4-2002

Price Dispersion and Differentiation in Online Travel: An Empirical Investigation

Eric. K. Clemons
University of Pennsylvania

Il-Horn Hann

Lorin. M. Hitt
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/oid_papers

Part of the [E-Commerce Commons](#), [Marketing Commons](#), and the [Other Economics Commons](#)

Recommended Citation

Clemons, E. K., Hann, I., & Hitt, L. M. (2002). Price Dispersion and Differentiation in Online Travel: An Empirical Investigation. *Management Science*, 48 (4), 534-549. <http://dx.doi.org/10.1287/mnsc.48.4.534>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/oid_papers/203
For more information, please contact repository@pobox.upenn.edu.

Price Dispersion and Differentiation in Online Travel: An Empirical Investigation

Abstract

Previous research has examined whether price dispersion exists in theoretically highly efficient Internet markets. However, much of the previous work has been focused on industries with low cost and undifferentiated products. In this paper, we examine the presence of price dispersion and product differentiation using data on the airline ticket offerings of online travel agents (OTAs). We find that different OTAs offer tickets with substantially different prices and characteristics when given the same customer request. Some of this variation appears to be due to product differentiation—different OTAs specialize by systematically offering different trade-offs between ticket price and ticket quality (minimizing the number of connections, matching requested departure and return time). However, even after accounting for differences in ticket quality, ticket prices vary by as much as 18% across OTAs. In addition, OTAs return tickets that are strictly inferior to the ticket offered by another OTA for the same request between 2.2% and 28% of the time. Overall, this suggests the presence of both price dispersion and product differentiation in the online travel market.

Keywords

electronic markets, price dispersion, product differentiation

Disciplines

E-Commerce | Marketing | Other Economics

Price Dispersion and Differentiation in On-Line Travel:

An Empirical Investigation

forthcoming, Management Science

Eric K. Clemons¹⁾, Il-Horn Hann²⁾, and Lorin M. Hitt¹⁾

¹⁾Department of Operations and Information Management
The Wharton School, University of Pennsylvania

²⁾Graduate School of Industrial Administration
Carnegie Mellon University

First Draft: April 1998
Current Draft: August, 2001

We would like to thank David Croson, Rachel Croson, Moti Levi, Eli Snir, Ho-Geun Lee, Rebecca Tsui, the participants of the Workshop on Information Systems and Economics, the Department Editor, Associate Editor, and five anonymous referees for helpful comments on earlier drafts of this paper. This paper has benefited from seminars at Carnegie Mellon University, MIT, New York University, Ohio State University, Stanford University, University of California at Berkeley, University of Maryland at College Park, University of Minnesota, University of Southern California, University of Texas at Austin, and University of Washington at Seattle, and the University of Pennsylvania. We would also like to thank Chad Ross for programming assistance. Funding for this work was provided by the Reginald Jones Center Project on Information: Industry Structure and Competitive Strategy.

Price Dispersion and Differentiation in On-Line Travel:

An Empirical Investigation

Abstract

Previous research has examined whether price dispersion exists in theoretically highly efficient Internet markets. However, much of the previous work has been focused on industries with low cost and undifferentiated products. In this paper, we examine the presence of price dispersion and product differentiation using data on the airline ticket offerings of online travel agents (OTAs). We find that different OTAs offer tickets with substantially different prices and characteristics when given the same customer request. Some of this variation appears to be due to product differentiation -- different OTAs specialize by systematically offering different tradeoffs between ticket price and ticket quality (minimizing the number of connections, matching requested departure and return time). However, even after accounting for differences in ticket quality, ticket prices vary by as much as 18% across OTAs. In addition, OTAs return tickets that are strictly inferior to the ticket offered by another OTA for the same request between 2.2% and 28% of the time. Overall, this suggests the presence of both price dispersion and product differentiation in the online travel market.

1. Introduction

Since the beginning of Internet commerce, the impact of the Internet on pricing and competition has been actively debated by both practitioners and academics. Some authors have argued that the increasing availability of price and product information in on-line markets will create highly competitive, “friction free” commerce. For example, Combes and Patel (1997) describe the customer environment for Internet-based travel services:

“...a whole new level of convenience and ubiquity to the shopping experience. Consumers are empowered with the ability to price and compare features with ease. They can inquire about various aspects of a travel destination without having to speak to a travel agent... *or they can quickly and simply find the lowest fare to Las Vegas.*” (italics added)

These types of arguments have also been supported theoretically and by case examples in the context of Internet commerce (Malone, Yates and Benjamin 1987; Bakos 1991, 1997; Clemons and Hitt, 2001) and are consistent with predictions of the economic theory of search as search costs converge to zero (see e.g., Stigler, 1961; Salop and Stiglitz, 1976).

In contrast, there is an emerging stream of research that suggests that there is some observed variation in prices across retailers (“price dispersion”) in Internet markets, even for commodities like books and music compact discs (CDs) (Smith and Brynjolfsson, 2001; Bailey, Brynjolfsson and Smith, 1997). Others have begun to investigate specific types of market imperfections that could lead to these results, such as customer learning (Johnson et. al., 2000), brand loyalty (Chen and Hitt, 2000), or systematic variations in the nature of products offered in online versus regular channels (Lee, 1998). Understanding the presence or absence of exploitable imperfections in Internet markets and their implications for pricing strategy is critical for the long-term viability not only for Internet retailers, but also for firms that must compete in environments with increasingly informed consumers.

In this paper, we contribute to this latter stream of research in several specific ways. First, we consider the presence or absence of price dispersion in a complex product with multiple quality attributes. Most previous work on price dispersion in Internet markets has been done on commodities such as books and CDs. Second, because we consider complex products, we are able to investigate the role of product differentiation, where producers offer different combinations of product features to attract different customer segments, in contributing to price variation.

We choose the online travel agent industry as the setting for study because travel is an important online market and the product is complex, yet fully describable, enabling products of differing qualities to be objectively and accurately compared. Using software agents, we made over 900 ticket requests (drawn from actual travel patterns at a traditional travel agent) identically and simultaneously to five online travel agents – four popular publicly available OTAs and one proprietary system. We then examine the characteristics of the ticket recommendations they provided in response to our requests.

Our data indicate that OTAs systematically provide ticket recommendations with different prices and qualities, and with average prices varying as much as 28% across OTAs for the same ticket requests. When we account for variation in prices due to differences in quality using a hedonic price model, the price variation drops to about 18% across OTAs, suggesting some effects of product differentiation. This is further supported by the observation that OTAs, given the same ticket request, return a ticket unambiguously inferior to one offered by another OTA between 2.2% and 28% of the time, depending on which two OTAs are compared. Our results suggest that this industry is characterized by both product differentiation and random error in product selection.

2. Pricing in Online Travel

2.1. The Online Travel Market

Online travel agents (OTAs) provide a point of contact via the World Wide Web (WWW) to enable customers to search for appropriate flights and fares and make a selection, which is then booked and ticketed by the OTA. There are dozens of OTAs representing on-line travel agents, airlines, traditional travel agents, and computerized reservation systems, although the top two agents accounted for 60% of all traffic in 2001. According to Forrester research, total on-line travel sales are growing 60% per year and are expected to total 12% of the travel market by 2003.

The operational process of an OTA is straightforward. It collects information from the customer, principally departure and arrival cities (airports) and preferred flight times. The OTA then takes this request and some additional parameters set by the OTA and submits these to a computerized reservation system (CRS), which searches for relevant flights from the collection of offerings from all airlines. The agent then takes the collection of flights returned by the CRS, selects one or more flights for presentation to the customer and sorts the final output. When the customer chooses to purchase a ticket, the OTA processes the booking with the CRS and receives a commission from the airline in return.¹ An important feature of this process is that OTAs have to pay the CRS a fee for each request,² but only collect revenue if the customer books a flight. Since only 1-5% of “browsers” actually make a reservation (Machlis, 1997), targeting segments of customers to increase purchase probability is crucial to maximizing profits.

3.2 Application of the Theory of Search and Product Differentiation to OTAs

Competing OTAs have to take into account the search behavior of their customers. It has been argued that greater market transparency in electronic markets would lead to greater price

competition (Bakos, 1991; Benjamin and Wigand, 1995). With vanishing search costs, competing firms offering undifferentiated products have to charge the same competitive price (Bertrand 1883). Undoubtedly, electronic markets on the Internet have made it far easier for consumers to search for services from various OTAs. However, initial evidence (Bailey et al., 1997; Bailey, 1998; Smith and Brynjolfsson, 2001) suggests that this does not hold for near-commodity products like books, CDs, and software. While we cannot determine whether all costs associated with online search have disappeared, we can certainly observe *if market participants behave as if search costs were zero*, by utilizing results from well-established economic models.

In addition, the OTA environment has several unique competitive characteristics. First, OTAs can only select tickets from the (possibly large) available set offered from the airlines – they do not have the ability to alter prices or other product features. OTAs compete for consumers by striving to select the best available tickets according to their preferences, attempting on the consumers' behalf to offset non-competitive pricing by airlines (Borenstein, 1991). However, serving potential customers is not without costs for OTAs. Exhaustive search through computer reservation systems may prove to be too costly, hence, OTAs may choose to specialize and serve a limited number of segments. Any systematic variation in ticket selection would imply a horizontal form of differentiation among OTAs.

Second, because we are considering a good with multiple characteristics and heterogeneous consumer preferences over these characteristics, all comparisons must be made relative to one or more specifications of preferences. While there can be a wide range of consumer preferences for airline tickets, we follow industry practice of considering two specific groups: price-sensitive leisure travelers and time-sensitive business travelers. By offering

multiple choices for a given request, OTAs can attempt to serve multiple customer groups simultaneously.

The central question we analyze in this paper is whether OTAs, given the same request, provide the same ticket recommendations. To the extent that this is not true, we further try to distinguish between two alternative explanations: 1) that differing ticket selection is due to systematic attempts at product differentiation, and 2) that differing ticket selection represents errors on the part of the OTAs. The economic theory of search suggests that when customers are well informed about available prices and face few barriers to searching multiple providers, equilibrium prices will converge to marginal cost, eliminating price dispersion, even in markets where goods are horizontally differentiated (Bakos, 1997).³ In our setting, given that there are few barriers to searching multiple travel sites and the good is sufficiently expensive to justify significant search effort, an analogous prediction is that all OTAs should provide a similar set of “efficient” recommendations. We define *efficiency* here as selecting a ticket that offers the bundle of attributes (price, connections, timeliness, etc.) that maximizes utility for a given consumer segment.

In this paper we focus on testing two hypotheses. First, to the extent that ticket selection is “efficient” in our sense and they share a common pool of tickets to choose from, we would generally expect:

Hypothesis 1 (H1): For a given customer request and specification of consumer preferences, the selected recommendation for each OTA will have the same characteristics (including price and non-price attributes)

To the extent this is not true, we can then explore whether these different recommendations are efficient in the sense of offering maximum utility for different segments. To make this comparison, it is useful to create a single index of utility in which we compute a

“quality-adjusted” price for a ticket following methods used in hedonic price analysis (Griliches, 1961; Chow, 1967). Thus, we also test:

Hypothesis 2 (H2): For a given customer request and specification of consumer preferences, all OTAs will return identical ticket prices after controlling for price variation due to ticket characteristics.

2.3 Relevant Attributes of Airline Tickets

Airline tickets are complex but unambiguously describable. A ticket for a direct flight can be fully described by specifying the airline, departure and arrival airport, departure and arrival time, price, class of service, and fare restrictions. Tickets for connecting flights have additional information related to each segment, but the salient features can be summarized by the number of connections, the total connection time, and the total duration of the trip. The principal concern of our analysis is isolating the relevant ticket attributes that can vary across the selections of various OTAs. Of the list of possible attributes, some are held constant across OTAs by study design, others are fixed due to technological constraints, and some can vary based on design choices by the OTA. In our ticket requests to the OTAs we specify departure and arrival airport (“citypair”), service class (all tickets are coach class), and travel dates and times, eliminating variations across OTAs from these sources. Technological limitations of the OTAs or legal restrictions on their behavior eliminate variations due to fare restrictions and airline.⁴ Our software agents make simultaneous reservations on all OTAs, eliminating variations in fare structure over time or time to departure. Thus, of the universe of ticket characteristics described above, the principal ticket characteristics that can vary by OTA in our study are:

- Timeliness: an inquiry for a specific flight includes the desired departure time for both legs of the flight. An OTA is not bound to report only flights that meet the time requirements specified by the traveler, but can also select other flights, which might be less expensive.

- Number of connections: connections often permit more options to be considered, which may make a flight less expensive, and may sometimes be necessary to ensure the timeliness of the flight.
- Length of connections: an OTA has to make the decision about the duration of a connection that is acceptable for the customer, given his or her priorities for time and price.
- Flight Duration: although principally a function of connections and connection time, an OTA can choose a more complex, longer routing in an attempt to save cost.
- Price: OTAs can select tickets with different characteristics even on the same routing which leads to differences in fares.

This list represents the exhaustive list of dimensions upon which tickets can vary by OTA and will form the basis for a set of comparisons that attempts to determine whether OTAs offer optimal or inferior flight offerings when provided the same request. It should also be noted that this list includes all available preferences that can currently be expressed by the consumer to the OTA as well as preferences that the OTAs can express to the CRS to guide ticket selection. We therefore believe that we are able to either control for or measure all the fundamental quality attributes that consumers are currently able to use to describe airline tickets.

3. Data and Methodology

3.1. Data Sources and Collection Methods

Our analysis begins with a data set of actual tickets written by a major corporate travel agent in the U.S. for five corporate clients in the month of April 1997.⁵ By using a base of actual tickets for the analysis we are able to make comparisons among OTAs under a scenario that closely matches how they would be used by a group of travelers. To ensure that all systems could process these tickets, we applied a number of screening rules to eliminate problematic tickets.⁶ The five online travel agents to be studied are selected based on recommendation of market participants and represent relatively large OTAs.⁷ Three of the five OTAs provide users

with the option to indicate a preference for either price or time. We treat each OTA that offer the ability to choose between time preference and price preference as two separate OTAs, which results in our studying eight OTAs.

Using this screened sample of tickets, the reservations are then replicated on each of the OTAs using intelligent agents. This ensures high data quality and allows simultaneous data collection across all agents, virtually eliminating the chance of fare changes influencing the results. The set of flight alternatives offered by OTAs are stored in a database. Based on the input of market participants, we apply two different decision rules to select a single flight from the set of alternative flights offered. In our first decision rule, we emulate the preferences of business travelers by making timeliness our top priority. Of all flights offered, we select those flights that depart in the time window of one hour before or after the specified departure time for the departure and return flight (these guidelines are those typically used by travel agents in determining timeliness). A ticket that meets the time window constraint in one leg is preferred to a ticket that does not meet the time window constraint at all. If a ticket meets the time window in only one leg, the tickets where the time window is met in the departure leg was given preference. In case of a tie, the cheaper ticket is given preference. The second decision rule reflects the consumer preferences of price sensitive leisure travelers. Hence, price is the top priority, with timeliness (meeting the time window) as the tiebreaker. For discussion purposes, we label these the *time-priority* and *price-priority* data sets. We also collect other data, such as the difference of the desired and actual departure time as well as the length of the connections.

We ran our intelligent agents for four consecutive days for 24 hours. Each day, our set of intelligent agents made reservations for flights that were originally made on that weekday with the same number of days to departure, seeking to replicate travel attributes such as trip duration

and departure times. On average, each of our agents made between 300 and 500 reservation requests per day. The total number of reservation requests that the intelligent software agents made depended highly on the availability of the OTAs⁸ and the response time of the slowest OTA. In order to obtain a consistent data set, we discard all reservations where one or more OTA made no recommendation. This reduces our data set to 939 unique tickets for which we have reservation recommendations from all eight OTAs.

3.2. Data Characteristics

Table 1 shows some basic statistics of our sample. Overall, we have a total of 7512 tickets with an average price of \$557 in the time priority data set and \$515 in the price priority data set.⁹ The tickets otherwise show similar characteristics between the two data sets except for the percentage of flights meeting the time window. The price priority data set contains more than twice as many tickets that violate the time window. There is large variation across the sample in prices, primarily due to route differences. The standard deviation of price is nearly \$400 and the tickets range in price from \$81 to \$2118.

4. Empirical Analysis and Hypothesis Testing

4.1. Model Formulation

The key challenge in empirically testing our predictions is properly modeling the relationship between ticket quality and price so that we can create a single index of ticket utility. The critical dimensions that affect ticket prices are the route (embodying characteristics such as distance, competition, and demand), timing (especially “Saturday night stay”, advance purchases, and constraints on departure and arrival time) and the characteristics of connections (including number and length, as well as overall flight duration). These characteristics span most

of the relevant characteristics and are consistent with our earlier discussion of factors the OTA can control, as well as all the features that the customer can request. To this base model of ticket prices we include dummy variables for each OTA – the coefficients on these dummy variables represent the average quality-adjusted price of the OTAs ticket selections, which can be interpreted as average utility levels.

There is substantial debate in the literature on hedonic price models as to the functional form of the price-characteristic relationship. For our context, we adopt the log-linear form, which has been extensively used in the airline economics literature (Borenstein 1989; 1991; Evans and Kessides 1993). In this formulation, characteristics lead to a percentage increase in the base price rather than an absolute increment. This appears more consistent with actual pricing behavior in the market (e.g. permitting a connection on a \$1000 flight could reduce the price by \$200, while permitting a connection to a \$150 flight may only save \$30). In addition, the combination of city-pair fixed effects and the log-linear model virtually eliminates the heteroskedasticity that is present due to the fact that higher price tickets have more price variation (the remaining heteroskedasticity is not economically meaningful and is addressed by the use of White Robust standard errors).¹⁰ Thus, the general model we estimate is:

$$\begin{aligned} \log(p_{OTA}) = & \sum_j \beta_j OTA_j + \delta_1 * TimeWindow_{departure} + \delta_2 * TimeWindow_{return} + \\ & + \delta_3 * Connection_{departure} + \delta_4 * Connection_{return} + \delta_5 * DD7 + \delta_6 * DD14 \\ & + \delta_7 * DD21 + \delta_8 * DD28 + \delta_9 * Saturday + \sum_i \gamma_i * Citypair_i + \varepsilon \end{aligned} \quad (1)$$

For our main formulation, all variables are binary. $TimeWindow_{departure}$ and $TimeWindow_{return}$ are 1 if the time window is met on departure and return respectively, and zero otherwise. $Connection_{departure}$ and $Connection_{return}$ are 1 if there is a connection on departure and return respectively, and zero otherwise. $Saturday$ is 1 if there is a Saturday night stay and zero

otherwise. The variables $Citypair_i$ represent a dummy variable for each of the 436 combinations of departure and return cities in our dataset (we estimate this as a fixed effect model). The variables of the form $DDnn$, represent a dummy variable capturing the number of days to departure (with nn representing the minimum number of days before departure in that bucket – for example, $DD7$ is 1 if days to departure is greater than 7 but less than 14). We also explore models in which we examine deviation from desired time window and length of connection (both measured in minutes) with virtually identical results.¹¹ We thus conduct all the analysis using the simpler model because it is more easily interpreted and compared with our frontier analysis in Section 4.5.

In interpreting the coefficients of Equation 1, it is important to consider the interaction between the quality characteristics and the city-pair characteristics. In this model, the quality characteristics do not represent the pure difference due to connections or time, but the differential when the number of connections is above the norm for that route. For example, if a route always has connections, then there will be no premium attributed to a connection on that route. In essence this formulation accurately removes variation in ticket prices due to these characteristics, which is our primary objective. However, these coefficients should not be interpreted as the simple price premium or discount due to these quality characteristics in isolation, although they usually are in the same general direction as the unconditional relationship.

4.2. Baseline Comparison

Our analysis begins with a relatively simple log-linear model that relates the log of the ticket price to the identity of the OTA and the city pair but omitting the quality variables for connection and timeliness, which is the appropriate model for investigating price in Hypothesis

1. We drop the dummy variable for OTA1 (on average, the lowest price OTA) to avoid multicollinearity with the intercept term so all coefficients can be interpreted as differences in average price compared to OTA1. Therefore, if H1 were true, we would expect

$\beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = \beta_7 = \beta_8 = 0$. In other words, on average, the OTAs should provide tickets with the same prices.

The first column of Table 2 depicts the coefficients for the time-priority data set, while the second column depicts the coefficients for the price-priority data set. In both models, we can clearly reject hypothesis H1, that prices are equal across OTAs (for time-priority data set, $F_{7, 7069} = 65.50$, $p < .0001$; for price-priority data set, $F_{7, 7069} = 79.31$, $p < .0001$).

OTA4 shows the greatest price variation from OTA1, with $\beta_4 = 27.8\%$ for the time-priority data set and $\beta_4 = 25.9\%$ for the price-priority data set. All coefficients are significant at a $p < .001$ level. Interestingly, OTA2 is the only OTA whose deviation from the average price noticeably changes when applying a different decision criterion. For the time-priority data set, OTA2 is about 17.2% more expensive than OTA1, while it is only 5.2% more expensive for the price-priority data set. This suggests that OTA2 is the only agent that returns substantially different offerings to appeal separately to time vs. price sensitive travelers. Moreover, we actually find that preference specification (price, time) on systems that allow this (OTA2, OTA3, and OTA4), actually has no effect on quoted prices – p-values for tests of equality across systems are all above .85, and there is no significant difference in the orderings of tickets offered.¹²

When coding the intelligent agents, we discovered that OTA2 and OTA3 appear to share a common search engine. This suggests that they could offer similar tickets. To verify this in our data, we tested whether OTA2 and OTA3 quoted identical prices (the actual test is $OTA2_{time} = OTA2_{price} = OTA3_{time} = OTA3_{price}$). The results suggest that these two systems and their

time/price options are virtually identical ($F_{3,7069}=.051$, $p>.98$ for the time priority data set; $F_{3,7504}=.002$, $p>.99$ for the price priority data set).

Because the time/price options and OTA2/OTA3 data appear to be redundant and including them in the analysis would only increase the power of our statistical results (perhaps artificially), we take a conservative approach and delete the redundant systems from the analysis. We then repeat the baseline analysis (Table 2, columns 3 and 4) on the four remaining OTAs and find similar results – a price dispersion of approximately 28% between the highest and lowest price OTA for both decision criteria.

4.3. Variations in Characteristics of Selected Tickets

Tickets in our sample have four key characteristics indicative of “quality”: meeting time window requirements on departure, meeting time window requirement on return, having no connection on departure, and having no connections on return. This yields 16 possible types of tickets. In Tables 3a and 3b, we count the number of ticket recommendations from each OTA of each quality. For example, the first row of Table 3a indicates that OTA2 quoted 5 ticket recommendations that had no connections but failed to meet the time window requirements on both departure and return. OTA1 quoted 54 ticket recommendations, OTA4 quoted 14 ticket recommendations and OTA5 quoted only 1 ticket recommendation with these characteristics. To examine whether there are systematic variations across OTAs in the type of tickets they issue, we employ the simple sign tests and the Wilcoxon signed rank tests to evaluate whether the distributions across ticket types are the same for all OTAs. This is a test of the latter part of hypothesis H1. Moreover, given our earlier results, we can check to see if the systems with higher prices also tend to yield tickets with higher quality, which would be indicative of differentiation.

Our results (Tables 4a and 4b) suggest that in both data sets, OTA5 and OTA4, the highest price OTAs, are quite similar ($p > .10$ for sign and Wilcoxon test for time and price priority data set) and that they are different from OTA1 (at least $p < .05$ for the sign test for both data sets and at least $p < .10$ for the Wilcoxon test for both data sets) which offers the lowest quality tickets. In some cases we can also distinguish OTA5 and OTA4 from OTA2. We also find that OTA1 and OTA2 are similar in the price priority data set but not in the time priority data set ($p > .10$ for the price priority data set and $p < .10$ for the time priority data set). Overall this is consistent with the idea that variation in prices is matched to variation in ticket quality. Moreover, the results appear to be fairly robust to alternative priority criteria (price/time) and statistical tests (Wilcoxon/simple sign test).

4.4. Testing for Price Variation Across OTAs Accounting for Ticket Quality

The previous results suggest that different OTAs appear to be targeting different price-quality segments among consumers. In this section we estimate the full hedonic price model including the terms for time window, connections, and Saturday night stay. The specific model is shown in Equation 1.

The results of the full model are shown in Table 2 columns 5 and 6. First, we should note that the signs on the quality control variables are consistent with prior expectations: refusing to accept flights outside the time window makes travel between 7.9% and 11.6% more expensive than accepting them; refusing to accept a connection can increase prices up to 17% on a route that normally has connections. Saturday night stays decrease the ticket price by up to 30.9%. Similar results hold if we do not include the OTA dummy variables and just estimate the hedonic model alone.¹³ Overall, this provides some confidence that the model is directionally correct.

Using the reduced set of OTAs as in the previous analysis, we are able to reject equality of the OTA effects in all models at $p < .0001$ level, suggesting that quality variation is not the only source of price variation in the sample. However, the estimates show that price variation across OTAs is reduced; for the time and price regressions, the price variation are 18.1% and 19% respectively as opposed to 27.8% for the base regression model. This suggests that product differentiation accounts for some of the previously observed price variation. To gauge the robustness of this result, we estimate an even more general model that includes interactions between the quality attributes and the city-pair variable – this allows the quality effects to vary by route (regression results not shown). Even in this model, price dispersion is still evident. For the time-priority data set, OTA2_{time} was 9.9% more expensive than OTA1, OTA4_{time} and OTA5 were 16.3% and 13.4% more expensive than OTA1. Similar results are found for the price priority dataset.

Overall, this suggests that product differentiation accounts for at least 10% of the variation in actual ticket price between OTAs. This represents a variation of about \$50 in the price of an average ticket, or about 35% of the overall price dispersion across OTAs. We cannot conclude for certain that these estimates represent the true contribution of product differentiation since this calculation requires that there is no specification error in the hedonic model that is correlated with a specific OTA. However, it does suggest that there is a significant contribution of product differentiation to overall price dispersion across OTAs as well as some price dispersion. In the next section, we consider an alternative approach to demonstrate the presence of the price dispersion that does not rely on the accuracy of the hedonic model.

4.5 Frontier Analysis

An alternative way of gauging inefficiency of OTAs is to compare the best ticket offerings for each request to see the extent to which some OTAs offer tickets that are demonstrably inferior to a selection provided by another OTA. While this relies on having a proper characterization of the underlying dimensions of quality, it is not subject to issues of functional form or specification error as in the earlier analysis. To implement this analysis, for each ticket request for each pair of OTAs, we categorize the relationships as either: 1) identical – the OTAs give the same ticket, 2) one OTA strictly dominates another – an OTA provides a ticket which was at least as good on all characteristics (time window, connection, price) and strictly better on at least one, or 3) non-comparable – one OTA gives a ticket that was superior in one dimension, while the other returns a ticket superior on another dimension. We focus particularly on dominated tickets since this is a clear example of inefficiency and independent of the decision rule (if a ticket is dominated, it is necessarily dominated for both decision rules).

An analysis of dominated tickets for each pair of OTAs is presented in Tables 5a and 5b. To interpret the results, consider the relationship between OTA2 and OTA1. OTA2_{time} strictly dominates OTA1 82 times, but OTA1 strictly dominates OTA2_{time} 122 times. When comparing the strict dominance of all OTA pairs, we observe that OTA1 is least often strictly dominated; this may be attributable to their strategy of finding low ticket prices making it unlikely that they are dominated on price. However, we cannot find any OTA pair for which we can clearly state that one clearly dominates the other.¹⁴

While this does not provide an indication of the presence of differentiation, it does strongly suggest the presence of inefficiencies. Moreover, since all OTAs offer a significant number of dominated tickets, it also suggests a component of randomness in the error. However,

systems that aggressively pursue high quality tickets tend to make these “mistakes” with greater frequency.

5. Discussion

Overall, we find that different OTAs offer different types of tickets and that they do so at substantially different prices. Our analysis suggests the presence of horizontal product differentiation in the sense that different providers offer tickets of systematically different quality and that price variation across OTAs in the hedonic model is reduced when we include controls for quality. Another way to view this result is to plot the average ticket price and quality (along different dimensions) for each OTA. The plot for price satisfaction (ratio of the prices of each OTA to OTA1) vs. connection satisfaction (percentage of flights without connections) is shown in Figure 1a and a similar plot for price versus timeliness is shown in Figure 1b. The graphs clearly show that OTAs tend to “specialize” on particular ticket characteristics – OTA1 is clearly the price leader with OTA4 and OTA5 focusing on connections and timeliness, respectively. We also note that no single OTA is clearly inferior in both analyses.

These graphs are consistent with our regression and frontier analyses that suggest that the variance in ticket prices across OTAs is partially due to quality variation and partially due to inefficiency. While one could generally rationalize a finding of product differentiation with the existence of entry barriers in an otherwise efficient market, it is harder to rationalize why firms offer inefficient recommendations. In the remainder of this section, we explore three possible explanations: technological constraints related to implementing differentiation, agency problems, and attempts at price discrimination through OTA design.

5.1. Technological Constraints

OTAs achieve product differentiation through parameterization of their requests to the CRS. Each OTA uses a specific set of decision rules to identify which tickets will or will not be considered in the search process – altering these criteria causes different tickets to be offered. Interviews with a system designer suggest that the key parameters relate to: permitted deviation from requested departure or return time, minimum savings required to justify a connection, maximum duration of a connection, and minimum distance before a connection is considered. While we cannot observe directly the choices different OTAs make on these dimensions, we attempt to infer the choices on these parameters through their selection behavior. To accomplish this examination, we extend the model in Equation (1) to include interaction terms between the quality attributes (connections, timeliness) and the OTA (we center the interaction terms so that the coefficients on other terms are comparable to our previous regressions). Intuitively, this is a crude measure of the “shadow cost” of deviations from ideal ticket characteristics imposed by the parameter settings chosen by the OTA.¹⁵ To reduce the number of coefficient estimates, we aggregate the *TimeWindow* and *Connection* variables to represent the meeting of both time windows (departure and return) or having any connections. The results are shown in Table 6.

Overall, there is a wide variation in the effects of connections and satisfaction of time window requirements across OTAs. For the time-priority data set, OTA4_{time}'s ticket prices decrease by only 4.4% when a connection is involved (not significant)¹⁶, while all other OTAs reduce the prices by 10.9% to 21.3% (all estimates are significant at the $p < .001$ level). OTA5's ticket prices are actually 11.7% cheaper when they leave on time than when they leave outside of the specified time window. OTA4_{time}'s prices increase by 6.9% (significant at $p < .01$ level) for flights that leave on time, while OTA2 time and OTA1 have an increase of 9.8% and 22.4%, respectively (all significant at $p < .0001$ level). Similar results are found for the price-priority data

set. The high variation in the coefficients for the interaction terms of OTA and ticket characteristics reflects the differences in the set of parameters that an OTA has chosen. For example, OTA5 requires the highest saving to justify a connection and OTA1 allows for the highest deviation from the specified time of departure. To the extent that these settings are not ideal for all flights, this may explain at least some of the price dispersion we observe.

5.2. Agency Problems

Given that agents are representatives of the airlines and not the consumer, it is possible that agents offer higher price tickets in an effort to increase their own revenue (or the revenue of their owners for those OTAs affiliated with airlines). While a fully efficient market would cause consumers to avoid agents that offer unnecessarily expensive tickets, limited search, advertising, brand loyalty, or other market imperfections may allow this behavior. To investigate this possibility we examined the various commission structures in place for OTAs at the time of our study. While there is a substantial variation across different airlines in commission structure, making such a comparison difficult, one OTA (OTA5) performs their services entirely on a flat fee basis for their customers and rebates all commissions. Thus, of all the systems, they should have no incentive to offer tickets with inflated prices. If agency problems were the only concern, then we would expect OTA5 to have lower ticket prices on average. However, our data suggest that OTA5 offers tickets with the second highest average price, even after controlling for quality. While this does not completely rule out possible agency effects, this provides limited evidence against the hypothesis that agency problems lead to the price dispersion we observe.

5.3. Price Discrimination

The strategy of the two OTAs that are operated by the same company, namely OTA4 and OTA1, is interesting. In this case, one company offers two different on-line services with

different interface characteristics and very different prices. OTA1 offers the cheapest tickets with lowest convenience and OTA4 offers the most expensive tickets with highest convenience. Simultaneously, the front-end for OTA1 can be best described as archaic, while the front-end for OTA4 is state-of-the-art. Since one would generally believe a better interface is preferable for the consumer, a possible explanation for this behavior is that the traveler's willingness to pay is correlated with their valuation of interface quality. For time-sensitive travelers it can certainly be argued that they do not have the patience to work through the unintuitive mainframe menu of OTA1, while price-sensitive travelers may endure the procedure. In other words, the difficulty in using OTA1's user interface serves as a screen to prevent the time-sensitive travelers from exercising personal arbitrage. Further, a traveler who is willing to learn OTA1's user interface can be assumed to be computer-savvy and therefore more likely to be knowledgeable about alternative OTAs and their interfaces. This type of behavior is widespread and has a long history. Perhaps the most famous example is discussed by Dupuit (1849) on price discrimination in the French railroad system:

It is not because of the few thousand francs which would have to be spent to put a roof over the third-class seats that some company or other has open carriages with wooden benches... What the company is trying to do is prevent the passengers who can pay the second-class fare from traveling third-class; it hits the poor, not because it wants to hurt them, but to frighten the rich... And it is again for the same reason that the companies, having proved almost cruel to third-class passengers and mean to second-class ones, become lavish in dealing with first-class passengers. Having refused the poor what is necessary, they give the rich what is superfluous (p. 23).

In the OTA context, whether this is a deliberate strategic choice or the result of historical accident is not clear. OTA1 is originally developed as proprietary software and is designed to connect to a mainframe. OTA4 is built with the interactivity in mind that the World Wide Web provides. This results in two systems that couple the ticket selection behavior with customer segmentation through interface design, consistent with a price discrimination explanation.

5.4. Conclusion

Our results suggest that even in a market with a potentially low degree of search costs and strong incentives for consumer search, there exists persistent price dispersion across service providers. One possible explanation is that search costs prevalent in traditional markets are replaced by new types of costs. For example, OTAs may create switching costs by requiring the customer to "sign-up" by entering personal information that later reduces the time to find and book flights. When non-zero switching and search costs are combined with uncertainty about which OTAs will truly provide the best flight for the consumers' preference, it may be optimal for consumers to limit their search to a single or to a few OTAs. If this is the case, then consumers would be unambiguously better off using a "Super OTA" that searches all available OTAs and returns their flight recommendations. One system indeed saw this opportunity and offered to search four different OTAs for the best fare. Of those four OTAs, two OTAs – OTA2 and OTA3 – are in our sample. The "Super OTA" itself was operated by the same company that operated OTA2 and OTA3. But even before this system was upgraded from a limited beta version to a full version, the two other systems were withdrawn, claiming they already offered the best prices. The two remaining systems offer, according to our analysis, the same flight recommendations. More recently, some Internet firms provide this service on an independent basis and appear to generate improved recommendations. This suggests that our findings of price dispersion are robust to a substantial amount of innovation in price search capability that has occurred in the two years following our data collection. That is, significant price dispersion appears to continue to exist.

Our results also show that service differentiation is a key strategic component of online sellers that offer access to heterogeneous goods. While this may appear unusual for markets that

should theoretically have greater information transparency, it mirrors behavior in non-electronic markets: by exploiting consumers' heterogeneity in tastes and uncertainty of vendor quality, vendors can ease price competition by segmenting the market. Moreover, with a product with little cost of differentiation, new strategies may emerge; for example, creating "Sister OTAs" strategically to capture different customer segments and utilizing user interface quality to segment the market. However, at least in this setting, differentiation strategies are not without costs – our results suggest that some of the inefficiency in selection may be due to imperfect implementation of differentiation strategies.

Tables

Table 1: Ticket Characteristics

Characteristic	Time Priority	Price Priority
Total Tickets	7,512	7,512
Average Price	\$557	\$515
Standard Deviation	\$395	\$376
Tickets with connections	38.0%	40.8%
Tickets not w/in time window	17.6%	38.8%

Table 2: Price Differences Among OTAs (Baseline Regression)

VARIABLE	Log-linear TIME Full Data Set	Log-linear PRICE Full Data Set	Log-linear TIME Reduced Data Set	Log-linear PRICE Reduced Data Set	Log-linear TIME	Log-linear PRICE
Intercept	5.768*** (0.097)	5.731*** (0.040)	5.851*** (0.168)	5.695*** (0.059)	5.764*** (0.146)	5.651*** (0.053)
DD7	0.468*** (0.018)	0.490*** (0.017)	0.460*** (0.025)	0.487*** (0.024)	0.460*** (0.026)	0.486*** (0.025)
DD14	0.175*** (0.019)	0.219*** (0.018)	0.206*** (0.027)	0.238*** (0.026)	0.205*** (0.027)	0.237*** (0.026)
DD21	0.077*** (0.020)	0.087*** (0.018)	0.085*** (0.028)	0.087*** (0.026)	0.084*** (0.027)	0.088*** (0.026)
DD28	0.090*** (0.022)	0.093*** (0.022)	0.078*** (0.031)	0.090*** (0.031)	0.083*** (0.032)	0.091*** (0.031)
OTA2 _{time}	0.172*** (0.014)	0.052*** (0.012)	0.172*** (0.014)	0.052*** (0.013)	0.112*** (0.014)	0.048*** (0.012)
OTA2 _{price}	0.172*** (0.014)	0.052*** (0.012)				
OTA3 _{time}	0.175*** (0.014)	0.052*** (0.012)				
OTA3 _{price}	0.178*** (0.014)	0.053*** (0.012)				
OTA4 _{time}	0.278*** (0.015)	0.259*** (0.012)	0.278*** (0.014)	0.259*** (0.014)	0.181*** (0.015)	0.190*** (0.014)
OTA4 _{price}	0.278*** (0.015)	0.259*** (0.014)				
OTA5	0.252*** (0.016)	0.208*** (0.015)	0.252*** (0.015)	0.208*** (0.014)	0.155*** (0.016)	0.153*** (0.014)
Time Window Departure					0.116*** (0.022)	0.096*** (0.019)
Time Window Return					0.130*** (0.020)	0.079*** (0.014)
Connection Departure					-0.006 (0.030)	0.005 (0.030)
Connection Return					-0.170*** (0.033)	-0.120*** (0.031)
Saturday Night Stay	-0.300*** (0.016)	-0.330*** (0.016)	-0.293*** (0.024)	-0.310*** (0.020)	-0.286*** (0.023)	-0.308*** (0.022)
N	7512	7512	3756	3756	3756	3756
F (Prob>F)	65.499 p<0.0001	79.315 p<0.0001	29.55 p<0.0001	34.66 p<0.0001	31.09 p<0.0001	35.93 p<0.0001
R ²	80.56	83.39%	79.81%	82.26%	80.77%	82.92%

White standard errors in parenthesis; *** - p<.001, ** - p<.01, * - p<.05

Table 3a: Ticket Characteristic Combinations by OTAs (Time Preference)

TWD	TWR	CND	CNR	OTA1	OTA2 time	OTA4 time	OTA5	Total
				54	5	14	1	74
			✓	13	1	2	0	16
		✓		6	0	0	0	6
		✓	✓	52	6	1	1	60
	✓			56	28	49	13	146
	✓		✓	20	3	27	0	50
	✓	✓		8	1	0	0	9
	✓	✓	✓	60	25	9	7	101
✓				87	14	32	25	158
✓			✓	13	3	0	0	16
✓		✓		19	2	25	0	46
✓		✓	✓	69	28	4	11	112
✓	✓			279	467	634	638	2018
✓	✓		✓	15	24	6	0	45
✓	✓	✓		19	21	9	0	49
✓	✓	✓	✓	169	311	127	243	850
				939	939	939	939	3756

✓ - Means that this criteria was met

Table 3b: Ticket Characteristic Combinations by OTAs (Price Preferences)

TWD	TWR	CND	CNR	OTA1	OTA2 time	OTA4 time	OTA5	Total
				79	44	32	18	173
			✓	13	3	2	0	18
		✓		6	3	1	0	10
		✓	✓	68	69	1	15	153
	✓			54	40	62	67	223
	✓		✓	22	3	22	0	47
	✓	✓		6	4	0	0	10
	✓	✓	✓	72	66	12	36	186
✓				84	76	61	68	289
✓			✓	13	5	1	0	19
✓		✓		9	10	15	0	34
✓		✓	✓	74	116	5	27	222
✓	✓			267	287	593	520	1667
✓	✓		✓	15	16	3	0	34
✓	✓	✓		16	11	6	0	33
✓	✓	✓	✓	141	186	123	188	638
			Total	939	939	939	939	3756

✓ - Means that this criteria was met

TWD (Time Window Departure) TWD is checked if the departure flight left within the time window (one hour before and after the desiredtime).
TWR (Time Window Return) TWR is checked if the return flight left within the time window (one hour before and after the desired_{time}).
CND (Connection Departure) CND is checked if the departure flight involved at least one connection.
CNR (Connection Return) CNR is checked if the return flight involved at least one connection.

Table 4a: Distribution Tests for Similarity of Characteristics Across Systems (Time Preference)

Sign Test TIME	OTA1	OTA2	OTA4 _{time}	OTA5
OTA1	-	n = 16 M=4*	n = 16 M = 13**	n = 16 M=14***
OTA2	-	-	n = 15 M = 8	n = 15 M = 13***
OTA4 _{time}	-	-	-	n = 12 M = 9
OTA5	-	-	-	-

Wilcoxon TIME	OTA1	OTA2	OTA4 _{time}	OTA5
OTA1	-	n = 16 T = 35*	n = 16 T = 21**	n = 16 T = 31*
OTA2	-	-	n = 15 T = 6.5***	n = 15 T = 23**
OTA4 _{time}	-	-	-	n = 12 T = 20.5
OTA5	-	-	-	-

***-p<.01, ** - p<.05, * - p<.1

Table 4b: Distribution Tests for Similarity of Characteristics Across Systems (Price Preference)

Sign Test PRICE	OTA1	OTA2	OTA4 _{time}	OTA5
OTA1	-	n = 16 M=5	n = 15 M = 12**	n = 16 M=13**
OTA2	-	-	n = 16 M = 12*	n = 16 M = 13**
OTA4 _{time}	-	-	-	n = 15 M = 9
OTA5	-	-	-	-

Wilcoxon PRICE	OTA1	OTA2	OTA4 _{time}	OTA5
OTA1	-	n = 16 T = 50	n = 15 T = 21.5**	n = 16 T = 33.5*
OTA2	-	-	n = 16 T = 42.5	n = 16 T = 29**
OTA4 _{time}	-	-	-	n = 15 T = 59
OTA5	-	-	-	-

***-p<.01, ** - p<.05, * - p<.1

Table 5a: Comparison of OTAs for Similar and Dominated Tickets (Time Preference)

		Y			
	X Strictly Dominates Y	OTA1	OTA2 _{time}	OTA4 _{time}	OTA5
X	OTA1	0	122	74	207
	OTA2 _{time}	82	0	112	241
	OTA4 _{time}	20	110	0	268
	OTA5	103	193	199	0

		Y			
	X is equal to Y	OTA1	OTA2 _{time}	OTA4 _{time}	OTA5
X	OTA1	939	297	337	161
	OTA2 _{time}	297	939	430	266
	OTA4 _{time}	357	430	939	316
	OTA5	161	266	316	939

Table 5b: Comparison of OTAs for Similar and Dominated Tickets (Price Preference)

		Y			
	X Strictly Dominates Y	OTA1	OTA2 _{time}	OTA4 _{time}	OTA5
X	OTA1	0	163	82	216
	OTA2 _{time}	88	0	80	185
	OTA4 _{time}	21	67	0	258
	OTA5	102	122	177	0

		Y			
	X is equal to Y	OTA1	OTA2 _{time}	OTA4 _{time}	OTA5
X	OTA1	939	350	337	161
	OTA2 _{time}	350	939	308	192
	OTA4 _{time}	337	308	939	293
	OTA5	161	192	293	939

Table 6: The Effect of Interactions between Ticket Characteristics and OTAs on Ticket Price

VARIABLE	Log-linear TIME	Log-linear PRICE
Intercept	5.961*** (0.145)	5.761*** (0.060)
DD7	0.455*** (0.026)	0.488*** (0.025)
DD14	0.202*** (0.027)	0.236*** (0.026)
DD21	0.080*** (0.027)	0.084*** (0.026)
DD28	0.077*** (0.032)	0.090*** (0.032)
OTA2 _{time}	0.105*** (0.014)	0.039** (0.012)
OTA4 _{time}	0.193*** (0.015)	0.204*** (0.015)
OTA5	0.176*** (0.017)	0.145*** (0.014)
OTA1 * (CN-CN _{avg})	-0.109*** (0.025)	-0.077*** (0.023)
OTA2 _{time} * (CN-CN _{avg})	-0.196*** (0.026)	-0.124*** (0.024)
OTA4 _{time} * (CN-CN _{avg})	-0.044 (0.034)	-0.023 (0.032)
OTA5 * (CN-CN _{avg})	-0.213*** (0.030)	-0.162*** (0.027)
OTA1 * (TW-TW _{avg})	0.224*** (0.021)	0.189*** (0.020)
OTA2 _{time} * (TW-TW _{avg})	0.098*** (0.033)	0.086*** (0.021)
OTA4 _{time} * (TW-TW _{avg})	0.069** (0.033)	0.080*** (0.030)
OTA5 * (TW-TW _{avg})	-0.117** (0.049)	0.076*** (0.025)
Saturday Night Stay	-0.285*** (0.024)	-0.301*** (0.023)
N	3756	3756
F (Prob>F)	31.14 p<.0001	35.82 p<.0001
R ²	80.96%	83.02%

White standard errors in parenthesis; ***-p<.001, ** - p<.01, * - p<.05

Figures

Figure 1a: Efficient Frontier - Price Satisfaction vs. Connections

Price Satisfaction = $\frac{\text{Price Price(OTA1)}}{\text{Price(OTA)}}$; Connection Satisfaction is the percentage of flights offered without connections.

Figure 1b: Efficient Frontier: Price Satisfaction vs. Timeliness

¹ At the time of data collection (1997), the commissions ranged from roughly 5-8% of ticket price with an overall cap of about \$50 to various flat fee schedules up to \$35. These commissions are similar but do vary by airline. By the year 2000, most airlines had moved to a fixed commission schedule of about \$15 for online bookings and some airlines have even explored eliminating commissions for online booking. OTAs also earn revenues from banner advertisements, special promotional programs and the sale of packaged travel, although at the time of our study, regular air commissions were the dominant source of revenue.

² Conversations with OTA managers indicate that the fee depends on the revenue generated. Since CRS's also compete with each other, travel agents, who produce high volumes of bookings, will obtain the CRS service for free or even be paid for the bookings. According to the interviewee, the OTAs at the point of the study did not cross the required threshold volume. The website of the airline typically does not have to pay the fee as these websites are typically connected to their internal host.

³ Horizontal differentiation is when different products are offered to target different consumer tastes (e.g. color) along non-quality dimensions. This contrasts with vertical differentiation, where different price-quality combinations are offered (Beath and Katsoulacos, 1991).

⁴ At the time of our study, OTAs could only offer a single fare class which was the lowest cost available at the time of ticketing (typically restricted Y-class tickets). In November, 1984, following the Sabre-American Airlines antitrust suit with the Department of Justice, new rules were introduced that prohibited agents from biasing recommendations by airline without customer request. Where available, our software agents specified "no preference" for airline choice in all requests.

⁵ The reservation data only provide the input to the requests. It is the specification of the traveler's preference that allows us to verify the hypotheses for price sensitive as well as time sensitive travelers.

⁶ Specifically, we required that all tickets to have either 2 or 4 segments, had the same departure and final destination, and represented travel entirely within the U.S.

⁷ OTAs did not provide any market share data, hence we relied on industry participants to select significant competitors.

⁸ It was not uncommon for one or several of the OTAs to be unavailable. In that case, the intelligent software agents stopped the reservation request after a specified waiting time and continued with the next reservation.

⁹ While the difference between these means is significant, the time priority dataset is only 8% more expensive, reflecting the fact that on many routes lowest price and most timely flights are the same due to limited choice.

¹⁰ The log-linear formulation reduced heteroskedasticity substantially compared to the linear formulation, resulting in a reduction of the R^2 of the White Test regression from 12% to 0.9%. Due to our large sample size, this is still statistically significant but not economically significant. We therefore report White heteroskedasticity-consistent standard errors (White, 1980) in all the tables (which are essentially identical to the OLS errors).

¹¹ The OTA coefficients were typically within one percent of model (1). Results are available from the authors upon request.

¹² We perform a t-test to explore whether or not the mean ranks of the systems with the option of indicating different preferences (e.g., $OTA2_{time}$ and $OTA2_{price}$) are the same. For the time-priority (price-priority) data set we obtain the following results: $OTA2$ $t=0.10$ ($t=0.15$), $OTA3$ $t=0.31$ ($t=0.15$) and $OTA4$ $t=0$ ($t=0$).

¹³ Results are available from, the authors by request.

¹⁴ We extended this analysis for all 16 possible subgroups of ticket categories to see whether or not an OTA is strictly dominant for specific ticket characteristics. Our results did not indicate any strict dominance of an OTA for any subgroup.

¹⁵ This measure is crude because it is only identified if the ticket set returned by an agent has substantial variation on these dimensions. For example, to estimate the shadow cost of missing connections, we need to observe the same agent returning tickets with and without connections for a similar request. For agents which aggressively optimize on a single dimension, there will be little variation in the sample set leading to potentially anomalous results. However, for agents which tend to trade off multiple dimensions, this analysis will more accurately reveal this tendency.

¹⁶ This result is probably due to the fact that $OTA4_{time}$ offered the fewest connections of all OTAs. After accounting for all ‘necessary’ connections through the city-pair variables, there is not enough data variation left to achieve statistical significance.

References

- Bailey, J., Brynjolfsson, E. and Smith M.D. "In Search of "Friction-Free Markets": An Exploratory Analysis of Prices for Books, CDs and Software Sold on the Internet," working paper, Technology, Management and Policy, Massachusetts Institute of Technology, 1997.
- Bailey, J. P. 1998. "Electronic Commerce: Prices and Consumer Issues for Three Products: Books, Compact Discs, and Software," Organization for Economic Co-Operation and Development, Working Paper OCDE/GD(98)4 (<http://www.oecd.org/dsti/sti/it/ec/prod/ie98-4.pdf>)
- Bakos, J. Y. "A Strategic Analysis of Electronic Marketplaces," *MIS Quarterly*, Vol. 15, No. 3, September 1991, pp. 295-310.
- Bakos, J. Y. "Reducing Buyer Search Costs: Implications for Electronic Marketplaces," *Management Science*, Vol. 43, No. 12, 1997, pp. 1676-1692.
- Beath, J. and Katsoulacos, Y. *The Economic Theory of Product Differentiation*. Cambridge University Press, Cambridge, 1991.
- Benjamin, R. and Wigand, R. "Electronic Markets and Virtual Value Chains on the Information Superhighway," *Sloan Management Review*, Winter 1995, pp. 62-72.
- Bertrand, J. "Théorie Mathématique de la Richesse Sociale," *Journal des Savants*, 1883, pp. 499-508.
- Borenstein, S. "Hubs and High Fares: Dominance and Market Power in the U.S. Airline Industry," *Rand Journal of Economics*, Vol. 20, Autumn 1989, pp. 344-65.

Borenstein, S. "The Dominant-firm Advantage in Multiproduct Industries: Evidence from the U.S. Airlines," *Quarterly Journal of Economics*, Vol. 106, November 1991, pp. 1237-66.

Brynjolfsson, E. and Smith, M., "Frictionless Commerce: A Comparison of Internet and Conventional Retailers," *Management Science*, April 2000 (volume 46, number 4), pages 563-585

Chen, P. and Hitt, L. M. "Switching Cost and Brand Loyalty in Electronic Markets: Evidence from On-Line Retail Brokers," *21st Annual International Conference on Information Systems*, Brisbane, Australia, 2000.

Chow, G. C. "Technological Change and the Demand for Computers," *American Economic Review*, Vol. 57, No. 5, December 1967, pp. 1117-1130.

Clemons, E.K. and Hitt, L.M. "The Internet and the Future of Financial Services: Price Transparency, Disintermediation and Differential Pricing," in *The Economic Payoff from the Internet Revolution* (R. Litan and A. Rivlin, eds.), pp. 97-128, Brookings Institution Press, 2001.

Combes, G. C. and Patel, J. J. "Creating Lifelong Customer Relationships: Why the Race for Customer Acquisition on the Internet is so Strategically Important," *iword*, Vol. 2, Issue 4, Hambrecht & Quist 1997.

Dupuit, J. "On Tolls and Transport Charges," in *International Economic Papers*, London, Macmillan, 1962, pp. 7-31; original version: "De l'influence des péages sur l'utilité des voies de communication," in *Annales des Ponts et Chaussées*, 2^e série, 1849, pp. 207-248.

Evans, W.N. and Kessides, I.N. "Localized Market Power in the U.S. Airline Industry," *The Review of Economics and Statistics*, Vol. 75, Issue 1, February 1993, pp. 66-75.

Griliches, Z. "Hedonic Price Indexes for Automobiles: An Econometric Analysis of Quality Change," in *The Price Statistics of the Federal Government*, General Series No. 73, New York, Columbia University Press for the National Bureau of Economic Research, 1961, pp. 137-196.

Johnson, E.J., Moe, W., Fader, P., Bellman, S., and Lohse, J., "On the Depth and Dynamics of Online Search Behavior," Working Paper, June 2000, The Wharton School, University of Pennsylvania.

Lee, H. G. "Do Electronic Marketplaces Lower the Price of Goods?" *Communications of the ACM*, January 1998, Vol. 41, No. 1, pp. 73-80.

Machlis, S. "Profits Elude Travel Sites," *Computerworld*, December 29, 1997, pp. 53-54.

Malone, T. W., Yates, J., and Benjamin, R. I. "Electronic Market and Electronic Hierarchies," *Communications of the ACM*, Vol. 30, No. 6, 1987, pp. 484-497.

Salop, S., Stiglitz, J. "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion," *Review of Economic Studies*, Vol. 44, 1976, pp. 493-510.

Smith, Michael D and Brynjolfsson (2001, forthcoming), Erik, Consumer Decision-Making at an Internet Shopbot, J. *Industrial Economics*, December.

Stigler, G. J. "The Economics of Information," *The Journal of Political Economy*, Vol. , No. 3, June 1961, pp. 213-225.

White, H. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," *Econometrica*, Vol. 48, pp. 817-838, May 1980.