Transactions on Engineering and Computing Sciences - Vol. 11, No. 3
Publication Date: June 25, 2023
DOI:10.14738/tecs.113.14739.
Battiti, R., Brunato, M., & Battiti, F. (2023). Pragmatic Estimation of the Elasticity of Demand in Hospitality from Noisy Reservations. Transactions on Engineering and Computing Sciences, 11(3). 26-47.

Pragmatic Estimation of the Elasticity of Demand in Hospitality from Noisy Reservations

Roberto Battiti
Universita di Trento, Dipartimento di Ingegneria e Scienza` dell'Informazione, Via Sommarive 14, I-38123 Trento, Italy

Mauro Brunato
Universita di Trento, Dipartimento di Ingegneria e Scienza` dell'Informazione, Via Sommarive 14, I-38123 Trento, Italy

Filippo Battiti
Ciaomanager Srl, Via Ragazzi del '99 17, I-38123 Trento, Italy

ABSTRACT
In hotels, future prices should be determined based on predictions of future demand and the responses of clients to price changes. Price changes will actively affect future demand. If this effect is neglected, opportunities for promptly and accurately setting prices will be lost, as well as potential increased profits. The role of elasticity of demand in setting optimal prices is well-developed in Economics, but the practical estimation of elasticity for a specific hotel from data about the pickup of reservations is worth investigating. In this paper, we highlight the risk of estimations based on a single A/B test and propose practical rules and pragmatic experiments. The analysis is based on statistics but simplified so that the results can be easily applied in single hotels without excessive disruption of daily operations. After defining rules to derive error bars on the estimates, we experiment in different situations, including estimations from scratch or by gradually tracking abrupt or seasonal changes, that are more realistic for a hotel in operation. Our methods can be useful to hotel managers who wish to decide about prices based on measuring the response of their potential customers, in a data-driven manner and based on statistically significant estimates.

Keywords Pickup elasticity of demand, Revenue management, Optimal room pricing, Rule of thumb

INTRODUCTION
Hotel sell their finite inventory of rooms (organized into sets of room types) over a booking window, that can be as long as one-two years in advence. Revenue Management Talluri *et al.* (2008) deals with deciding what to sell, when to sell, and for which price, in order to increase revenue and profit. Hotel Revenue Management is a complex area characterized by nonlinearities, many parameters and constraints, and stochasticity, in particular in the demand by customers. It uses models of the demand for periodic forecasts, models of how customer respond to prices (elasticity of demand) and optimization methods. In large hotels and real-

world contexts it suffers from the *curse of dimensionality* Bellman (2015): when the number of variables increases (number of rooms, number possible prices, number of reservation rules) exact solutions by dynamic programming or by alternative global optimization techniques cannot be used and one has to resort to intelligent heuristics, i.e., methods which can improve current solutions but without formal guarantees of optimality. Because learning about the context and optimizing are interrelated in this area, methods integrating machine learning from data and intelligent optimization are a viable option Battiti *et al.* (2008), Battiti and Brunato (2018).

Defining prices for selling accommodations is a critical managerial activity in hospitality. Finding proper prices can be the fastest and most effective way to increase profits (Marn *et al.* 2003). Atkinson *et al.* (1995) argue — with some dramatization — that pricing is frequently the only active revenue-creating decision made within a business. From the origin of Economics, a basic "input signal" for pricing is given by the response of customers: the market price reflects the interaction between supply and demand, with the additional complication for the hospitality sector of a *perishable inventory*: if a room night is not sold at check-in time, its potential revenue is lost forever. Furthermore, the rooms inventory is fixed in the short term: hoteliers cannot respond to demand surges by building more rooms to sell.

Despite its relevance and the abundance of academic studies on pricing, anecdotal evidence and research in hospitality management manifest a significant gap between the interest of hoteliers and the contributions of academics, and persistent difficulties in the daily activity of hotel managers. Walker (1997) argues that in pricing decisions hoteliers cannot have advanced knowledge of how much demand will change in response to a given change in the room price. Middleton and Clarke (2012) feel that for most tourism businesses, "pricing is more art than science." Schmidgall (2002) agrees that establishing prices to maximize revenues is extremely difficult. While normative models developed by academics "are impressive in their mathematical sophistication and claims to internal validity, few efforts are marked by the pragmatism necessary to impact on managerial practice" (Bonoma *et al.* 1988). For example, common practices include establishing prices based largely on competitors' prices (*going-rate pricing*), i.e., assuming that somebody else (arguably with better methods) already solved the problem, or dynamically changing prices via rigid rules based on occupancy levels, or setting prices based on some experience accumulated from the previous years.

From the eighties, more principled ways of adjusting prices and/or room allocation are studied in Revenue Management (RM) (Talluri *et al.* 2008, Brunato and Battiti 2020) which adopts demand forecast and mathematical optimization to identify revenue- and profit-maximizing prices. In the large area of Revenue Management, which cannot be reviewed in this short paper, two issues are often not given sufficient attention in standard methods: the explicit treatment of estimation errors and the circular nature of the process. After estimating the demand, prices are changed so that the realized demand will be *modified* actively and purposefully w.r.t. the prediction. For example, if overbooking is predicted for a future high-season period (Fig. 1), and an opportunity for raising prices is identified, one should determine an optimal price increase, one that will reduce the demand but not in an excessive manner that will leave rooms unsold.

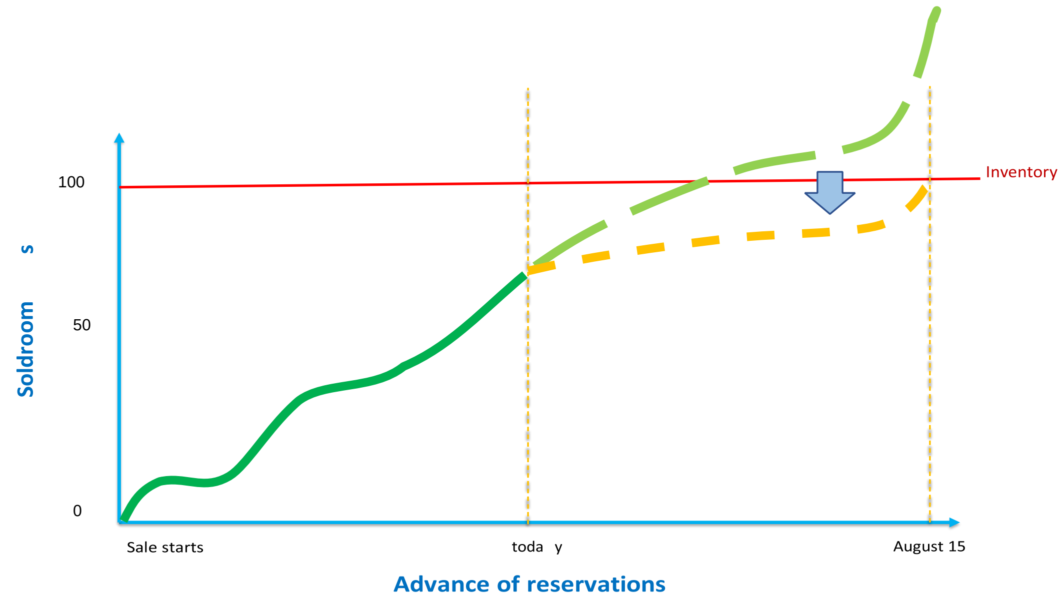

Figure 1: Booking curve for a check-in date (Aug 15). Solid line: Reservations received up to today (OTB - on the book). Dashed curve: forecast of future reservations. Dotted curve: forecast after increasing prices (modified by the price elasticity of demand). The optimal price increase should lead to a full hotel without overbooking at the check-in date.

Over-reacting can be dangerous and lead to reduced profits. To accomplish an optimal reaction, how the demand will be changed by price changes has to be measured, an area related to the concept of *demand elasticity*. Without a careful estimation of the elasticity, the concept of *optimal price* is of mathematical interest but is difficult to apply in practice. We do not underestimate the huge impact of current RM practices on profitability but we argue that proper estimation of elasticity can improve results even further.

In this paper, we assume the point of view of the hotel owner or manager, and we aim at defining simple but principled rules which can be used in practice to measure the effect of pricing on the demand. In particular, we focus on the problem of estimating the (own-) *elasticity of demand w.r.t. prices ceteris paribus* ("other things being equal or held constant"). The fact that lowering prices will stimulate demand and raising prices will quench it is part of the standard knowledge base of hoteliers but measuring the effect in a scientific and *statistically significant manner* is a subject worth investigating.

Intrinsic difficulties in estimating the elasticity are related to various facts. First, many hotels are dealing with numbers of daily reservations that are far from infinite, often they are small (and integer) numbers in the tens. The variance of the estimates can be huge as well as the risk in using estimates for correcting prices. Then the demand is dynamic and changes rapidly, depending on the season, the day of the week, etc., in ways that may interfere with the estimation process. As an extreme example, if the elasticity is evaluated by aggregating high- and low-season data, one may wrongly conclude that high prices (in high season) are "causing" a larger demand and occupancy. The possible confusion between correlation and causation and the effect of hidden variables (like the season) are lurking around. Price endogeneity occurs when prices are influenced by demand, i. e., higher prices are observed when demand is high and lower prices are observed when demand is low. Failure to correct for price endogeneity is critical, as it will result in biased estimates and incorrect elasticity calculations.

In addition, estimates are *censored* in a statistical sense: often only the partial demand leading to accepted reservations is observed. E.g., if only reservations are counted and the hotel reaches full occupancy, no further signal about demand changes can be measured. Improvements in measuring the entire potential demand can be accomplished by analyzing client searches on the hotel website, but a growing portion of searches happen by intermediaries like OTA's, making a solid estimation very difficult. Similarly, many customers who are potentially interested but avoid reserving at the current price remain invisible.

A final difficulty is related to *designing experiments* (Cox and Reid 2000) in hospitality and executing them with the explicit goal of estimating the elasticity, an activity which should not interfere too much with the daily operation. Experimenting is costly and the design of the experiment should aim at estimating elasticity while controlling the experimental costs. On the other hand, estimating the elasticity from price changes that occur "spontaneously" as a side-effect of daily price-setting activity is subject to several variables and to a level of stochastic noise that may easily hide the relevant signal.

While following statistics, we present the core of the analysis so that practitioners can easily understand and apply the results without advanced knowledge of differential calculus or statistics.

In the following part of this paper, Section 2 summarizes the state of the art in measuring elasticity, in particular for the hospitality sector, Section 3 highlights the danger and risk of naive estimations and introduces empirical laws derived from statistics with associated confidence intervals, Section 4 analyzes experimental results obtained by using the proposed estimation methods, and Section 5 presents concluding remarks.

ELASTICITY OF DEMAND IN THE HOSPITALITY SECTOR

One cannot set optimal prices without understanding how customers respond to them, by refusing offers, or by accepting and buying stays. We summarize basic formulations from classical Economics and then address complications arising when estimating customer response in hospitality from the partial, noisy, and limited reservations received. Let's assume that the quantity of rooms Q bought in a day (without considering the issue of finite and perishable inventory) is described by a function $Q(P,\xi)$, in which P is the price and ξ is an array of variables summarizing the additional relevant factors beyond the price that are influencing the customers' decision. The first radical assumption is to compare situations for the same (or very similar) value for all remaining factors ξ.

A dimensionless coefficient, not depending on units of prices, to measure how the number of rooms sold varies as price changes is the "price elasticity", which measures how *percent changes* in volume are related to percent changes in price.

Definition of Elasticity of Demand

For a small variation of price *dP* causing a change *dQ* in rooms bought, the *price elasticity of demand* E_d can be expressed as:

$$\text{Elasticity}=E_d=(Dq/Q)/(dP/P) = (dQ/dP)/(Q/P) \qquad (1)$$

or the ratio of (*dQ/dP*) to the value of the average function (*Q/P*).

Because quantities tend to decrease when prices increase, theoretical price elasticities are almost always negative, although it is customary to ignore the sign (with a slight risk of ambiguity).

It is easy to understand the relevance of elasticity on revenue changes. If one increases the price by *dP* and measures a change *dQ* in quantity, the revenue, given by quantity times price, will be changed as follows:

$$(Q+dQ)(P+dP) = Q\,P+dQ\,P+dP\,Q+dQ\,dP \qquad (2)$$
$$= Q\,P+dP\,Q\,(E_d+1)+dQ\,dP \qquad (3)$$
$$\approx Q\,P+dP\,Q\,(E+1) \qquad (4)$$

where the quadratic term *dQ dP* becomes negligible for small changes and is omitted in the last step. One concludes that the revenue after the change is (approximately) equal to the previous revenue if the elasticity is equal to −1. This is the case if one decreases the price by 10% and observes a quantity increased by 10%. If the elasticity is −0.5 and one decreases the price by 10%, one observes a quantity increased by 5%, which is not sufficient to offset the reduced revenue for all reservations. Therefore, the change has an overall negative impact on revenue. In general, the demand is *inelastic* when the elasticity is less than one (in absolute value): price changes have a relatively small effect on the quantity demanded. The demand is *elastic* when the elasticity is bigger than one.

More formally, by taking the derivatives of revenue w.r.t. quantity *Q*, one can easily demonstrate that revenue is maximized when the price is set so that the elasticity is exactly one. The following equation holds:

$$R' = P(1+1/E_d) \qquad (5)$$

where R' is the *marginal revenue* and *P* is the price

A seminal analysis of the role of elasticity in setting optimal prices is in Nash (1975). Nash's analysis (equations 2-5), can be used to determine whether a price increase or decrease would be appropriate in a given situation. A large number of studies in Economics are devoted to demand elasticity. Tellis (1988) is a meta-analysis of econometric studies in hundreds of markets starting from the sixties. In general, it is recognized that studies that use overt intervention to elicit customer response may distort estimates because they elevate price consciousness.

Studies About Elasticity in Hospitality

Because the inventory of rooms in hotels is limited, hoteliers study the arrival of reservations for future days (the s.c. *booking curves* or *pickup curves*) to estimate the overall demand for future check-in days. If a situation of overbooking (more rooms requested and possibly booked than the available inventory) is predicted for a future period, and if one aims at profit maximization, prices should be increased for that period (Fig. 1). This reaction corresponds to the rule of thumb of "avoiding selling out too soon", and of keeping precious rooms for late-arriving requests, often by customers prepared to pay more, like businessmen.

Knowledge of elasticity allows determining the *optimal price increase* for that future day so that the entire inventory will be sold without overbooking and without being left with empty rooms at check-in time. Every surge of demand for future days, when identified in advance, becomes an opportunity to increase profit by rapidly increasing prices, low-hanging fruit to be picked by Revenue Management. The response has to be prompt, but an excessive increase will quench requests too abruptly and leave the hotel with unsold capacity, this is why price adjustments need to consider elasticity.

Very precise estimations of the elasticity are not needed in this case. If the booking window is large, and therefore the arrival of reservations is distributed over many days in advance w.r.t. the check-in day, an approximated setting of the price in the early days can be corrected and fine-tuned later on, when more accurate estimates become available. On the contrary, more careful estimations are needed if the hotel is in the low or middle season, far from full occupancy. In this situation, the hotel management may be tempted to lower prices to stimulate demand, but this is not always the optimal strategy. Keeping the prices stable or even increasing prices could be a better strategy. In this case, knowledge of elasticity and the variable and fixed costs are critical to deciding about increasing or decreasing prices, or even closing the hotel if the predicted revenue is not even sufficient to cover the costs.

Hiemstra and Ismail (1993) found that the price elasticity of demand varies across hotel segments' room rates. Lewis and Shoemaker (1997) explore room rate setting in connection with a consideration of what the market can bear. A macro-economic study of the U.S. lodging industry including estimates of the elasticity of night stays w.r.t. GNP or w.r.t. room rates is for example Wheaton and Rossoff (1998). The hotel industry's discounting philosophy has uncharitably been described as "being similar to negotiations in a flea market or on a used car lot" in Hanks *et al.* (2002). Fibich *et al.* (2005) derive an expression for the price elasticity of demand in the presence of reference price effects ("the price consumers have in mind and to which they compare the shelf price of a specific product"). Canina and Carvell (2005) study the demand for urban hotels and found that demand is rather price-inelastic and price elasticity measures vary across market segments. It argues that estimating the price elasticity of demand for the overall lodging industry can be different from the price elasticity of demand at the property level.

Researchers studied various methods to estimate the price elasticity of demand in various contexts (Chung 2006, Skuras *et al.* 2006). The estimates produced in many studies have large confidence intervals and are of limited utility for practical use. Chung (2006) concludes that ignoring quality adjustment in either prices or quantities can cause biased price elasticities.

By analyzing aggregated data from the hotel industry at the beginning of this century and considering the context of a hotel's competitive set Enz *et al.* (2009) argue that the effectiveness of pricing strategies critically depends on the price elasticity of demand. It also mentions that empirical studies on the price elasticity of demand often produce disappointing results, that estimates based on aggregate demand are of limited practical use, and that the difficulty and uncertainty in the estimates cause most managers to "steer clear of estimating demand curves when making pricing decisions". They advocate a more principled approach based on demand data that are specific to a given property.

The effect of advertising on elasticity is investigated in Chen *et al.* (2015). Aziz *et al.* (2011) explicitly represent demand elasticity in a simulator. Elasticity is present in Bayoumi *et al.* (2013), which considers a case study of a specific hotel and estimates elasticity by fitting a probit function to the historical data. To be more robust w.r.t. errors in computing elasticity, simulations with different elasticity values are adopted. Vives *et al.* (2019) consider a log-linear function form ($\log Q = \alpha \log P + ...$), derived the own-price point elasticity via regression for different seasons, dates of stay, etc., and tests the model on a specific hotel. Apart from exceptions like the previous paper, the majority of the demand models in the literature estimate general market price elasticities (Canina and Carvell 2005, Song *et al.* 2011, Cross *et al.* 2009, Tran *et al.* 2015).

A critical step for a specific hotel is to estimate the demand response of that hotel with its characteristics and target customers to price variations, particularly for touristic hotels that are increasingly impacted by the emerging online transient customers. Most elasticity studies involve demand data that are aggregated and that, to a first approximation, can be considered real numbers. The fact that data are quantized (e.g., approximated by integer numbers) can render such methods dangerous and misleading (Vardeman and Lee 2005). Changes in the daily requests tend to be small numbers, in particular for medium-sized properties. The demand is *highly stochastic*: even if the average demand is constant the variation from day to day can be large. The large standard deviation that is characteristic of hotel demand can almost hide the "weak signal" of variations caused by price changes.

The demand is *intrinsically dynamic*: the reservations received in time for a given check-in date are characterized by a *booking curve* with a pattern depending on the hotel, the season, the room type, the customer segment, etc. Again, the number of reservations received may depend more on the booking curve pattern than on price variations and this effect has to be properly discounted before recovering changes caused by price.

A final issue has to do with the possibility to experiment. One could theoretically *design* appropriate experiments to estimate the elasticity (e.g. by randomizing higher and lower prices for the same future day) but experiments risk disrupting normal operations, annoying customers, reducing profit, and must be administered with care.

ESTIMATING PICKUP ELASTICITY: EMPIRICAL LAWS DERIVED FROM STATISTICS

To obtain pragmatic methods which can be used by medium-size hotels one must simplify the problem and consider affordable techniques, based on data that is easily collected as part of the daily hotel operation. We assume that the only data observed are the arrival of reservations for future days (including modifications and cancellations). In addition, we assume that the external context (like the competition) is sufficiently stable over the estimation period. This measure can be defined as “pick-up elasticity”, in which pick-up refers to the process of collecting reservations in time for future check-in days. In “pick-up elasticity” the quantity Q is therefore given by the quantity of reservations received for future days.

It is well known in science and statistics that raw comparisons of data are not sufficient for sound decisions. For example, if a hotel receives 12 reservations today for a future day X (when rooms are selling for 90$), it raises prices to 100$ and tomorrow it receives 15 reservations for the same future day, of course, it cannot conclude that higher prices cause more reservations. If yesterday a hotel received one request for a future day, it increases the price by 1$ (from 100 to 101), and then receives two requests on the next day, the raw measurement of elasticity is going to be an enormous 100, while typical theoretical and practical values are close to the equilibrium value of 1. There can be tens of “hidden variables” explaining the result, like different people searching with different budgets, different tastes, different impressions derived from the hotel advertisement, a large group reserving for day X, etc. Before taking a decision we should always make sure that the observed difference is *statistically significant*, i.e., not likely to occur by chance, by the stochastic “noise” implicit in many variables which cannot be controlled (Kahneman *et al.* 2019), but instead likely to be attributable to a specific cause.

For marketing and management reasons modern hotels do not sell single rooms but accommodation types (like “Basic”, “Superior”...). Estimating elasticity in hospitality for an individual business with tens of rooms of a specific type is a kind of “perfect storm” with small and integer numbers counted every day and complications to measure the effect of prices while being immersed in a variety of concurrent changes (day of the week, weather conditions, season, advertisement, occupation levels, groups, multiple-duration reservations, etc.).

On The Risk of A/B Testing with Small Reservation Numbers

A basic source of knowledge for estimating customers’ responses is to change prices and measure the resulting reservations received (for comparable periods and similar overall contexts), also called A/B testing (Kohavi and Longbotham 2017). The result is stochastic. Let’s assume that the experiment considers two prices p^+ and p^- and let λ^+ and λ^- be the “true” rates of arrival of reservations corresponding to the prices. Let’s assume for the current analysis that the distribution of daily reservations is described by a Poisson distribution (Papoulis 1990) After fixing a period of analysis, in our case a day, if λ is the average daily number of reservations received and k is the actual number observed, the probability $\Pr(X=k)$ of observing k reservations is given by:

On The Risk of A/B Testing with Small Reservation Numbers

A basic source of knowledge for estimating customers' responses is to change prices and measure the resulting reservations received (for comparable periods and similar overall contexts), also called A/B testing (Kohavi and Longbotham 2017). The result isstochastic. Let's assume that the experiment considers two prices p^+ and $p^-$and let $\lambda^+$and λ^- be the "true" rates of arrival of reservations corresponding to the prices. Let's assume for the current analysis that the distribution of daily reservations is described by a Poisson distribution (Papoulis 1990) After fixing a period of analysis, in our case aday, if λ is the average daily number of reservations received and k is the actual number

observed, the probability Pr(X=k) of observing k reservations is given by:

$$\Pr(X=k) = \lambda^k e^{-\lambda} / k! \quad (6)$$

The distribution has a single parameter λ which is the expected number of reserva- tions and also its variance (the standard deviation is, therefore,$\sqrt{\lambda}$). A Poisson model is justified when reservation events are independent (the occurrence of a reservation does not affect the probability that subsequent reservations arrive, that is not the case

— as an example — for group reservations), with a constant arrival rate λ. The Pois- son distribution is also the limit of a binomial distribution, for which the probability ofsuccess for each trial equals λ divided by the number of trials, as the number of trials approaches infinity. A way for a hotel manager to visualize the process is that all peo-ple in the world — a good approximation of an infinite number — throw a dice every day (or every time interval of reference), with a negligible individual probability of re-serving a specific hotel but so that the overall average reservation rate is λ . In practice, hotel reservations may show over-dispersion and/or multi-modality (Song 2021), witha different variation w.r.t. Poisson, but for the sake of brevity, we consider only Poissondistributions in this paper.

Given a sample of n measured values $k_i \in \{0,1,...\}$, for i = 1,...,n, the maximum likelihood estimate is Papoulis and Saunders (1989)

$$\widehat{\lambda}_{\mathrm{MLE}} = \frac{1}{n}\sum_{i=1}^{n} k_i. \quad (7)$$

A simple experiment is to count the number of reservations (in a day) obtained for two different prices, to identify the price leading to the bigger number of reservations. This context is called *Poisson races*: the "horses" in the race are the two prices, and the winning price is the one leading to more reservations. Of course, single estimates can be seriously wrong, e.g.,$\widehat{\lambda^+}$can be bigger than$\widehat{\lambda^-}$ while λ^+ is smaller than λ^-. This can be the case if more reservations are observed for a bigger price (ceteris paribus).

Let's imagine that a hotelier is counting reservations for two different prices, modeled by two different Poisson distributions $X \approx Poi(\lambda)$ and $Y \approx Poi(\mu)$, with $\lambda < \mu$, and let's evaluate the probability of an incorrect judgment, i.e. of obtaining a count from a single experiment $X > Y$. The distribution of $K = X - Y$ is described by the Skellam distribution (Irwin 1937):

$$p(K;\lambda,\mu) = \Pr\{K = k\} = e^{-(\lambda+\mu)} \left(\frac{\lambda}{\mu}\right)^{k/2} I_k(2\sqrt{\lambda\mu}) \quad (8)$$

where Ik(z) is the modified Bessel function of the first kind.

The plots in Fig. 2 show the probability of incorrect judgment from a single A/B test, i.e., the probability of concluding that the average number of reservations is bigger or equal when the truth is that it is less. The results are very similar (although always less than 50%) if one excludes the case of an equal result. Because one is estimating a derivative, the percent change in price has to be small (let's say less than 10%) so that the corresponding true change in quantity will be comparably small (if the absolute

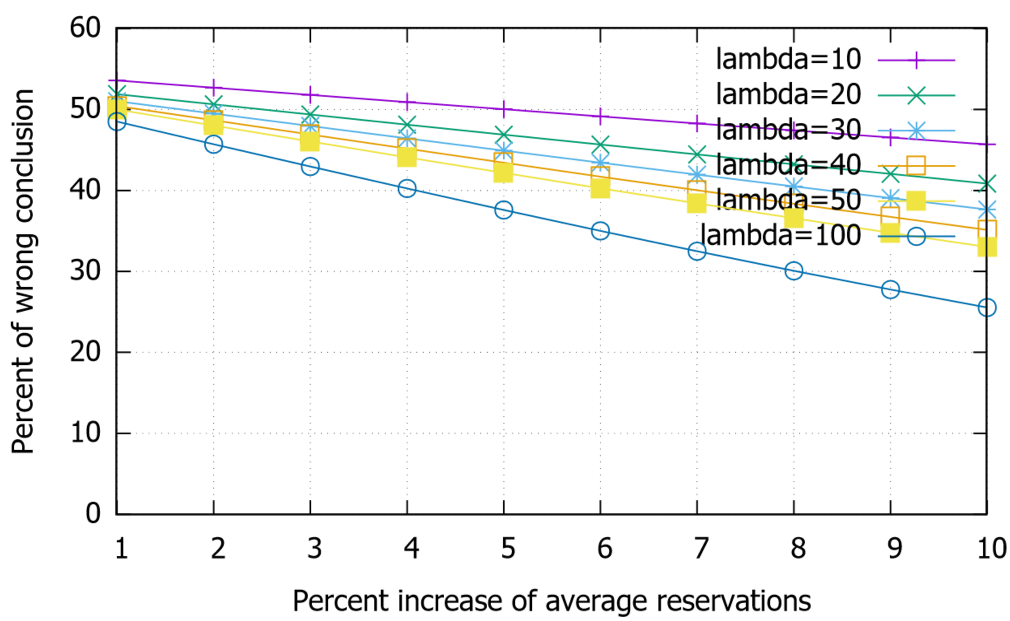

Figure 2: Percent of wrong conclusions as a function of the percent increase in the true average reservations, based on a single experiment about reservation counts. Curves are for different values of the initial average reservations λ

elasticity is not too far from 1). In addition, excessive price changes will disrupt daily operations, surprise customers, and cause potential loss of profit. This is why we limit consideration to a maximum "real" increase of 10% in Fig. 2.

When the initial average reservations range from 10 to 100 (reasonable values for the number corresponding to a specific room type of standard hotels) the probability of incorrect judgment is very large — close to 50% — and in any case bigger than 25% if the true change is less than

10% of the initial number of reservations. For example, with 20 average reservations and a price increase inducing a 10% reduction, one will measure no change or an increase in reservations with a probability of about 40%.

Therefore, deriving a conclusion from a single A/B test — with a limited number of reservations — about increasing or decreasing prices is not too far from deriving conclusions by flipping a coin, a fact that should convince also the most reluctant hoteliers of the need for repeated statistical estimations which consider also error bars and the probability of incorrect conclusions.

Confidence Intervals for Elasticity

In statistics, one accepts a certain probability of wrong estimation (specified as a target α value) and one can derive an interval around the estimated mean in which the true value should be (*interval estimation*). One introduces a confidence level ($1-\alpha$). If one runs many experiments and uses the confidence level to derive confidence intervals (CI), the confidence level can be interpreted as the long-run proportion of experiments so that the computed confidence interval contains the true value.

We summarize the recent comparison of nineteen confidence intervals for the Poisson mean (Patil and Kulkarni 2012) to derive some pragmatic estimates which can be used by hotel managers. It is useful to distinguish the case of a "large" average (λ) value, say more than 4, from the case of "small" values. For large values, the seminal estimate of Garwood (1936) is still among the recommended choices. From Garwood (1936), given a single observation k from a Poisson distribution with mean λ (a single count of reservations in a day), the confidence interval for λ with a confidence level ($1-\alpha$) is:

$$\tfrac{1}{2}\chi^2(\alpha/2;2k) \le \mu \le \tfrac{1}{2}\chi^2(1-\alpha/2;2k+2), \qquad (9)$$

where χ is the chi-squared distribution. For small average values, according to Barker (2002) — a case which can be relevant if the elasticity is measured for a single apartment, a small bed and breakfast, or even for a specific room type of a hotel with a small number of rooms — one among the best-suggested estimates is the modified Wald: the interval is between the lower limit:

for $k = 0$: 0, for $k > 0$: Wald limit (10)

and the upper limit:

for $k = 0$: $-\log(\alpha/2)$, for $k > 0$: Wald limit (11)

in which Wald limit is

$$k \pm Z\alpha/2\sqrt{k} \qquad (12)$$

where $Z_{\alpha/2}$ is the $(1-\alpha/2)100$-th percentile of the standard normal distribution. The formulas can look intimidating for a hotel manager but are easily realized with statistical software.

Let's now analyze and simplify the formula to give more intuition and practical indications. A first observation is related to repeating the experiment (e.g., for multiple equivalent days), obtaining n measured values k_i each drawn from a Poisson distribution with mean λ. Luckily, the sum of independent Poisson distributions is still Poisson, and the λ's are simply summed to obtain the expected value for the sum. One can therefore start from: $k = \sum_{i=1}^{n} k_i$, calculate an interval for $\mu = n\lambda$, and then derive the interval for λ.

If the standard deviation $\sqrt{\lambda n}$ is adopted as the error bar for μ, the error bar on the estimate $\hat{\lambda}$ $=1/n\sum_{i=1}^{n} k_i$ is $\sqrt{\lambda/n}$.

To estimate the elasticity, we need to estimate a difference between two quantities of reservations (for two different prices), let's call them Q and Q'. According to standard rules for the propagation of uncertainty, the approximation for the standard deviation of a difference $Q-Q'$ of two stochastic variables, from their respective standard deviation σ_Q and $\sigma_{Q'}$, and from their covariance $\sigma_{QQ'}$, is:

$$\sigma^2_{Q-Q'} \approx \sigma^2_Q + \sigma^2_{Q'} - 2\sigma_{QQ'} \quad (13)$$

$$\approx 2\sigma^2_Q \quad (14)$$

in the reasonable assumption that the two quantities are approximately uncorrelated (zero covariance) and that the quantities have the same standard deviation. A demonstration can be found in Ku *et al.* (1966) and is summarized in the footnote [11]

Because the price difference is known (it is not a stochastic variable), one, therefore, obtains an error bar on dQ/Q equal to $\sqrt{2/(n\lambda)}$ and an error on the elasticity equal to:

[1] Any non-linear differentiable function, $f(a,b)$, of two variables, a and b, can be expanded with Taylor's formula as:

$$f \approx f^0 + \frac{\partial f}{\partial a}a + \frac{\partial f}{\partial b}b$$

and, after using the formula for variance of a linear combination of variables:

$$Var(aX + bY) = a^2 Var(X) + b^2 Var(Y) + 2ab * Cov(X,Y)$$

one obtains:

$$\sigma_f^2 \approx \left|\frac{\partial f}{\partial a}\right|^2 \sigma_a^2 + \left|\frac{\partial f}{\partial b}\right|^2 \sigma_b^2 + 2\frac{\partial f}{\partial a}\frac{\partial f}{\partial b}\sigma_{ab}$$

where σ_f is the standard deviation of the function f, σ_a is the standard deviation of a, σ_b is the standard deviation of b and $\sigma_{ab} = \sigma_a\sigma_b\rho_{ab}$ is the covariance between a and b. The application to the case of a difference can be obtained by using the formula for $f = a\,A - b\,B$.

$$\text{Elasticity} = \widehat{\text{Elasticity}} \pm \sqrt{\frac{\frac{2}{n\lambda}}{dP/p}}$$

$$= \widehat{\text{Elasticity}} \pm dp/p \sqrt{\frac{2}{n\lambda}} \quad (16)$$

where $\widehat{\text{Elasticity}}$ is the estimate.

The error bar is going to zero in a manner inversely proportional to the square root of the product λn. If one fixes the desired error bar (e.g. of 0.2, that is reasonable for elasticity estimates) and solves the equation, one obtains:

$$dp/p \sqrt{\frac{2}{n\lambda}} = 0.2 \quad (17)$$

$$\lambda n = 2/(0.2\ dP/P)^2 \quad (18)$$

With some color, we may call the above relationship the "Rule of patience for elasticity estimation". Finally, if we assume a 10% change in the price ($dP/P = 0.1$), the ballpark estimate for the number of experimental measures required to have an elasticity estimate with an error bar of 0.2 is $n \approx 5000/\lambda$. The first important observation is that the error bar on the elasticity depends on the product λn. Hotels with smaller average reservation rates λ will need to compensate with a larger number of tests n to get equivalent relative errors. For example, 50 tests are required if the number of daily reservations is 100, while about 200 tests are required for 25 daily reservations.

A second observation and approximation are related to the number of rooms. If a hotel is working at reasonable occupation levels and if each reservation is for a single night, one expects an average total number of reservations received in a day (for all future possible check-in dates) similar to the total number of rooms. There can be exceptions (lucky days in which the number of reservations received for future days is

bigger than the number of rooms), but they will need to be counterbalanced by days with fewer reservations otherwise, overbooking will be reached. We can therefore substitute to obtain:

$$(\text{number of rooms})\times(\text{number of experiments}) \approx 2/(0.2\ dP/P^2) \quad (19)$$

There are no free meals. Linearly, if the hotel has a smaller number of rooms, it needs to increase the number of experiments to compensate and to get equivalent relative error bars.
The results of some experiments in measuring elasticity are shown in Fig. 3. In these experiments, the true elasticity is one, one considers different initial rates of daily reservation arrivals (10,20, 100), changes prices by 10%, and measures the effect on reservation arrivals to estimate elasticity. The estimations are repeated for a given number of experiments *n*. Let's underline that the assumption of true elesticity equal to one is used in the experiments but it does not limit in any way the validity of the approach for arbitrary real values of the elasticity. As expected, the initial estimates are very noisy, then only very gradually do the estimates

converge towards the true value of 1 (with a standard deviation which decreases like $1/\sqrt{\lambda n}$. After observing these concrete cases, the amount of stochasticity in the estimate and the difficulties of naive estimations should become evident. In three different series of experiments, as expected, the relevant parameter is the product λn. The different evolutions become similar when plotted as a function of λn on the x-axis.

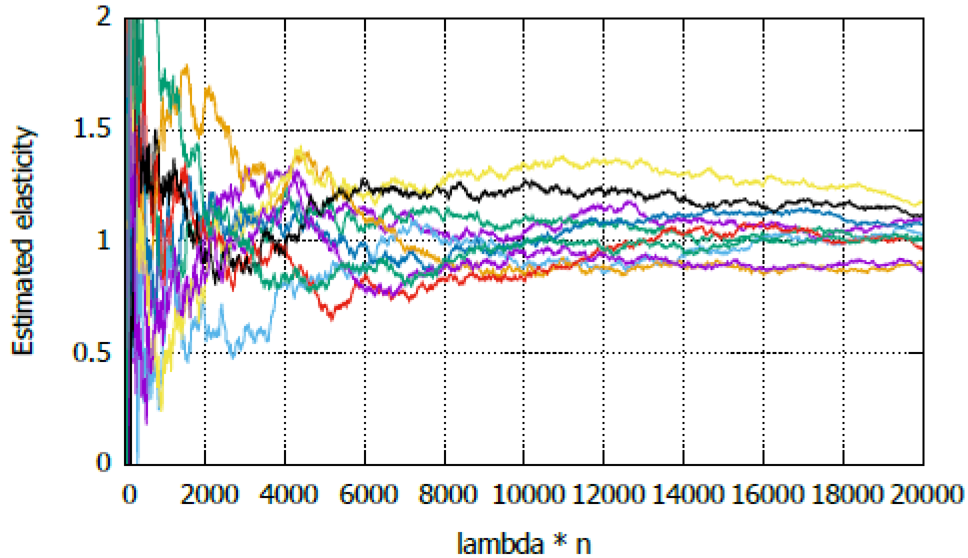

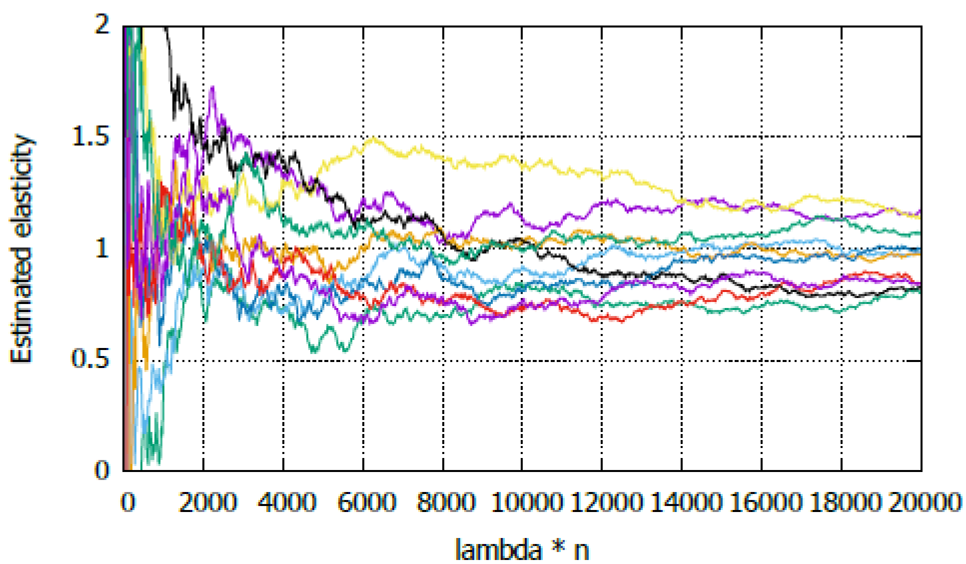

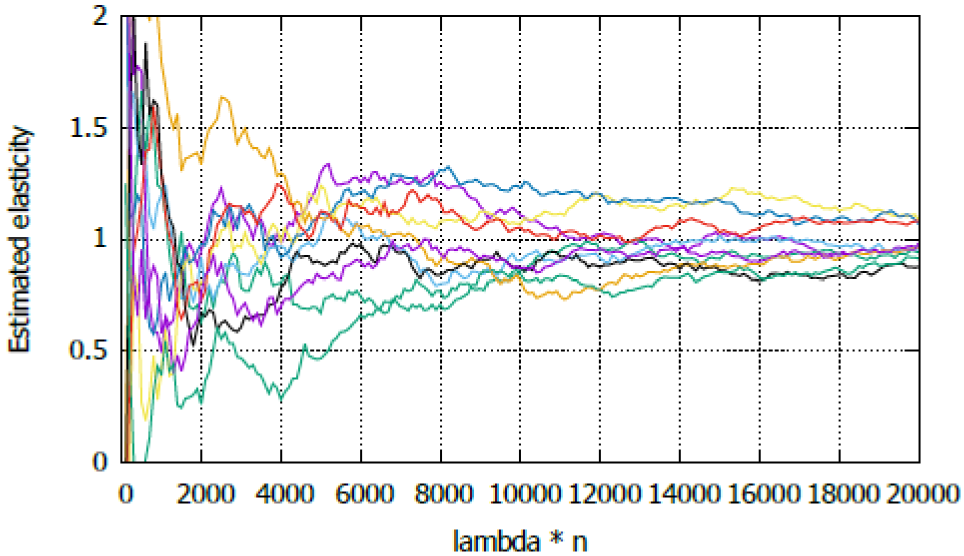

Figure 3: Estimating elasticity: n is the number of tests, λ is the average number of daily reservations received (from top to bottom: 10,20, 100). Ten estimation sequences for each plot.

EXPERIMENTAL RESULTS

The above results about estimates and error bars can be used in different contexts. In this section, we propose a possible pragmatic experiment to estimate elasticities. We assume that current prices are already in the proper ballpark for the future booking window and that elasticity is not too far from the theoretically optimal value of −1 (for the case of unlimited inventory).

The experiment that we consider is the following one. For a given experiment period (of *n* days, a multiple of 7), the entire set of prices for future days is modified with an alternating "comb" pattern (for the entire hotel booking window, with typical values which can range from three months to one year). On even days prices are increased by 5% w.r.t. current prices, on odd days prices are decreased by 5%. For each day, the number of received reservations (for all future days in the booking window) is counted. Experimenting for a multiple of seven days will give equal weight to all days in the week in the two pricing cases (e.g., it would be improper to have more weekend days in the high-price experiments than in the low-price ones). The alternating pattern is intended to minimize the effect of additional factors influencing the elasticity beyond the price. The 5% value is a compromise between ensuring a non-negligible signal but avoiding excessive disruption that may annoy customers. Larger modification values can of course be used if deemed appropriate by management.

A second possibility would be to randomize the pricing, by deciding each day (independently) about an increase or decrease with 50% probability. But this randomized choice can introduce imbalances caused by different days of the week represented in different ways in the high- and low-price cases. Because the estimates are censored, if some future days are fully booked, the counts for that specific future day (referring to the increased and decreased price) are omitted from the totals in the estimates.

In the following part, we consider the first possibility (λ is the average number of reservations received in a day) and n is the number of experiments (estimates of elasticity that require two days each) so that the number of days is $2n$.

An assumption that can be made is that the elasticity is not very far from −1, the equilibrium value for the case of normal, non-perishable goods. If the absolute value of the elasticity is close to one, a given small percentage change in the prices will be translated into a similar percentage change in the reservation arrival rates. We distinguish two situations. In the first one, a hotel is starting from scratch, without any knowledge of the elasticity. We assume that the true elasticity is equal to 1. Fig. 4 shows the results of ten different experiments (with different seeds for the random number generator). As expected, the initial estimates are very noisy. Only after some months of experiments do the estimates begin to stabilize around the ground-truth value of 1.0. Given the amount of noise and the slow convergence, this "brute-force" estimation should be avoided.

A different situation is that of a hotel that already has a current estimate of the elasticity, so the task is to gradually update it to reflect changes in customers' responses in time. Even if a hotel is completely new, it could start by considering an initial value derived from global studies about the location, or the kind of hotel. A method to gradually update a current estimate is the exponential moving average(EMA), also known as an exponentially weighted moving average (EWMA) (Gardner 1999). It applies

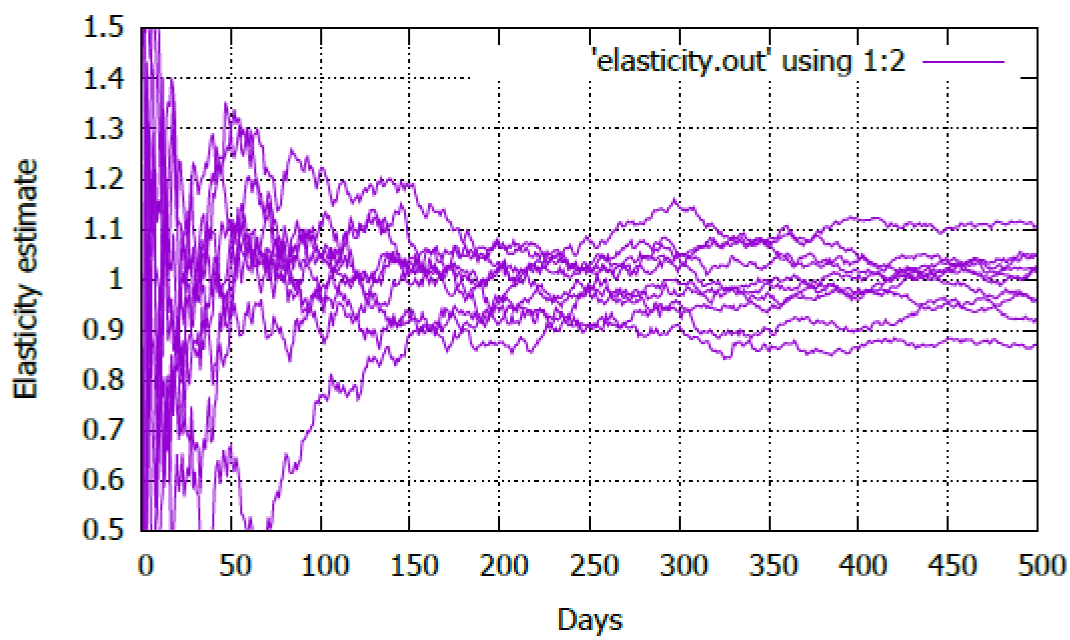

Figure 4: Estimating elasticity from scratch. 100 rooms, 100 average reservations received per day.

weighting factors for the average which decrease exponentially for measurements executed in days in the past. The intuition is that recent estimates have a bigger influence on determining the current elasticity w.r.t. older values. The EMA for a series E may be calculated recursively:

$$S_t = \begin{cases} E_0, & t = 0 \\ (1-\alpha)E_t + \alpha S_{t-1}, & t > 0 \end{cases} \quad (20)$$

Where E_t is the value at a time t. S_t is the value of the EMA at any time t. The coefficient α represents a constant smoothing factor between 0 and 1. Because the new EMA estimate is a weighted average of the previous estimate multiplied by α and the new measure multiplied by $(1-\alpha)$, a smaller α reduces the contribution of the previous estimate and, therefore, discounts older observations faster. In general, the α value should reflect the expected period in which significant changes in the elasticity may happen. A rule of thumb is that if changes happen in about n days, a reasonable value is $\alpha \approx (1-1/n)$. E.g., if α is 0.99, the average will reflect evaluations averaged over (approximately) 100 days.

In the experiment reported in Fig. 5, the hotel starts with an elasticity $E_0 = 0.7$. We assume that the real elasticity jumps from 0.7 to 1.0, α is 0.995, E_t is evaluated by aggregating reservation counts for the two prices setup from day 0 to the current day. As it can be seen from the results, the initial noise is now reduced to a more acceptable level, and the change in the elasticity is gradually incorporated into the estimate. The same EWMA method can be used to track seasonal patterns. Let's imagine that the elasticity changes with a sinusoidal pattern between 0.8 and 1.2, for one year.

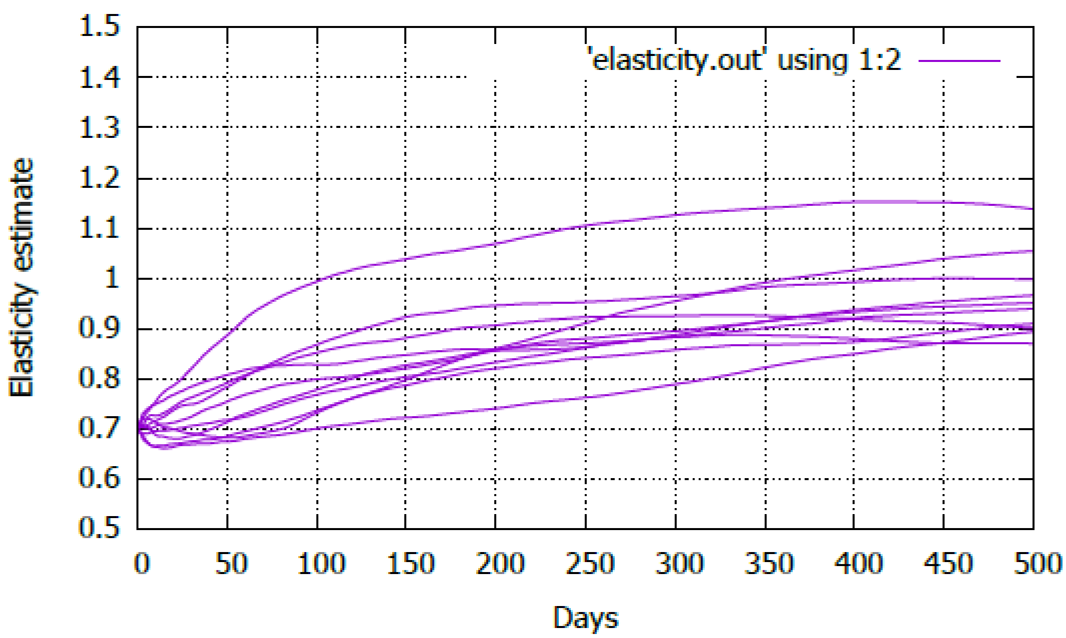

Figure 5: Estimating elasticity with daily updates. 100 rooms, 100 average reservations received per day.

The experiment in Fig. 6 is intended to model a situation where there is a high season with a smaller elasticity and a low season with a large one. In this E_t is evaluated

from scratch after each experiment (lasting two days). With the same α value of 100, the elasticity estimate follows the seasonal pattern, allowing better tuning of prices in the two seasons.

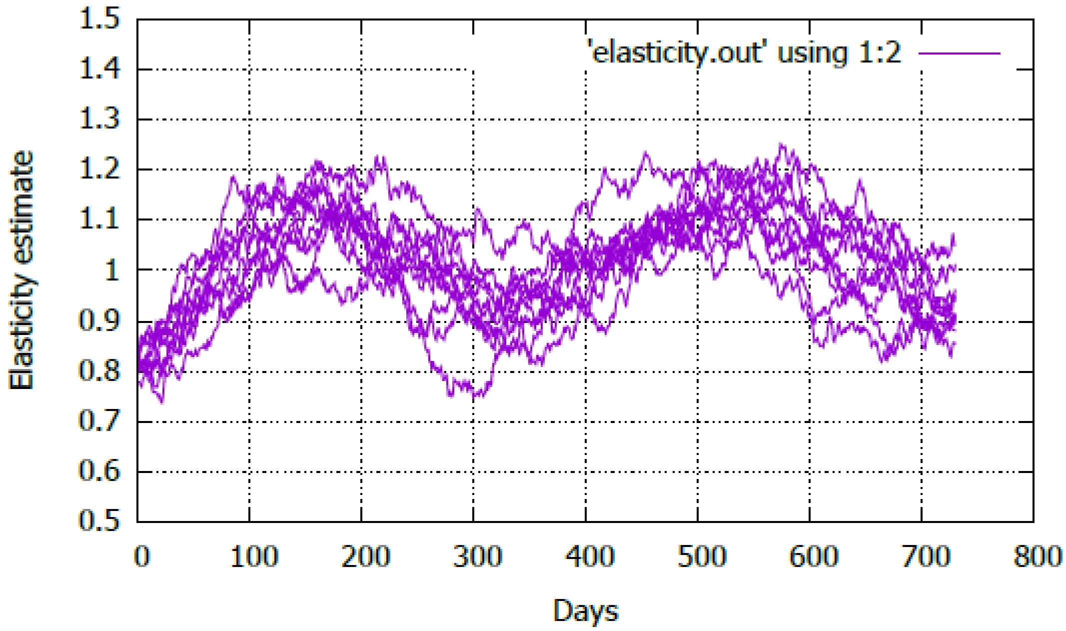

Figure 6: Tracking a seasonal elasticity pattern. 200 average reservations received per day.

The above presented experiments and additional ones in realistic sistuations (which are not reported for the limits on the size of this paper) demonstrate the practical applicability of the proposed estimation methods. The required inputs are aggregated data from the stream of reservations that are easily collected by the hoteliers, or by their supporting Property Management System (PMS) software.

CONCLUSIONS AND FUTURE EXTENSIONS

The building blocks of advanced and scientific RM (Cross 1997) based on dynamic prices are i) models to forecast demand (based on historical patterns, received reservations on the book, and external factors) and ii) optimization schemes to determine better prices for future days. In particular, reservation-based forecasting methods are widely used (Andrawis *et al.* 2011). This process is not one-shot: when prices are changed, the demand forecast itself will need to be updated. To automate the above process one needs therefore to estimate the so-called *price elasticity of demand*. Knowing how customers will react to price changes will permit us to close the above estimation-andoptimization loop.

While the issue is tremendously relevant, the abstract theory is well known, and the concept of elasticity was introduced more than one hundred years ago, estimating the elasticity in a robust scientific manner in the practical context of hotels is far from trivial even with current technology. The task is so delicate, complex, and error-prone to deserve detailed investigations. We concentrated on a very pragmatic and realistic context: that of estimating the “pickup elasticity”, by measuring how the arrival rate of reservations for future days (a quantity that is easily measured by hoteliers) is affected by a price change (ceteris paribus, everything else is — approximately — held constant). Other evaluations by questionnaires or macro-economic studies are almost useless because the real client behavior can be very different from answers given in an artificial context and because target customers of a specific hotel can respond in their way.

We quantified by a statistical analysis that: i) evaluations based on simplistic A/B testing are so error-prone to resemble a random coin toss ii) there is no alternative to patiently collecting a sufficient number of estimations at two price levels to render the estimation error so low that the estimate can be safely used to tune prices. An inferior alternative is to base the pricing on feeling, hopes, tradition, . . . without profiting from serious and objective experimentation.

We identified pragmatically two quantities that are relevant in the error estimation: the size of the room inventory and the number of days of experimentation. The relevant factor is their product and the error decreases like one over the square root of their product. This rule of thumb is sufficiently simple and robust to be used by most hotel managers, while more detailed estimations can be done with the support of software. It is to be remembered that a very precise evaluation is not needed because the pricetuning activity is repeated frequently when new reservations and new predictions are available. On the other hand, estimations with small error bars will permit a faster adaptation of the prices to the revenue-maximizing levels, without losing opportunities.

The difficulties and patience involved in estimating elasticities by monitoring customers’ responses should underline the importance of simple pricing schemes and sensible design of accommodation types. If pricing rules “split the hair in four” with too many parameters (restrictions, minimum stay, children discounts...) and accommodation types contain only a few rooms, the elasticity estimation problem can very rapidly become so difficult that a scientific estimation will have to be abandoned.

A limitation of the current study is the consideration of a simple Poisson model (more complex models of demand may show over-dispersion, a large variation w.r.t. Poisson, which will render the estimation even slower).

References

Andrawis, R. R., Atiya, A. F. and El-Shishiny, H. (2011), ‘Combination of long term and short term forecasts, with application to tourism demand forecasting’, *International Journal of Forecasting* 27(3), 870–886.

Atkinson, H., Berry, A., Jarvis, R. *et al.* (1995), *Business accounting for hospitality and tourism*, Chapman & Hall Ltd.

Aziz, H. A., Saleh, M., Rasmy, M. H. and ElShishiny, H. (2011), ‘Dynamic room pricing model for hotel revenue management systems’, *Egyptian Informatics Journal* 12(3), 177–183.

Barker, L. (2002), 'A comparison of nine confidence intervals for a Poisson parameter when the expected number of events is less than or equal to 5', *The American Statistician* 56(2), 85–89.

Battiti, R. and Brunato, M. (2018), *The LION way. Machine Learning* plus *Intelligent Optimization*, LIONlab, University of Trento, Italy.
URL: *http://intelligent-optimization.org/LIONbook/*

Battiti, R., Brunato, M. and Mascia, F. (2008), *Reactive Search and Intelligent Optimization*, Vol. 45 of *Operations research/Computer Science Interfaces*, Springer Verlag.

Bayoumi, A. E.-M., Saleh, M., Atiya, A. F. and Aziz, H. A. (2013), 'Dynamic pricing for hotel revenue management using price multipliers', *Journal of Revenue and Pricing Management* 12(3), 271–285.

Bellman, R. E. (2015), *Adaptive control processes: a guided tour*, Vol. 2045, Princeton university press.

Bonoma, T. V., Crittenden, V. L. and Dolan, R. J. (1988), 'Can we have rigor and relevance in pricing research', *Issues in Pricing–Theory and Research, Lexington Books, Lexington* pp. 333–360.

Brunato, M. and Battiti, R. (2020), 'Combining intelligent heuristics with simulators in hotel revenue management', *Annals of mathematics and artificial intelligence* 88(1), 71–90.

Canina, L. and Carvell, S. (2005), 'Lodging demand for urban hotels in major metropolitan markets', *Journal of Hospitality & Tourism Research* 29(3), 291–311.

Chen, C.-M., Lin, Y.-C. and Tsai, Y.-C. (2015), 'How does advertising affect the price elasticity of lodging demand? Evidence from Taiwan', *Tourism Economics* 21(5), 1035–1045.

Chung, C. (2006), 'Quality bias in price elasticity', *Applied Economics Letters* 13(4), 241–245.

Cox, D. R. and Reid, N. (2000), *The theory of the design of experiments*, Chapman and Hall/CRC.

Cross, R. G. (1997), 'Revenue management: Hard-core tactics for market domination', *Cornell Hotel and Restaurant Administration Quarterly* 2(38), 17.

Cross, R. G., Higbie, J. A. and Cross, D. Q. (2009), 'Revenue management's renaissance: A rebirth of the art and science of profitable revenue generation', *Cornell Hospitality Quarterly* 50(1), 56–81.

Enz, C. A., Canina, L. and Lomanno, M. (2009), 'Competitive pricing decisions in uncertain times', *Cornell Hospitality Quarterly* 50(3), 325–341.

Fibich, G., Gavious, A. and Lowengart, O. (2005), 'The dynamics of price elasticity of demand in the presence of reference price effects', *Journal of the Academy of Marketing Science* 33(1), 66–78.

Gardner, E. S. (1999), 'Note: Rule-based forecasting vs. damped-trend exponential smoothing', *Management Science* 45(8), 1169–1176.

Garwood, F. (1936), 'Fiducial limits for the Poisson distribution', *Biometrika* 28(3/4), 437–442.

Hanks, R. D., Cross, R. G. and Noland, R. P. (2002), 'Discounting in the hotel industry: A new approach', *Cornell hotel and restaurant administration quarterly* 43(4), 94– 103.

Hiemstra, S. J. and Ismail, J. A. (1993), 'Incidence of the impacts of room taxes on the lodging industry', *Journal of Travel Research* 31(4), 22–26.

Irwin, J. O. (1937), 'The frequency distribution of the difference between two independent variates following the same Poisson distribution', *Journal of the Royal Statistical Society* 100(3), 415–416.

Kahneman, D., Lovallo, D., Sibony, O., Torrance, A. and Von Hippel, C. (2019), *A structured approach to strategic decisions*, MIT Sloan Management Review.

Kohavi, R. and Longbotham, R. (2017), 'Online controlled experiments and A/B testing', *Encyclopedia of machine learning and data mining* 7(8), 922–929.

Ku, H. H. *et al.* (1966), 'Notes on the use of propagation of error formulas', *Journal of Research of the National Bureau of Standards* 70(4), 263–273.

Lewis, R. C. and Shoemaker, S. (1997), 'Price-sensitivity measurement: A tool for the hospitality industry', *Cornell Hotel and Restaurant Administration Quarterly* 38(2), 44–54.

Marn, M. V., Roegner, E. V. and Zawada, C. C. (2003), 'The power of pricing'.

Middleton, V. T. and Clarke, J. R. (2012), *Marketing in travel and tourism*, Routledge.

Nash, J. F. (1975), 'A note on cost-volume-profit analysis and price elasticity', *The Accounting Review* 50(2), 384–386.

Papoulis, A. (1990), *Probability and statistics*, Prentice-Hall, Inc.

Papoulis, A. and Saunders, H. (1989), 'Probability, random variables and stochastic processes'.

Patil, V. and Kulkarni, H. (2012), 'Comparison of confidence intervals for the Poisson mean: some new aspects', *REVSTAT–Statistical Journal* 10(2), 211–227.

Schmidgall, R. S. (2002), *Hospitality industry managerial accounting*, Educational Institute, American Hotel & Motel Association.

Skuras, D., Petrou, A. and Clark, G. (2006), 'Demand for rural tourism: the effects of quality and information', *Agricultural economics* 35(2), 183–192.

Song, H., Lin, S., Witt, S. F. and Zhang, X. (2011), 'Impact of financial/economic crisis on demand for hotel rooms in Hong Kong', *Tourism Management* 32(1), 172–186.

Song, K.-S. (2021), 'Simultaneous statistical modelling of excess zeros, over/underdispersion, and multimodality with applications in hotel industry', *Journal of Applied Statistics* 48(9), 1603–1627.

Talluri, K. T., Van Ryzin, G. J., Karaesmen, I. Z. and Vulcano, G. J. (2008), Revenue management: Models and methods, *in* '2008 Winter Simulation Conference', IEEE, pp. 145–156.

Tellis, G. J. (1988), 'The price elasticity of selective demand: A meta-analysis of econometric models of sales', *Journal of marketing research* 25(4), 331–341.

Tran, X. V. *et al.* (2015), 'Effects of economic factors on demand for luxury hotels rooms in the US', *Advances in Hospitality and Tourism Research (AHTR)* 3(1), 1– 17.

Vardeman, S. B. and Lee, C.-S. (2005), 'Likelihood-based statistical estimation from quantized data', *IEEE Transactions on Instrumentation and Measurement* 54(1), 409–414.

Vives, A., Jacob, M. and Aguilo, E. (2019), 'Online hotel demand model and own-´ price elasticities: An empirical application in a mature resort destination', *Tourism Economics* 25(5), 670–694.

Walker, T. (1997), 'Calculating breaking point in the hotel pricing game', *Hotels* 31(11), 105.

Wheaton, W. C. and Rossoff, L. (1998), 'The cyclic behavior of the US lodging industry', *Real Estate Economics* 26(1), 67–82.