usually surpasses supply for the majority of the weeks, most of the available units will be sold out. Therefore, it is reasonable to expect that the consumers have a lower price sensitivity in August than other non-peak months. A sample output of the price sensitivity for one of the tested regions by arrival months is shown in Figure 3. While price sensitivity may vary by arrival month, it may also vary by when the reservation was made - i.e. differentiating the impact of price for last-minute shoppers versus early planners. However, bookings may be concentrated during a few periods rather than evenly distributed across all booking windows. Owing to potential differences in price sensitivity by arrival month as well as booking lead time, and the resulting sparsity that results from consumers making reservations up to a year in advance, we adjust for the differences in price sensitivity using a rolling window analysis. The rolling window analysis provides for smoothed, yet heterogeneous estimates of price sensitivity across the portfolio of units. We later assume the products across different weeks in the same month are homogeneous in terms of their (monthly) utility with weekly differences (e.g. for school breaks or holidays) added through use of the booking curve heuristic.

Rolling window analysis

The price sensitivity for vacation rentals may depend on various factors, from check-in season, booking lead time to product availability. For example, for the vacation rental business studied, customers from most regions will plan their vacations in January and

February, recognized as a popular "shopping season", for many different summer vacation stay dates. During the shopping season, the price sensitivity might be different from other seasons. One of the key assumption underlying traditional statistical models is that the model parameters are constant over time. In order to relax this assumption, a common technique called "rolling window" is introduced to sequentially compute parameter estimates over a fixed time window across the data set.

The rolling window technique is used to estimate the utility of product j for a check-in month m with lead time (days prior to arrival that booking was made) t in the form:

$$U_{i,j,m,t} = \alpha_{m,t} + \gamma_t \hat{\eta}_{j,m} + \beta_{5,m,t} PR_{j,m} + \varepsilon_{i,j,m,t} \tag{5}$$

where $\gamma_t \hat{\eta}_{j,m}$ is the scaled utility from the first estimation.

If the price parameter (PR) is indeed constant over lead times, then we should not observe significant parameter coefficient variations over the time windows. Otherwise, the rolling window should be able to detect a non-stationary pattern in the parameters. The method segments reservations into different booking lead times and estimates the offset utility for teach time window. The first component is the expected utility across individuals for a given arrival month; the second component is the offset utility or deviation from the expected utility - estimated from the remaining (time-varying) price sensitivity across lead times. In essence we assume price sensitivity is time-varying by both arrival month as well as how far in advance (of the arrival date) the booking was made. A similar setting may be an airline with different prices (owing to different supply-demand characteristics) for direct versus connecting flights (for the same origin-destination pair) along with different sensitivity to price changes for these two flights over the 180 days that consumers are able to book each flight owing to when leisure versus business travellers may be seeking to book.

Even though the parameters are estimated simultaneously, one can think of the rolling window analysis as a two-step process where $\eta_{j,m}$ is first estimated across the entire sample (for each arrival month m), then in the second stage the algorithm utilize observations within the range from t_0 to $t_0 + h$ (along with the original estimate of $\eta_{j,m}$) to estimate the parameter at time t_0, note that m represents the arrival month and the time window t is used to assess differences in price sensitivity over time leading up to the stay . As illustrated in Figure 4, the sample is then rolled forward with a given step s, and the same model fitting process is repeated across all time (booking) windows. As described above, for a certain lead time d, we take the booking samples from day d to $d+h$ to estimate the utility of product j for arrival month m from lead time d in the form:

$$\hat{\eta}_{j,m,t} = \hat{\alpha}_{m,t} + \hat{\eta}_{j,m} + \hat{\beta}_{5,m,t} PR_{j,m,t} \qquad (6)$$
$$t \in [d, d+h)$$

Resulting in the corresponding choice probabilities used in the optimization computed as follows:

$$\hat{p}_{i,j,m,t} = \frac{e^{\hat{\eta}_{j,m,t}}}{1 + e^{\hat{\eta}_{j,m,t}}} \qquad (7)$$

The expected product utility across individuals from iteration 1 in previous session is directly fed into the second iteration. Then the projected utility of product j for arrival month m and lead time d is estimated and fed into the optimization. In the first iteration, the baseline model is fitted to estimate the expected utility ($\eta_{j,m}$) across all individuals and lead times; in the second iteration the utility from bookings with different lead times is projected simultaneously for the corresponding rolling windows. It should be noted that while no operations (addition or subtraction) are performed on η during estimation, in the two stage estimation process we assume utilities are of the same scale during each estimation stage, i.e. $\gamma_t = 1$ in (5).

Ideally, the two iterations in our estimation process should be completed in a single step with a single likelihood function, but owing to sparse data (as we are estimating utilities at the vacation unit level across a booking window) we estimate the utilities in a two-step process. Further, assuming $\gamma_t = 1$ may induce unknown bias into our estimates which may distort our optimization results, which is hopefully moderated by our booking curve heuristics.

An example of the final smoothed curve from the remaining (time-varying) price sensitivity across different lead days for one of the tested regions is shown in Figure 5.

It should be noted, an alternative to the rolling horizon estimation process would be the inclusion of lead time as an attribute, along with the interaction of lead time and price within the choice model. The non-linearity of price sensitivities over the booking window as illustrated in Figure 5 along with concentration of bookings during certain time periods (resulting in sparse bookings during parts of the booking window) favored use of the rolling window analysis.

Figure 4 **Rolling window algorithm partitions.**

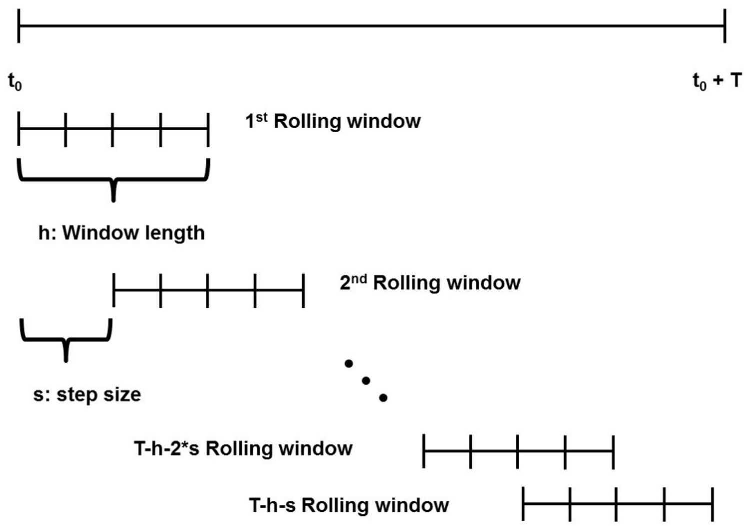

Figure 5 **Time-varying price sensitivity.**

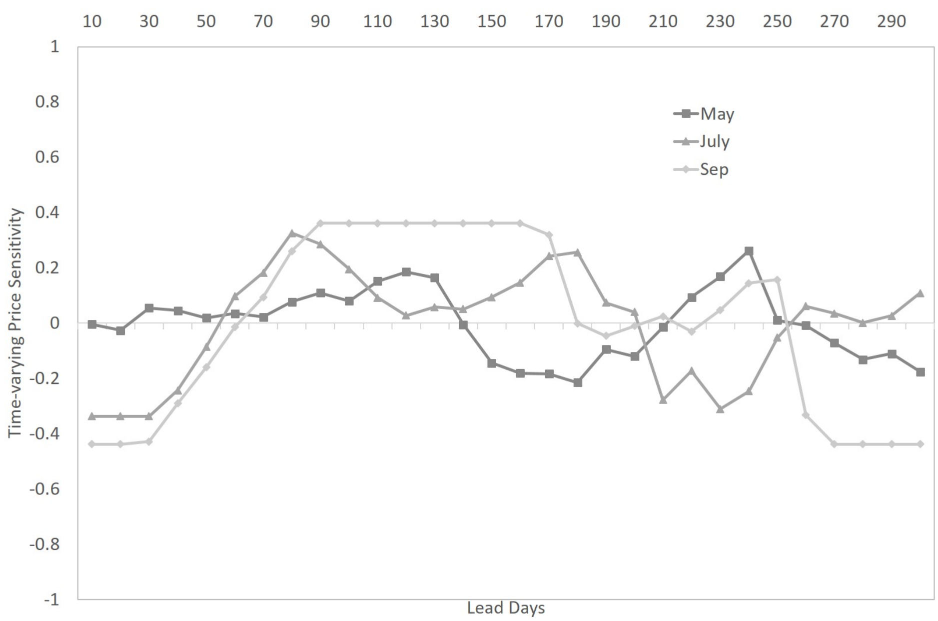

Price optimization

In the following we will discuss the optimization of vacation rental unit prices. There are two challenges we will face in this problem:

- Nonlinearity: Because the predicted purchasing probability from the choice model is a nonlinear function, the objective function will also be nonlinear.

• Multi-stage decision process: The pricing problem for vacation rentals is a multi-stage decision process with Wyndham being able to change prices over time for a given rental unit.

Throughout the booking window the opportunity exists to update selling prices for each unsold rental unit. Generally, there are two different ways to solve this pricing optimization:

• Dynamic programming (DP): a systematic methodology to solve this multi-stage problem in a global way by looking at all price decision in the whole process together;

• Myopic optimization: a greedy algorithm to make decisions for each period in isolation, not considering (potential) future price changes if inventory remains unsold.

While clearly a vacation rental company with a long selling horizon and multiple opportunities to change prices would benefit from a full DP implementation to set the optimal prices, such a formulation in this context would quickly suffer from the curse of dimensionality owing to the choice probabilities for a particular rental unit being a function of the other units currently available.

Unlike DP, a myopic optimization only focuses on maximizing the gain from the next stage given the current rental units available. Logically, myopic or greedy approaches are best suited to settings where prices are monotonically decreasing during the decision window, in our case monotonically decreasing during the selling horizon or booking window. If prices are dropping over the selling window (for a particular stay date), selling this piece of inventory earlier is better than selling later, and under this circumstance optimizing the next period (in isolation) becomes an approximation to the global optimum. Therefore, the original problem of multi-stage decision process is simplified to be a static problem, which determines the optimal price stage by stage in order to maximize the global system revenue.

The formula for this myopic optimization to maximize the revenue is given as follows:

$$\sum_{w}\sum_{j} P_{j,w}\frac{e^{\eta_{j,w}}}{\sum_{j=1}^{N} e^{\eta_{j,w}}} \tag{8}$$

where $P_{j,w}$ is the decision variable of the price for product j and arrival week w; the detailed description for the other notations is outlined in the following sections.

A myopic approach is suitable for the rental company's pricing problem as the majority of arrival weeks have supply in excess of demand with about 70% of the rental company's inventory full risk (Type I). Lastly, as shown in Figure 6, most rental unit-stay period

Figure 6 Sample booking and price curves.

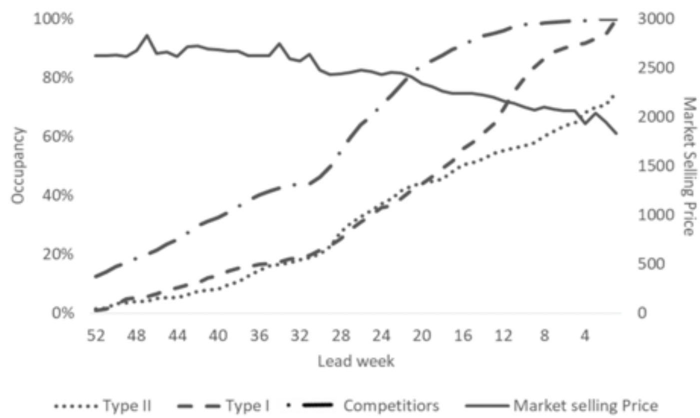

combinations have decreasing prices during the selling horizon. Figure 6 shows sample booking curves of the two product types from the rental company as well as its main competitor in July for a representative region. The x-axis is weeks prior to stay with the booking curve displayed as occupancy (% of units rented) on the primary y-axis and the solid line corresponding (secondary y-axis) plotting average transacted prices. The combination of at risk inventory, surplus capacity and decreasing prices over the selling horizon naturally lend themselves to an objective of maximizing instantaneous profit and may not require the full complexity of a dynamic programming approach.

For some peak demand periods the price curve may not be monotonically decreasing (owing to elevated demand and tighter supply) and as a result may be better suited for a full DP approach. To fix this limitation, a booking curve heuristic is proposed to provide feedback into the optimization and in essence ensure that prices are not decreased dramatically early within the selling horizon. As all rental units are rented out on a weekly basis, the optimization runs weekly to optimize revenues for the coming stay weeks using weekly inventory and pricing information from both the rental company and its competitors.

Maximization objective:

$$\sum_{w}\sum_{j_1=1}^{N_1}\pi_{j_1,w}\cdot P_{j_1,w}+\theta\sum_{j_2=1}^{N_2}\pi_{j_2,w}\cdot P_{j_2,w} \tag{9}$$

Where

$$\pi_{j,w}=\frac{e^{\eta_{j,w}}}{\sum_{j=1}^{J}e^{\eta_{j,w}}}$$

$$P_{j,w} = x_{j,w} \cdot MP_w$$

Where N_1 and N_2 represent the number of the available properties from Type I and Type II products at the current lead time for arrival week w; θ is the revenue discounting factor for the contract-based product, determined by the profit ratio between Type I and Type II products; MP_w is the weighted market selling price for arrival week w from the similar products; $x_{j,w}$ is the decision variable for the optimization model, indicating the price ratio for the product j in the arrival week w relative to the market selling price of all the similar products.

Cost constraints:

For Type II product,

$$P_{j_2,w} \geq (1 + PM) \cdot C_{j_2,w} \tag{10}$$

Where PM is a managerial set minimum profit margin for non-risk inventory , with $C_{j_2,w}$ the contracted cost associated with the product j_2.

Discount off rack rate constraints:

For both Type I and Type II product,

$$(1 - q_l) \cdot RR_{j,w} \leq P_{j,w} \leq (1 - q_u) \cdot RR_{j,w} \tag{11}$$

Where $RR_{j,w}$ is the rack rate for the product j of arrival week w. Rack rate is the ceiling or brochure price of the product. This constraint sets the boundary for the maximum discount (r_l) below and maximum markup (r_u) above the rack rate. The boundaries can be varied by different lead time and product types.

Discount off market selling price constraints:

For Type I product,

$$(1 - r_l) \cdot MP_w \leq P_{j,w} \leq (1 - r_u) \cdot MP_w \tag{12}$$

Where MP_w is the weighted market selling price for arrival week w for similar products as product j. This constraint sets the boundary for the maximum discount (r_l) below and maximum markup (r_u) above the market selling price for similar products. These boundaries can be varied with lead time and arrival weeks (peak vs. non-peak). For a typical peak week, a wider boundary can provide the algorithm more flexibility to adjust the price in order to achieve a higher revenue potential, while for a non-peak week a relatively narrow boundary can constrain the current price to a reasonable price range aligned with the market price to hedge against the potential unsold risk.

Booking curve constraints:

As described earlier, there are different controlling bands with different confidence intervals for the booking curve heuristic to provide feedback into the optimization. For each product type in a certain arrival week, a booking curve is projected separately by time series analysis to control its own booking pace.

If the utilization is in the lower alarming range of the booking curve, it will feedback a price reduction signal into the optimization. The price reduction (s_1) is determined by two factors - price sensitivity from the choice model and the gap between the current utilization and the targeted utilization.

$$AP_w \leq (1 - s_1) \cdot AP'_w \tag{13}$$

Where AP_w is the average price of the available inventory after adjustment for arrival week w, and AP'_w is the average system price before adjustment. This constraint makes sure that after the price adjustment the average price of the available inventory should be at least marked down by s_1.

However, because on both sides there are equivalent numbers of products, we can simplify the constraint to be:

$$\sum_j P_{j,w} \leq (1 - s_1) \cdot \sum_j P'_{j,w} \tag{14}$$

Where $P_{j,w}$ is the price after adjustment, and $P'_{j,w}$ is the current system price before price reduction.

If the utilization is in the upper alarming range of the booking curve, it will feedback a price increase signal into the optimization. Similar to (s_1,) the price increase (s_2) is a function of price sensitivity from the choice model and the gap between the current utilization and the targeted utilization.

$$\sum_j P_{j,w} \geq (1 + s_2) \cdot \sum_j P'_{j,w} \tag{15}$$

If the utilization is in the lower warning range of the booking curve, a signal of "Do not increase price" will feedback into the optimization from the booking curve heuristic:

$$\sum_j P_{j,w} \leq \sum_j P'_{j,w} \tag{16}$$

Table 2 Year-over-Year comparisons of Model versus Control Performance.

	Occupancy		Average Daily Rate		Revenue per Unit	
	TEST	CONTROL	TEST	CONTROL	TEST	CONTROL
REGION A	1.012	0.989	1.064	1.036	1.077	1.024
REGION B	1.178	1.079	1.056	1.001	1.116	1.078

If the utilization is in the upper warning range of the booking curve, a signal of "Do not decrease price" will feedback into the optimization from the booking curve heuristic:

$$\sum_j P_{j,w} \geq \sum_j P'_{j,w} \tag{17}$$

In order to keep the optimization feasible, some of the constraints are set as soft constraints with a violation penalty in the objective function. The penalty weights associated with different constraints are determined by the priority of the business requirements. For ease of presentation, these soft constraints are omitted in the above description.

Model Implementation and Results

The choice model is estimated in SAS PROC MDC, while the optimization is implemented in Xpress with the nonlinear optimization package. In an effort to assess the model performance, vacation rental units were priced for part of 2018 season in two major markets (we disguise these markets as Region A and Region B). In each region units were priced via the proposed model for 3 out of 7 main arrival months for the entire booking horizon (typically 6 months). To benchmark the model, we compare year-over-year (2018 versus 2017) performance. Table 2 displays Average Daily Rate, Occupancy and Revenue per Unit (year-over-year) growth for the two regions from the model tested period and the 4 other control months.

Table 2 indicates strong model performance across all metrics and both regions - if we look at Revenue Per Unit and compare the Test versus Control relative growth rates we see that the model outperformed the control in Region A by 5.2% (1.077/1.024) and 3.5% in Region B (1.116/1.078).

Using web scraping we are also able to estimate competitor sales and their resulting occupancy during both the Test and Control periods. Table 3 summarizes occupancy lifts for Wyndham Destinations versus their competitors. The 1.024 in Table 3 represents the relative occupancy growth (year-over-year) for the Test versus Control periods (1.024 =

Table 3 Wyndham Destinations Occupancy Performance versus Competition.

	WYND	Competition	Relative
REGION A	1.024	0.983	1.042
REGION B	1.091	0.944	1.156

1.012/0.989 from Table 2) for Wyndham with 0.983 the same metric for the competitor - indicating that Wyndham Destinations grew occupancy 4.2% faster (relative to its competition) in Region A and 15.6% faster in Region B.

In an ideal setting we would have used the model to price a subset of rental units (for a subset of the booking window), comparing the performance of these units over time to those not priced with the model - in essence comparing KPI growth for 2018 (versus 2017) for region/month pairs where the model was used relative to those where the model was not. We realize KPI improvement could have been caused by factors other than our model, we attempt to account for this by comparing occupancy growth for tested region/months versus those estimated from competitors. So while we were unable to perform an isolated test-versus-control study to validate model performance the combination of Tables 2 and 3 indicate strong model performance. Table 2 benchmarks unit performance across periods of model use (Test) versus those without model use (Control), indicating gains from model use, whereas Table 3 benchmarks/controls for overall market condition changes. Together, Table 2 and 3 indicate that not only did the model generate gains (Test versus Control) but that these gains can be attributed to the model as the gains were above and beyond that of the underlying demand as measured by competition occupancy gains. This comparison of relative (to the market) growth in occupancy is encouraging given it came with a solid increases in ADR as well. We are unable to illustrate the similar results to Table 3 for daily rate and revenue per unit as we don't know prices at which the competition transacted as the scrapped prices are retail prices for vacation units only whereas units are usually sold as packages with air.

Summary and Limitations

In this paper we develop and document the implementation of a choice-based approach to revenue management for vacation rentals. The heterogeneous nature of vacation rental inventory lends to use of choice-based approaches to setting prices. We utilize a myopic optimization approach to setting prices, while this may be a limitation of our approach, the re-optimization of these prices, combined with both the monotonically decreasing nature of rental unit prices over the selling horizon and booking curve heuristics, decreases the effect of this limitation as evidence by the revenue per rental unit lift (as high as 5.2%) attained through model use.

While our approach has greatly improved pricing decisions it is not without limitations. The current approach assumes a similar (to its own) price change frequency by Wyndham's competition as price scraping is performed on a weekly basis and as a result may miss price changes. Given that most price changes are downward in nature the impact of capturing a price change a few days late, while minimal, may result in having slightly higher prices in the interim. The model, is in essence a competitive pricing tool (versus a revenue management system) owing to the lack of incorporation of a demand forecast. The use of a booking curve heuristic to constrain/signal price changes assumes current booking patterns are close to optimal, and while this may seem a strong assumption, the myopic nature of the optimization (preference for early sales at higher prices) should improve the booking curve over time. The use of a booking curve heuristic will also limit suggested prices and in essence assumes stable price and booking patterns, so while this approach limits future prices it also serves to stabilize potential competitive reactions that may result from aggressive price moves.

As discussed earlier, the rolling window approach for estimation of heterogeneity in price effects assumes utilities are of the same scale during each estimation stage as we don't re-scale utilities during the second stage of estimation. While utilities may be different peak versus off-peak, our two-stage estimation process is re-estimating utilities over the booking horizon for a fixed stay period and should be consistent with this assumption, but nonetheless may still result in unknown bias for our price effects.

While model results and field tests were very encouraging, full scale deployment of the model has been hindered by the spin-off of Wyndham Worldwide into two separate publicly traded companies - Wyndham Hotels and Resorts and Wyndham Destinations.

References

Anderson, CK. And X Xie. 2016. Dynamic pricing in hospitality: overview and opportunities. $International Journal of Revenue Management$ 9 (2-3), 165-174.

Coretjens, M.L. and D.A. Gautschi. 1983. Formal choice models in marketing. $Marketing Science$, 2 pp. 19-56.

Elmaghraby, W. and Keskinocak, P. (2003). Dynamic pricing in the presence of inventory considerations: research overview, current practices, and future directions. $ManagementScience$, Vol. 49, No. 10, pp.1287–1309.

Gibbs, C., Guttentag, D., Gretzel,U., Morton, J., and A. Goodwill. (2018). Pricing in the sharing economy: a hedonic pricing model applied to Airbnb listings. $JournalofTravel\&TourismMarketing$, 35:1, 46-56.

Gallego, G., Ratliff, R.M. and Shebalov, S. (2015). A general attraction model and sales based linear program for network revenue management under customer choice. $OperationsResearch$, Vol. 63, No. 1, pp.212–232.

Garrow, L. and Koppelman, F. (2004). Multinomial and nested logit models of airline passengers' no-show and standby behavior. $JournalofRevenueandPricingManagement$, Vol. 3, pp.237–253.

Kwok, L. and Xie, K.L. (2018). Pricing strategies on Airbnb: are multi-unit hosts revenue pros? $InternationalJournalofHospitalityManagement$, 67, 174–84.

Magno, F., Cassia, F. and Ugolini, M. (2018). Accommodation prices on Airbnb: effects of host experience and market demand. $TheTQMJournal$, Vol. 30 No. 5, pp. 608-620.

McFadden, D.L. 1973. Conditional Logit Analysis of Qualitative Choice Behavior, In: Frontiers in Econometrics, ed. by P. Zarembka, Academic Press: New York. pp. 105-142.

Orkin, E.B. (1988). Boosting your bottom line with yield management. $CornellHospitalityQuarterly$, February. pp. 52-56.

Talluri, K.T. and van Ryzin, G.J. (2004). The Theory and Practice of Revenue Management, Springer Science and Business Media, Inc., Boston, ISBN 0-387-24376-3.

Ye P., Qian J., Chen J., Chen-hung W., Zhou Y., De Mars S., Yang F., and Zhang L. (2018) Customized regression model for Airbnb dynamic pricing. Proc. 24th ACM SIGKDD Internat. Conf. Knowledge Discovery Data Mining (ACM, New York), 932–940.