Authority for Consumers & Markets

Working paper

Price effects of non-brand bidding agreements in the Dutch hotel sector

Stefan Haasbeek, Jan Sviták and Jan Tichem

June 7th, 2019

Markets that work well for people and businesses

Abstract
Recently competition authorities have enforced against agreements among competitors to refrain from bidding on each other's brand-related keywords on search engines. In this paper, we investigate the effect of these so-called "non-brand bidding agreements" on hotel prices in the Netherlands. Hotels sell their product on their own website and through Online Travel Agents (OTAs). Some hotels restrict OTAs in bidding on their brand name on search engines. We use data on hotel pricing and the presence of an advertisement restriction on the hotel level. We apply a data-driven trajectory balancing approach to correct for unobserved heterogeneity between hotels that do and do not impose advertising restrictions on OTAs. The analysis shows that NBBAs increase price on hotel websites relative to the price on OTAs. We conclude that the advertising restrictions are likely to lead to higher prices on hotel websites, and that potential ad spend savings are not passed on to consumers in the form of lower prices.

Keywords: advertising, competition, hotel, non-brand bidding agreements, Online Travel Agents, price effects, search advertising

Version: June 2019

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

1. Introduction

Competition authorities have recently enforced against restrictions on online search advertising that firms impose on competitors. The U.S. Federal Trade Commission found that *1-800 Contacts*, the biggest U.S. retailer of contact lenses, violated Section 5 of the FTC Act by entering into an agreement of this kind with fourteen of its competitors. The firms agreed not to bid for paid links in case the user's search phrase contains a competitor's brand name (FTC, 2018). Competitors of *1-800 Contacts* thus agreed not to participate in ad auctions when consumers use keywords containing the brand "*1-800 Contacts*", and vice versa. The European Commission recently imposed a fine on *Guess*, a manufacturer of clothing apparel and accessories, for, amongst others, forbidding its distributors to bid on *Guess*' brand name in sponsored search auctions (EC, 2018). Following the market study on digital comparison tools by the UK Competition and Markets Authority (CMA, 2017), we refer to this type of restriction as a "non-brand bidding agreement" (NBBA).

NBBAs seem quite common. Besides the examples mentioned above, the CMA encountered NBBAs in the UK markets for broadband, credit cards, energy, and home insurance (CMA, 2017). This paper studies the hotel sector in the Netherlands, where NBBAs are widespread. As NBBAs have only recently come to the attention of competition authorities, empirical evidence regarding their effects on consumer welfare can help to shape competition policy in this area. Whereas there exists considerable evidence on the effects of NBBAs on web traffic (see our section 2, where we review the related literature), this paper is, to the best of our knowledge, the first to study the price effects of NBBAs.

In the Dutch hotel sector, suppliers sell their product through their own website and through OTAs. Hotels determine the terms of trade on the OTA, so they set their own price. OTAs charge a commission per transaction. There is considerable variation between hotels regarding NBBA-status, which allows us to estimate effects over a relatively large sample. Our theory of harm is that hotels, by negotiating NBBAs with OTAs, can to some extent shield their own online sales channel against competition from other hotels present on the OTAs. A customer using a branded keyword on a search engine is clearly considering the hotel. The NBBA reduces the probability that the consumer will visit the OTA, where she is confronted with many other brands. The NBBA may thus reduce the extent to which consumers compare different brands, which allows the hotel to charge a higher price on its own website. Note that we do not posit that the NBBAs affect the price of the hotel on OTAs. We consider this possibility implausible, since once a customer reaches an OTA, the hotel has to compete there with many other hotels. An NBBA may however prevent the consumer from considering the offers on the OTA altogether. In sum, our theory of harm is that the NBBAs of hotels essentially induce market segmentation by consumer groups, as OTAs cannot target the customers that use a branded keyword with search ads.

Importantly, NBBAs may also generate efficiencies. CMA (2017) notes that NBBAs may, amongst others, reduce advertising costs. In the context we study, the possibility for ad spend savings seems a priori plausible. NBBAs can reduce advertising costs as these clauses imply that fewer bidders join the sponsored search auction. The restricting hotel can therefore refrain from bidding altogether, or win the

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

auction at a lower bid. Moreover, the presence of competitors' ads may crowd out organic (free) clicks to the supplier, especially for well-known brands (Blake et al., 2015; Simonov et al., 2017; Simonov and Hill, 2018). Bidding on one's own brand name may thus be quite costly. From a competition policy point of view, however, an important question is whether hotels pass on any savings on advertisement costs to consumers. This is a necessary requirement for firms to successfully invoke efficiency-based exceptions to a finding of cartel behaviour. Under the efficiency explanation, we expect that hotels reduce price on their own website.

To distinguish between these conflicting hypotheses regarding price, we test whether the difference in price on the hotel's website and the price on the OTA is positively or negatively affected by the presence of an NBBA. We use the OTA price level as a benchmark, since we consider it unlikely that NBBAs affect OTA pricing (we elaborate on this assumption in section 7). For this analysis we use data from a meta-search site on hotel prices in the Netherlands. We select all hotels that advertise their own website on the meta-search site for at least 20% of the dates within the observed period. In addition to this, we collected data from two OTAs on the presence of an NBBA between the hotels in our dataset and the OTA. This yields a non-random dataset of 183 hotels, of which 131 have an NBBA with at least one OTA (in almost all cases, a hotel has an NBBA with both OTAs or none).[1]

As the hotels in our dataset are not assigned to a treatment randomly, identifying the causal effect of NBBAs on the price difference between the hotel's website and the OTA is non-trivial. A greater or smaller price difference between NBBA and non-NBBA hotels may well be caused by other differences between the hotels than NBBA-status. To address this problem, we balance NBBA hotels and non-NBBA hotels on the basis of characteristics that we do observe. We compare hotels on the basis of their pricing on one of the OTAs (recall that we assume that the OTA price is unaffected by NBBAs). The approach we adopt is called 'trajectory balancing' (Hazlett & Xu, 2018; Hazlett, 2018). By using the OTA price, we match hotels on a core strategic parameter that they choose.

Before we describe the results, we note that hotels in our sample do not have complete freedom in setting the price on their own website. In the Netherlands, OTAs are allowed to impose so-called Price Parity Clauses (PPCs) on hotels, which imply that hotels cannot post a lower price on their own website than on the OTA. PPCs are implemented to prevent hotels from advertising their hotel on OTAs but conducting the transaction on their own website at a lower price, thereby saving on OTA commissions. Next, hotels may be reluctant to price their website higher than the OTA out of fear of cannibalizing their own channel (EU Competition Authorities, 2016). In this market context, one may therefore doubt whether NBBAs will lead to higher prices on hotels' websites at all, since hotels cannot price lower than the OTA and generally avoid to price higher. Be this as it may, in our dataset hotels price their website on average 3 percent lower than OTAs. Hotels thus consistently choose to price their own website cheaper than OTAs, which is consistent with other research.

[1] We do not share our data, nor reveal the identities of the firms who supplied the data, for confidentiality reasons.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

Our estimations consistently show that NBBAs lead to a higher price on the hotel's website relative to the OTA price, almost closing the gap between the OTA price and the price on the hotel website. In our specification including all hotels, we find that NBBAs lead to a price increase on the hotel website relative to the OTA price of roughly 2 percent. We also run the analysis for a subset of hotels who relatively often violate the price parity clause OTAs impose on hotels. For these hotels we find a stronger positive price effect of NBBAs, roughly a 5 percent price increase on the hotel website relative to the OTA price. Finally, in a number of cases the hotel is sold out on OTAs, in which case the PPC is completely irrelevant. We control for this situation in our model and interact it with NBBA-presence. The interaction variable is strongly related to higher prices on the hotel website for the subsample of hotels that are often out-of-parity. We conclude that in the Dutch hotel sector, i) NBBAs are likely to lead to higher prices on hotels' websites, and ii) hotels do not pass on possible cost savings on advertising to consumers in the form of lower prices.

The remainder of this paper is organised as follows. Section 2 discusses the related literature, section 3 develops our hypotheses regarding the price effects of NBBAs in the Dutch hotel sector, section 4 discusses our empirical approach, section 5 discusses the dataset, section 6 presents our results, section 7 discusses our assumption that NBBAs do not affect OTA pricing, and section 8 finishes with concluding remarks.

2. Related literature

Our research is related to the literature on the role of (search) advertising in consumers' search and purchase decisions, and how this affects competition. Haan and Moraga-González (2011) develop a theoretical model where advertising increases the saliency of the firm. A more salient firm is more likely to be searched first by consumers. Being searched first is profitable since given the existence of some search costs, firms have some market power over visiting consumers. This provides a rationale to invest in advertising. Firms are in a classic prisoner's dilemma: they would be better off if no one advertised, but each firm has an incentive to become prominent. Desai, Shin and Staelin (2014) show theoretically that firms may bid on a competing firm's brand name in order to siphon off traffic. The competing firm may optimally respond by bidding on its own name, even if it would not be optimal to do so in case competitors refrain from bidding. Bidding on one's own brand name can be costly because it crowds out free organic traffic. The defensive motive, however, may outweigh this cost. In these papers, advertising can harm consumers because it does not alter preferences for firms in equilibrium, but its costs are passed on in the price. Sayedi et al. (2014) consider that besides search advertising, firms may advertise to raise general awareness which leads to both direct sales for the firm but also more searches for the product in general. In this setting, competitors can choose to free-ride on the firm's expenses to increase general awareness of the product by bidding on the firm's brand name. If the competitor does not invest in awareness itself, it has more means to bid aggressively in the paid search auction. In response, the firm may choose to allocate more of its budget to traditional advertising, which also yields some direct

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

sales. The authors show that this mechanism can explain why search engines rank ads not only on the basis of bids, but also by relevance.

Studying experimentally the case of *eBay*'s search advertising, Blake et al. (2015) find that if *eBay* does not advertise on keywords containing *eBay*'s brand name, virtually all paid clicks are recovered through organic clicks. Paid advertising on its brand name thus almost completely cannibalizes on *eBay*'s organic traffic. During the experiment, *eBay* did not face competitors' ads. The authors also find in another experiment that paid advertising on non-branded keywords crowds out organic traffic too much to be profitable for *eBay*. The results of Blake et al. (2015) may be explained by the enormous strength of *eBay*'s brand name, but there seems to be more to the story. Golden and Horton (2018) study an experimental setting where two closest competitors bid on each other brand names, and also find defensive bidding to be ineffective for a brand that is not well-known. In their study, when one of the firms experimentally shuts down advertising on its own brand name, the competitor does not receive more clicks. However, for another lesser known brand Coviello et al. (2017) do find that defensive bidding is effective.

Simonov et al. (2017) take on the issue of brand bidding for a very large sample. They study experimentally the effectiveness of brand bidding for a large sample of 2,500 brands on Bing. The authors find that, on average, firms are in a prisoner's dilemma when considering to bid on each other brand name, as predicted by the models cited above. The authors also find that when a firm bids on its own brand name it virtually always wins the top ad position, but the presence of competitors significantly increases the advertising costs of the focal brand. The cost per click for focal brands increases from $0.23 to $0.60 - $1.03, depending on the number of competitors bidding. Compared to no ads shown at all, placing one ad for the brand searched for increases the probability of a click on the brand (either paid or organic) from 77% to 79%. This effect almost disappears for very strong brands, which is consistent with Blake et al. (2015). Although paid advertising leads to more clicks for smaller brands, it strongly cannibalizes organic clicks: advertising leads to 60% of clicks being paid. Simonov et al. (2017) also study the impact of competitors' ads on the effectiveness of brand advertising. Compared to when only an ad is shown from the brand searched for, the addition of competitors' ads reduces traffic to the focal brand by 1-5% of clicks, depending on the number of competitors bidding. The addition of competitors' ads leads to an even greater share of paid results for the focal brand, up to 84%. Finally, when the focal firm does not advertise, competitors' ads lead to a large loss of traffic. When four competitors are present, the focal brand loses 42% of the total traffic it would have had if no ads are present at all. The authors explain this strong 'click-stealing' effect by referring to users' inclination to click on top-ranked results, that is, the position effect (see e.g. Baye et al., 2016, and Ursu, 2018).

Simonov and Hill (2018) study brand bidding for almost 1,500 brands on Bing by using randomized experiments on Bing. The authors find that competitors in top ad positions receive 6-20% of the clicks. If the focal brand is in the top position competitors only receive 1-3% of the clicks. Simonov and Hill (2018) also study the quality of traffic, as measured by the share of searchers that quickly return to the search engine after clicking an ad. When competitors are in the top paid position, this 'quick back rate' lies

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

around 43%, which is significantly higher than the 6% 'quick back rate' of clicks on focal brands' links. For more relevant competitors, the quick back rate lies 5 percentage points lower. Simonov and Hill (2018) conclude that their results imply a mixed verdict on the question whether brand-bidding should be prohibited. Clearly, for many consumers competitor's bids are valuable as the majority of consumers clicking these ads does not quickly return to the search engine, and more relevant competitors are able to retain more clicks. This implies that brand searches are not purely navigational (that is, an alternative for typing in the URL). However, a significant share of consumers clicking these ads lose time if a competitor's ad is in the top position compared to when the focal brand's ad is in the top position, as evidenced by the higher quick back rate. These costs are not significant in practice because many brands successfully defend their brand, but this is costly for firms, as shown by Simonov et al. (2017).

Finally, some papers study the effect of ads on organic clicks for wider sets of keywords. Yang and Ghose (2010) study the effectiveness of paid advertising on Google by a large nationwide retailer for hundreds of different keywords. By estimating a structural model of consumer click- and purchase behaviour, the authors find that an organic listing on average positively affects clicks and conversion through paid links, and vice versa. This interdependence is strongest for the least competitive keywords such as branded keywords (containing the retailer's brand name), and weakest for the most competitive keywords such as keywords containing the brand name of suppliers of the retailer. The authors also find in a field experiment that the presence of ads reduces organic clicks for some keywords, but that for a majority of keywords the presence of ads increases the average click-through from organic results. Studying a dataset of non-branded searches, Baye et al. (2016) also find that the presence of an ad increases the number of organic clicks. These results suggest that when a keyword does not contain firm's brand name, bidding for ad space can also increase the number of organic clicks.

3. Possible price effects of NBBAs

In this section we formulate our hypotheses regarding the price effects of NBBAs negotiated by hotels. We start by outlining the general anti- and pro-competitive price effects based on the economics literature, and then turn to the specific circumstances of the sector we study.

3.1 Pro- and anti-competitive effects of NBBAs

Our first hypothesis is that an NBBA shields the hotel's website from competition from other hotels on the OTAs. The (empirical) literature on brand bidding shows that competitors can get some clicks by bidding on a firm's brand name, especially if the focal brand chooses not to participate in the ad auction. In our setting, OTAs are highly relevant competitors for hotels' websites, as next to the hotel itself, they contain a large number of other hotel offers. Simonov and Hill (2018) show that a substantial share of consumers value competitors' ads when performing a brand search, and this share is higher if the competitor is more relevant. NBBAs imposed on OTAs thus reduce the amount of clicks OTAs receive from brand-

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

searching consumers. In turn, this may enable the brands to raise the price on their own website. By reducing the amount of information on competitive offers readily available to consumers, NBBAs increase the costs of inspecting competing offers. Higher search costs typically lead to a lower elasticity of demand and therefore higher prices, as in e.g. Stahl (1989), Wolinsky (1986), and Anderson and Renault (1999). Our hypothesis thus treats advertising as informative to consumers as opposed to persuasive.[2] We consider this a plausible view on paid search advertising by the OTAs, as one of the core functions of OTAs is to help consumers compare competing offers.

More recent literature on consumer search has identified situations where higher search costs may lead to *lower* prices. We do not think these possibilities overturn our theory of harm in the present context. First, Armstrong and Zhou (2011) show that firms may optimally set a low price to gain a prominent (top) position in the ranking of platforms. This yields more sales and profits, but the firm must set a sufficiently low price to persuade the consumer not to search any further. If search costs increase, consumers are more reluctant to search further which increases the value of being prominent. Firms therefore set lower prices as search costs increase. This mechanism does not affect our theory of harm, as we postulate that NBBAs prevent consumers reaching OTAs altogether. Second, Moraga-González et al. (2017) show that an increase in search costs may decrease the number of searchers and hence demand, which in turn may imply lower prices. We also consider this mechanism unlikely to reverse our theory of harm, since NBBAs imposed on OTAs affect only consumers using a branded keyword, who clearly are already searching. At the same time, consumers considering whether to start searching have many alternatives to start their search which are not affected by the NBBAs (we discuss the importance of alternative ways of searching in more depth later on in this section). Hence, our first hypothesis is:

HYPOTHESIS 1 NBBAs shield hotels' websites from competing offers, which enables hotels to charge higher prices on their own website.

NBBAs may also lead to lower prices. Given that OTAs cannot bid on hotels' brand names, hotels no longer need to bid defensively on their own brand name to counter OTA ads. Hotels can then choose to lower their bids or refrain from bidding altogether. This mechanism has its theoretical underpinning in the models by Haan and Moraga-González (2011) and Desai et al. (2014), which show that firms face a prisoner's dilemma in advertising. That cost savings due to NBBAs are a real possibility is shown by Simonov et al. (2017) who, in their sample of 2,500 brands, find economically significant increases in the cost-per-click for the focal brand when competitors also bid on the focal brand name. For competition policy purposes, however, the question is whether potential costs savings are passed on to consumers to such an extent that it compensates for any competitive harm. That is why we focus our research on the price effect of NBBAs. Our second hypothesis is:

[2] The informative and persuasive view on advertising have contrasting normative implications on advertising. Under the former view, advertising may improve market outcomes whereas under the latter outcome advertising leads to changes in consumer tastes due to spurious product differentiation and creating brand loyalty. See Bagwell (2007) for an overview of the economic literature on advertising.

HYPOTHESIS 2 NBBAs yield savings on search advertising costs, which hotels can pass on to consumers in the form of lower prices on their website.

3.2 Details of the Dutch hotel sector

In this sub-section we provide more details on the sector we study, and discuss how these impact our hypotheses.

Hotels distribute their product in a number of ways. First, hotels sell directly to consumers, both offline and online. Second, hotels may sell through OTAs (e.g. *Agoda, Booking.com*, *Expedia, Hotels.com*, *Hotels.nl*, *Hotelspecials.nl, HRS*). In that case the hotel determines the sales price on the OTA and typically pays a commission as a percentage of the sales price to the OTA for each transaction. Third, hotels may use meta-search sites (MSSs) to increase their sales (e.g. *Google Hotel Finder, Kayak, TripAdvisor, Trivago*). MSSs do not conduct transactions but refer customers to a sales channel. The sales channel usually pays a cost per click to MSSs for referrals.

From the above, it is apparent that consumers can use several ways to search before purchase. Besides using a branded keyword on a search engine and clicking one of the paid or organic links, consumers may use non-branded keywords on search engines such as “compare hotels”. Consumers may also navigate directly to hotels’ websites, OTAs, and meta-search sites. Sometimes, search engines also show their own meta-search service in response to (branded) keywords. The extent to which NBBAs may lead to higher prices on hotels’ websites clearly depends on the alternatives for branded searches, and the extent to which they are used by consumers. We do not have data on the relative importance of using branded keywords for consumer search in this market, and we are not able to observe the intermediate steps in the causal chain running from the presence of an NBBA to higher prices on the websites of hotels. Instead, we evaluate the price effect of NBBAs on price directly.

In its digital comparison tools study, the CMA (2017) identified different types of NBBAs. The first, so-called ‘narrow NBBA’, implies that one advertiser agrees not to bid on another firm’s brand name when the search phrase equals that brand name. Second, under a ‘wide NBBA’, one advertiser agrees not to bid on another firm’s brand name whenever the search phrase includes that firm’s brand name. The last type, referred to as ‘negative matching agreements’, is where the restricted advertiser agrees to add another firm’s brand name to its list of ‘negative keywords’, which prevents the ads from appearing at all times when the search phrase includes the firm’s brand name. In our setting, most of the NBBAs are negative matching agreements (123 hotels from all 131 with some form of NBBA). This means that from all the known types of search advertising, Dutch hotels mostly apply the type with the strongest possible effects, both pro- and anti-competitive.

The Dutch hotel sector is characterised by the presence of narrow PPCs. Under these clauses hotels cannot post a lower price on their own website relative to the price hotels set on the OTA imposing the

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

PPC, although hotels can differentiate between OTAs.[3] The rationale behind the PPCs is that hotels cannot use the OTAs to advertise their rooms, but then cut out OTAs from transactions (and thereby preventing commission payments) by posting a slightly lower price on the hotel's website. In practice, however, hotels do not perfectly comply with PPCs. A study by 10 EU Competition Authorities found that in Member States with PPCs, 35 percent of hotels reported that they did undercut OTAs on their own website, 48 percent of which said they did so most of the time (EU Competition Authorities, 2016, p. 14). According to Centre for Market Insights (2019), 38 percent of hotels undercut the OTAs always or most of the time while 40 percent always abide the PPCs. The present study also finds that hotels do undercut OTAs. In our dataset, after correction for hotel characteristics, hotels without NBBA price on average 3 percent lower on their own websites compared to OTAs, whereas hotels with NBBA price 1 percent lower on their websites compared to OTAs. For a subsample of out-of-parity (OOP) hotels[4], this price difference is 9 percent for non-NBBA hotels and 4 percent for NBBA hotels. We'll discuss the implication of PPCs, and the fact that they are not perfectly adhered to, for the interpretation of our results in the next section.

Finally, brands with NBBAs that are active on the Dutch market are in many cases franchise formulas. The 131 NBBA-hotels in our dataset together represent only ten brands. Although brands negotiate on NBBAs with OTAs, individual franchisees make their own pricing decisions. We therefore run our analysis of the price effect of NBBAs at the hotel-level, as each hotel of the same brand may price differently. Indeed, we find strong variation in price strategy of hotels within the same brand. Hotels of the same brand differ strongly in the percentage of cases where they abide by the PPC (see section 5).

4. Methodology

In order to assess whether NBBAs lead to higher or lower prices on hotels' websites, we develop a difference-in-differences (DiD) model with NBBA adoption as the treatment. There is no variation in treatments across time in our dataset so we cannot construct a classic DiD model. Instead, we estimate the treatment effect as an increase in price on the hotel website relative to the price on the OTAs in our dataset, controlled for analogical price differences used by the hotels without NBBA. Therefore, instead of the typical difference-in-differences across time and control/treatment groups, we estimate the treatment effect using different dimensions, namely sales channels and control/treatment groups. By using this approach we employ a crucial assumption in testing our hypotheses, namely that NBBAs do not affect OTA pricing. In section 7, where we interpret our results, we provide a detailed discussion of

[3] Previously, major OTAs imposed so-called Wide PPCs (also called Across Platform Parity Agreements) in the Netherlands. Under this clause hotels are obliged to give the lowest price to OTAs compared to any other distribution channel, including other OTAs. Major OTAs narrowed the scope of PPCs to hotel websites in the EU since 1 July 2015, following commitment decisions by the French, Italian, and Swedish competition authorities. In some European countries, narrow PPCs in the hotel sector are prohibited either by courts, competition authorities or due to national sector-specific legislation, such as in Belgium, France, Germany, Italy, and Sweden. Economists have studied both wide and narrow PPCs, and their effects on competition, pricing, and consumer welfare. See e.g. Edelman and Wright (2015), Boik and Corts (2016), Johansen and Vergé (2017), and Hunold et al. (2018).

[4] Out-of-parity is defined as the hotel being more than 1 percent cheaper on the own website for at least 30 percent of the time. From our full sample of 183 hotels, 50 hotels satisfy this definition.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

the plausibility of this assumption.

Our empirical model can be written as follows:

$$\log\left(P_{ijtg}\right) = \alpha_i + \beta_j + \gamma T_{NBBA} + \delta_d + f(t,g) + \varepsilon_{ijtg} \tag{1}$$

where P_{ijtg} is price for hotel i, in distribution channel j, for check-in date t, searched g days ahead of the stay. α_i and β_j are fixed effects for hotels and sales channels, respectively. T_{NBBA} is a dummy variable equal to 1 for observations corresponding to websites of hotels with a NBBA. Parameter γ is the average treatment effect on the treated (ATT).[5] δ_d is a set of fixed effects for each day of the week and $f(t,g)$ is a joint function of check-in date t and number of days g between the search date and the stay. This function controls for hotels' dynamic pricing strategies with respect to search and check-in dates. ε_{ijtg} is the model error term.

To allow $f(t,g)$ to be flexible enough to capture a wide range of interactions between search date and check-in date we estimate the above model semi-parametrically as a generalized additive model (GAM) (Hastie & Tibshirani, 1990; Wood, 2006), with $f(t,g)$ being modelled as a tensor product smooth. A GAM model allows us to correct for the complex time trends regarding search and check-in dates in a flexible way while preserving the linear structure of the model, which aids interpretation. The model is estimated using the library mgcv in statistical software R.

In the following, we denote $p_{ijtg} = \log\left(P_{ijtg}\right)$ for ease of exposition. For channels $j \in \{D, OTA\}$, denoting the direct channel and an OTA, respectively, T_i being equal to 1 for NBBA-suppliers and to 0, otherwise, and X_i denoting observation characteristics being conditioned upon in the model, we can write the ATT estimate as:

$$\hat{\gamma} = \left(\hat{\mathbb{E}}[p_i^D | T_i = 1, X_i] - \hat{\mathbb{E}}[p_i^{OTA} | T_i = 1, X_i]\right) - \left(\hat{\mathbb{E}}[p_i^D | T_i = 0, X_i] - \hat{\mathbb{E}}[p_i^{OTA} | T_i = 0, X_i]\right) \tag{2}$$

where $\hat{\mathbb{E}}$ denotes an empirical average implied by the model. The unbiasedness of the estimate can then be derived as:

$$\begin{aligned} \mathbb{E}[\hat{\gamma}|X] = (\beta_D + \gamma + \mathbb{E}[p_i^{OTA} | T_i = 1, X_i] - \mathbb{E}[p_i^{OTA} | T_i = 1, X_i]) \\ - (\beta_D + \mathbb{E}[p_i^{OTA} | T_i = 0, X_i] - \mathbb{E}[p_i^{OTA} | T_i = 0, X_i]) = \gamma \end{aligned} \tag{3}$$

The above result is crucially dependent on β_D (the difference between the price on the direct channel and an OTA) being constant for both treated and untreated suppliers absent the agreements. This is analogical to the usual parallel trends assumption. Strict exogeneity of the disturbance term ε_{ijtg} is also necessary for the above to hold.

[5] We follow *inter alia* Athey & Imbens (2006) and Heckman & Vytlacil (2005) in interpreting the DiD parameter as an ATT rather than an average treatment effect (ATE). ATT is basically an average difference between the outcome (treated) and the counterfactual (untreated) for the treated units whereas the ATE is an average difference between treated and untreated outcomes/counterfactuals for all units.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

The NBBA hotels, which are large internationally operating brands, may, however, price differently from hotels that were not able to negotiate such agreements with OTAs regardless of the NBBA. Hotels with NBBAs are therefore not necessarily comparable to hotels without NBBAs. As a remedy to this possible source of heterogeneity, researchers typically utilize some version of synthetic control group or propensity score matching methods which ensure that units in treatment and control groups are on average comparable with each other based on some observed characteristics. For an incomplete but useful treatment of the topic see e.g. Abadie et al. (2010, 2015), Doudchenko & Imbens (2016). Athey & Imbens (2017), Athey et al. (2018a), Xu (2017), Hainmueller (2012) and references therein.

As hotel pricing (including β_D) is likely to be heavily influenced by unobserved characteristics of the hotels (location, catering to specific type of customers, management choices) we compare hotels on the basis of their pricing on one of the OTAs (OTA1), (recall we assume the price on this sales channel to be unaffected by NBBAs). We use an approach called trajectory balancing (Hazlett & Xu, 2018; Hazlett, 2018). In what follows, we mimic the steps in Hazlett & Xu (2018) with minor adjustments for our application.

First, we make a standard conditional ignorability assumption implying that the non-treatment outcomes $p_{igt}^{D,0}$ are independent of the treatment status T_i when controlled for the prices on the untreated channel p_i^{OTA}:

ASSUMPTION 1 Conditional ignorability

$$p_{igt}^{D,0} \perp T_i | p_i^{OTA} \tag{4}$$

This assumption basically states that potential differences in unobserved characteristics influencing $p_{igt}^{D,0}$ between treated and control units are accounted for by p_i^{OTA} and the identified treatment effect can thus be attributed to the treatment status.

Next, we assume that the unobserved non-treatment outcome for the NBBA hotels, $p_{igt}^{D,0}$, is linear in a feature expansion of the observations in the non-treated channel p_i^{OTA}: This assumption is the basis for our construction of counterfactuals.

ASSUMPTION 2 Linearity of non-treatment outcome in $\varphi(p_i^{OTA})$

$$\mathbb{E}\left[p_{igt}^{D,0} | p_i^{OTA}\right] = \varphi(p_i^{OTA})'\theta_{gt} \tag{5}$$

where φ is some feature mapping with f-dimensional output and θ_{gt} is a vector of f parameters.

As noted by Hazlett & Xu (2018) a number of popular methods for treatment effect estimation such as

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

DiD and two-way fixed effects models, structural time-series cross-sectional models and interactive fixed effects models implicitly relies on linearity of the non-treatment outcomes in untreated observations. Assumption 2 allows for dependence on a richer set of features of p_i^{OTA}.

In our application, we basically assume that the trajectory of prices on OTA1 captures the unobserved characteristics of the hotels that also determine the pricing on the direct channel.

We aim to ensure that the features implied by φ are on average the same for the treated and control hotels. Achieving mean balance on $\varphi(p_i^{OTA})$ between the control and treatment groups will imply mean balance on $\mathbb{E}[p_{igt}^{D,0}|p_i^{OTA}]$. Comparison of this counterfactual with the observed treated outcomes then yields a consistent estimate of the ATT. To achieve the balance on $\varphi(p_i^{OTA})$ we make the following feasibility assumption:

ASSUMPTION 3 Feasibility of φ-balance
There exists a set of non-negative weights w_i for the control hotels such that $\sum_{T_i=0} w_i = N_0$ and features of the untreated outcomes are balanced between the treatment and reweighted control groups:

$$\frac{1}{N_1}\sum_{T_i=1}\varphi(p_i^{OTA}) = \frac{1}{N_0}\sum_{T_i=0} w_i\varphi(p_i^{OTA}) \tag{6}$$

Where N_0 is a number of control hotels and N_1 is a number of treated (NBBA) hotels.

Seeking balance on the higher order representation of the untreated trajectories $\varphi(p_i^{OTA})$ instead of balancing directly on p_i^{OTA} ensures that control hotels more similar to the NBBA hotels will get larger weight. The approach balancing only means of p_i^{OTA} could lead to very dissimilar hotels being given large weights only because their average matches the average NBBA hotels' prices. This appears crucial for our application as we will be dealing with a highly unbalanced panel and the mean price trajectory will thus be driven not only by hotels' strategies but also by availability of observations for specific hotels at different combinations of check-in and search dates. Our goal is, however, to capture the unobserved characteristics influencing the pricing strategies.

Hazlett & Xu (2018) and Hazlett (2018) propose a kernel-based choice of φ where they construct a Gram matrix K with elements $k_{i,i'} = k(p_i^{OTA}, p_{i'}^{OTA})$ with function $k()$ being a kernel. In general, kernels are functions evaluating similarity between the inputs in a symmetric manner such that $k(a,b) = k(b,a)$. Each row of K is thus a vector of similarities with respect to other hotels. Hazlett (2018) demonstrates that mean balance on rows of K for treated and control subsets implies mean balance on the features given by φ using the Mercer's theorem (Mercer, 1909) which shows that $k_{i,i'}$ can be seen as an inner product of $\varphi(p_i^{OTA})$, $\varphi(p_{i'}^{OTA})$. Hence, we look for weights satisfying the following:

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

$$\frac{1}{N_1}\sum_{T_i=1} K_i = \frac{1}{N_0}\sum_{T_i=0} w_i K_i \tag{7}$$

where K_i is the i-th row of K.

An approximate balance will, however, suffice as an exact balance on all N dimensions of K appears rarely feasible. The entropy balancing algorithm (Hainmueller, 2012) from the kbal package (Hazlett, 2018) is used for minimizing the l_1-norm of the differences between the (weighted) mean vectors corresponding to treated and controlled hotels, respectively.

Many choices of the kernel function $k(p_i^{OTA}, p_{i'}^{OTA})$ are possible. Nevertheless, we need to take into account that we are dealing with a highly unbalanced panel. For each hotel, we observe prices for different combinations of check-in and search dates. We solve this by estimating the following GAM model for each hotel:

$$p_{itg}^{OTA} = \delta_d + f(t,g) + u_{itg} \tag{8}$$

where δ_d denotes intercepts for each day of the week and $f(t,g)$ a tensor product smooth analogical to the one used in (1).

Consequently, we can use the estimates to capture the price trajectory as a function of the covariates:

$$\hat{f}_{p_i^{OTA}}(t,g) = \hat{\delta}_d + \hat{f}(t,g) \tag{9}$$

Next, we compare the estimated trajectories between the two hotels evaluated at check-in and search dates corresponding to the hotel with less observations. Hereby, we obtain vectors of the same lengths and fill in eventual gaps between observations. Evaluating the functions at dates corresponding to the hotels with less observations then ensures that we don't rely on unreliable model projections in regions that are not populated with sufficient observations[6]. We use the Gaussian kernel for the comparison:

$$k\left(\hat{f}_{p_i^{OTA}}, \hat{f}_{p_{i'}^{OTA}}\right) = \exp\left(-\frac{\left\|\hat{f}_{p_i^{OTA}}(X_k) - \hat{f}_{p_{i'}^{OTA}}(X_k)\right\|^2}{b}\right) \tag{10}$$

where X_k is a matrix of covariates for hotel $k = i$ if $n_i^{OTA} < n_{i'}^{OTA}$ and $k = i'$ otherwise with n_i^{OTA} being the number of observations on OTA1 for hotel i. The scale parameter b measures how close the compared

[6] The number of observations for all suppliers increases as the search date approaches the check-in date and in the summer months. Hence, a supplier with less observations will typically generate searches within a region well covered by the observations of the supplier with more observations. This implies we rely mainly on interpolation and do not wish to base our analysis on extensive extrapolation.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

vectors should be in order to be deemed “similar”. As noted by Hazlett (2018), Hainmueller & Hazlett (2014) and Schölkopf & Smola (2002) in practice it is useful to choose b proportional to n_k^{OTA}. We use $\frac{n_k^{OTA}}{100}$ as this choice produces values spread across the unit interval, and consequently, lead to satisfactory balance on K_i and p_i^{OTA}.

The Gaussian kernel satisfies the Mercer condition and the similarity matrix is thus equal to the inner product of the features of $\hat{f}_{p_i^{OTA}}$ and $\hat{f}_{p_{i'}^{OTA}}$ captured by φ. As we do not compare the prices directly but use their trajectories smoothed by the GAM model we formally adjust ASSUMPTION2 to:

ASSUMPTION 2A Linearity of non-treatment outcome in $\varphi\left(\hat{f}_{p_i^{OTA}}\right)$

$$\mathbb{E}\left[p_{igt}^{D,0}\middle|\hat{f}_{p_i^{OTA}}\right] = \varphi\left(\hat{f}_{p_i^{OTA}}\right)'\theta_{gt}$$
$$\mathbb{E}\left[u_{itg}\right] = 0 \tag{11}$$

The $\mathbb{E}\left[u_{itg}\right] = 0$ assumption ensures that the predictions from model (8) are unbiased and hence, the balancing also leads to mean balance on p_i^{OTA}.

The unbalanced nature of our panel also leads us to running a weighted DiD GAM model rather than compare average treated and control trajectories as changes over time can be influenced by a change in pricing as well as by the set of hotels with available price at the given time point. A GAM model allows us to include additional controls such as hotel level fixed effects to correct for the composition effects.

Hence, in addition to Assumption 1, we wish to ensure unbiasedness of the DiD GAM model, i.e. mean independence $\mathbb{E}\left[w_i\varepsilon_{ijtg}\middle|T_i, X_i\right] = 0$, and thus, we assume:

ASSUMPTION 4 Strict exogeneity of $w_i\varepsilon_{ijtg}$

$$w_i\varepsilon_{ijtg} \perp T_i, X_i \tag{12}$$

We run the model (1) with weights obtained from optimizing (7) divided by the number of observations for hotel i, i.e. $\frac{\widehat{w}_i}{n_i}$[7] and obtain an ATT estimate $\hat{\gamma}_{GAM}$. Dividing by n_i ensures that the ATT is estimated as an average across hotels, where under Assumptions 1, 2A, 3 and 4 this estimate is unbiased. We also calculate the weighted DiD estimate without running (1) using only searches for which there are observations available for both OTA1 and the hotel website:

[7] Dividing by the number of observations ensures that each supplier has *ceteris paribus* the same weight in determining the ATT.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

$$\hat{\gamma}_{DiD} = \frac{1}{N^1}\sum_{i:T_i=1}\frac{1}{n_i}\sum_{tg}\left(p^D_{itg} - p^{OTA}_{itg}\right) - \frac{1}{N^0}\sum_{i:T_i=0}\frac{\widehat{w}_i}{n_i}\sum_{tg}\left(p^D_{itg} - p^{OTA}_{itg}\right) \tag{13}$$

This estimator is unbiased under Assumptions 1, 2A and 3[8]. Hazlett & Xu (2018) compare the outcome and the counterfactual directly. In our case, however, we stick to the DiD structure because the difference relative to the OTA channel is crucial in our application and correcting for potential differences in price level enhances the precision of our estimate[9]. Formulating Assumption 2A in terms of the difference $p^D_{itg} - p^{OTA}_{itg}$ leads to the same estimator (13) and follows the approach of Hazlett & Xu (2018) more closely. It makes the use of the GAM model (1) less intuitive, however, and it restricts our ability to gain additional insights from searches not listing prices on both sales channels. Additionally, as noted by Arkhangelsky et al. (2019), methods relying both on balancing the control and treated units and modelling the outcomes typically outperform approaches relying only on one technique. The program evaluation literature refers to a double robustness property where misspecification of only the balancing weights or only the conditional outcome model does not violate consistency of the treatment effect estimate (Athey et al., 2018b; Belloni et al. 2014, Chernozhukov et al., 2018; Hirshberg & Wager, 2018; Imbens & Rubin, 2015; Newey et al., 2004; Scharfstein et al., 1999).

Due to the 2-step nature of the procedure and the obvious within cluster (i.e. hotel) dependence we cannot rely on the default asymptotic GAM standard errors for inference. Therefore, we employ a non-parametric bootstrap and calculate test statistics from the distribution of coefficient estimates based on bootstrapped re-samples (see e.g. Efron, 1979; Horowitz, 2001). To account for the within-hotel correlation, we utilize a block bootstrap by sampling the whole hotel vectors with replacement (Hazlett & Xu, 2018; Cameron et al. 2018). In addition to bootstrapping standard errors, we construct second-order accurate bias-corrected accelerated (BC_α) confidence intervals (Efron, 1987; Efron & Tibshirani, 1986; Diciccio & Efron, 1996) and use these intervals for statistical significance testing. The interval endpoint corresponding to a two-sided test with confidence level α is given by:

$$\hat{\gamma}_{BC_\alpha}\left[\frac{\alpha}{2}\right] = \hat{G}^{-1}\Phi\left(z_0 + \frac{z_0 + z^{\left(\frac{\alpha}{2}\right)}}{1 - a\left(z_0 + z^{\left(\frac{\alpha}{2}\right)}\right)}\right) \tag{14}$$

where $\hat{G}$ is the empirical distribution of the bootstrap estimates, Φ is the standard normal cumulative distribution function, $z^{\left(\frac{\alpha}{2}\right)}$ is the critical point $\Phi^{-1}\left(\frac{\alpha}{2}\right)$[10], z_0 is the bias correction parameter and a the acceleration parameter.

[8] See Appendix A for the proof.
[9] After balancing the price level on the OTA will be approximately equal for both treated and control suppliers. Potential remaining differences may, however, impact the precision of the estimate even though they do not impact unbiasedness.
[10] Hence, for a two-sided test at a 10% significance level we use $z^{\left(\frac{0.1}{2}\right)} = -1.645$ and $z^{\left(1-\frac{0.1}{2}\right)} = 1.645$.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

The parameters z_0 and a are estimated as:

$$\hat{z}_0 = \Phi^{-1}\left(\frac{\#\{\hat{\gamma}_b < \hat{\gamma}\}}{B}\right)$$
$$\hat{a} = \frac{\sum_{b=1}^{B}\left(\hat{\gamma}_b - \overline{\hat{\gamma}_b}\right)^3}{6\left\{\sum_{b=1}^{B}\left(\hat{\gamma}_b - \overline{\hat{\gamma}_b}\right)^2\right\}^{\frac{3}{2}}} \tag{15}$$

where $\overline{\hat{\gamma}_b} = \frac{1}{B}\sum_{b=1}^{B}\hat{\gamma}_b$ and B is the number of bootstrap samples.

Thus, the bias-correction parameter is dependent on the share of bootstrap estimates lower than the sample estimate $\hat{\gamma}$ and the acceleration parameter is equal to the sixth of the skewness of the bootstrap distribution.

5. Data

To estimate the model, we use data from two sources:

i) We use information of two OTAs on the presence (and type) of NBBA between the OTA and hotels, and hotel characteristics, such as star rating, chain affiliation and number of beds for hotels in the Netherlands;

ii) A meta-search site has provided us with data on prices of hotels in the Netherlands on the hotel website and the two OTAs. The meta-search site returned this information to consumers searching on the site for one night stay in a double room in hotels in the Netherlands. The searches were carried out in the period February 18, 2017 – March 31 2018.

The two datasets are merged based on the hotel names. We remove some outliers[11] and only keep hotels with sufficient observations for their website to allow useful comparison. We therefore drop hotels whose website is listed for less than 20% of the searches. We also remove three hotels with too few observations[12] to allow a flexible GAM estimation and two hotels without any available prices for OTA1. The summary statistics for the remaining hotels are presented in Table 1.

Almost all NBBAs are negative matching agreements (123 out of 131). Regarding observed characteristics, the NBBA- and non-NBBA-hotels are on average similar except for the size – NBBA-hotels have on average a larger number of rooms. The NBBA hotels promote their website more often on the metasearch site. The NBBA hotels also seem to abide to the PPC with the OTAs more carefully, undercutting the OTA price only in 15% of the cases compared to 42% of the searches being out-of-

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

[11] This concerns 0.15% of price points deemed likely faulty, i.e. price more than three times higher than hotel average.
[12] Less than 500 searches.

parity (OOP) in case of the non-NBBA hotels.

Table 1: Summary statistics of hotels by NBBA-type

	All hotels		NBBA hotels		No-NBBA hotels	
	Mean	St.dev	Mean	St.dev	Mean	St.dev
Average price OTA1	167.09	92.91	167.34	90.58	166.46	100.04
Average price OTA2[1]	164.80	93.38	164.77	90.19	163.88	102.72
Average price direct channel	167.05	94.41	168.27	91.15	163.97	103.60
Share of searches OOP[2] on OTA1	0.24	0.25	0.17	0.18	0.41	0.32
Share of searches OOP[2] on OTA2[1]	0.22	0.24	0.14	0.14	0.43	0.29
Star rating	3.85	0.67	3.85	0.67	3.87	0.66
Number of rooms	169.92	117.30	188.46	121.75	123.21	90.62
Share of searches without OTA price	0.09	0.09	0.09	0.09	0.08	0.10
Share of searches with direct channel	0.63	0.21	0.70	0.18	0.47	0.19
#hotels	183		131		52	

[1] Based on 180 hotels with room listings on OTA2

[2] Out-of-parity price is defined as the price on the hotel website lower than on the OTA by more than 1% and less than 20% (the latter condition intends to limit the influence of product differences and the former condition intends to limit rounding discrepancies).

The lower OOP rate of NBBA-hotels might suggest that these hotels abide the parity clauses because of the NBBA. However, this seems to be unlikely given that most hotels determine their pricing strategy on lower level than at the brand level. Indeed, we observe significant variation in OOP rate within hotel groups, that have negotiated an NBBA with the OTAs as demonstrated by Figure 1. The boxes span the interquartile range of the OOP rates while the vertical lines indicate the full support of OOP values from minimum to maximum. The horizontal lines denote median values. This variation in OOP rates within hotel groups makes it likely that prices are set at the hotel level rather than at the brand level. The notion of NBBA's negotiated at hotel group level directly determining prices set on hotel level appears less plausible in light of the heterogeneity in pricing strategies within most NBBA hotel groups.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

Figure 1 Boxplot of OOP rates per hotel within a hotel group (NBBA hotels)

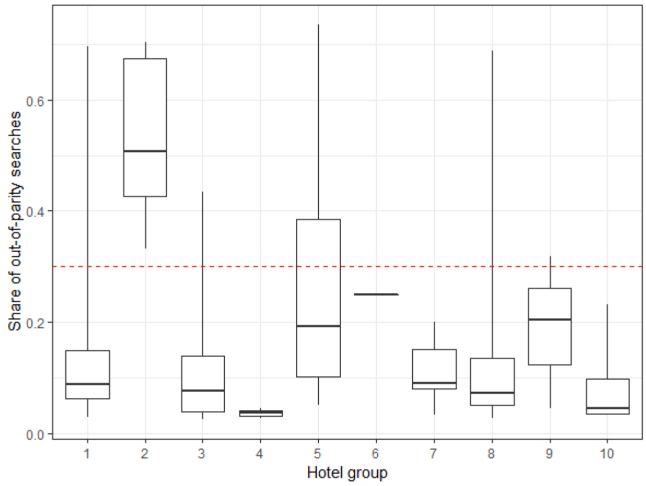

6. Empirical results

In this section we present the results from the estimation of the GAM models weighted using the results from the trajectory balancing procedure described in section 4 and demonstrate the quality of the balance.

As discussed above, the procedure entails constructing a similarity parameter valued between 0 and 1 for each pair of hotels in the dataset. The similarity is assessed using predictions from a GAM model (Equation (9)) for each hotel's prices on OTA1, which we assume to be a sales channel unaffected by NBBAs. Comparing the smoothed price trajectories according to (10) leads to a measure of similarity valued between 0 and 1 for each pair of hotels. Appendix B provides illustration of the constructed similarity measure.

Balancing the parts of the resulting similarity matrix corresponding to NBBA- and non-NBBA-hotels such that the weighted sample of control hotels resembles the composition of the NBBA hotels leads to importance weights that are used for further estimation. To check whether the resulting weights indeed lead to improved trajectory balance we compare the mean kernel vectors for treated and control hotels with and without weights. Recall the matrix K with elements specified by equation (10). The mean kernel vector is then a simple average over the rows corresponding to the group of hotels in question. Figure 2 presents scatter plots of the mean kernel vectors for the control hotels with and without weights against the mean kernel vector for the treated. The two panels correspond to different subsamples of hotels based on their propensity to price out-of-parity on OTA1. In all cases we observe major improvements compared to the unweighted comparison and see the balancing as successful.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

Figure 2 Balancing the mean kernel vectors of the treated and control hotels

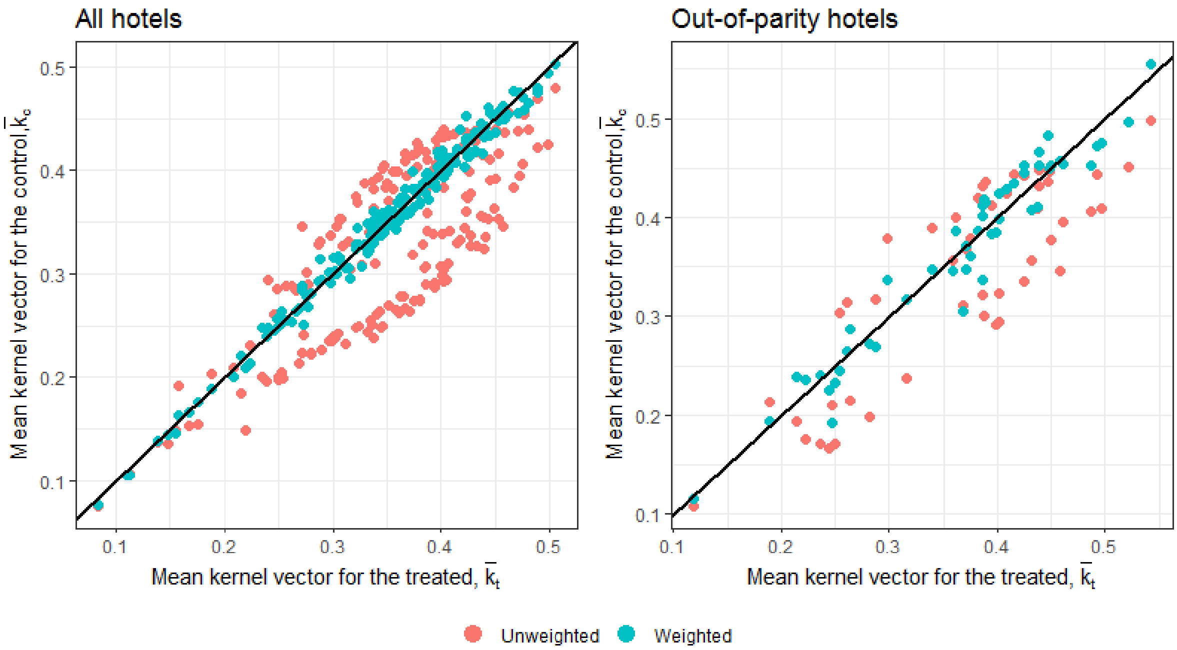

Although we do not balance on the average trajectories, it is interesting to see whether we see improvement also in that regard. Table 2 presents the root mean squared errors (RMSE) between the actual average price trajectories for both groups for the baseline approach without reweighting and also for the trajectory balancing approach. The average trajectories are calculated using a prediction from a GAM model specified similarly to (8). We observe a significant reduction in RMSE in all cases and consider the balancing successful. Furthermore, Appendix C provides graphical illustrations of the similarity in price trends across two dimensions, check-in date and the number of days prior to the check-in date.

Table 2: RMSE between trajectories of NBBA and non-NBBA hotels

	Subset of hotels	
	All	OOP
Average without weights	0.06	0.06
Weighted average	0.01	0.02

Table 3 presents the results from the GAM model (1). In addition to variables and coefficients mentioned in the previous sections, we also include a dummy for searches where the hotel website is the only channel available (i.e. no price is available on either OTA), and we interact this term with the NBBA indicator. The interaction term is interesting because it measures the effect in a situation where there is no intra-brand competition between the direct and OTA channels. In this case the PPC is not a limitation on hotel pricing and so the NBBA is more likely to increase the price on the hotel website. We only present the results from the weighted specification here as we consider the balancing crucial for correct causal inference. The results for the unweighted DiD GAM are reported in Appendix D.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

As discussed before, NBBAs may not lead to higher prices in the context of online hotel booking markets because i) (some) hotels (some of the time) abide by the PPCs with OTAs, and ii) hotels avoid being more expensive than any OTA. However, some hotels do violate the PPC relatively often (out-of-parity hotels ("OOP-hotels")). In column (2) we present results from the specification focusing on OOP-hotels, which violate the PPC at least in 30 percent of the observations. Unsurprisingly, the difference in price between the hotel website and OTA is larger for this subset of hotels. Furthermore, the positive effect of the NBBA on price is greater in magnitude and is statistically significant at lower significance level in case of the OOP hotels.

Table 3: Estimation results: GAM with weights based on trajectory balancing (bootstrapped standard errors in parentheses)

	Subset of hotels	
	All	OOP
	(1)	(2)
direct channel	-0.031***	-0.090***
	(0.015)	(0.023)
OTA2	-0.004	-0.015***
	(0.003)	(0.009)
No OTA listed	0.069**	0.063*
	(0.037)	(0.043)
NBBA effect ($\hat{\gamma}$)	**0.022****	**0.050*****
	(0.015)	**(0.024)**
NBBA x no OTA listed	**0.055**	**0.101****
	(0.039)	**(0.048)**
Hotel FE	YES	YES
Weekday FE	YES	YES
$f(t,g)$	YES	YES
(adjusted) R^2	0.754	0.684
#hotels	183	50
#searches	560 488	129 220
#price quotes	1 266 443	277 929

Statistical significance: $p < 0.1$ *, $p < 0.05$ **, $p < 0.01$ ***
Testing based on BC_{α} confidence intervals
Based on 1000 bootstrap samples

The signs and sizes of all estimated coefficients appear logical and consistent with our theory of harm. They are also roughly similar to the unweighted DiD estimates presented in Appendix B. The control hotels are on average 3% cheaper on their direct channel compared to the OTAs and do not price differentiate significantly between the two OTAs. In case of no availability on the considered OTAs, the prices on the direct channel increase on average by 6-7%. The treatment effects estimates suggest that the NBBA hotels' direct channels are 2% more expensive virtually closing the gap to the OTA price. In case of no OTA listing the price in the direct channel increase by additional 5%. The results in column (2)

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

confirm our intuition that the effects should get stronger for the OOP hotels as the coefficients increase in magnitude.

To check the robustness of the results to potential model misspecification we also present the estimator from the simpler weighted DiD estimator (13) in Appendix E. The magnitudes of the effects and their statistical significance are very similar which gives us confidence in the model and the quality of the balancing as we are able to arrive at comparable estimates even without corrections such as hotel fixed effects and $f(t,g)$. On the other hand, the GAM specification provides us with additional insights regarding pricing on OTA2 (for which we would need to run a separate model otherwise), price level in the direct channel relative to the OTAs and in cases of OTA unavailability.

The significance tests in Table 3 are based on BC_{α} confidence intervals, which are second-order accurate as opposed to t-tests based on bootstrapped standard errors, which rely on stronger distributional assumptions and are only first-order accurate. Figure 3 offers some additional insights from the bootstrap distributions of the treatment effect estimates. The vertical lines depict the two types of confidence intervals and the normal densities are centred around the sample estimate and indicate how realistic the assumptions behind standard t-test are in our case. We can notice that the confidence intervals get wider for the specification with OOP hotels, which is logical as we reduce the number of observations substantially. The estimates shift, however, enough to the right that the estimates retain statistical significance. The bias correction and acceleration makes the confidence intervals shift away from zero. The comparison of the normal densities with the histograms demonstrates the importance of correcting for bias and skewness when constructing confidence intervals.

Figure 3 Bootstrap distribution of the treatment effect estimates

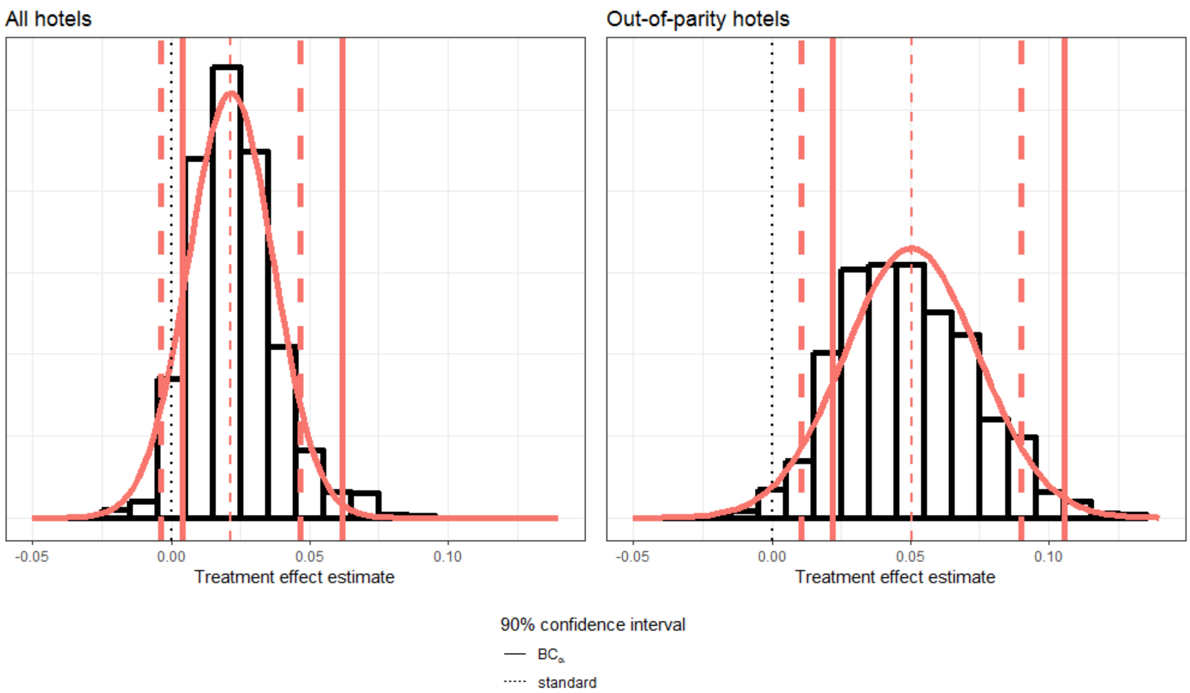

The treatment effect is statistically significant at least at 5% significance level for both specifications and the magnitude of the estimates are similar to the robustness checks in Appendix E.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

7. Interpretation of the estimates

Our empirical analysis consistently shows that NBBAs increase the price hotels set on their own website relative to the price hotels set on OTAs. The dataset we have does not allow us to determine whether or not NBBAs affect the level of the OTA price. We therefore consider this possibility in this subsection in order to be able to draw any conclusions on whether NBBAs are harmful to competition (leading to higher prices) or generate cost savings that are passed on to consumers (leading to lower prices). In the following we argue that NBBAs are unlikely to affect OTA pricing. In light of this assumption, we interpret our empirical results such that NBBAs i) make hotels increase price on their websites, and ii) do not lead to ad spend savings that hotels pass on as lower prices.

NBBAs are unlikely to enable hotels to profitably *increase* their price on the OTA. The reason is that once a customer reaches the OTA, the hotel faces competition from many other hotels on the OTA. This competition is not relaxed by the NBBA. Once a customer reaches the OTA, an individual hotel must still make an effort to make a sale there (by setting a low price, providing good quality, getting good reviews, obtaining a recommendation from the OTA if possible, etc.). In other words, NBBAs may prevent customers using branded queries on search engines from reaching the OTA, but NBBAs do not relax competition on the OTA. At this point one may counter that if NBBAs successfully prevent customers that use branded queries to reach OTAs, the number of OTA users decreases. This may reduce hotels' incentives to compete aggressively on the OTA, leading to higher prices on the OTA. However, given that only 131 hotels have an NBBA with OTAs and OTAs list thousands of hotels, we do not think NBBAs do substantially affect the number of OTA visitors. Moreover, even if NBBAs would significantly reduce the number of OTA visitors, this would affect *all* hotels. Hence the treatment effect we estimate cannot be attributed to this mechanism. Another possibility is that if NBBAs allow hotels to raise price on their website, at some point the hotel price may be higher than the OTA price. This is something most hotels seem to avoid (see the evidence for this discussed later in this paragraph). Therefore, in order to reap the full potential of the NBBA, hotels may increase the OTA price too. We leave this issue aside because our assumption that NBBAs do not increase the OTA price is a conservative one. After all we conclude that NBBAs lead to higher prices on hotel websites. If, contra our assumption, NBBAs do lead to higher OTA prices as well, our approach underestimates the adverse price effects of NBBAs but our conclusion would still be valid.

The possibility that NBBAs make hotels *decrease* their price on OTAs potentially overturns our conclusion. Theoretically, if NBBAs yield cost savings on advertising and hotels choose to pass these on as lower prices, hotels may (partly) do so by lowering the OTA price. Furthermore, if the PPC is a binding restraint on a hotel's pricing, the hotel can only pass on the cost saving to the extent it lowers the OTA price. However, as mentioned before, hotels do not adhere perfectly to PPCs in our sample (which is consistent with other research, see section 3.2). So practically the PPCs do not imply that hotels must reduce OTA prices if they decide to pass on cost savings due to NBBAs. In addition to this, hotels

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

advertise on their brand name on search engines to attract traffic to their own website. If NBBAs reduce these costs, NBBAs thus reduce the cost of sales of the hotel's website. We also note that search advertising on one's own brand name is different from general means of advertising. General advertising is directed towards creating brand awareness, whereas search advertising on one's brand name aims at generating conversions on the hotel's website. This leads us to expect that if hotels pass on any savings on brand bidding, they do so through their own website price rather than through the OTA price.

Next, hotels have a strong preference for having the lowest price on their website because this leads to a greater share of transactions on the hotel website. This yields at least two benefits for hotels: avoiding OTA commissions and having access to more consumer data. In line with this, a study by 10 EU Competition Authorities found that 80 percent of hotels indicate that they don't price their website higher than OTAs. The reasons most frequently given are that hotels "don't want (their) hotel website to be more expensive than OTAs" and "don't want to divert sales away from direct channels" (EU Competition Authorities, 2016, pp. 14, 15). Using this fact, we can at least establish the following conclusion: our results cannot be explained by the efficiency rationale alone. The reason is that even if hotels do pass on part of the savings through the OTA price, they would pass on the same amount or more through the website price. This implies that the price difference between OTAs and hotel websites would increase due to NBBAs. What we find empirically is the opposite: NBBAs decrease the price difference.

Further, hotels differ in the extent to which they price their website lower than the OTA. For a subsample of hotels that are relatively often out-of-parity, we find (not surprisingly) that they give a greater discount on their website compared to the full sample (hotel website is 9 percent cheaper than the OTA, compared to 3 percent for the full sample). For this sub-sample we also find that NBBAs lead to a higher website price relative to the OTA price. Out-of-parity hotels demonstrably care little about abiding by the PPC. So for these hotels we find it implausible they would choose to pass on any cost saving through the OTA price. For out-of-parity hotels we do not only find the qualitatively same result as for the full sample, we also find that the NBBA-effect on the website price is stronger for out-of-parity hotels compared to the full sample (5 percentage points, compared to 2 percentage points for the full sample). This finding is not well explained by the efficiency hypothesis. First, we see no reason why out-of-parity hotels enjoy greater savings and/or would decide to pass on more of the efficiency. Second, for out-of-parity hotels it is even less convincing than for the full sample of hotels that they would pass on the greater part of the efficiency through the OTA price. The theory of harm, by contrast, can explain why we find a greater NBBA-effect for OOP hotels. OOP hotels are cheaper on their website to begin with, and so they can raise their website price more than other hotels before reaching the point where the website price is higher than the OTA price (which hotels are keen to avoid). Hence we infer that NBBAs do raise hotel website prices but do not lead to cost savings that are passed on to consumers through prices.

Lastly, for a small percentage of searches on the meta-search site in our dataset, hotels have no rooms available anymore on the OTA for the date of stay. In this case, the PPC is completely irrelevant as there is no OTA price. We control for this situation in our model and interact it with NBBA-presence. The interaction variable is strongly related to higher prices on the hotel website for the sub-sample of out-of-

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

parity hotels (the interaction variable is of the same sign but smaller in size and not statistically significant for the full sample). This result can only be explained by our theory of harm. The reason is that a hotel having no availability on the OTA for a particular date is unrelated to savings on brand bidding due to NBBAs. When consumers use branded keywords, hotels do not know the intended date of stay of the consumer (because this is not included in the search), implying that hotels cannot make their brand advertising contingent on OTA availability. However, hotels do know when they're sold out on OTAs. In these cases hotels apparently use the protection provided by NBBAs to increase price on their website.

We are aware of an alternative explanation for our finding that NBBAs reduce price differences between OTAs and hotel websites. It can be conceived that hotels are more inclined to adhere to the PPC because of the OTA's willingness to accept an NBBA. Under this *quid pro quo* view, hotels and OTAs come to a mutual understanding that they do not 'steal' each other's customers: OTAs do not target ads to consumers with a revealed preference for some hotel (as indicated by the use of a branded keyword on the search engine), in return for which hotels do not cut out the OTA from transactions by posting lower prices on their websites. Note that even if this is the underlying mechanism, it does not imply anything about how hotels choose to adhere to the PPC (by lowering the OTA price, increasing the website price, or a combination of both). There is a number of arguments that go against this interpretation, however. First, as noted in section 3.2, NBBAs are negotiated at the level of brands whereas prices are set by individual hotels (which is understandable given the fact that hotels adjust price over time to optimize occupancy rates). The ten brands with NBBAs together account for 131 hotels. We find that hotels of the same brand exhibit strong variation in out-of-parity rates. This makes it less plausible that NBBA hotels reciprocate to the NBBA-status by abiding with the PPC. Second, also with an NBBA, hotels are on average price out-of-parity. Third, this hypothesis cannot explain our result, for OOP hotels, that NBBAs have a strong effect on the website price if the hotel has no availability on the OTA, in which case the PPC is irrelevant.

8. Concluding remarks

We consistently find that NBBAs reduce the price differential between OTAs and hotel websites, where hotels price their websites lower than OTAs to begin with. In light of the industry characteristics, we consider it unlikely that NBBAs affect OTA pricing. Therefore we interpret our estimates in such a way that NBBAs in the Dutch hotel sector are i) likely to lead to higher prices on hotels' websites, and ii) do not make hotel pass on possible cost savings on advertising to consumers in the form of lower prices.

Finally, we did not analyse all possible theories of harm nor all possible efficiencies. In particular, our analysis ignores the possibility that NBBAs reduce the match value for consumers by decreasing the chance they find a better match at OTAs. Our analysis also ignores the possible efficiencies arising from a reduction of consumer confusion during search, increased hotel investments in their brands because of a reduction of free-riding by OTAs on brand names, and cost savings that are passed on to consumers in the form of more quality.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

9. References

Abadie, A., Diamond, A., and Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American statistical Association*, 105(490), 493-505.

Abadie, A., Diamond, A., and Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, 59(2), 495-510.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., & Wager, S. (2019). Synthetic difference in differences (No. w25532). National Bureau of Economic Research.

Anderson, S.P. and Renault, R. (1999). Pricing, product diversity, and search costs: a Bertrand-Chamberlin-Diamond model. *The RAND Journal of Economics*, 30(4), 719–735.

Armstrong, M. and Zhou, J. (2011), Paying for prominence. *The Economic Journal*, 121(556), 368-395.

Athey, S., Bayati, M., Doudchenko, N., Imbens, G., and Khosravi, K. (2018a). Matrix completion methods for causal panel data models (No. w25132). National Bureau of Economic Research.

Athey, S., and Imbens, G.W. (2006). Identification and inference in nonlinear difference-in-differences models. *Econometrica*, 74(2), 431-497.

Athey, S., and Imbens, G.W. (2017). The state of applied econometrics: Causality and policy evaluation. *Journal of Economic Perspectives*, 31(2), 3-32.

Athey, S., Imbens, G. W., & Wager, S. (2018b). Approximate residual balancing: debiased inference of average treatment effects in high dimensions. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 80(4), 597-623.

Baye, M.R., De los Santos, B., and Wildenbeest M.R. (2016), What's in a name? Measuring prominence and its effect on organic traffic from search engines. *Information Economics and Policy*, 34, 44-57.

Belloni, A., Chernozhukov, V., & Hansen, C. (2014). High-dimensional methods and inference on structural and treatment effects. Journal of Economic Perspectives, 28(2), 29-50.

Blake, T., Nosko, C., and Tadelis, S. (2015), Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment. *Econometrica*, 83(1), 155-174.

Boik, A. and Corts, K.S. (2016). The effects of platform most-favored-nation clauses on competition and entry. *Journal of Law and Economics*, 59(1), 105-134.

Cameron, A. C., Gelbach, J. B., and Miller, D.L. (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90(3), 414-427.

Centre for Markets Insights (2019), *ShoppingTomorrow – Online Market Places Report,* https://public.tableau.com/profile/cmihva#!/vizhome/ShoppingTomorrow-Online-Marketplaces-Report-2019/Online-Marketplaces-in-the-Netherlands-Experiences-of-Retailers-and-Hotels

Chernozhukov, V., Escanciano, J. C., Ichimura, H., Newey, W. K., & Robins, J. M. (2016). Locally robust semiparametric estimation. arXiv preprint arXiv:1608.00033.

Competition and Markets Authority (2017), *Digital comparison tools market study. Paper E: Competitive landscape and effectiveness of competition*, market study.

Coviello L, Gneezy U., and Goette L. (2017). A large-scale field experiment to evaluate the effectiveness of paid search advertising. CESifo Working Paper Series No. 6684.

Desai, P.S., Shin, W. and Staelin R. (2014). The Company That You Keep: When to Buy a Competitor's Keyword. *Marketing Science*, 33(4): 485-508.

DiCiccio, T. J. and Efron, B. (1996). Bootstrap confidence intervals. *Statistical Science*, 11(3), 189-212.

Doudchenko, N. and Imbens, G. W. (2016). Balancing, regression, difference-in-differences and synthetic control methods: A synthesis (No. w22791). National Bureau of Economic Research.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

European Commission (2018), Case AT.40428 – *Guess*.

Edelman, B. and Wright, J. (2015). Price coherence and excessive intermediation. *The Quarterly Journal of Economics*, 130(3), 1283-1328.

Efron, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics*, 1-26.

Efron, B. (1987). Better bootstrap confidence intervals. *Journal of the American statistical Association*, 82(397), 171-185.

Efron, B. and Tibshirani, R. (1986). Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical science*, 54-75.

EU Competition Authorities (2016). Report on the monitoring exercise carried out in the online supplier booking sector by EU competition authorities in 2016.

Federal Trade Commission (2018). Docket No. 9372 – *1-800 Contacts, Inc.*

Golden J.M. and Horton, J.J. (2018). The effects of search advertising on competitors: an experiment before a merger. Working paper.

Haan M.A. and Moraga-González, J.L. (2011). Advertising for attention in a consumer search model. *The Economic Journal*, 121(552), 552-579.

Hainmueller, J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis*, 20(1), 25-46.

Hainmueller, J., and Hazlett, C. (2014). Kernel regularized least squares: Reducing misspecification bias with a flexible and interpretable machine learning approach. *Political Analysis*, 22(2), 143-168.

Hastie, T. J. and Tibshirani, R. J. (1990). *Generalized additive models, volume 43 of Monographs on Statistics and Applied Probability*.

Hazlett, C. (2018). Kernel balancing: A flexible non-parametric weighting procedure for estimating causal effects. Forthcoming*, Statistica Sinica*.

Hazlett, C. and Xu, Y. (2018). Trajectory Balancing: A General Reweighting Approach to Causal Inference with Time-Series Cross-Sectional Data.

Heckman, J.J. and Vytlacil, E. (2005). Structural equations, treatment effects, and econometric policy evaluation. *Econometrica*, 73(3), 669-738.

Hirshberg, D. A., & Wager, S. (2017). Augmented minimax linear estimation. arXiv preprint arXiv:1712.00038.

Horowitz, J. L. (2001). The bootstrap. In *Handbook of econometrics* (Vol. 5, pp. 3159-3228). Elsevier.

Hunold, M., Kesler, R., Laitenberger, U., and Schlütter, F. (2018). Evaluation of best price clauses in online hotel bookings. *International Journal of Industrial Organization*, 61(x), 542-571.

Imbens, G. W., & Rubin, D. B. (2015). Causal inference in statistics, social, and biomedical sciences. Cambridge University Press.

Johansen, B.O. and Vergé, T. (2017). Platform price parity clauses with direct sales. Working Papers in Economics 01/17, University of Bergen.

Moraga-Gonzalez, J.L., Sándor, Z., and Wildenbeest, M.R. (2017). Prices and heterogeneous search costs. *The RAND Journal of Economics* 48(1), 125-146,

Newey, W. K., Hsieh, F., & Robins, J. M. (2004). Twicing kernels and a small bias property of semiparametric estimators. Econometrica, 72(3), 947-962.

Sayedi, Jerath and Srinivasan (2014). Competitive poaching in sponsored search advertising and its strategic impact on traditional advertising. *Marketing Science*, 33(4), 586-608.

Scharfstein, D. O., Rotnitzky, A., & Robins, J. M. (1999). Adjusting for nonignorable drop-out using semiparametric nonresponse models. Journal of the American Statistical Association, 94(448),

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

1096-1120.

Schölkopf, B., Smola, A. J., & Bach, F. (2002). *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press.

Simonov, A. and Hill, S (2018). Competitive advertising on brand search: traffic stealing, adverse selection and customer confusion. Columbia Business School Research Paper No. 18-59.

Simonov, A., Nosko, C., and Rao, J.M. (2017). Competition and crowd-out for brand keywords in sponsored search. *Marketing Science*

Stahl, D.O. (1989). Oligopolistic pricing with sequential consumer search. *American Economic Review*, 79(4), 700-712.

Ursu, R. (2018). The power of rankings: quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science*, 37(4), 530-552.

Wolinsky, A. (1986). True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics* 101(3), 493–511.

Wood, S.N. (2006). *Generalized additive models: an introduction with R*. Chapman & Hall/CRC.

Xu, Y. (2017). Generalized synthetic control method: Causal inference with interactive fixed effects models. *Political Analysis*, 25(1), 57-76.

Yang, S. and Ghose, A. (2010). Analyzing the relationship between organic and sponsored search advertising: positive, negative, or zero interdependence? *Marketing Science*, 29(4), 602-623.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

Appendix A – Unbiasedness of the weighted DiD estimator

Under assumptions 1, 2A and 3:

$$\mathbb{E}\left[\hat{\gamma}|\hat{f}_{p_i^{OTA}}\right] = \frac{1}{N^1}\sum_{i:T_i=1}\frac{1}{n_i}\sum_{tg}\mathbb{E}\left[p_{itg}^{D} - p_{itg}^{OTA}|T_i = 1, \hat{f}_{p_i^{OTA}}\right]$$
$$-\frac{1}{N^0}\sum_{i:T_i=0}\frac{1}{n_i}\sum_{tg}\widehat{w}_i\mathbb{E}\left[p_{itg}^{D} - p_{itg}^{OTA}|T_i = 0, \hat{f}_{p_i^{OTA}}\right] \quad \text{(A1)}$$

$$\text{[A-3]} \quad = \frac{1}{N^1}\sum_{i:T_i=1}\frac{1}{n_i}\sum_{tg}\gamma_{itg} + \mathbb{E}\left[p_{itg}^{D,0}|T_i = 1, \hat{f}_{p_i^{OTA}}\right] - \frac{1}{N^0}\sum_{i:T_i=0}\frac{1}{n_i}\sum_{tg}\widehat{w}_i\mathbb{E}\left[p_{itg}^{D,0}|T_i = 0, \hat{f}_{p_i^{OTA}}\right] \quad \text{(A2)}$$

$$\text{[A-1]} \quad = \gamma + \frac{1}{N^1}\sum_{i:T_i=1}\frac{1}{n_i}\sum_{tg}\mathbb{E}\left[p_{itg}^{D,0}|\hat{f}_{p_i^{OTA}}\right] - \frac{1}{N^0}\sum_{i:T_i=0}\frac{1}{n_i}\sum_{tg}\widehat{w}_i\mathbb{E}\left[p_{itg}^{D,0}|\hat{f}_{p_i^{OTA}}\right] \quad \text{(A3)}$$

$$\text{[A-2A]} \quad = \gamma + \frac{1}{N^1}\sum_{i:T_i=1}\frac{1}{n_i}\sum_{tg}\varphi\left(\hat{f}_{p_i^{OTA}}\right)'\theta_{gt} - \frac{1}{N^0}\sum_{i:T_i=0}\frac{1}{n_i}\sum_{tg}\widehat{w}_i\varphi\left(\hat{f}_{p_i^{OTA}}\right)'\theta_{gt}$$
$$= \gamma + \frac{1}{N^1}\sum_{i:T_i=1}\varphi\left(\hat{f}_{p_i^{OTA}}\right)'\bar{\theta}_i - \frac{1}{N^0}\sum_{i:T_i=0}\widehat{w}_i\varphi\left(\hat{f}_{p_i^{OTA}}\right)'\bar{\theta}_i \quad \text{(A4)}$$

$$\text{[A-3]} \quad = \gamma \quad \text{(A5)}$$

In step (A2) we split the differences into separate summations and use the fact that balance on the $\varphi\left(\hat{f}_{p_i^B}\right)$ implies balance on p_i^B. Hazlett (2018) shows that kernel balancing ensures balance in means. Together with unbiasedness of $\hat{f}_{p_i^B}$ this is sufficient for the expected means of p_i^B for treated and control suppliers to cancel out.

In step (A3) we write the ATT γ as a weighted average of the individual time-specific treatment effects γ_{itg}. The weighting ensures that we measure ATT as an average across suppliers.

Proof for the GAM estimator using model specification (1) follows the same logic. The difference is that the expectations will be conditioned on the regressors and Assumption 4 is needed to proceed as shown above.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

Appendix B – Similarity illustration

Figure B1 illustrates the smoothing produced by the GAM model (9) for a particular hotel. We can see that the extreme (dark red and dark blue) observations from the raw data in the left panel do not appear in the model prediction in the right panel, but the overall trends in the dynamic pricing strategy of the hotel are captured well.

Figure B1 Smoothing demonstration

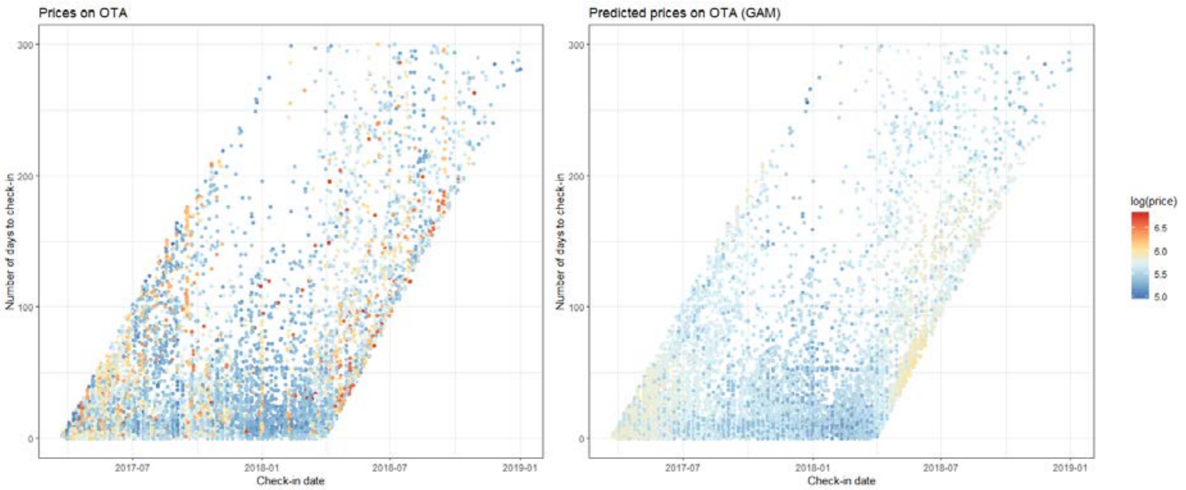

Based on the smoothed trajectories (9) we calculate similarity for all pairs of hotels according to the kernel function (10). Figure B2 illustrates the resulting similarity measure with an example for a large 5 star hotel in Amsterdam (top left panel). The top right panel shows a hotel with a high similarity (k = 0.88). Indeed, we can observe substantial resemblance. The hotels are also similar with respect to the number of stars, size and location, which explains why the dynamic pricing strategies follow the same trends. The hotel in the bottom left panel is located on the other side of the country and is smaller and cheaper overall. We can also notice that the prices do not always rise for the same combinations of check-in and search dates as the hotels in the top row. The moderate difference in overall price level and significant price increase in the summer implying some similarity in relative pricing trends lead to a moderate similarity parameter value of 0.45. Finally, the hotel in the bottom right panel is much cheaper than the three afore-mentioned hotels, 2 stars suggest lower quality and it is located outside large cities. Price increases appear in similar places as in Amsterdam which is understandable, given the geographical distance from Badhoevedorp. Nevertheless, this is true mainly for the summer months and not the rest of the observed period and there appear to be no significant differences between days of the week, which are clearly visible in the other three panels. Therefore, the value of 0.06 suggests very low similarity with the hotel in the top left panel. Overall, a simple eyeball test corresponds to the calculated similarity values.

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

Figure B2 Similarity illustration

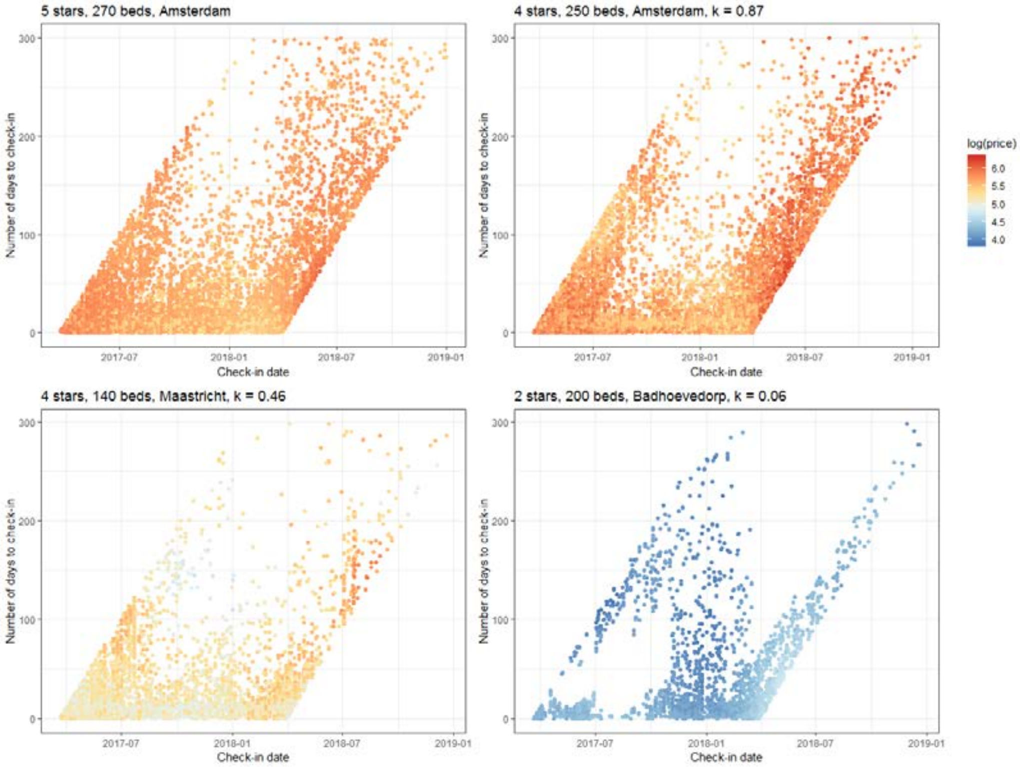

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

Appendix C – Balanced price trajectories

Figure C1 shows graphically the similarity in price trends across two dimensions, check-in date and the number of days prior to the check-in date. The black line depicts the smoothed actual prices of the NBBA hotels on OTA1, the red line represents the unweighted non-NBBA hotels' prices and the green line shows the price trajectories for the weighted non-NBBA hotels. We can notice that in most cases the green line follows the black line rather closely while the red one deviates on multiple occasions. Finally, the right panel shows that the trajectories are less similar for the OOP hotels. Note that this does not mean that the balancing has been unsuccessful in this case as Figure 2 and Table 2 show that the kernel representation of the hotels in the two groups is well balanced. As noted above, the remaining discrepancies are due to hotel composition effects that are more prominent because of the lower number of available observations.

Figure C1 Comparison of price trajectories

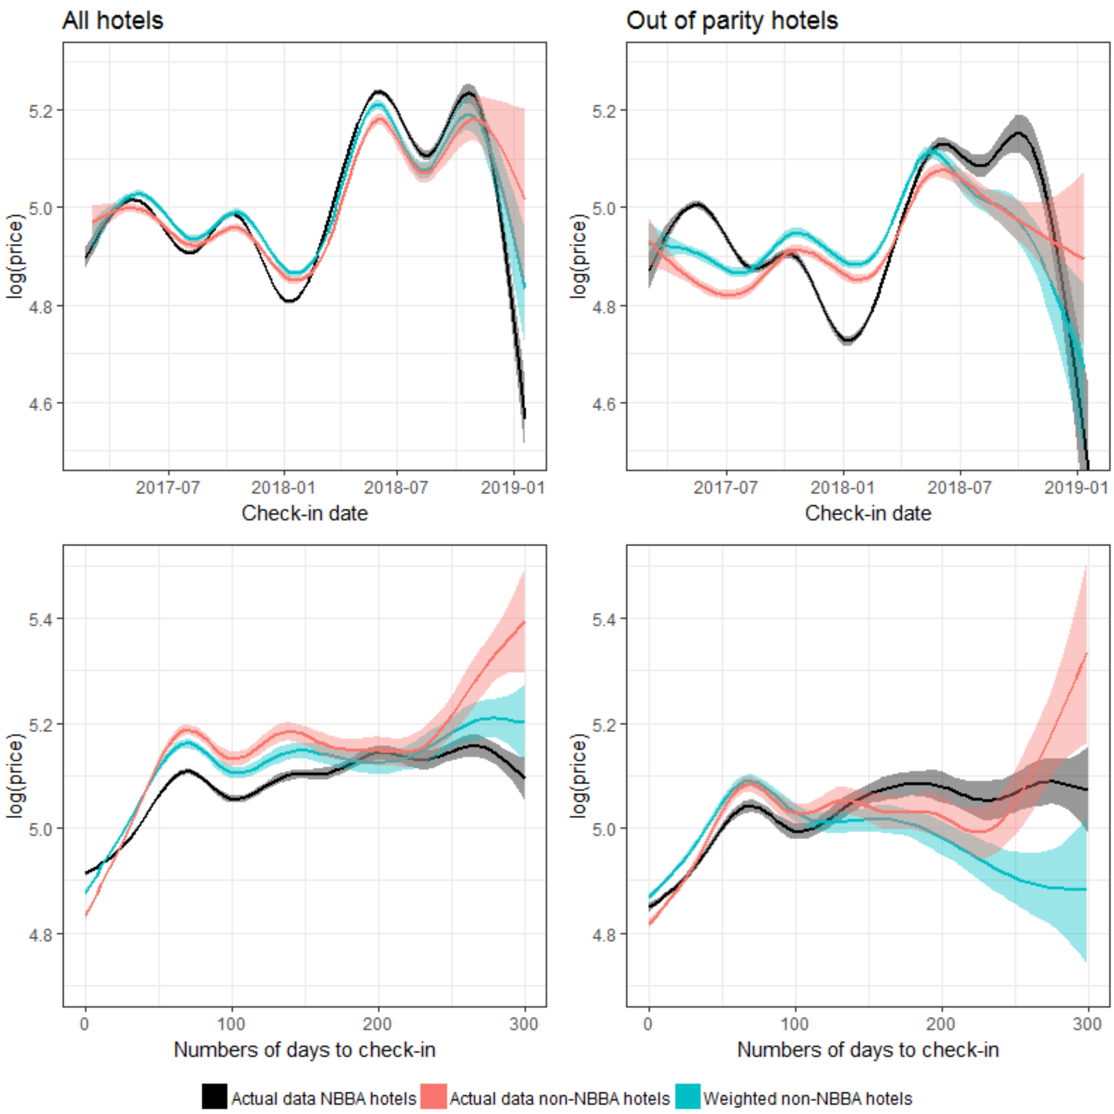

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

Appendix D – Estimation results for unweighted DiD GAM

Table D1: Estimation results: GAM
(bootstrapped 95% t-confidence intervals in parentheses)

	Subset of hotels	
	All	OOP
	(1)	(2)
direct channel	-0.037***	-0.069***
	(0.008)	(0.011)
OTA2	-0.005**	-0.011*
	(0.002)	(0.006)
No OTA listed	0.052***	0.044***
	(0.026)	(0.022)
NBBA effect	**0.028*****	**0.030*****
	(0.008)	**(0.013)**
NBBA x no OTA listed	**0.074****	**0.130*****
	(0.028)	**(0.033)**
Supplier FE	YES	YES
Weekday FE	YES	YES
$f(t,g)$	YES	YES
(adjusted) R^2	0.771	0.730
#suppliers	183	50
#searches	1 266 443	277 929

Statistical significance: p < 0.1 *, p < 0.05 **, p < 0.01 ***
Testing based on BC_{α} confidence intervals
Based on 500 bootstrap samples

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague

Appendix E – Estimation results for simple weighted DiD

Table E1: Estimation results: DiD with weights based on trajectory balancing (bootstrapped standard errors in parentheses)

	Subset of hotels	
	All	OOP
	(1)	(2)
NBBA effect ($\hat{\gamma}$)	0.022*	0.044*
	(0.018)	(0.022)
#hotels	183	50
#searches	292 975	57 027
#price quotes	585 950	114 054

Statistical significance: $p < 0.1$ *, $p < 0.05$ **, $p < 0.01$ ***

Testing based on BC_α confidence intervals

Based on 1000 bootstrap samples

Figure E1: Bootstrap distribution of treatment effect estimates

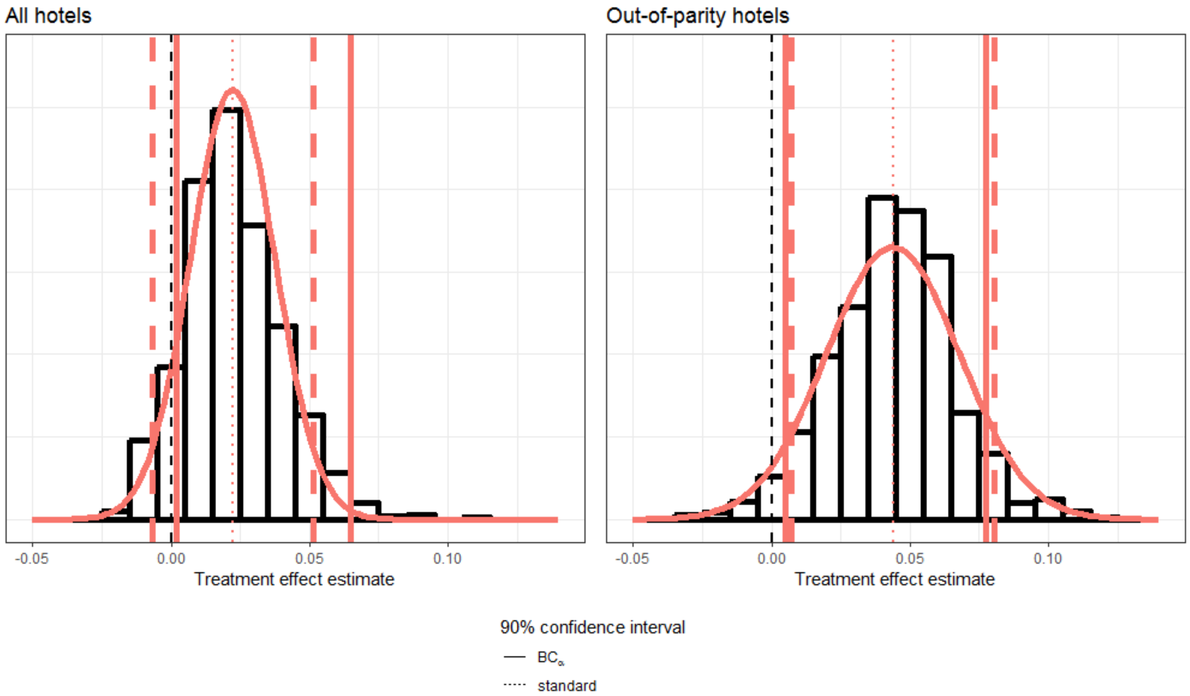

Muzenstraat 41 | 2511 WB The Hague
Postbus 16326 | 2500 BH The Hague