Universitat de les Illes Balears

DEA WP no. 92
Working Paper Series

Dynamic pricing and revenues of Airbnb listings: estimating heterogeneous causal effects

Veronica Leoni
Universitat de les Illes Balears
E-mail: veronica.leoni@uib.es

Jan Olof William Nilsson
Universitat de les Illes Balears
E-mail: william.nilsson@uib.es

May 2020

Departament
d'Economia
Aplicada

Dynamic pricing and revenues of Airbnb listings: estimating heterogeneous causal effects.

Veronica Leoni[1]

Jan Olof William Nilsson[2]

Abstract

This paper investigates the extent to which the implementation of intertemporal price discrimination affects Airbnb listings' revenue. We found that on average, a price surge (i.e., increasing the price as we approach the date of service consumption) has an adverse effect on revenue. However, the magnitude of such effect exhibits significant heterogeneity among listings. Through the application of generalized random forests, a causal machine learning technique, we identify exacerbating and moderating treatment modifiers and shed light on the listing dimensions that cause price surges to be particularly detrimental for hosts' revenues.

[1] Corresponding Author
Department of Applied Economics, Universitat de les Illes Balears (UIB), Cra. de Valldemossa, km 7.5, 07122, Palma de Mallorca, Spain, veronica.leoni@uib.es

[2] Department of Applied Economics, Universitat de les Illes Balears (UIB), Cra. de Valldemossa, km 7.5, 07122, Palma de Mallorca, Spain, william.nilsson@uib.es

1. Introduction

The role played by price is frequently debated in economics, particularly in the field of hospitality and airline transportation management. The importance attributed to pricing decisions in this area mainly arises from the intrinsic nature of the tourism product as a perishable and non-storable service.

Tourism establishments require efficient inventory management of their fixed capacity and a reliable forecast of uncertain tourism demand. As these companies are generally characterized by high fixed costs and low marginal costs, they tend to maximize revenues via price adjustments. However, this is usually complex. Setting an inadequate price, for instance, could lead to immediate forgone revenue and also have a negative effect on customers' perception of quality and positioning in the market: A bold discounting practice could send the wrong signal and give the impression that the good or service is deteriorating. In addition, fare adjustments could affect the so-called "reference price" (i.e., the subjective value against which consumers evaluate current product prices to assess their attractiveness) and influence future purchase intentions.

The dramatic evolution of information and communication technology (ICT) has allowed tourism managers to be more responsive and quicker in adjusting prices based on market conditions. The term *dynamic pricing* refers to upward or downward price adjustments over time.

Most previous literature on the adoption of dynamic pricing strategies by airline carriers (particularly low-cost carriers) and hotels finds evidence of two main types of price discrimination: intertemporal price discrimination (hereafter IPD) and dynamic adjustments to stochastic demand. Companies typically use a combination of these strategies; i.e., they set the price as a function of the time horizon between the booking time and the travel date (IPD) and as a function of the current level of capacity (dynamic adjustments to stochastic demand). Empirical evidence (Koenigsberg et al., 2008; Escobari, 2012; Malighetti et al., 2009) highlights a tendency toward price increases as the travel date approaches and as the number of rooms (seats) available decreases.

The adoption of dynamic pricing strategies by traditional tourism firms has attracted great interest, but there is still scant knowledge of the pricing strategies adopted by the emerging *peer-to-peer* segment. In addition, the effect of pricing strategies on performance deserves substantially more attention.

Only a few papers have explored the adoption of revenue management techniques by nonprofessional tourism agents (Gibbs et al., 2018b; Kwok and Xie, 2018; Magno et al., 2018; Oskam, 2018). The pricing strategies analyzed converge on two main types: (i) price positioning—i.e., the price level compared with that of close competitors; and (ii) price volatility—i.e. the absolute overall variations in the price during a specific time period.

The Airbnb market is characterized by highly differentiated supply, which allows hosts to actively use the price mechanism. The heterogeneity of supply is also reflected in how hosts adjust their listings' prices over time. Within the same market, and at the same time, it is possible to observe both price discounts and price surges, and both policies may be optimal. This is difficult to evaluate, because the effect of price changes can have different effects for different hosts, and we have not found any empirical study that takes this into account.

Therefore, this paper has three main goals. First, we will examine the adoption of IPD in the peer-to-peer market. Using a panel dataset of daily prices set by Airbnb hosts in Ibiza, we go one step further than a mere volatility estimation by discerning the direction of the price adjustments as the service consumption date approaches. Second, we will identify the effect of the implementation of dynamic pricing on the listings' revenue. The framework we adopt uses potential outcomes (often referred to as the Rubin causal model) to define the causal effects of the treatment (pricing strategy) on the final outcome (revenue). And third, we will analyze the heterogeneity in treatment effects—i.e., by examining the differences across subjects that moderates the positive/negative effect of the pricing strategies on the final revenue. To accomplish these goals, we carefully define a "price surge" variable, which can measure dynamic pricing even with restricted availability of data. We implement a recent machine learning technique, *Generalized Random Forest*, which is able to identify heterogeneous treatment effects, even in cases in which the treatment is endogenous.

To the best of our knowledge, this is the first attempt to examine IPD in the peer-to-peer market and the first to thoroughly assess the effect of this strategy on agents' revenue. Whereas most of the literature focuses on the estimation of average effects, in this study we also analyze the heterogeneity in treatment effects, since listings might yield different results in response to the same strategy. This paper provides evidence of the differences in effects across subjects and sheds light on the features that render a pricing strategy particularly favorable or detrimental.

2. Literature review

This study is related to two streams of the literature: dynamic pricing strategies in both the standard accommodation industry and peer-to-peer marketplaces and a broader literature on the estimation of heterogeneous treatment effects.

2.1. Dynamic pricing in traditional markets

General economic theory claims that the price, in the absence of market failures, should be a valid economic signal that helps customers make informed decisions, since it includes information regarding the quality of a good or a service. A rich branch of the literature (de Olivera Santos, 2016; Espinet et al., 2003; Thrane, 2005; Abrate and Viglia, 2016; Chen and Rothschild, 2010; Magno et al., 2018; Gibbs et al., 2018; and Wang and Nicolau, 2017, among others) uses hedonic price modeling, in which the price of a tourism good is examined as function of its own characteristics by disentangling the effect of each feature on the final price (Lancaster 1966; Rosen, 1974).

Customers' reactions to price variations (price elasticity) do not uniquely link to the product's features but also to the "reference price" against which they evaluate the actual price (Monroe, 1973). A growing body of literature investigates the existence of different reference prices and their effect on purchasing intention (Oh, 2003; Viglia et al., 2016; Nieto-Garcia et al., 2017, among others). For further details, we refer the reader to Viglia et al. (2016).

The price of a tourism product works as a tool for operation management. Setting a price implies taking the following conditions into consideration: The tourism product is (i) de facto a service, (ii) perishable, and (iii) non-storable. Pricing decisions are further complicated by the fact that since most tourism suppliers (accommodation and transportation sectors) have a limited capacity (number of rooms per day or number of seats on a flight) and a limited time to sell (advanced purchase), they generally face high fixed costs and an uncertain demand. Price can help stimulate demand when demand is sluggish (via price discounts) or control it if it exceeds the remaining capacity (price surging).

Therefore, the most straightforward way to maximize profit is through the maximization of revenues, since most of the costs are non-recoverable and marginal costs tend to be very small. A growing body of literature examines optimal pricing when goods have the above-mentioned characteristics, most of which concern accommodation (Abrate et al., 2012; Abrate and Viglia, 2016; Melis and Piga, 2017) and air transport management (Malighetti et al., 2009; Koenigsberg et al., 2008; Escobari, 2012; Willliams, 2018).

From a theoretical standpoint, Gallego and Ryzin (1994) suggest that it is possible to choose an optimal price and maximize revenues subject to two constraints: the length of the time horizon and the available capacity. Price discrimination that is related to the time horizon between the booking date and the travel date—in industry jargon, the “lead time”—is known as intertemporal price discrimination (IPD). Most previous research on IPD examines airline management (Escobari, 2012; Malighetti et al., 2009; Williams, 2018; Koenisberg, 2006; Alderighi et al., 2016; Morlotti et al., 2017; Alderighi et al., 2018). Less attention has been paid to the accommodation industry (Abrate and Viglia, 2016; Melis and Piga, 2017; Abrate et al., 2012; Abrate et al., 2019).

Intertemporal price adjustments can be revealed by either upward or downward variations during the booking period that precedes the service consumption (Melis and Piga, 2017). The sign and the magnitude of the price fluctuation depend on two factors: the expectation regarding the customer’s willingness to pay and the capacity stock (Alderighi, 2016). We could therefore describe the price as a function of time-invariant characteristics (e.g., room facilities or flight route), the date of travel (seasonality), the lead time, and the remaining capacity. Previous research suggests that the form of the optimal price could be represented as a hyperbola, with the price increasing as the travel date nears and the remaining capacity reduces (Koenigsberg et al., 2008; Escobari, 2012, Malighetti et al., 2009).

Customers with lower willingness to pay tend to book further in advance, which guarantees a certain level of occupation (or load factor in the case of airlines) for the company, which is necessary for the coverage of variable costs. However, the most desirable customers are those who present a lower elasticity to price and who tend to book later (Bergantino and Capozza, 2015). Thus there is a trade-off between ensuring a threshold level of occupation and the desire to allocate part of the capacity to customers with the highest willingness to pay. These decisions involve a certain degree of risk, since accepting the "wrong booking" or setting an inadequate fare means that although any booking generates new revenue, this increase may prove to be lower than what could have been obtained through a different sales decision. This is normally referred to as "revenue dilution." How companies adjust their pricing dynamically could also depend on the characteristics of their customer population. Potential buyers may exhibit heterogeneity in their valuation of the service and in their degree of patience, based on the waiting cost (Su, 2007). In this regard, the optimal strategy also depends on customers' strategic behavior.

Despite the effectiveness of revenue management techniques in improving both the firm's profitability and social welfare (by allocating the product to customers who value it the most), there are several reasons for more uniform pricing. One is the above-mentioned reference price. Companies implementing price discounts could permanently affect the reference price of their customers—i.e., once a client receives a lower fare, it will be difficult to revert to a higher price without affecting their purchasing intention. Moreover, price fluctuations might be perceived as unfair, especially by loyal customers, and disincentivize future purchasing.

Another side effect of dynamic pricing concerns the ability to produce reliable forecasts. If management is unable to correctly estimate demand, it is more suitable to adopt uniform pricing (Melis and Piga, 2017). Melis and Piga (2017) find evidence that the use of dynamic pricing is actually less widespread than what is normally suggested by the literature.

Despite the vital importance of revenue management in the tourism sector and widespread academic interest in this topic, there is still a lack of knowledge on the real quantitative effect of dynamic pricing strategies on revenue, compared with more uniform pricing. Recent work by Abrate et al. (2019) explores the effect of IPD on revenue maximization. Their main finding is that higher price variability leads to higher revenue. In addition, the authors state that the positive effect of price variation, together with adequate inventory management, outweighs the negative effect of perceived unfairness.

2.2 Dynamic pricing in peer-to-peer markets

The term "peer-to-peer economy" has a range of meanings that can be summarized as the activity of sharing access to goods or services, for free or for a fee, between nonprofessional suppliers (peers) that is often mediated through two-sided online platforms.

The ICT revolution has enabled the rise of these activities, which have particularly flourished in the tourism sector (Ert et al., 2016). The most popular peer-to-peer platform in the accommodation segment is Airbnb, and Uber has been the pioneer in the road transportation segment.

One of the main factors influencing the decision to buy from a nontraditional provider is the economic gain obtained through cost reduction, since these services are typically cheaper than canonical ones (Guttentag, 2015). Also, people who participate in the sharing economy often cite concerns about sustainability and enjoyment of the activity itself—i.e., having more authentic experiences—as reasons to choose this option (Kim et al., 2015; Hamari et al., 2016).

Price is a core competitive advantage for the sharing economy (Tussyadiah and Pesonen, 2016; Gibbs et al., 2018), and therefore pricing strategies are integral to its success. In this respect, Uber and Airbnb have adopted different approaches. Uber sets the price of a ride using a "price surge" mechanism: The more people request rides in a given area (and during a specific time), the more the price will increase and, in turn, ensure that the service is assigned to the customer who values it the most. By contrast, Airbnb has developed a highly sophisticated algorithm, called Smart Pricing, that suggests the optimal price (Hill, 2015; Ye et al., 2018) but leaves the host free to decide whether to accept the recommendation. In two-sided marketplaces, pricing is a crucial lever to better match supply and demand (Ye et al., 2018). Despite the importance Airbnb ascribes to its pricing policies (e.g., enough to develop the Smart Pricing algorithm), there is still a lack of empirical analysis of the real effect of pricing strategies listings' profitability.

Gibbs et al. (2018) investigate the adoption of dynamic pricing strategies by Airbnb hosts and compare these methods with the strategies of nearby hotels. To compare price fluctuations, they analyze the coefficient of variation of Airbnb monthly prices. The results show an important heterogeneity between Airbnb hosts and relatively lower adoption of dynamic pricing strategies compared with the hotel segment.

Kwock and Xie (2018) examine the effect of two different pricing strategies on Airbnb listings' revenue. The first is "price positioning"—i.e., the gap between the price set by a specific host and the average price set by its competitors. The second strategy examines dynamic pricing, which is captured by the standard deviation of listings' prices in a given month. Their results favor the adoption of both strategies, which, ceteris paribus, would generate higher revenues. In addition, they show significantly better performance of multi-listing hosts (hosts who manage more than one property) compared with single-listing hosts.

Magno et al. (2018) investigate the price of Airbnb listings in Verona (Italy) and find evidence that hosts are gaining revenue management competencies by adjusting prices dynamically depending on market conditions.

Finally, Oskam et al. (2018) explore the effect of dynamic pricing on revenue performance. The dynamic pricing is proxied by the number of different prices and the mean of positive and negative price changes. Using multiple OLS, they show that hosts who adjust their prices more frequently outperform others in terms of both revenues and occupancy rates.

Most of the above-mentioned studies measure dynamic pricing through the coefficient of variation (Gibbs et al., 2018; Kwock and Xie, 2018; Oskam et al., 2018) or the standard deviation of price (Kwock and Xie, 2018). These measures afford insight into the adoption of nonuniform pricing by peer-to-peer agents, but do not explore the directions of such price changes. A step toward assessment of the direction of price changes has been accomplished through the use of two other measures: the count of negative and positive price changes and the average positive and negative price variations (Oskam et al., 2018). However, including these variables in a regression model and making a ceteris paribus interpretation departs from the aim of analyzing dynamic pricing in the form of IPD. In general, a high average positive price change will be accompanied by a high average negative price change, unless the price actually is changed due to an adjustment in the long-run market situation. When these variables, as well as the revenue, are measured over a long period, it is difficult to maintain the assumption of exogeneity; for example, low revenue could lead to larger average negative price changes. Accordingly, while these variables include the sign of the change, the analysis still concerns variation in the price over time *for different dates*, rather than the price variation over time for the *same date*.

2.3 Heterogeneous treatment effects

One of the aims of this work is measuring the extent to which the implementation of dynamic pricing affects listings' revenue. The causal relationship between price and revenue is studied in the framework of potential outcomes, generally known as the Rubin causal model (Shekhon, 2008). A causal effect is defined as the difference between two potential outcomes, but only one of the two outcomes is observed in practice. This type of model is used in different disciplines, including biomedical research (e.g., to test a drug's effect on patients) or policy evaluation (to assess a policy's effect). The average treatment effect (ATE) measures the expected effect from treatment, expressed by a binary explanatory variable, for a randomly chosen unit in the population. When the treatment has a continuous interval, as in this study, we talk about average partial effect.

Over the last few years, we have explored the heterogeneity of treatment effects (HTE) with growing interest. Any given treatment—for instance, a price fluctuation—might affect subjects in different ways. Studying HTE implies examining the differences across subjects to reveal the conditions under which treatments are especially effective (or ineffective). As suggested by Lechner (2019), it is valuable to uncover underlying heterogeneity at the "finest possible level of granularity"; this yields better understanding of the causal mechanisms at work. It also sheds light on the distributional aspects of a treatment by identifying, for instance, groups that win and groups that lose. The analysis of HTE can help in designing and deploying policies to maximize their effectiveness.

Methodologies used to estimate HTE vary across the literature. Classical nonparametric approaches include nearest-neighbor matching and kernel methods (Crump et al., 2008; Xie et al., 2012; Verhofstadt and Maertens, 2014; and Chen et al., 2019, among others). Recently, there have been rapid advances in the development of causal machine learning techniques to study treatment effects in both experimental and observational studies. A number of papers (Green and Kern, 2012; Athey and Ibens, 2016; Wendling et al., 2018; Athey and Wager, 2019; Daoud and Johansson, 2019; Farbmacher et al., 2019; Lechner, 2019) adopt tree-based methods for estimating HTE, building on Breiman's random forest algorithm (Breiman, 2001). Athey et al. (2019) introduced the generalized random forest, which allows for an endogenously determined treatment as long as an instrumental variable is available. Their model is, accordingly, particularly suitable for use in this application; more details can be found in Section 4.

Causal machine learning could be an alternative to classical regression models for the study of heterogeneous causal effects. Conversely from linear regression, which can lose statistical power and suffer from computational issues when adding many interactions, *causal forests* can detect high-dimensional covariate-based treatment effect heterogeneity. Moreover, these models allow for more flexibility in incorporating nonlinear and interaction terms automatically (Strittmatter, 2018). Finally, by adopting causal forests, we reduce cherry-picking behaviors because we learn about the moderators deductively from the data, rather than inductively, by adding interactions. Ad hoc searches for particularly responsive subgroups can mistake noise for a true treatment effect (Davis and Heller, 2017).

3. Data

3.1 Sample

The data set we use in this study is provided by Airdna[i] and consists of daily information on listings published on Airbnb's platform for Ibiza (Balearic Islands, Spain) during the 2016 summer season. In this analysis, we consider only those listings with at least one active booking for the analyzed month and another active booking for the previous month. Our sample contains 2,599 listings for the month of July and 3,049 for the month of August.

3.2 Defining and measuring dynamic pricing ($\boldsymbol{W_i}$)

One of the key concerns in dynamic pricing strategies is to set the correct price at the correct time. Despite this, most of the measures currently used for dynamic pricing focus on price variation or price changes for different dates. Using such measures to evaluate whether dynamic pricing affects revenues can be misleading, because observing variation does not reveal the quality of the changes. Ideally, we would like to observe the complete history of the price for a specific listing until the date of reservation or until the date has passed without any reservation having occurred. However, obtaining such data is highly demanding. We propose an alternative, which requires much less information. Our idea is to compare the price for similar dates of arrival that were reserved with different lead times. The first booking to occur implies that the price at that time is settled, and the dates still available are open to price adjustments. When the second date is booked (or when the check-in date is reached without any booking), another price is recorded. It is then easy to compare the two prices and evaluate whether the host has increased or decreased the price as the consumption date approaches. Our measure of dynamic pricing is based on a difference of log prices—or, equally, the log of a price ratio. This is our treatment variable, which we label a *price surge* variable, because positive values correspond to an *increase* in the price as the service consumption date approaches. Negative values of the variable are found if the price was decreased. To ensure easier comprehension of how the price surge variable (i.e., the treatment, W_i) is constructed, Table 1 includes a calendar for a randomly chosen property during the month of August 2016.

[Insert Table 1 about here]

Table 1 reports the property identity number (Column 1), the calendar date (Column 2), the status (Column 3), the price (Column 4), the booking date (Column 5), and the booking reference number (Column 6).

The variable *status* (Column 3) assumes the value “R” if that date has been booked (reserved), “B” if the host has deliberately blocked it to avoid unintended bookings, and “A” if that date has not been booked (i.e., available). The variables *bookeddate* (Column 5) and *reservationid* (Column 6) record a missing value if the status is “blocked” or “available.”

We generated the variable *lead* that corresponds to the time lapse (in days) between the date of service consumption and the booking date (i.e., lead time). For each booking reference (*reservationid*), we created a variable, *average_price* (Column 8), which is the average price of that specific booking.

To define a measure of dynamic pricing, we consider two types of price variations: (i) the difference in price between two active bookings and (ii) the difference between active bookings and “non-booked” dates. In the first case, we evaluate the price direction (increasing or decreasing as the lead time reduces) by looking at the price difference and the lead time (advance in purchase). In the second case, we do not have the purchase advance of a non-booked date, but we expect that the price of the non-booked day is the last price the host set (e.g., before the end of the potential consumption date) in order to achieve a booking; in this sense, we consider the price for a non-booked date as being subsequent to the price of a booked date[ii].

We created Δa_{ijk} (i=1, …,I, where i refers to the *propertyid*; j=1,…,J, where j refers to the month; k=1,…,K, where k refers to the number of bookings made by listing i in period j), which corresponds to the difference between the average price of non-booked days[iii] (transformed in natural logarithm) and the above-mentioned *average_price* (transformed in natural logarithm). $\overline{\Delta}a_{ij}$ corresponds to the average of Δa_{ijk} by *propertyid* and for a specific month. These variables assume a positive value if the host increases the price of days still not booked and a negative value if the host is setting a price discount.

In addition, we consider the price difference between two bookings (in chronological order according to the *bookeddate*). The variable Δr_{ijk} corresponds to the difference between the *average_price* of the two bookings[iv] (both transformed into natural logarithm). The price differences we record are restricted to bookings with consumption dates of 15 days or less, and local holidays are excluded. Following the example in Table 1, we compare the *average_price*[v] of reservation #65214714 and the previous one, #65214712. As we can see, the first booking has an average price ($303), which is higher than the second booking ($293). As the first booking (#65214712) has been booked with a higher advance, the variable Δr_{ijk} will therefore assume a negative value: The host is decreasing the price as we get closer to the check-in date. Again, if we compare the later booking reference #65214717 to the previous one, #65214714, we can see that the host is still decreasing the price as we get closer to the check-in date, and hence stimulates demand via price discounts[vi].

In order to have a non-missing value for Δa_{ijk}, we need at least one non-booked day in that specific month, whereas for Δr_{ijk} we require at least two active bookings. $\bar{\Delta} r_{ij}$ corresponds to the average of Δr_{ijk}, by *propertyid* and for a specific month. The final variable used in the econometric model is the average of $\bar{\Delta} a_{ij}$ and $\bar{\Delta} r_{ij}$,, hereafter referred to as the treatment W_{ij}. The treatment is hence the average of these two measures ($\bar{\Delta} a_{ij}$ and $\bar{\Delta} r_{ij}$) if both are available. If a listing does not have two active bookings during the specific month, we use $\bar{\Delta} a_{ij}$ as the final measure. If no days were left available (fully booked), $\bar{\Delta} a_{ij}$ cannot be calculated and $\bar{\Delta} r_{ij}$ is used instead of the average. Averaging these two variables ($\bar{\Delta} a_{ij}$ and $\bar{\Delta} r_{ij}$) means that for each listing, the treatment will assume only one monthly value. The final model will thus be implemented on a cross-sectional dataset.

In Figure 1, we show the distribution of the W_{ij} for the months of June, July, and August 2016.

[Insert Figure 1 about here]

During all three periods, about 60% of the listings pursued a price decrease and about 6% opted for more uniform pricing[vii]. The remaining listings exhibit a price increase for bookings closer to the service consumption date, but the magnitude of such price variation is weaker than the one registered for price drops.

[Insert Figure 2 about here]

Figure 2 shows the relationship between W_{ij} and the lead time, using the local polynomial[viii] function. The graph shows a tendency for positive price changes for hosts recording higher average lead times.

In this study, we consider the price variation in terms of the booking advance—that is, the distance (in days) between the booking date and the date of consumption. Our treatment variable measures the magnitude of the price variation as well as its sign, to capture how Airbnb hosts apply changes to the price as we approach the specific date. However, our measure is not free of limitations, since the price for each calendar date corresponds to the last price that was made available on the platform. The price variable in our panel dataset is the price at which a specific date was booked or the last attempted price of a non-booked room. Most of the literature on airlines and a few studies in the accommodation sector dispense with the use of more data on the price trajectory for specific dates (e.g., 30, 20, 10, 5, and 1 day in advance); this approach allows for wider comprehension of the pricing variation over time, compared to our restricted data, which could overlook the intermediate prices changes. It is important to bear in mind that airlines and hotels can continue to perform price changes until they are fully booked for the specific flight or date. In contrast, an Airbnb host will stop varying the price for a specific date once it is booked, unless the booking is later cancelled. Hence, this difference in the data is quite natural, but obtaining the price trajectory until a reservation is made would enable more detailed analysis of intertemporal price variations.

Another potential limitation is that our measure of dynamic pricing could overlook seasonality or days characterized by a strong demand (local events), which could lead to price jumps independent of the lead time. In essence, the idea is to get as close as possible in order to compare the intertemporal price variation on the same date, in cases in which the price trajectory is absent. Accordingly, it is important to use dates that are as similar as possible. This is the reason for using consecutive bookings for the same *month*, when the effect of the season should be fairly small, excluding bank holidays and bookings separated by more than 15 days. In this way, the price variation due to time elapsed between the booking date and the service consumption is more directly targeted.

Despite these limitations, which are both related to the absence of price trajectories, our analysis aims to overcome the main issues of the measures adopted so far in the peer-to-peer literature, which were discussed at the end of Subsection 2.2. Notice that the method that will be introduced in Section 4 is also applicable to situations with price trajectories. The key difference is only the construction of the price surge variable, which offers avenues for future research on dynamic pricing in traditional markets.

3.3 Variables

After discussing the treatment, we introduce our dependent variable, the listings' revenue, and the set of control variables that will be used in our econometric model.

The outcome Y_i is the listings' (log) revenue (nominal value, expressed in USD) during a specific month. In Figure 3, we show the distribution of the revenue for the months of July and August 2016.

[Insert Figure 3 about here]

In Table 2, we show descriptive statistics for the variables. The first subgroup contains time-invariant characteristics and the second specific variables for each month.

[Insert Table 2 about here]

The first group contains variables related to the size of the property, such as the number of bedrooms (*bedrooms*), bathrooms (*bathrooms*), the maximum number of guests allowed (*maxguest*), and a dummy for the type of property listed (*entire_property*), where the excluded category is *room* (i.e., the host is renting a private or shared room rather than the entire dwelling); variables defining the location, such as distance to the closest beach (*distance_beach*), distance to the closest club (*distance_disco)*, a set of dummy variables for the municipality (Sant Antoni de Portmany, San Joan de Labritja, San Josep de sa Talaia, Santa Eulària des Riu), with Eivissa as an excluded category; and information on the online reputational capital of the property, such as the number of reviews left by previous guests (*numberofreviews*) and a set of dummy variables for the score rating (*high_rated*, *medium_rate*, *never_rated*),[ix] where the excluded category is *low_rate*. In addition, we include information about the professionality of the host: the number of listings managed by each host (*multiple_host*), the host's response rate to guests' queries (*responserate*), instant booking acceptance (*instant_booking*), longevity of the account (*longevity*) and the cancellation policy (*cxl_mod*; *cxl_strict*), with a flexible cancellation policy as an excluded category. Finally, we include two proxies for the property's quality: the amount of the deposit required upon booking (*securitydeposit*) and the first price set by the host upon entering the market (*lnpublishednightlyrate*).

The second group of control variables consists of information on the listing's performance, such as the occupancy rate (*OR_June; OR_July*), the average local performance of listings located in the same area *(OR_June_local*; *OR_July_local*), the average lead time-—i.e., the average number of days between the booking date and the service consumption (*leadtime_July*; *leadtime_August*), and the percentage of days in a specific month that were booked with an advance of at least 30 days (*perc30_June*; *perc30_July*); these variables provide insight on the pace of bookings.

4. Methodology

The outcome, Y_i, is the log revenue; the treatment variable, W_i, is a continuous measure of dynamic pricing—i.e., the price surge variable; and $\boldsymbol{X}_i$ is a set of covariates. W_i is explained in detail in the previous section, but it is important to bear in mind that higher values correspond to increasing the price as the date gets closer. The variable can also take negative values if the price has been decreased as the date gets closer. The data are represented by $(Y_i, W_i, \boldsymbol{X}_i)$, for $i = 1, \dots, n$, which are independently and identically distributed cases. The purpose is to estimate how W_i affects the outcome while allowing for heterogeneity in such effect; hence $\tau_i(x)$ is expressed with an indicator on that it depends on $\boldsymbol{X}$. This is an important feature of the model because it is not possible in advance to discard either positive or negative effects of the variable. For example, a host who expects that they will be able to rent on all days could increase the price to (possibly) increase the revenue. Of course, increasing the price comes with a higher risk of a lower occupancy rate. Another host, who does not expect to be able to rent on all days, could decide to lower the price to increase the occupancy rate, which could compensate for the lower price and accordingly imply an increased revenue. We used generalized random forest (Athey et al., 2019) for the empirical analysis, which was performed using the R package *grf*, version 0.10.4 (Tibshirani et al., 2019). Our review of the method includes the key ideas; we refer the reader to Athey et al. (2019) for technical details.

Consider the model:

$$Y_i = \mu(\boldsymbol{X}_i) + \tau_i(x)W_i + \varepsilon_i \ .$$

We initially assume *unconfoundedness* or *exogeneity*, $Y_i^{(t)} \perp W_i | \boldsymbol{X}_i$, which implies that the magnitude (and sign) of the treatment is random when conditioned on all observed covariates. It is important to underline that the assumption rules out additional unmeasured confounding variables that are related to the potential outcomes ($Y_i^{(t)}$) and degree of the treatment t.

Athey et al. (2019) consider random forest as an adaptive kernel method, in which weights, $a_i(x)$, capture the frequency of occasions on which observation i is found in the same leaf as the test point x when B trees are grown. The treatment effect is estimated according to

$$\hat{\tau}(x) = \frac{\sum_{i=1}^{n} a_i(x)(Y_i - \widehat{m}^{(-i)}(X_i))(W_i - \hat{e}^{(-i)}(X_i))}{\sum_{i=1}^{n} a_i(x)(W_i - \hat{e}^{(-i)}(X_i))^2},$$

and conditionally centered outcomes are used. Random forest—any other procedure can also be used where the variables $\boldsymbol{X}$ are used to obtain fitted values for Y and W. $\hat{m}^{(-i)}(X_i)$ and $\hat{e}^{(-i)}(X_i)$ are *leave-one-out* estimates—i.e., observation i is not used to estimate the regression function for itself. Conditionally centering the outcomes implies that the effects of the covariates are regressed out from all of the outcomes; i.e., it is a residual that is finally used in the growing of the causal forest. In our application, the variable *leadtime* is excluded from $\boldsymbol{X}$ when conditional centering is performed, whereas it is included to detect potential heterogeneity. The reason is that the average number of days between the booking date and the service consumption is, at least partly, a consequence of the performed dynamic pricing. Price discounts can, for example, reduce the average *leadtime*.

The general idea for splitting rule in the growing of the trees in causal forests is to greedily choose the variable and where to make the split, to maximize treatment effect heterogeneity when partitioning the data. Causal forest uses a resampling method with the data that are labeled *honest*. The default option is to split the data into two parts: Observations in the first half are used to grow the tree, while those in the second part are used to calculate within-leaf estimations (Athey et al., 2018). Thus there is no "waste" in the procedure, because each data point is, in fact, used in both roles.

The assumption of *unconfoundedness* is important, but not always fulfilled. Consider, for example, that unobserved factors related to treatment also affect the outcome, and ε would be correlated with W. In such a situation, it is necessary to use an alternative method to consistently identify the treatment effect. In our case, the treatment is not randomly assigned: The concern is that the *price surge* variable is driven by the decision of the host, although it is filtered through the market process. A host chooses their own price policy, but the difference in price (the sign and magnitude of the price surge) depends on when, and at what price, the listing was rented, and this is out of the host's control (ex ante, the host does not know at what price level that specific date will be booked). If unobserved factors are related to W and they are also related to the revenue, there would be an endogeneity problem. Considering that hosts make a forecast on their probability to rent on every (or most) days, and if they consider that probability to be high, they can increase the price in an attempt to increase revenues. Other hosts, who consider the probability to be low, could reduce their price to increase the occupancy rate and (hopefully) increase revenues. The first "solution" for overcoming the endogeneity problem is to ensure that the set of control variables, X , is sufficiently complete and fulfills the assumption of *unconfoundedness*. If hosts base their decision on their occupancy rate from the previous month, this is a natural control variable. For this reason, we also include in $\boldsymbol{X}_i$ variables related to the individual property's past performance (performance of the previous month).

Another key performance measure is how early reservations, in general, are made. For instance, at the end of June, a host knows how early the reservations were made for that month, and this could be used in setting the price for July. The percentage of reservations that were made at least 30 days in advance for the month preceding the month for the outcome is a formalization of the idea. The published nightly rate (*publishednighlyrate*) the host first introduced on the platform can be considered a subjective indicator of the quality, where the host took into consideration factors related to location, or listing-specific characteristics, that could be partially unobserved by an analyst.

Despite a careful review of the variables to include in X , there may still be unobserved factors that imply an endogeneity concern. The second solution is finding a variable, Z , that is related to the treatment but is independent of ε conditional on X. This allows us to relax the assumption of *unconfoundedness*. The relevant data now consist of $(Y_i, W_i, \boldsymbol{X}_i, Z_i)$ and the treatment effect is identified according to

$$\hat{\tau}_{IV}(x) = \frac{Cov(Y_i, Z_i | X_i = x)}{Cov(W_i, Z_i | X_i = x)}$$

where the procedure of centering also includes Z. We use the *price surge* variable from the previous month as the instrumental variable.

When interpreting $\tau(x)$ (or $\tau_{IV}(x)$)—i.e., the effect of the treatment, W_i, on the final outcome—we should remember that the treatment, the price surge variable, is essentially the difference between two log-prices, ($W_i = lnP_1 - lnP_2$), which corresponds to the log-ratio, $W_i = \ln \frac{P_1}{P_2}$. In this log-log model, in which the dependent variable is log-revenue, an increase in the price ratio will produce a (positive or negative) effect on the revenue, all other variables kept constant. An increase in the price ratio by 1% will multiply Y_i by $e^{\tau(x)*\ln(1.01)}$. $\tau(x)$ is the coefficient of the treatment effect in the corresponding model, in which differences are allowed across differences in the control variables. The measured treatment effect is hence the difference in the revenue (Y_i) as a consequence of an increase in the price surge variable(W_i), conditioned on a set of covariates ($\boldsymbol{X}_i$). The model can, of course, also be used to evaluate the effect of a reduction in the price surge variable—i.e., the host is reducing the price of available dates as the calendar date is approaching. The results section includes estimates from both models.

5. Results

A key issue for the empirical analysis is choosing a model that allows for heterogeneity in the effect of dynamic pricing, because the same policy may not fit all of the hosts in the market. Despite this, we start the empirical analysis by estimating a linear regression (OLS) and an instrumental variable regression, from which we obtain an average partial effect. The instrumental variable we use is the same measure of dynamic pricing but calculated from the previous month. As discussed in the methodology section, the existence of factors that affect both the final outcome and the treatment would lead to endogeneity concerns. In Table 3, we show the output of the instrumental variable regression. We can see that, controlling for all others factors, the effect of an increase in W_i (the treatment) by one unit would decrease the final revenue; this is true for both July (-0.9761) and August (-0.7133), with the effect being slightly stronger for July. This is equivalent to saying that, ceteris paribus, an increase of 10% in the price ratio would produce a revenue decrease by 8.88% in July and 6.57% in August. We point out that a weaker effect of the treatment on the revenue would have occurred if the endogeneity issue had been overlooked. The coefficient of W_i, using a simple OLS model, would have been -0.5463 for July and -0.3769 for August. Summarizing the effect of dynamic pricing with a single coefficient for all observations could, however, hide relevant results. We found strong heterogeneity in the treatment effects, and hosts could face effects that are far from what the average effect indicates. For this reason, we extend our analysis by estimating the causal forest, as explained in the previous section, both with and without using an instrumental variable—i.e., the lagged dynamic pricing variable from the previous month. Our preferred model is causal forest using an instrumental variable, and our presentation of the results focuses on that model. As suggested by Basu et al. (2018) and demonstrated by Athey et al. (2019), we first trained a pilot causal forest on all of the features. We then trained another forest on only those features with reasonable variable importance[x]. Notice that all control variables are used in the centering procedure, which helps to reduce the confounding effects; in contrast, a restricted number of variables are used to detect heterogeneity, which allows us to focus on the real treatment modifiers.

Figure 4 shows the distribution of *out-of-bag* heterogeneous treatment effect predictions. As the model uses resampling, we can obtain out-of-bag predictions, which are obtained from a model that did not use that particular observation in building it. Notice that we refer to the prediction of the individual treatment effect and not the outcome itself, since our aim is to assess the effect of the pricing strategy on revenue rather than predict the revenue itself.

Using causal forests, with an instrumental variable, we found a negative treatment effect for all of the Airbnb listings in July, whereas in August we found a positive treatment effect, but only for a negligible number of observations. The conditional average partial effect (CATE) for the month of July is -1.048 (std. err. 0.307) and for August is -0.759 (std. err. 0.305). This means that in July, a 10% increase in the price ratio (W_i), would (on average) decrease the revenue by 9.50%, and in August by 6.98%. A *decrease* in the price ratio by 10% would, on the other hand, *increase* the revenue on average by 11.67% in July and by 8.32% in August. As widely discussed in the Methodology section, one of the key advantages of our model is that we do not limit the analysis to the average effects, but also investigate the heterogeneity in such effects.

In Figure 5, we show the ranking of variables' importance for the months of July and August 2016. These lists allow us to examine the nature of the heterogeneity and understand which variables maximize $\hat{\tau}$ variance. The ranking suggests the more relevant modifiers—i.e., along which dimensions the variance of the effect is higher. However, it does not provide any intuition on the specific characteristics that cause the price surge to be particularly detrimental. For both periods, the variables that rank higher are those related to the listing's quality (security deposit and price), location (distance from the beach or from a nightlife area), booking pace (percentage of bookings made at least 30 days in advance and leadtime), and the host's experience (longevity).

In Figures 6a and 6b, we plot the relationships between the treatment effect and the most relevant treatment modifiers. The purpose of these graphs is to shed light on the relationship between the covariates and the treatment effect. However, we should be aware that the variables do not exist independently, and their interactions are potentially complex in the estimated model.

The functions shown in the graphs represent predictions of the treatment effect. In Figure 6a, the solid black line shows the relationship between the treatment effect (y-axis) and the lead time (x-axis), keeping all other features at their median level. We can see that for both months, the effect is more negative the closer to the consumption date. This is equivalent to saying that the effect of the price surge could be particularly detrimental if, on average, the bookings come with a lower anticipation.

In the same graph (Figure 6a), we represent the treatment effect along the lead time and with the features at their median level, but for the low (10^{th} percentile) and the high (90^{th} percentile) levels of the variable *distance_beach.* The graph suggests that listings closer to the beach (dotted green line) register less negative treatment effects, compared with listings within a higher distance from the beach (dotted red line). This is valid for both periods, with the negative effect of the unfavorable location (being far from the beach) being slightly stronger in July.

In order to better elucidate the interactions between the variables and provide insights into the heterogeneity of the treatment effect, we present two extreme cases. In Figure 6b, we have increased the restrictions on the values assumed by a specific variable. The black dotted line represents the treatment effect of listings closer to the beach (this is the same as reported in Figure 6a), the blue line represents the listings close to the beach and belonging to experienced hosts (90^{th} percentile of the variable *longevity*), and the pink line represents only those listings with the previously mentioned specifications and with most of the bookings made with a large advance (90^{th} percentile of the variable *perc30June*). The graph demonstrates that although the effect of the price surge is still negative, this subgroup of listings appears to be less penalized by the price surge. This suggests that these specific characteristics might cause customers to be less sensitive to a price increase.

By contrast, listings far from the beach (dotted red line), belonging to less experienced hosts (dotted gray line), and with a high ratio of late bookings (dotted pink line) register a stronger loss in revenue as a consequence of the price surge.

These graphs represent extreme cases, whereas the majority of the listings might present a combination of the above-mentioned characteristics. However, this simplification illustrates the magnitude of the variation of the treatment effect along the covariate space and helps to elucidate which characteristics cause customers to be more or less sensitive to price changes.

6. Limitations and Suggestions for Future Research

As with all research, our study is not free of limitations. The first limitation was discussed in the Methodology section. Ideally, we would have preferred to have the entire price trajectory for a specific date (e.g. 60, …,25, etc. days prior to the date of consumption), but obtaining such data is highly demanding. Although our price surge index provides a measure of IPD with less available information, meticulous implementation is required. We make an important assumption—i.e., that two consecutive bookings that are close in time (restricted to be less than 15 days apart) are qualitatively similar. Hence, the difference in the price is mostly attributable to the difference in the purchase advance (lead time), and not to differences in season or particular circumstances, such as local events. The qualitative similarity between dates is a nontrivial assumption, since comparing two dates that clearly belong to different seasons, and hence to periods with "a temporal imbalance in the phenomenon of tourism" (Butler, 1998), could lead to myopic estimates. However, this concern does not seem relevant for our analysis of July and August[xi].

Another important consideration concerns the possible qualitative difference between weekdays and weekends, on which price dynamics could sensibly be diverse. Comparing the price of two bookings, one for the weekend and one for weekdays, could result in a price variation that is mainly due to the qualitative difference of the days booked, rather than to the difference in the lead time. Since in our analysis we compare the average of bookings' prices (which could include both types of day), this issue is less relevant. In addition, in the absence of a systematic pattern for weekend vs. weekday pricing, the difference will appear as noise and be dealt with through the IV procedure. Despite this, we ran a sensitivity check by performing the entire analysis on two separate samples: one including only weekdays and the other only weekends. The sensitivity analysis confirms that previous results are robust. This outcome could be linked to the type of tourism in Ibiza, which is mainly sea and sun; on the other hand, this robustness could not be guaranteed for cultural tourism destinations (or urban tourism destinations), where the heterogeneity of days (and weekly seasonality) could be stronger. Future extensions of our model should consider this aspect when building the price surge variable.

Another caveat, which opens avenues for further research, concerns the unbooked days between two reservations. For instance, if one booking ends on Sunday and the following begins on Wednesday, the days in between become more difficult to book, independent of their price. This is particularly relevant for "booking gaps" that occur during weekdays. In this regard, hosts should be careful in managing their spare capacities and set the most appropriate minimum stay requirements to minimize unattractive gaps.

The second limitation is related to the scope of the research question. While we can identify how different IPD policies would affect revenues—and for which kind of listings the effect is larger or weaker—our method does not specify how much the price should be changed. For almost all listings, we find that a price discount would increase the revenues, but we do not determine how much the discount should be. Naturally, the estimated effects are identified with the price variation found in the data, and in practice the discount could become too large. Defining an optimal pricing policy requires answering both when and by how much the price should be changed. This is an important research area that requires appropriate analytical techniques, which allows the targeting of optimality for the specific decision-making unit.

7. Conclusion

In our research, we define IPD as the increase (or decrease) in the prices of two close dates, which present different booking horizons. Our measure of dynamic pricing reveals both the magnitude and the sign of the price adjustment, and hence is more informative than simple price dispersion or the number of price changes—methods that have previously been used in the literature. This implies that our results are much more closely linked to the concept of IPD and, as such, have more relevant managerial implications because they provide valuable information on the effect of a price increase or decrease, rather than the simple absolute fare variation. Our treatment variable is the price surge—i.e., the increase in the price as we approach the date—while the outcome is the monthly revenue generated by the home rental activity. On average, we found that an increase in the price ratio by 10% leads to a decrease in revenue of 9.50% in July and 6.98% in August. Su's (2007) model describes how suppliers should adapt their pricing strategies depending on their customers' composition. Our results suggest that high-valuation guests in Ibiza tend to be more eager to book buy early than low-valuation consumers. For this reason, in line with Su (2007), it would be preferable to decrease the price over time, because high-value customers would buy early at high prices, whereas low-value customers are willing to wait. The tendency toward price decrease over time is also confirmed by Melis and Piga (2017), who find that hotels in the Balearic Islands demonstrate a high propensity for markdown pricing.

The applied econometric model allows use to detect heterogeneous effects. This is important, because different hosts can face different effects on revenues from the same price change. Listings for which revenue is more sensitive to changes in the price also have more to gain if a price discount on unrented days is applied. In our study, we identified the treatment modifiers—i.e., the exacerbating or moderating variables that render the price surge particularly detrimental to revenue. Our results are in line with prior studies of the accommodation sector, since a closer distance to the main point of interest (in this case, the beach) is generally associated with higher demand (or performance) (Arbel and Pizam, 1977; Tussyadiah, 2016; Yang et al., 2017; Gunter and Önder, 2018). The experience of the host (or hotel manager) is found to positively affect performance (Brouder and Heriksson, 2013; Gémar et al., 2016; Xie and Mao, 2017), and the booking pace is a key determinant of revenue. Releasing an excessive percentage of the capacity with a high advance could have a negative effect on performance, whereas releasing restrictions contributes avoiding revenue dilution (Legohérel et al., 2013).

Our model includes instrumental variables, and we find that the effect on revenues would be biased toward zero without incorporating this feature in the estimation technique. Addressing the endogeneity problem implies a more negative effect of a price surge. This indicates that unobserved factors that are good for revenues are also related to a higher price surge variable. This is quite natural, considering that hosts can be aware of characteristics that are missing in our data, and positive factors could incline hosts to be more confident about a price surge. Unfortunately, as it turns out, this was a suboptimal policy for most of the listings if we consider maximizing revenues as the only target.

Reference

Abrate, G., and Viglia, G. (2016). Strategic and tactical price decisions in hotel revenue management. *Tourism Management*, *55*, 123-132.

Abrate, G., Fraquelli, G., and Viglia, G. (2012). Dynamic pricing strategies: Evidence from European hotels. *International Journal of Hospitality Management*, *31*(1), 160-168.

Abrate, G., Nicolau, J. L., and Viglia, G. (2019). The impact of dynamic price variability on revenue maximization. *Tourism Management*, *74*, 224-233.

Alderighi, M., Gaggero, A. A., and Piga, C. A. (2018). Static, dynamic and discriminatory pricing.

Alderighi, M., Nicolini, M., and Piga, C. A. (2016). Targeting leisure and business passengers with unsegmented pricing. *Tourism Management*, *54*, 502-512.

Arbel, A., & Pizam, A. (1977). Some determinants of urban hotel location: The tourists' inclinations. *Journal of Travel Research*, *15*(3), 18-22.

Athey, S., and Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, *113*(27), 7353-7360.

Athey, S., and Wager, S. (2019). Estimating Treatment Effects with Causal Forests: An Application. *arXiv preprint arXiv:1902.07409*.

Athey, S., Tibshirani, J., and Wager, S. (2019). Generalized Random Forests. *Annals of Statistics*, 47(2), 1148-1178.

Basu, S., Kumbier, K., Brown, J. B., & Yu, B. (2018). Iterative random forests to discover predictive and stable high-order interactions. *Proceedings of the National Academy of Sciences*, *115*(8), 1943-1948.

Bergantino, A. S., & Capozza, C. (2015). One price for all? Price discrimination and market captivity: Evidence from the Italian city-pair markets. *Transportation Research Part A: Policy and Practice*, *75*, 231-244.

Breiman, L. (2001). Random forests. *Machine learning*, *45*(1), 5-32.

Brouder, P., & Eriksson, R. H. (2013). Staying power: what influences micro-firm survival in tourism? *Tourism Geographies*, *15*(1), 125-144.

Butler, R. (1998). Seasonality in tourism: Issues and implications. *The Tourist Review*.

Chen, C. F., and Rothschild, R. (2010). An application of hedonic pricing analysis to the case of hotel rooms in Taipei. *Tourism Economics*, *16*(3), 685-694.

Chen, P., Dong, W., Lu, X., Kaymak, U., He, K., and Huang, Z. (2019). Deep Representation Learning for Individualized Treatment Effect Estimation using Electronic Health Records. *Journal of biomedical informatics*, 103303.

Crump, R. K., Hotz, V. J., Imbens, G. W., and Mitnik, O. A. (2008). Nonparametric tests for treatment effect heterogeneity. *The Review of Economics and Statistics*, 90(3), 389-405.

Daoud, A., and Johansson, F. (2019). Estimating Treatment Heterogeneity of International Monetary Fund Programs on Child Poverty with Generalized Random Forest.

Davis, J., & Heller, S. B. (2017). Using causal forests to predict treatment heterogeneity: An application to summer jobs. *American Economic Review*, 107(5), 546-50.

De Oliveira Santos, G. E. (2016). Worldwide hedonic prices of subjective characteristics of hostels. *Tourism Management*, *52*, 451-454.

Ert, E., Fleischer, A., & Magen, N. (2016). Trust and reputation in the sharing economy: The role of personal photos in Airbnb. *Tourism Management*, *55*, 62-73

Escobari, D. (2012). Dynamic pricing, advance sales and aggregate demand learning in airlines. *The Journal of Industrial Economics*, *60*(4), 697-724.

Espinet, J. M., Saez, M., Coenders, G., and Fluvià, M. (2003). Effect on prices of the attributes of holiday hotels: a hedonic prices approach. *Tourism Economics*, *9*(2), 165-177.

Farbmacher, H., Kögel, H., and Spindler, M. (2019). Heterogeneous Effects of Poverty on Cognition.

Gallego, G., and Van Ryzin, G. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management science*, *40*(8), 999-1020.

Gémar, G., Moniche, L., & Morales, A. J. (2016). Survival analysis of the Spanish hotel industry. *Tourism Management*, *54*, 428-438.

Gibbs, C., Guttentag, D., Gretzel, U., Morton, J., and Goodwill, A. (2018). Pricing in the sharing economy: a hedonic pricing model applied to Airbnb listings. *Journal of Travel and Tourism Marketing*, *35*(1), 46-56.

Gibbs, C., Guttentag, D., Gretzel, U., Yao, L., & Morton, J. (2018). Use of dynamic pricing strategies by Airbnb hosts. *International Journal of Contemporary Hospitality Management*.

Green, D. P., and Kern, H. L. (2012). Modeling heterogeneous treatment effects in survey experiments with Bayesian additive regression trees. *Public opinion quarterly, 76*(3), 491-511.

Gunter, U., & Önder, I. (2018). Determinants of Airbnb demand in Vienna and their implications for the traditional accommodation industry. *Tourism Economics, 24*(3), 270-293.

Guttentag, D. (2015). Airbnb: disruptive innovation and the rise of an informal tourism accommodation sector. *Current issues in Tourism*, *18*(12), 1192-1217.

Hamari, J., Sjöklint, M., and Ukkonen, A. (2016). The sharing economy: Why people participate in collaborative consumption. *Journal of the association for information science and technology*, *67*(9), 2047-2059.

Hill, D. (2015). The secret of Airbnb's pricing algorithm. *IEEE Spectrum*, *20*.

Kim, J., Yoon, Y., and Zo, H. (2015, July). Why People Participate in the Sharing Economy: A Social Exchange Perspective. In *PACIS* (p. 76).

Koenigsberg, O., Muller, E., and Vilcassim, N. J. (2008). easyJet® pricing strategy: Should low-fare airlines offer last-minute deals? *QME, 6*(3), 279-297.

Kwok, Linchi, and Karen L. Xie. "Pricing strategies on Airbnb: Are multi-unit hosts revenue pros?" *International Journal of Hospitality Management* 82 (2019): 252-259.

Lechner, M. (2019). Modified Causal Forests for Estimating Heterogeneous Causal Effects.

Legohérel, P., Fyall, A., & Poutier, E. (Eds.). (2013). *Revenue management for hospitality and tourism*. Woodeaton: Goodfellow Publishers.

Magno, F., Cassia, F., and Ugolini, M. M. (2018). Accommodation prices on Airbnb: effects of host experience and market demand. *The TQM Journal*, *30*(5), 608-620.

Malighetti, P., Paleari, S., and Redondi, R. (2009). Pricing strategies of low-cost airlines: The Ryanair case study. *Journal of Air Transport Management*, *15*(4), 195-203.

Melis, G., and Piga, C. A. (2017). Are all online hotel prices created dynamic? An empirical assessment. *International Journal of Hospitality Management*, *67*, 163-173.

Monroe, K. B. (1973). Buyers' subjective perceptions of price. *Journal of marketing research*, *10*(1), 70-80.

Morlotti, C., Cattaneo, M., Malighetti, P., and Redondi, R. (2017). Multi-dimensional price elasticity for leisure and business destinations in the low-cost air transport market: Evidence from easyJet. *Tourism Management*, *61*, 23-34.

Nieto-García, M., Muñoz-Gallego, P. A., and González-Benito, Ó. (2017). Tourists' willingness to pay for an accommodation: the effect of eWOM and internal reference price. *International Journal of Hospitality Management*, *62*, 67-77.

Oh, H. (2003). Price fairness and its asymmetric effects on overall price, quality, and value judgments: the case of an upscale hotel. *Tourism management*, *24*(4), 387-399.

Oskam, J., van der Rest, J.P., Telkamp, B., 2018. What's mine is yours—but at what price?

Revenue Pricing Manag. 1–18 (OnlineFirst).

Sekhon, J. S. (2008). The Neyman-Rubin model of causal inference and estimation via matching methods. *The Oxford handbook of political methodology*, *2*, 1-32.

Strittmatter, A. (2018). What is the Value Added by using Causal Machine Learning Methods in a Welfare Experiment Evaluation? *arXiv preprint arXiv:1812.06533*.

Su, X. (2007). Intertemporal pricing with strategic customer behavior. *Management Science, 53*(5), 726-741.

Thrane, C. (2005). Hedonic price models and sun-and-beach package tours: the Norwegian case. *Journal of Travel Research*, *43*(3), 302-308.

Tibshirani, J., S. Athey, and S. Wager. "grf: Generalized Random Forests (Beta). R package version 0.10. 3." (2019).

Tussyadiah, I. P. (2016), “Factors of satisfaction and intention to use peer-to-peer accommodation”, *International Journal of Hospitality Management*, Vol. 55, pp. 70-80.

Verhofstadt, E., and Maertens, M. (2014). Can agricultural cooperatives reduce poverty? Heterogeneous impact of cooperative membership on farmers' welfare in Rwanda. *Applied Economic Perspectives and Policy, 37*(1), 86-106.

Viglia, G., Mauri, A., and Carricano, M. (2016). The exploration of hotel reference prices under dynamic pricing scenarios and different forms of competition. *International Journal of Hospitality Management*, *52*, 46-55.

Wager, S., and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523), 1228-1242.

Wang, D., and Nicolau, J. L. (2017). Price determinants of sharing economy based accommodation rental: A study of listings from 33 cities on Airbnb. com. *International Journal of Hospitality Management*, *62*, 120-131.

Wendling, T., Jung, K., Callahan, A., Schuler, A., Shah, N. H., and Gallego, B. (2018). Comparing methods for estimation of heterogeneous treatment effects using observational data from health care databases. *Statistics in medicine*, *37*(23), 3309-3324.

Williams, K. (2018). Dynamic airline pricing and seat availability.

Xie, Y., Brand, J. E., and Jann, B. (2012). Estimating heterogeneous treatment effects with observational data. *Sociological methodology*, *42*(1), 314-347.

Xie, K., & Mao, Z. (2017). The impacts of quality and quantity attributes of Airbnb hosts on listing performance. *International Journal of Contemporary Hospitality Management*.

Yang, Y., Mao, Z. and Tang, J. (2017), "Understanding guest satisfaction with urban hotel location", *Journal of Travel Research*, Vol. 57 No. 2, pp. 243-59.

Ye, P., Qian, J., Chen, J., Wu, C. H., Zhou, Y., De Mars, S., ... and Zhang, L. (2018, July). Customized Regression Model for Airbnb Dynamic Pricing. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 932-940). ACM.

Table 1. **Example**

propertyid	date	status	price[xii]	bookeddate	reservationid	average_price
(1)	(2)	(3)	(4)	(5)	(6)	(7)
1999468	01/09/2016	A	266			
1999468	02/09/2016	R	297	10/08/2016	65214712	303
1999468	03/09/2016	R	297	10/08/2016	65214712	303
1999468	04/09/2016	R	297	10/08/2016	65214712	303
1999468	05/09/2016	R	308	10/08/2016	65214712	303
1999468	06/09/2016	R	308	10/08/2016	65214712	303
1999468	07/09/2016	R	308	10/08/2016	65214712	303
1999468	08/09/2016	R	308	10/08/2016	65214712	303
1999468	09/09/2016	R	292	25/08/2016	65214714	292
1999468	10/09/2016	R	292	25/08/2016	65214714	292
1999468	11/09/2016	R	292	25/08/2016	65214714	292
1999468	12/09/2016	R	292	25/08/2016	65214714	292
1999468	13/09/2016	A	289			
1999468	14/09/2016	R	288	01/08/2016	65214717	305
1999468	15/09/2016	R	311	01/08/2016	65214717	305
1999468	16/09/2016	R	311	01/08/2016	65214717	305
1999468	17/09/2016	R	311	01/08/2016	65214717	305
1999468	18/09/2016	A	312			
1999468	19/09/2016	A	358			
1999468	20/09/2016	A	358			
1999468	21/09/2016	R	363	25/06/2016	65214720	363
1999468	22/09/2016	R	363	25/06/2016	65214720	363
1999468	23/09/2016	R	363	25/06/2016	65214720	363
1999468	24/09/2016	R	363	25/06/2016	65214720	363
1999468	25/09/2016	A	173			
1999468	26/09/2016	A	173			
1999468	27/09/2016	A	173			
1999468	28/09/2016	A	173			
1999468	29/09/2016	A	173			
1999468	30/09/2016	R	230	22/09/2016	65214722	230

Table 2. **Descriptive statistics**

variable name	meaning	mean	sd	min	max
Time invariant					
numberofreview	number of reviews	11.259	14.451	0	173
bedrooms	number of bedrooms	2.136	1.504	0	10
bathrooms	number of bathrooms	1.838	1.228	0	8
maxguests	maximum number of allowed guests	5.056	3.053	1	16
responserate	host response rate to customers' requests	85.197	30.224	0	100
securitydeposit	amount of deposit required	429.416	598.971	0	11749
distance_disco	distance (in km) to the closest club	4.832	4.841	0.010	16.250
distance_beach	distance (in km) to the closest beach	1.674	1.826	0.016	9.321
entireproperty	entire home/apartment (Dummy variable)	0.752	0.432	0	1
cancelpolicy_mod	moderate cancelation policy (Dummy variable)	0.089	0.285	0	1
cancelpolicy_strict	strict cancelation policy (Dummy variable)	0.781	0.414	0	1
instantbooking	bookings instantly accepted (Dummy variable)	0.854	0.353	0	1
munic3	listing located in Sant Antony de Portmany (Dummy variable)	0.086	0.280	0	1
munic4	listing located in Sant Joan de sa Labritja (Dummy variable)	0.057	0.232	0	1
munic5	listing located in Sat Josep de sa Talaia (Dummy variable)	0.279	0.448	0	1
munic6	listing located in Santa Eularia des Riu (Dummy variable)	0.217	0.413	0	1
longevity	number of months since account creation	14.751	13.909	0	73.833
multiplehost	entire home/apartment (Dummy variable)	0.687	0.464	0	1
medium_rate	score rating [4;4,5) (Dummy variable)	0.205	0.404	0	1
high_rate	score rating [4,5;5] (Dummy variable)	0.552	0.497	0	1
never_rated	listings with no score rating (Dummy variable)	0.166	0.372	0	1
ln_price	natural logarithm of the price set per night (by default)	5.225	0.865	2.485	7.948
OR_June	occupancy rate in June	0.521	0.272	0.033	1
OR_July	occupancy rate in July	0.657	0.259	0.032	1
OR_June_local	average occupancy rate in June in the same area	0.195	0.026	0.140	0.239
OR_July_local	average occupancy rate in July in the same area	0.299	0.034	0.199	0.357
perc30_june	percentage of days of June, booked with at least 30 days in advance	0.622	0.393	0	1
perc30_july	percentage of days of July booked with at least 30 days in advance	0.695	0.349	0	1
leadtimen_July	average lead time in July	100.448	50.506	0	181
leadtimen_August	average lead time in August	101.536	56.107	0	212

Table 3. **Instrumental variable regression**

	Revenue_July		Revenue_August	
	β	s.e.	β	s.e.
W_{ij} (price surge)	-0.914***	[-5.99]	-0.730***	[-5.88]
bedrooms	0.080***	[3.93]	0.108***	[5.14]
bathrooms	0.075***	[3.89]	0.049**	[2.46]
maxguests	0.001	[0.10]	0.003	[0.32]
responserate	0.001***	[2.64]	0.002***	[3.89]
securitydeposit	0.000***	[3.23]	0.000***	[3.15]
distancedisco	-0.013***	[-3.50]	-0.012***	[-3.20]
distancebeach	-0.000	[-0.06]	0.013	[1.52]
entireproperty	0.364***	[9.63]	0.404***	[10.50]
cancelpolicymod	0.108**	[1.98]	0.040	[0.72]
cancelpolicystrict	0.085**	[2.09]	0.094**	[2.23]
instantbooking	0.020	[0.58]	-0.029	[-0.86]
munic3	-0.088	[-1.44]	-0.108*	[-1.91]
munic4	-0.133**	[-2.05]	-0.141**	[-2.04]
munic5	-0.004	[-0.11]	-0.021	[-0.55]
munic6	-0.054	[-1.39]	-0.108***	[-2.81]
longevity	-0.009***	[-7.84]	-0.09***	[-8.12]
Or_{t-1}	0.465***	[9.21]	0.580***	[9.73]
$Or_{t-1}local$	0.860	[1.35]	0.306	[0.63]
multiplehost	0.0417	[1.53]	0.098***	[3.56]
$leadtimen_t$	0.002***	[7.60]	0.002***	[6.82]
medium_rate	0.060	[1.17]	0.046	[0.93]
high_rate	0.092*	[1.91]	0.096**	[2.02]
never_rated	0.048	[0.82]	0.146**	[2.38]
numberofreviews	0.007***	[6.78]	0.005***	[4.37]
ln_publishednightlyrate	0.452***	[18.00]	0.394***	[15.16]
$perc30_{t-1}$	-0.160***	[-4.36]	-0.108***	[-3.02]
_cons	4.524***	[24.05]	4.834***	[22.91]
N	2599		2226	
R^2	0.596		0.609	
adj. R^2	0.592		0.605	

t-statistics in brackets
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 1. **Price Surge (*W*$_i$)**

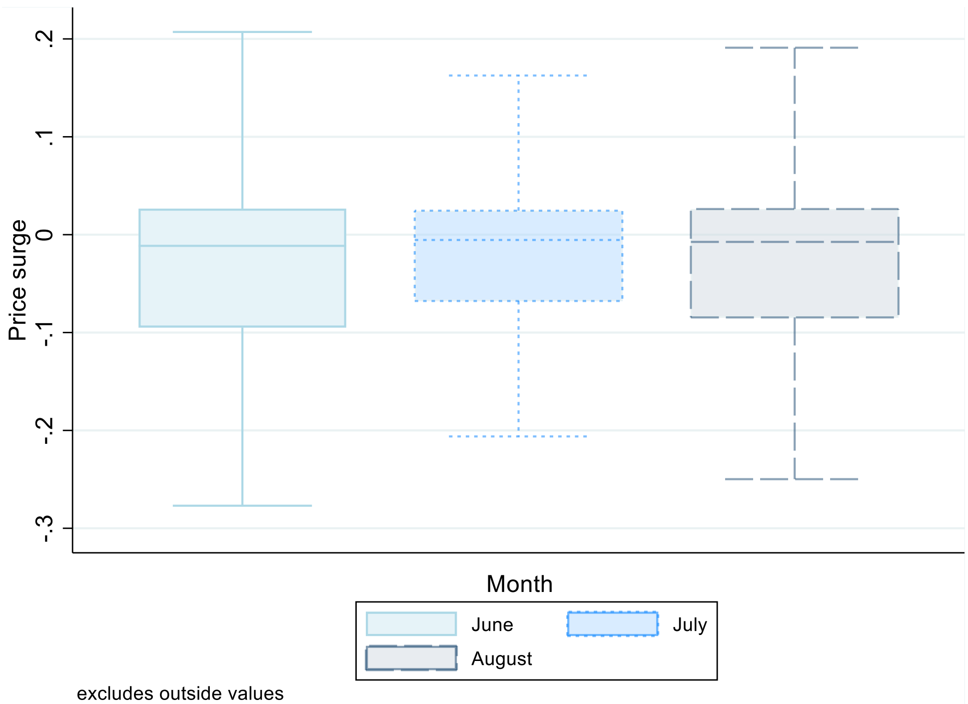

Figure 2. **Price surge (*W*$_i$) over lead-time (local polynomial)**

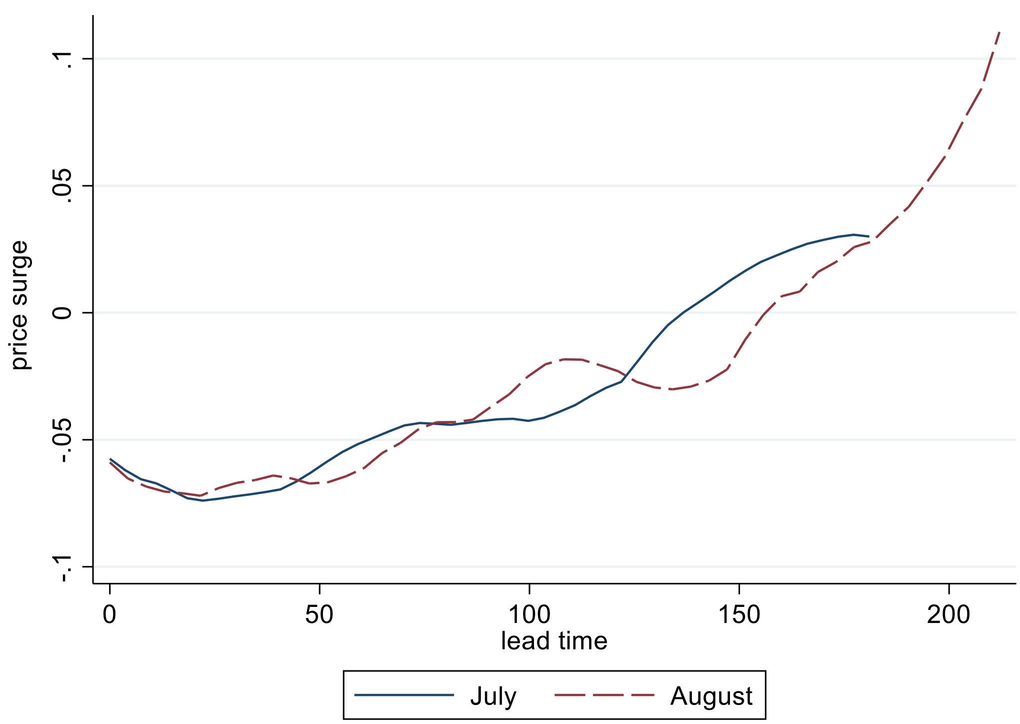

Figure 3. **Distributions of listings' log-revenue: July and August 2016**[xiii]

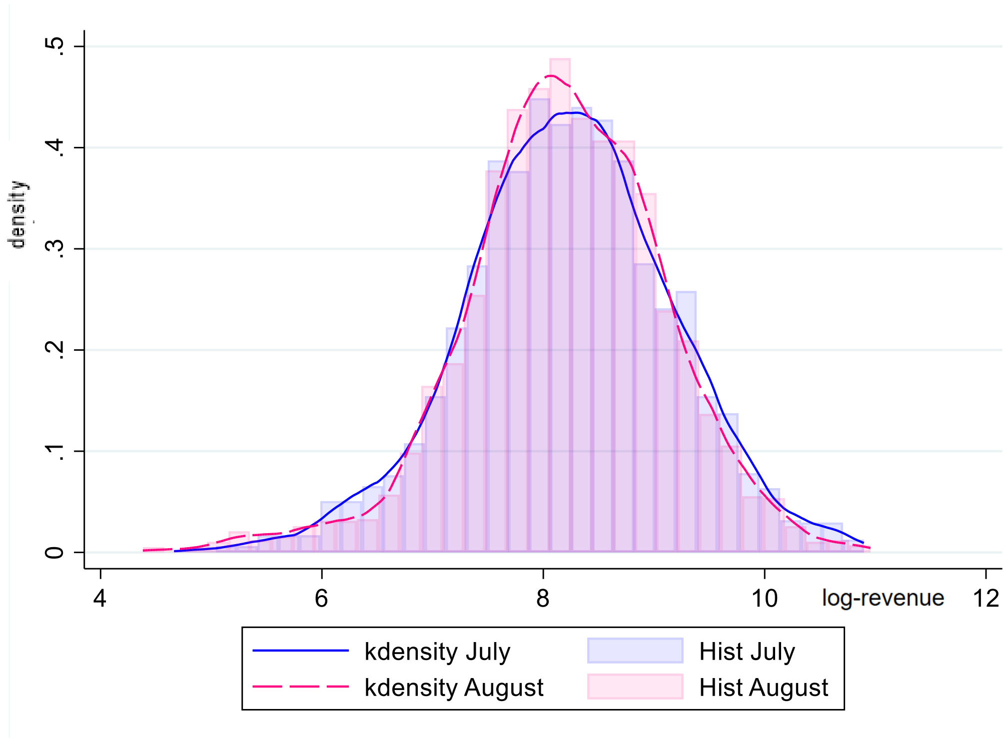

Figure 4. **Distribution of the heterogeneous treatment effect (out-of-bag) July and August 2016**

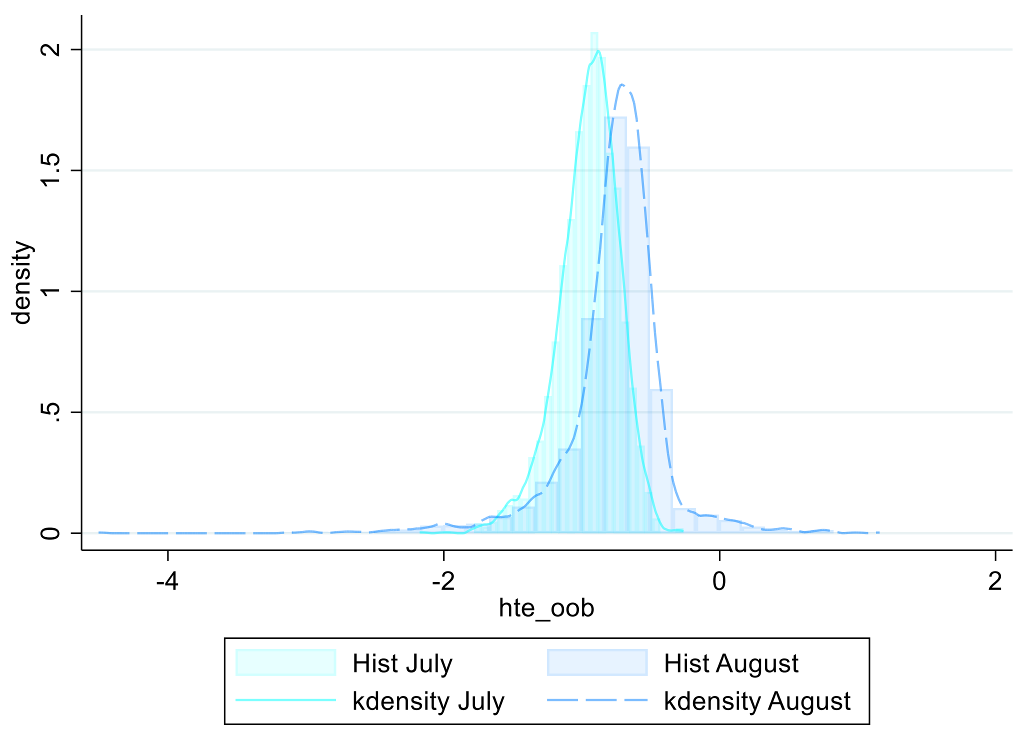

Figure 5. **List of variables that maximize the variance in the treatment effect**

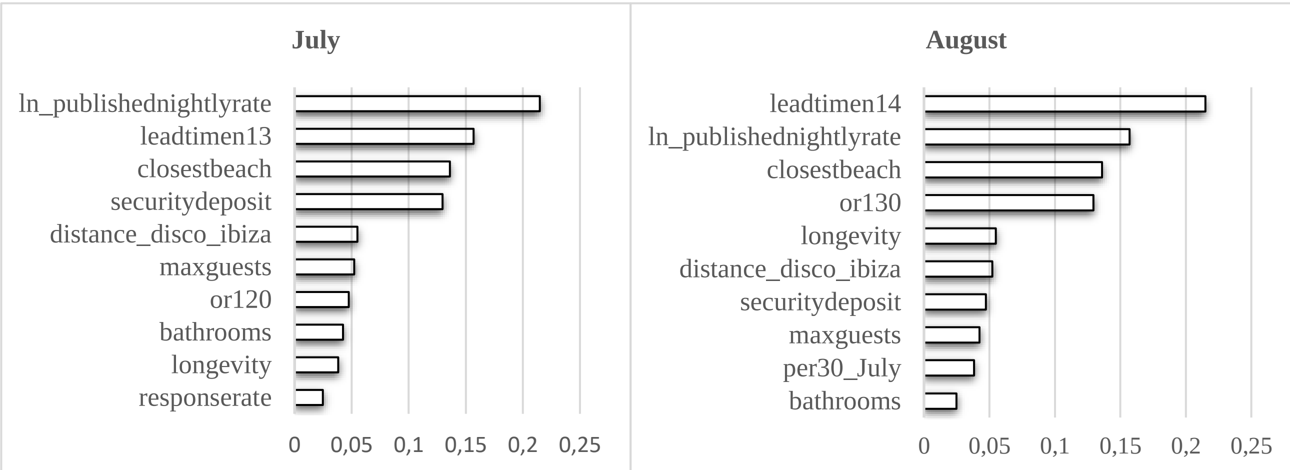

Figure 6a. **Treatment effects predictions**[xiv]

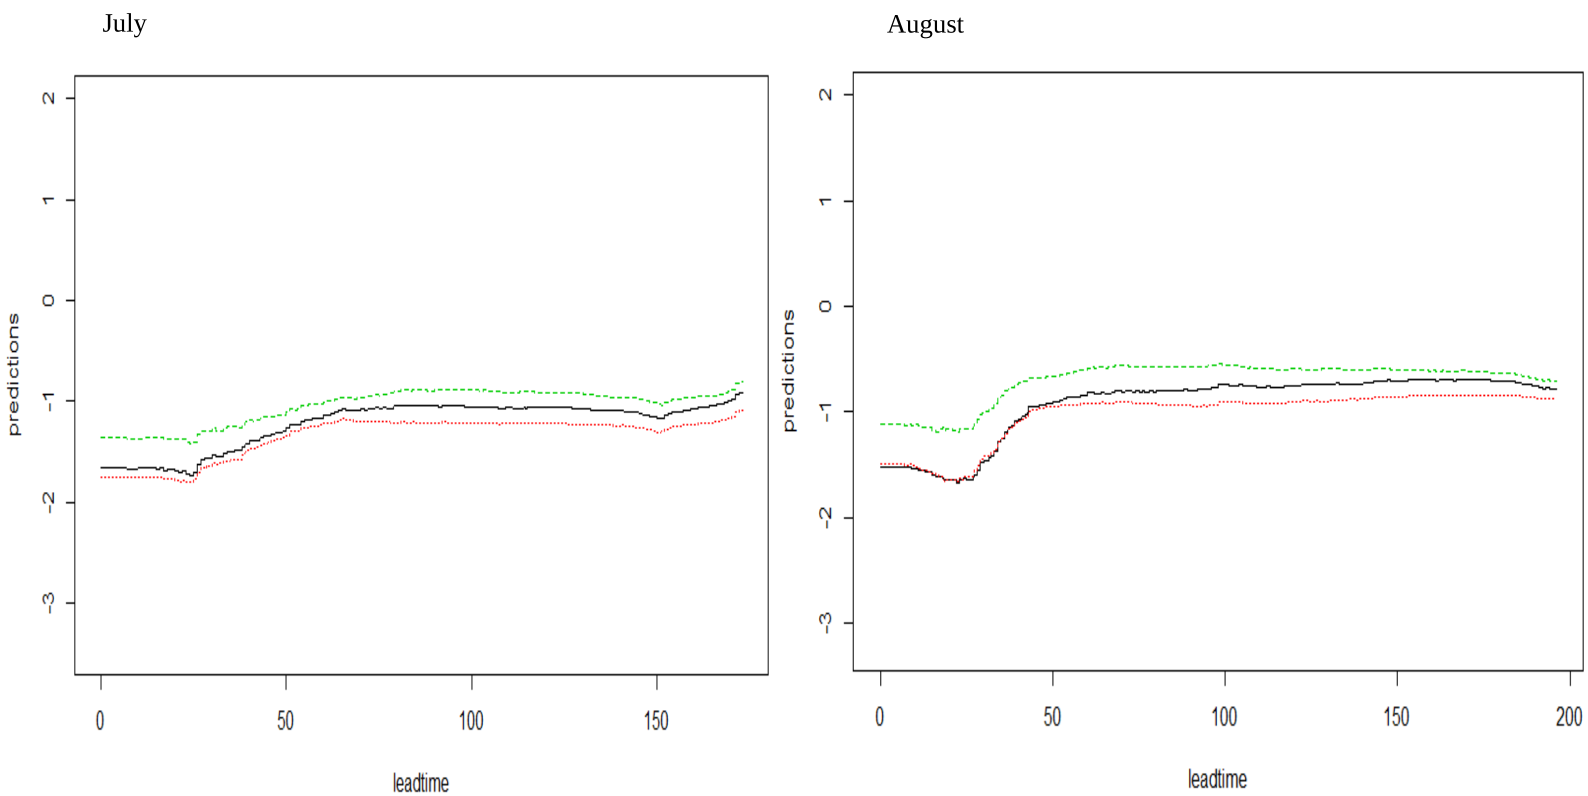

Figure 6b. **Treatment effects predictions**

July

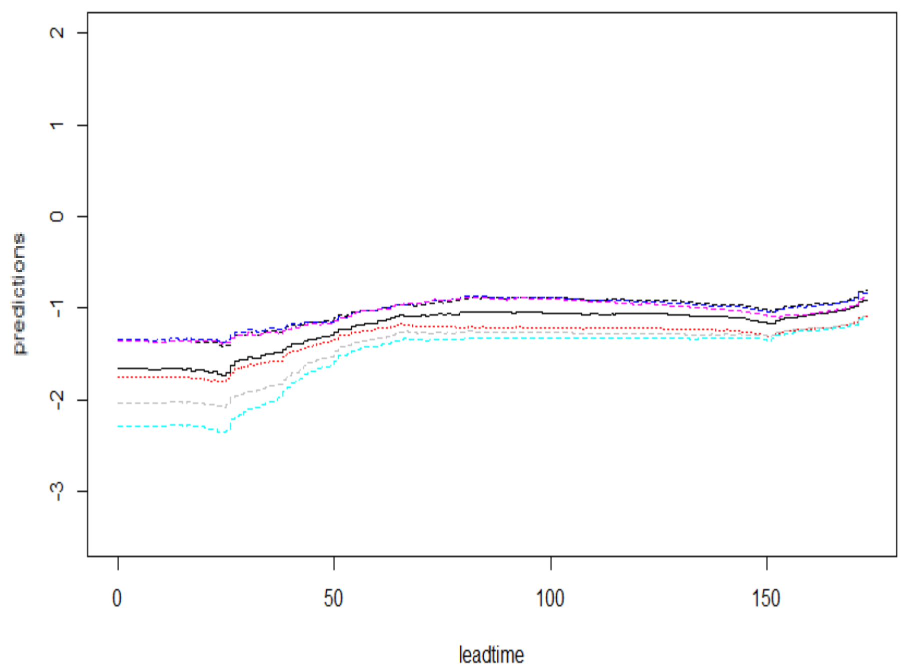

August

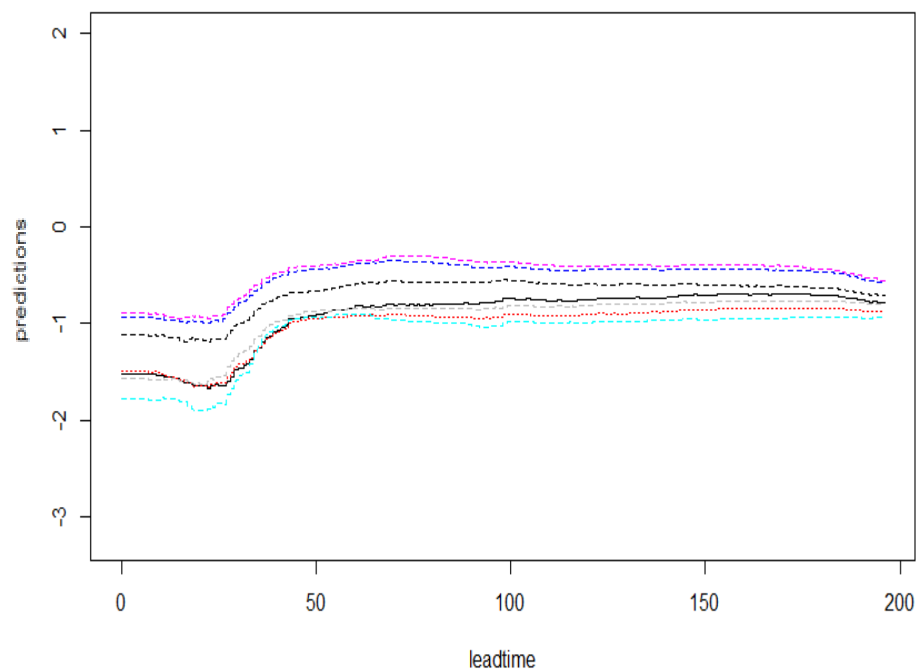

[i] https://www.airdna.co/pricing

[ii] Unless the booked date has been reserved with a zero leadtime (booked on the same date of service consumption), but we do not have such cases in our dataset.

[iii] Status=A; in this case, the average price of non-booked days is 244.8USD

[iv] We compare bookings of the same property and during the same month

[v] Price are expressed in US dollars

[vi] Note that #65214720 (in: 21-09-16/ out: 24-09/16) has been booked with more advance (25-06-2016); in computing the sign of the price variation, we consider the BOOKEDDATE (not the calendar date).

[vii] We consider uniform pricing when the value of the price surge is between -0.001 and +0.001.

[viii] It is a graphical representation of smoothed values deriving from a kernel-weighted local polynomial regression.

[ix] This set of dummy variables originates from a continuous variable (*overallrating*) containing the score rate (from 0 to 5) left by previous guests. *low_rate* equals one if *overallrating* is in the range [0, 4); *medium_rate* for the range [4, 4,5) and *high_rate* for the range [4,5, 5]. Listings with no available score are classified as "*never_rated*".

[x] We selected only those features with a variable importance above 0.2, which corresponds to the average value of the index "variable importance." For technical details, we refer to Athey et al. (2019).

[xi] Through descriptive statistics analysis, we found homogeneity in the demand (proxied by the occupancy rate) within the month of July and within the month of August, with the latter recording a slightly higher level of demand, which justifies our choice of the two-period analysis.

[xii] Prices (columns 4 and 7) are expressed in USD.

[xiii] *Hist*= histogram; *kdensity*= kernel density

[xiv] Lead time is expressed in days