The Power of Rankings:
Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions

- - ONLINE APPENDIX - -

Raluca M. Ursu*

*Stern School of Business, New York University, New York, New York 10012, rursu@stern.nyu.edu

Appendix A: Data Cleaning

The data set made available contains 9,917,530 observations. For my analysis, I filtered out the following three categories of observations, which reduced the data set to 4,503,043.

First, the data set contains some errors in the way price information was stored. For example, both very high (more than $19 million per night) and very low ($0.01 per night) prices appear in the data set. I corrected these errors for those consumers who made purchases, by removing search impressions that contain at least one observation for which the total amount spent exceeded the price paid multiplied by the length of the trip and the number of rooms booked plus taxes.[1]

Second, I chose to focus on "typical" search impressions and removed those that include prices lower than $10 or higher than $1,000 per night. By eliminating these search impressions, I also mitigated the first problem above for queries not ending in a transaction.

Third, the original data set contains observations on more than 20,000 destinations, with a median of two search impressions per destination. To observe enough variation in the position of a hotel under the Random ranking, I focused my attention on destinations with at least 50 search impressions.

Appendix B: Further Evidence for Section 3

Ranking Algorithm

In this section, I describe a basic *learning to rank* algorithm.[2] *Learning to rank* algorithms are machine learning algorithms used to rank documents based on relevance. In the current application, a document is a hotel and its characteristics. The system (e.g., the search intermediary) maintains a collection of these hotels, and when a consumer makes a search query, it proceeds to rank them. More precisely, the system uses a so-called training data set, which contains past consumer queries, hotel characteristics x_{ij} for each hotel j as observed by consumer i, and associated clicks and purchases made by the consumer. This latter information is interpreted as an indication of how relevant a hotel was for the consumer query and is known as the relevance score s_{ij} of a hotel. This relevance score is highest if a purchase occurred, and lowest if no click occurred. The purpose of the *learning to rank* algorithm is to use the data on hotel j's characteristics x_{ij} for a query performed by i and i's clicks/purchases s_{ij} for each j and to learn a function $f(x_{ij}, \gamma) = \hat{s}_{ij}$, so that the ranking order of predicted scores $\hat{s}_{ij}$ is exactly equal to that of observed s_{ij}. More concretely, this means the goal is to find a function that will rank at the top (the most relevant) the hotel the consumer will purchase, followed by hotels the consumer will click, and finally, by those the consumer will not consider. Importantly for this paper, a *learning to rank* model exploits the availability of data

[1] In the United States, hotel taxes range from 7% to almost 20% according to http://www.consumerreports.org/cro/news/2014/06/booking-a-hotel-these-cities-have-the-highest-hotel-taxes/index.htm, motivating my conservative approach of dropping observations where the implied tax is larger than 30%.

[2] A commonly used *learning to rank* algorithm is known as LambdaMART. It is described in more detail by Yoganarasimhan (2016).

on relevance scores of consumers, which makes the ranking of the search intermediary endogenous, because the position of a hotel in the ranking is set in response to predicted consumer choices.

Randomization Check

In this section, I show (i) consumers were randomly assigned to the two types of rankings, and (ii) in constructing the Random ranking, the position of the hotel was randomly generated.

To show consumers were randomly assigned to each type of ranking, I perform two tests. First, I test whether the time of arrival to Expedia's website is related to the type of ranking the consumer saw. One concern may be that different types of consumers visit the website at different times of the day, which would bias the results if the probability of observing one type of ranking is different at different times of the day. For example, suppose business travelers search after 5pm and they have a higher probability of purchase. Then, if after 5pm, the probability of observing Expedia's ranking were higher (instead of the Random ranking), a correlation would be present between consumer choices and Expedia's ranking in the data that is not due to the ranking observed, but rather to the way consumers were assigned to the two rankings. However, Figure 1 shows that this issue is not a concern in the data. More precisely, the left panel plots the number of search impressions made by the time of the day, and shows more search impressions occur in the afternoon and evening than in the morning. The right panel plots the fraction of search impressions seeing the Random ranking every 30 seconds during the course of one day in the entire data set. It shows the fraction of search impressions seeing the Random ranking is constant throughout the day, so that even though more queries happen in the second part of the day, they have the same probability of seeing the Random ranking. Thus, these figures suggest consumers were randomly assigned to see either type of ranking.

Figure 1: Number of search impressions occurring every 30 seconds during a day and the fraction seeing the Random ranking

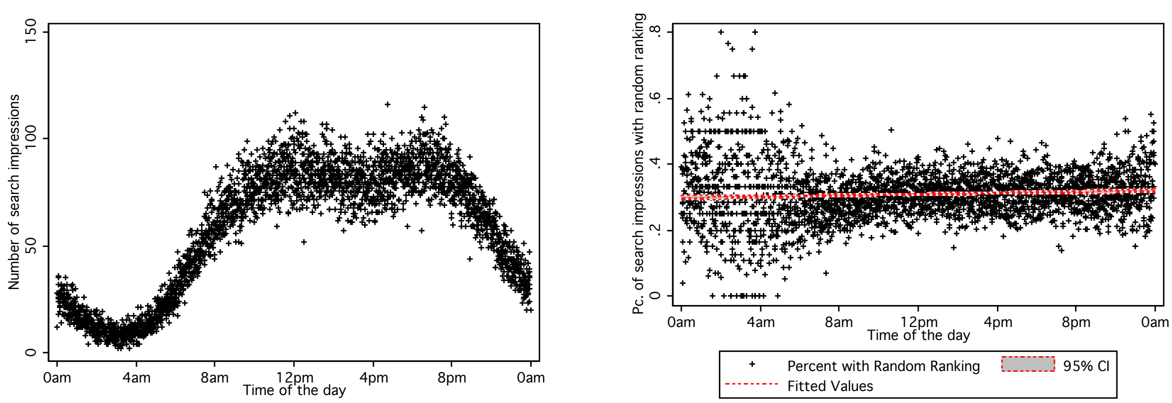

Number of search impressions

Fraction seeing the Random ranking

Second, I test whether consumer characteristics observed by Expedia prior to displaying a ranking are different for the two rankings. When the consumer arrives at Expedia's website, she reveals details about her upcoming trip, such as her destination, the length of the trip, how long in advance she is searching for, the number of travelers and rooms requested, as well as whether her trip includes a Saturday night. For some consumers, Expedia also has historical information on the average price and the number of stars of hotels previously purchased. One concern might be that this information affects consumers' probability of seeing either ranking. Table 1 shows this issue is also not a concern. Comparing separately search impressions ending and not ending in a transaction across the two rankings by means of a t-test, I find that consumers seeing Expedia's ranking have very similar characteristics to those seeing the Random ranking.[3] In particular, the difference in most characteristics is not statistically significant, whereas for the booking window and trip length, which show a significant effect, the magnitude of the effect is small. For example, the difference in trip length is less than 0.2 days (less than 7.4% difference), whereas for booking window, it is less than two days (the average booking window in the sample is 39 days, implying a difference smaller than 5%). Combined and given the large data samples, these findings suggest no systematic differences exist in consumer observables across the two rankings, confirming that consumers were randomly assigned to either ranking.

Table 1: T-test: Consumer observables (Expedia-Random)

	(1) No Tran	(2) Tran
Trip Length (days)	-0.1792^{***}	-0.0857^{***}
	(-6.5268)	(-3.3568)
Booking Window (days)	-1.9284^{**}	0.4653
	(-2.7148)	(0.6224)
Saturday Night	0.0126^{*}	-0.0018
	(2.2582)	(-0.2168)
Adults	-0.0034	-0.0475^{***}
	(-0.3268)	(-3.3687)
Children	-0.0051	-0.0205
	(-0.5477)	(-1.6416)
Rooms	-0.0034	-0.0058
	(-0.6458)	(-0.8529)
Observations	57,133	108,903
Consumer Hist. Stars	0.0684	0.0665
	(1.2424)	(1.3226)
Consumer Hist. Price	-9.3606	-8.0906
	(-1.2045)	(-0.9814)
Observations	1,286	6,264

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

To show hotels were randomly ordered under the Random ranking, as stated by the data provider in the description of the data,[4] I compare the average characteristics of the ranked hotels under the two rankings. My results can be found in Figure 2 below. I find that under the Random ranking, the quality of the ranked hotels is fairly constant across position, consistent with no correlation being present between the quality of the hotel and its position. This is to be expected if hotels were randomly ranked. By contrast, under Expedia's ranking, a strong correlation is present

[3] I separate search impressions ending and not ending in a transaction to avoid biasing results since converting search impressions were oversampled.

[4] See https://www.kaggle.com/c/expedia-personalized-sort/forums/t/5772/meaning-of-random-bool.

Figure 2: Characteristics of the ranked hotels under the two rankings

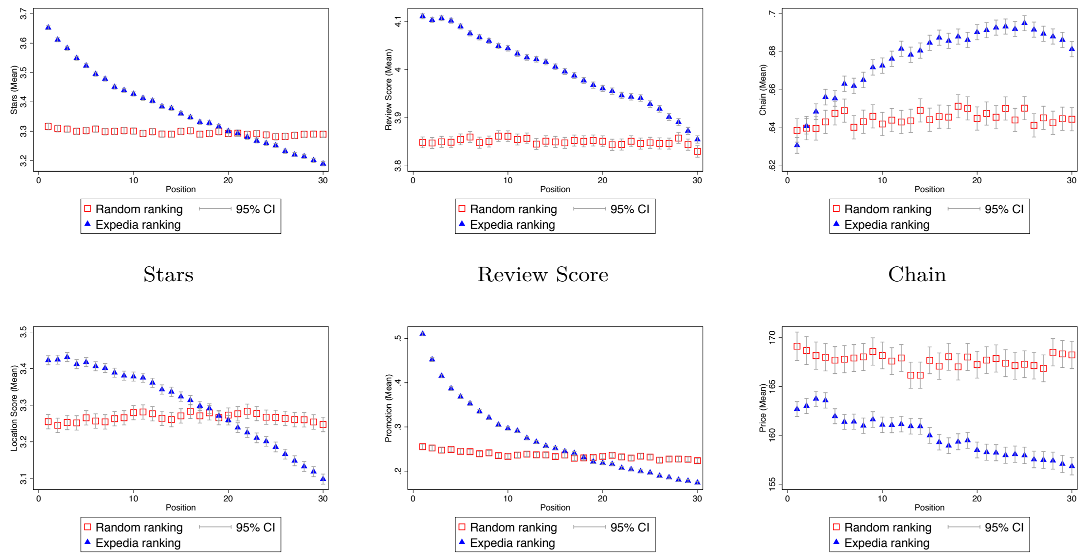

Stars Review Score Chain

Location Score Promotion Price

Note: Restrict attention to search impressions with opaque offers of similar length and focus on organic listings by not including the first and last two positions in each search impression. Position is the order in which hotels appeared on the page, not the actual position (difference may arise due to opaque offers). The difference in price level across the two rankings is due to the oversampling of transactions under Expedia's ranking.

between a hotel's characteristics and its position, as is expected of a curated ranking. Moreover, this result provides insights into the construction of Expedia's ranking: it favors non-chains that are more expensive, that have more stars, better reviews, better location, and that are running a promotion.

The WCAI Data Set

In this section, I describe the companion data set from the Wharton Customer Analytics Initiative (WCAI). It provides information on consumer queries for hotels on a popular online travel agent's (OTA) website between October 1 and 9, 2009.[5] This data set cannot be used to study the causal effect of rankings, because it does not include search impressions performed under a random ranking. However, it has information on some of the consumer groups that are excluded from the Expedia data set I use. For example, it contains queries where consumers do not click, where they sort/filter search results, as well as where they consider the additional result pages (beyond the first page of results). Thus, I use the WCAI data to explore the impact of the absence of these groups from the analysis. I find the groups that are missing generally constitute a small fraction of search queries, allowing me to state that their absence will not affect the representativeness of the Expedia data.

The WCAI data set includes search impressions in Paris, Budapest, Cancun, and Manhattan, but for the current analysis, I focus on the Manhattan data set. This data set contains 431,820

[5]I cannot disclose the name of the OTA that provided the data.

Table 2: Summary statistics comparison with the WCAI data set (Mean)

	Expedia	WCAI	WCAI Click
Hotel level			
Price	159.71	281.50	281.63
Stars	3.32	3.08	3.16
Review Score	3.89	4.01	4.03
Chain	0.66	0.35	0.37
Promotion	0.25	0.35	0.28
Observations	*4,503,043*	*431,820*	*108,066*
Search impression level			
Number of Hotels Displayed	27.12	25.00	25.00
Total Clicks	1.12	0.42	1.80
Total Transactions	0.66	0.01	0.03
Sort/filter	0%	32%	33%
First Page	100%	67%	68%
Random ranking	31%	0%	0%
Observations	*166,036*	*15,171*	*3,560*
Number of hotels	54,877	301	300
Number of destinations	788	1	1

Note: The WCAI data set does not contain a location score variable, so I omit it from the table for the Expedia data.

observations and 15,171 search impressions made by 4,752 consumers. Table 2 compares three data sets: the Expedia data, the WCAI data, and the WCAI data set containing search impressions with *at least* one click (similar to the Expedia data set), which I call WCAI-Click. The hotels displayed in the three data sets have similar characteristics (except that those in the Expedia data are cheaper, most likely due to the fact that it contains a much greater pool of destinations (788) and hotels (54,877) than those found in Manhattan).

Comparing the three data sets, I draw several conclusions. First, only 23% (3,560 out of 15,171) of the queries in the WCAI data set have at least one click. However, the Expedia data set, which only contains search impressions with at least one click, has a similar average number of clicks as the WCAI-Click data set (1.12 compared to 1.80).

Second, in the Expedia data set, two thirds of search impressions end in a transaction, whereas in the WCAI (WCAI-Click) data set, only 1% (3%) lead to a transaction. However, as discussed in section 3.2.3 (in the paper), search impressions ending and those not ending in a transaction in the Expedia data set were sampled randomly, thus not affecting the main analysis.

Third, the Expedia data set contains no queries in which consumers sort or filter. However, in the WCAI or WCAI-Click data sets, the minority of consumers actually make these choices: in the WCAI data set, out of all search impressions, only 4,860 contain filtered results (32%), whereas only 1,601 consumers filter (34%). Therefore, most consumers make choices from the ranking that is displayed, instead of sorting/filtering the hotels observed.

Finally, in the Expedia data set, I only observe the first page of results consumers saw in a search impression. However, the WCAI and the WCAI-Click data sets show the majority of consumers consider only the first page of results: out of all search impressions, 10,228 (67%) choose to consider only the first page of results, whereas out of all consumers, 3,027 (64%) consider only the first page of results. Thus, only observing consumer choices from the first page of results captures most consumers' behavior on an OTA.

The comparison of the Expedia and the WCAI data sets allows me to state that the former data set is representative of most consumers searching on an OTA.

Transaction Unconditional on Click

Figure 3: The effect of position on the transaction rate

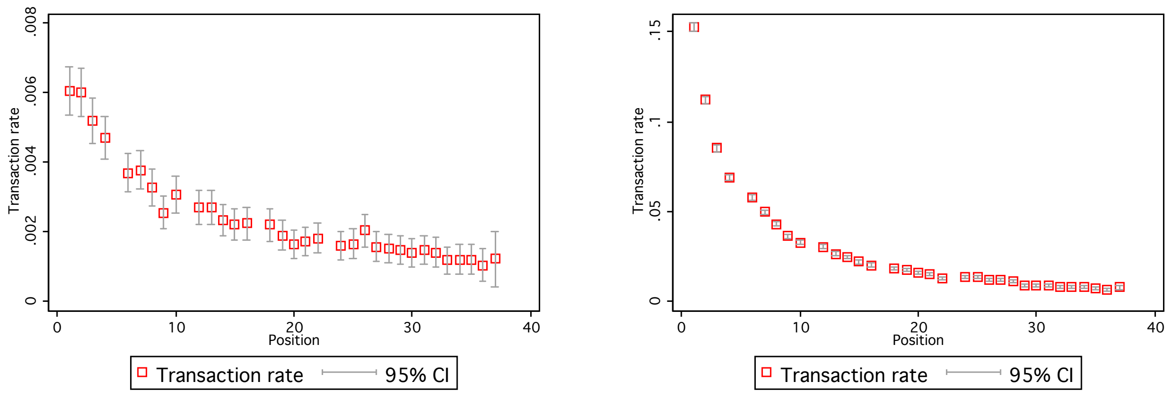

Random ranking

Expedia ranking

Summary Statistics for Clicked and Purchased Hotels

Table 3: Summary statistics on hotels displayed, clicked and purchased (Mean)

	Displayed	Clicked	Purchased
Price	159.71	146.84	137.32
Stars	3.32	3.43	3.40
Review Score	3.89	3.96	3.99
Chain	0.66	0.65	0.67
Location Score	3.17	3.29	3.20
Promotion	0.25	0.34	0.36
Observations	4,503,043	186,171	108,903

Robustness Checks for the Click Through Rate and the Conversion Rate Patterns under the Random Ranking

This section shows the same patterns as in Figure 1 in the paper hold when restricting attention to (i) small destinations with less than the median number of hotels in the sample, to control for the fact that in large destinations, sponsored ads are more likely (see Figure 4), (ii) search impressions with more than 30 hotels displayed, to control for the fact that some search impressions have few hotels displayed (see Figure 5), and (iii) search impressions ending in a transaction (see Figure 6).

Figure 4: Destinations with less than the median number of hotels: Random ranking

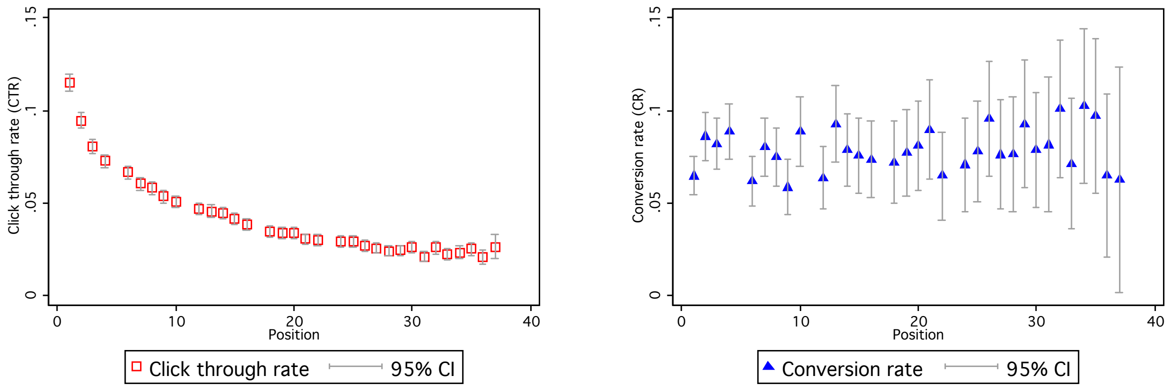

Figure 5: Search impressions longer than 30 displayed hotels: Random ranking

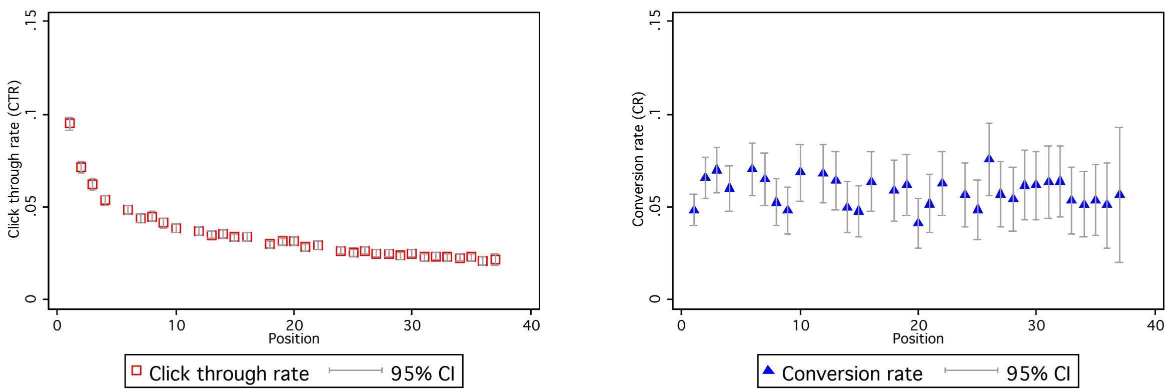

Figure 6: Search impressions ending in a transaction: Random ranking

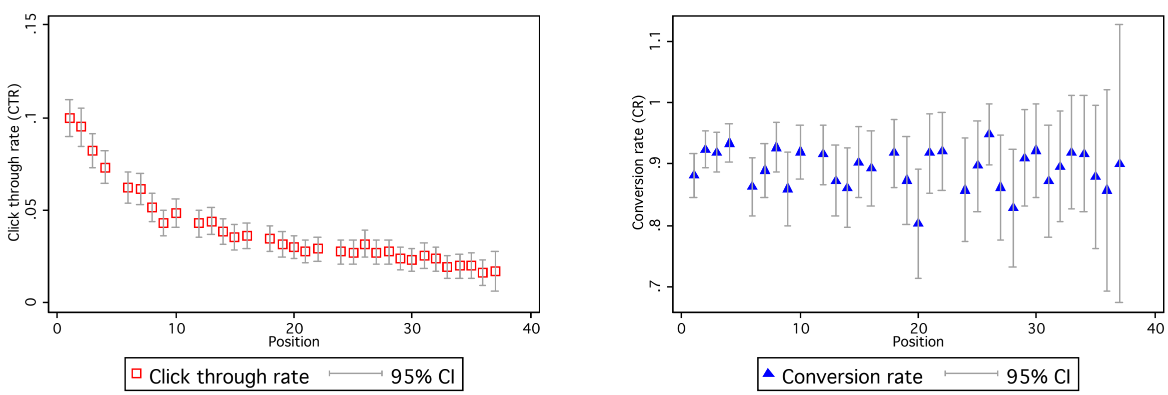

Note: The scale of the y-axis in the two figures is different.

Position Effect under the Random Ranking

Table 4: Estimates of click, transaction, and transaction conditional on click (Logit model): Random ranking

	Click			Transaction		Transaction conditional on click		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Position effect								
Position	-0.0494^{***}	-0.0481^{***}		-0.0467^{***}		-0.0016	-0.0002	
	(0.0005)	(0.0005)		(0.0020)		(0.0018)	(0.0019)	
Position 1			1.7231^{***}		1.6968^{***}			0.0830
			(0.0811)		(0.3392)			(0.3534)
Position 2			1.4629^{***}		1.7137^{***}			0.3452
			(0.0814)		(0.3392)			(0.3536)
Position 3			1.2969^{***}		1.5696^{***}			0.3453
			(0.0816)		(0.3399)			(0.3541)
Position 4			1.1713^{***}		1.4715^{***}			0.3778
			(0.0818)		(0.3405)			(0.3550)
Position 6			1.0519^{***}		1.2314^{***}			0.2006
			(0.0821)		(0.3423)			(0.3568)
Position 7			0.9554^{***}		1.2646^{***}			0.3612
			(0.0823)		(0.3423)			(0.3573)
Position 8			0.8970^{***}		1.1121^{**}			0.2204
			(0.0825)		(0.3439)			(0.3588)
Position 9			0.8205^{***}		0.8756^{*}			0.0471
			(0.0826)		(0.3464)			(0.3612)
Position 10			0.7207^{***}		1.0625^{**}			0.3730
			(0.0829)		(0.3447)			(0.3595)
Hotel characteristics								
Price		-0.0036^{***}	-0.0036^{***}	-0.0066^{***}	-0.0066^{***}		-0.0023^{***}	-0.0024^{***}
		(0.0001)	(0.0001)	(0.0004)	(0.0004)		(0.0003)	(0.0003)
Stars		0.4127^{***}	0.4129^{***}	0.4529^{***}	0.4520^{***}		0.0092	0.0090
		(0.0078)	(0.0078)	(0.0316)	(0.0316)		(0.0326)	(0.0326)
Review Score		0.0450^{***}	0.0458^{***}	0.1910^{***}	0.1918^{***}		0.1685^{***}	0.1673^{***}
		(0.0058)	(0.0059)	(0.0301)	(0.0301)		(0.0331)	(0.0331)
Chain		0.0527^{***}	0.0548^{***}	0.1341^{**}	0.1351^{**}		0.0932^{*}	0.0918^{*}
		(0.0109)	(0.0109)	(0.0447)	(0.0446)		(0.0466)	(0.0467)
Location Score		0.1177^{***}	0.1188^{***}	0.2711^{***}	0.2724^{***}		0.1666^{***}	0.1668^{***}
		(0.0046)	(0.0046)	(0.0208)	(0.0208)		(0.0212)	(0.0213)
Promotion		0.2351^{***}	0.2327^{***}	0.4230^{***}	0.4204^{***}		0.2184^{***}	0.2190^{***}
		(0.0105)	(0.0105)	(0.0411)	(0.0411)		(0.0433)	(0.0433)
Query characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Destination FE	No	Yes	Yes	Yes	Yes	No	Yes	Yes
Log-likelihood	-217,971	-211,254	-210,565	-20,871	-20,838	-12,486	-11,226	-11,209
Observations	1,245,455	1,220,917	1,220,917	1,205,168	1,205,168	54,614	53,172	53,172

Standard errors in parentheses

Note: Standard errors clustered at the search impression level. Regressions in columns (3), (5), and (8) include position FE for positions 1 through 37, and coefficients should be interpreted with respect to the omitted category: the last position. For visibility, I only display the coefficients for positions 1-10. I restrict the sample to search impressions with opaque offers, which means no hotel was displayed in positions 5, 11, 17, and 23 in the ranking observed by consumers. Query characteristics included in all regressions: trip length, booking window, number of adults and children traveling, number of rooms, and an indicator for a Saturday-night stay.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Position Effect under Expedia's Ranking

Table 5: Estimates of click, transaction, and transaction conditional on click (OLS): Expedia ranking

	Click			Transaction		Transaction conditional on click		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Position effect								
Position	−0.0031***	−0.0029***		−0.0025***		−0.0039***	−0.0030***	
	(0.0000)	(0.0000)		(0.0000)		(0.0002)	(0.0002)	
Position 1			0.1622***		0.1412***			0.0801***
			(0.0013)		(0.0013)			(0.0215)
Position 2			0.1184***		0.1010***			0.0643**
			(0.0012)		(0.0011)			(0.0215)
Position 3			0.0888***		0.0744***			0.0490*
			(0.0011)		(0.0010)			(0.0216)
Position 4			0.0715***		0.0588***			0.0390
			(0.0010)		(0.0009)			(0.0217)
Position 6			0.0586***		0.0481***			0.0402
			(0.0010)		(0.0009)			(0.0217)
Position 7			0.0493***		0.0399***			0.0279
			(0.0009)		(0.0008)			(0.0219)
Position 8			0.0413***		0.0335***			0.0294
			(0.0009)		(0.0008)			(0.0219)
Position 9			0.0335***		0.0271***			0.0268
			(0.0009)		(0.0008)			(0.0220)
Position 10			0.0303***		0.0239***			0.0117
			(0.0008)		(0.0007)			(0.0222)
Hotel characteristics								
Price		−0.0002***	−0.0002***	−0.0001***	−0.0001***		−0.0002***	−0.0002***
		(0.0000)	(0.0000)	(0.0000)	(0.0000)		(0.0000)	(0.0000)
Stars		0.0085***	0.0065***	0.0071***	0.0052***		0.0002	−0.0012
		(0.0002)	(0.0002)	(0.0002)	(0.0002)		(0.0024)	(0.0024)
Review Score		0.0020***	0.0021***	0.0021***	0.0022***		0.0159***	0.0158***
		(0.0002)	(0.0002)	(0.0001)	(0.0001)		(0.0023)	(0.0023)
Chain		0.0029***	0.0041***	0.0022***	0.0033***		−0.0016	−0.0004
		(0.0003)	(0.0003)	(0.0003)	(0.0003)		(0.0031)	(0.0031)
Location Score		0.0026***	0.0030***	0.0026***	0.0029***		0.0159***	0.0162***
		(0.0001)	(0.0001)	(0.0001)	(0.0001)		(0.0018)	(0.0018)
Promotion		0.0050***	−0.0003	0.0053***	0.0006		0.0166***	0.0135***
		(0.0003)	(0.0003)	(0.0003)	(0.0003)		(0.0029)	(0.0030)
Query characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Destination FE	No	Yes	Yes	Yes	Yes	No	Yes	Yes
Adjusted R^2	0.026	0.030	0.043	0.027	0.039	0.008	0.050	0.051
Observations	2,522,288	2,467,000	2,467,000	2,467,000	2,467,000	104,874	104,042	104,042

Standard errors in parentheses
Note: Standard errors clustered at the search impression level. Regressions in columns (3), (5), and (8) include position FE for positions 1 through 37, and coefficients should be interpreted with respect to the omitted category: the last position. For visibility, I only display the coefficients for positions 1-10. I restrict the sample to search impressions with opaque offers, which means no hotel was displayed in positions 5, 11, 17, and 23 in the ranking observed by consumers. Query characteristics included in all regressions: trip length, booking window, number of adults and children traveling, number of rooms, and an indicator for a Saturday-night stay.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 6: Estimates of click, transaction, and transaction conditional on click (Logit model): Expedia ranking

	Click			Transaction		Transaction conditional on click		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Position effect								
Position	-0.0944^{***}	-0.0900^{***}		-0.0947^{***}		-0.0240^{***}	-0.0194^{***}	
	(0.0005)	(0.0005)		(0.0005)		(0.0009)	(0.0010)	
Position 1			2.9731^{***}		3.0307^{***}			0.5504^{**}
			(0.0716)		(0.0798)			(0.1792)
Position 2			2.6408^{***}		2.6859^{***}			0.4135^{*}
			(0.0717)		(0.0799)			(0.1792)
Position 3			2.3561^{***}		2.3886^{***}			0.2911
			(0.0719)		(0.0801)			(0.1797)
Position 4			2.1537^{***}		2.1721^{***}			0.2164
			(0.0720)		(0.0803)			(0.1798)
Position 6			1.9748^{***}		1.9959^{***}			0.2257
			(0.0722)		(0.0805)			(0.1801)
Position 7			1.8247^{***}		1.8343^{***}			0.1397
			(0.0724)		(0.0807)			(0.1805)
Position 8			1.6778^{***}		1.6894^{***}			0.1495
			(0.0726)		(0.0809)			(0.1808)
Position 9			1.5096^{***}		1.5213^{***}			0.1347
			(0.0728)		(0.0812)			(0.1819)
Position 10			1.4296^{***}		1.4223^{***}			0.0356
			(0.0730)		(0.0815)			(0.1822)
Hotel characteristics								
Price		-0.0057^{***}	-0.0056^{***}	-0.0060^{***}	-0.0059^{***}		-0.0010^{***}	-0.0010^{***}
		(0.0001)	(0.0001)	(0.0001)	(0.0001)		(0.0002)	(0.0002)
Stars		0.2319^{***}	0.2040^{***}	0.2339^{***}	0.2034^{***}		0.0023	-0.0078
		(0.0059)	(0.0059)	(0.0063)	(0.0063)		(0.0163)	(0.0163)
Review Score		0.0851^{***}	0.0868^{***}	0.1087^{***}	0.1106^{***}		0.0994^{***}	0.0989^{***}
		(0.0059)	(0.0059)	(0.0064)	(0.0064)		(0.0140)	(0.0140)
Chain		0.0905^{***}	0.1068^{***}	0.0947^{***}	0.1120^{***}		-0.0152	-0.0072
		(0.0079)	(0.0079)	(0.0085)	(0.0086)		(0.0206)	(0.0206)
Location Score		0.0835^{***}	0.0896^{***}	0.1009^{***}	0.1078^{***}		0.1016^{***}	0.1042^{***}
		(0.0039)	(0.0040)	(0.0041)	(0.0042)		(0.0109)	(0.0109)
Promotion		0.0238^{**}	-0.0385^{***}	0.0480^{***}	-0.0180^{*}		0.1143^{***}	0.0908^{***}
		(0.0077)	(0.0079)	(0.0083)	(0.0085)		(0.0203)	(0.0205)
Query characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Destination FE	No	Yes	Yes	Yes	Yes	No	Yes	Yes
Log-likelihood	-399,658	-389,934	-387,048	-331,495	-328,920	-50,172	-47,111	-47,026
Observations	2,522,288	2,467,000	2,467,000	2,467,000	2,467,000	104,874	103,908	103,908

Standard errors in parentheses

Note: Standard errors clustered at the search impression level. Regressions in columns (3), (5), and (8) include position FE for positions 1 through 37, and coefficients should be interpreted with respect to the omitted category: the last position. For visibility, I only display the coefficients for positions 1-10. I restrict the sample to search impressions with opaque offers, which means no hotel was displayed in positions 5, 11, 17, and 23 in the ranking observed by consumers. Query characteristics included in all regressions: trip length, booking window, number of adults and children traveling, number of rooms, and an indicator for a Saturday-night stay.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Appendix C: Further Evidence for Section 4

Opaque Offer Illustration

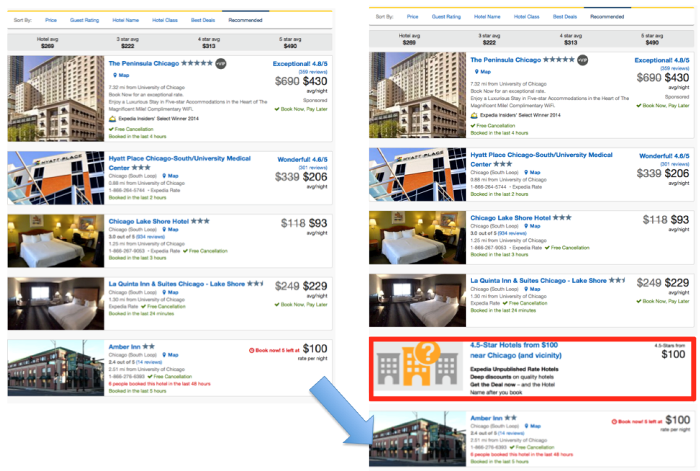

No opaque offer   With opaque offer

Exogeneity of Opaque Offers

The test of the effect of position on search cost is valid if opaque offers are exogenous search cost shifters, that is if consumers who see opaque offers are comparable to those who did not see such offers. Three criteria affect whether an opaque offer is displayed: (i) Expedia's revenue-management algorithm that might target certain consumers, (ii) location, and (iii) availability of rooms sold as opaque offers.[6] In what follows, I will provide evidence in support of availability as being the most likely criterion affecting the display of an opaque offer in the data. In this case, the display of an opaque offer is unrelated to consumer characteristics (exogenous), allowing me to use such offers to test for the effect of position on search costs.

First, and most importantly, Expedia's revenue-management algorithm might show opaque offers to those consumers who are more likely to respond positively to them. This would be a concern for exogeneity, because consumers who see opaque offers and those who do not see them would not be comparable. For example, given that opaque offers typically sell cheaper hotels, the primary customers of such offerings would likely be more price-conscious consumers that are less concerned about their travel dates. To show this possibility is not a concern in the data, I provide two pieces of evidence. First, in the data, 98.8% of consumers see opaque offers, making it unlikely that the algorithm used any particular feature to discriminate which consumers saw offers and which did not. Second, in Table 7, I show results from a t-test comparing observables of consumers who saw the Random ranking to check whether these influence who sees an opaque offer (consumers reveal these characteristics before they see the list of hotels satisfying their query and thus before seeing opaque offers). The table has three columns, considering first all search impressions, and

[6] According to `https://www.tnooz.com/article/expedia-integrates-hotwire-distressed-inventory-\in-hotel-booking-path/`.

then those ending and those not ending in transaction. A negative difference implies a higher value for the measured quantity in search impressions with opaque offers. I find that consumers who saw an opaque offer purchased more expensive hotels in the past (paid on average $74 more per hotel) than those who did not see an opaque offer, making them less price conscious. In addition, those consumers who saw an opaque offer searched for longer trips further in advance, making them more likely to be concerned about travel plans. Thus, the results show the opposite of what would be optimal if Expedia's revenue-management algorithm influenced the decision of which consumers saw an opaque offer.

Second, certain locations on Expedia may not display any opaque offers, because a time lag occurs in instituting the policy of displaying these offers. However, all destinations in my data contain opaque offers, making the second criterion not a discriminatory one here.

As a result, these data patterns provide evidence for the third criterion, availability, as the most likely criterion affecting whether a consumer in my data saw an opaque offer. Because availability is unrelated to consumer characteristics, this evidence allows me to use opaque offers to test whether position affects search costs.

Table 7: T-test: Consumer observables (search impressions without-with opaque offers): Random ranking

	(1) All	(2) Tran	(3) No Tran
Trip Length (days)	-0.3371^{***}	-0.2183	-0.3284^{**}
	(-3.3007)	(-0.9752)	(-3.0115)
Booking Window (days)	-16.3315^{***}	-1.3956	-17.3690^{***}
	(-6.1967)	(-0.2359)	(-6.1758)
Saturday Night	0.0481^{*}	-0.0793	0.0604^{**}
	(2.2835)	(-1.1584)	(2.7278)
Adults	0.1448^{***}	0.0808	0.1535^{***}
	(3.6665)	(0.6637)	(3.6843)
Children	0.0751^{*}	0.0546	0.0781^{*}
	(2.1636)	(0.5151)	(2.1314)
Rooms	0.0281	0.4500^{***}	-0.0155
	(1.4373)	(7.8867)	(-0.7472)
Observations	51,510	3,930	47,580
Consumer Hist. Stars	-0.5241^{**}	-0.5936^{*}	-0.2271
	(-2.6246)	(-2.5395)	(-0.6527)
Consumer Hist. Price	-73.9426^{*}	-88.4345^{*}	-57.6602
	(-2.5506)	(-2.2709)	(-1.1681)
Observations	1,296	195	1,101

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Appendix D: Further Evidence for Section 5

Price Endogeneity Concerns

Table 8 shows that observable hotel and query characteristics explain most of the variation in prices. More precisely, in the first column, I regress price on hotel and trip-date fixed effects in the largest destination in the data and obtain an adjusted R^2 of 0.725. Therefore, specific dates command different prices, but all consumers searching for the same trip-date see the same price for a particular hotel. In the next column, I add additional query characteristics, and include information about the average prices of similar hotels for the same trip and obtain a larger adjusted R^2 of 0.784. In the last three columns of Table 8, I show that a similar pattern holds across different destinations (the four largest destinations in the data).

Table 8: Estimates of price (OLS): Random Ranking

Destination	1	1	2	3	4
Hotel and trip date FE	Yes	Yes	Yes	Yes	Yes
Query characteristics					
Trip Length (days)		1.5323***	1.3647***	0.1758	2.3280***
		(0.1039)	(0.1188)	(0.1298)	(0.1427)
Booking Window (days)		−0.0748***	0.0305***	−0.0239***	−0.0521***
		(0.0038)	(0.0064)	(0.0072)	(0.0061)
Saturday Night		17.0277***	−0.5411	1.2326	−2.3260*
		(0.8444)	(0.9616)	(1.0052)	(1.1620)
Adults		1.2877***	9.0981***	4.6265***	2.7982***
		(0.1808)	(0.2883)	(0.2410)	(0.2951)
Children		1.4265***	14.2707***	3.5658***	5.5825***
		(0.2328)	(0.3711)	(0.2253)	(0.3107)
Rooms		−1.5471***	−5.8421***	−3.8202***	−3.3859***
		(0.3802)	(0.6420)	(0.6599)	(0.6430)
9am-6pm		0.0807	−0.7358	−2.1268***	−2.6452***
		(0.4201)	(0.5968)	(0.6197)	(0.6186)
6pm-midnight		−0.9365*	−2.4495***	−2.9490***	−2.2819***
		(0.4415)	(0.6570)	(0.7053)	(0.6537)
Weekend		0.9205**	−0.3109	−0.0833	−1.1583*
		(0.3309)	(0.5049)	(0.5542)	(0.4904)
Competition					
Avg. prices of similar hotels		−5.1619***	−1.8534***	−1.7380***	−2.0706***
		(0.0338)	(0.0169)	(0.0264)	(0.0207)
Promotion		−23.1960***	−23.6435***	−18.5009***	−27.5182***
		(0.3765)	(0.5669)	(0.7598)	(0.6052)
Adjusted R^2	0.725	0.784	0.815	0.836	0.803
Observations	111,685	111,247	67,425	34,878	61,022

Standard errors in parentheses
Note: Time of day of the search is with respect to the left-out variable: searches performed between midnight and 9am (local time). The average price of similar hotels is computed as the average price of hotels of the same type (chain vs. independent), with the same number of stars and reviews as the focal hotels for the same trip date (excluding the focal hotel). The variable "Weekend" identifies a query that happened on a weekend, whereas the variable "Saturday Night" identifies a trip that includes a Saturday-night stay. I restrict attention to hotels that are displayed at least 15 times, in order to include hotel fixed effects in all specifications above.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Evidence on Click Order

In this section, I show that the position of a hotel in the ranking is a good predictor of the order in which consumers click. To do so, I use the companion data set from WCAI, which contains information on consumers' click order in the form of time stamps associated with each click. I then ask what fraction of search impressions with at least two clicks had a click order that matched the position of the hotels clicked. I also compare this fraction with that ordered by price. Table 9 shows my results. I find that in 35% of all search impressions with at least two clicks and in the majority of search impressions (65%) with exactly two clicks, the position of the hotel exactly matches the click order of the consumer. By contrast, the price of the hotels clicked matches only the order of 20% of the clicks. This finding allows me to approximate consumers' click order with data on the position in which their clicks occurred.

Table 9: Evidence on consumer click order from the WCAI data set

	Percentage		
Search impressions	Position	Price	Position or Price
With at least two clicks	35%	20%	40%
With exactly two clicks	65%	49%	77%

Appendix E: Further Evidence for Section 6

Summary Statistics for the Four Largest Destinations in the Data

Table 10: Hotel and search impression level summary statistics: Four largest destinations (Mean)

Destination	1	2	3	4
Hotel level				
Price	146.30	254.51	133.62	197.56
Stars	4.00	3.65	3.30	3.43
Review Score	4.06	3.98	3.94	4.04
Chain	0.79	0.60	0.73	0.68
Location Score	4.04	4.66	2.81	4.13
Promotion	0.60	0.37	0.42	0.30
Observations	*111,822*	*85,497*	*51,681*	*68,453*
Search impression level				
Number of Hotels Displayed	28.52	28.98	28.79	30.29
Trip Length (days)	2.75	3.20	3.76	2.60
Booking Window (days)	51.95	48.30	45.52	43.08
Saturday Night (percent)	0.48	0.45	0.38	0.48
Adults	2.18	1.89	2.21	1.96
Children	0.23	0.32	0.93	0.36
Rooms	1.16	1.12	1.15	1.12
Total Clicks	1.14	1.15	1.14	1.12
Two or More Clicks (percent)	0.07	0.07	0.06	0.07
Total Transactions	0.54	0.55	0.48	0.65
Random ranking (percent)	0.44	0.37	0.46	0.30
Observations	*3,921*	*2,950*	*1,795*	*2,260*

Table 11: Estimation Results: Search impressions with at least two clicks

Destination	1	2	3	4
Panel A: Coefficients				
Utility (u)				
Price ($100)	-0.1982^{**}	-0.1823	-0.2866^{*}	-0.4359^{***}
	(0.0858)	(0.1466)	(0.1571)	(0.1561)
Stars	0.2201^{***}	-0.0338	-0.0122	0.0992
	(0.0257)	(0.0994)	(0.0981)	(0.0712)
Review Score	-0.1871^{*}	0.0021	-0.0550	-0.0341
	(0.1008)	(0.0659)	(0.0717)	(0.0734)
Chain	0.0003	-0.1178	0.0150	-0.0474
	(0.1051)	(0.1564)	(0.1672)	(0.1906)
Location Score	-0.2235^{**}	0.0765	-0.0247	0.0791
	(0.1040)	(0.0537)	(0.0850)	(0.0575)
Promotion	0.0950	-0.0194	0.2310	0.0401
	(0.0987)	(0.2355)	(0.1520)	(0.1887)
Outside option	0.1492	0.8611^{***}	0.3561	0.4883
	(0.1562)	(0.3084)	(0.3076)	(0.2807)
Search Cost (c)				
Position	0.0027	0.0073	-0.0110	-0.0022
	(0.0073)	(0.0127)	(0.0130)	(0.0109)
Constant	-2.4322^{***}	-2.6018^{***}	-2.0512^{***}	-1.8664^{***}
	(0.0024)	(0.0059)	(0.0069)	(0.0049)
Log-likelihood	-362	-137	-77	-104
Observations	3,127	1,158	666	912
Panel B: % Change cf. Full Sample ($)				
Baseline Search Cost	-64.40%	-60.16%	-82.97%	-81.09%

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Magnitude of Search Costs

In this section, I show evidence that the fact that search impressions contain mostly one click, together with the fact that the data set does not provide enough information to link search impressions made by the same consumer, leads to large search cost estimates. To this end, I re-estimate the model on two data sets. First, in Table 11, I estimate the model on the subset of search impressions that contain at least two clicks. Compared to the baseline search cost coefficient of -1.0305 in Table 8 (in the paper) for the Random ranking in the first destination, the coefficient is twice as large and equals -2.4322. Equivalently, in dollar terms, the baseline level of search costs decreases by as much as 83%, making these estimates comparable to those in the literature.

Second, being able to link different searches made by the same consumer considerably increases the number of clicks made by the unit of observation, and thus decreases the search cost estimate. The Expedia data set is provided at the level of a search impression thus, multiple searches performed by the same consumer cannot be linked. In the WCAI companion data set for Manhattan, I observe that most search impressions (77%) contain no clicks and that the average (median) number of clicks in search impressions with at least one click is 1.80 (1), comparable to the average (median) of 1.12 (1) in the Expedia data set. However, this number is increased if information at the consumer level is available. More precisely, consumers who make at least one click, make on average 3.28 clicks (median is two clicks). Thus, estimating search costs at the level of a search or at the consumer level may affect the results. Indeed, estimating the model on the WCAI companion data set for Manhattan, I find relatively large search costs of $71.37 when estimating at the search impression level, whereas at the consumer level, estimated search costs are 63% lower, equaling $30.20.[7] I conclude that the magnitude of search costs estimated in Table 8 (in the paper) is inherited from the limitations of the data set.

References

[1] Yoganarasimhan, H. 2016. Search Personalization using Machine Learning. Working paper.

[7] Analysis available upon request.